

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

SPINEX-TimeSeries: Similarity-based Predictions with Explainable Neighbors Exploration for Time Series and Forecasting Problems

Ahmed Z. Naser¹, M.Z. Naser^{2,3}

¹Department of Mechanical Engineering, University of Manitoba, Canada, E-mail: a.naser@umanitoba.ca

²School of Civil & Environmental Engineering and Earth Sciences (SCEEES), Clemson University, USA

³Artificial Intelligence Research Institute for Science and Engineering (AIRISE), Clemson University, USA
E-mail: mznaser@clemson.edu, Website: www.mznaser.com

Abstract

This paper introduces a new addition to the SPINEX (Similarity-based Predictions with Explainable Neighbors Exploration) family, tailored specifically for time series and forecasting analysis. This new algorithm leverages the concept of similarity and higher-order temporal interactions across multiple time scales to enhance predictive accuracy and interpretability in forecasting. To evaluate the effectiveness of SPINEX, we present comprehensive benchmarking experiments comparing it against 18 algorithms and across 49 synthetic and real datasets characterized by varying trends, seasonality, and noise levels. Our performance assessment focused on forecasting accuracy and computational efficiency. Our findings reveal that SPINEX consistently ranks among the top 5 performers in forecasting precision and has a superior ability to handle complex temporal dynamics compared to commonly adopted algorithms. Moreover, the algorithm's explainability features, Pareto efficiency, and medium complexity (on the order of $O(\log n)$) are demonstrated through detailed visualizations to enhance the prediction and decision-making process. We note that integrating similarity-based concepts opens new avenues for research in predictive analytics, promising more accurate and transparent decision making.

Keywords: Algorithm; Machine learning; Benchmarking; Time series; Forecasting.

1.0 Introduction

Time series analysis involves the study of data collected or recorded sequentially over time to extract meaningful patterns, trends, and insights [1]. Such a temporal ordering of data distinguishes time series from cross-sectional data and necessitates specialized techniques to understand the underlying mechanisms that generate the observed data and to forecast future values based on historical patterns. In other words, time series analysis inherently focuses on temporal data, with applications aiming to understand past behaviors, predict future trends, and identify cyclical patterns as well as anomalies. Such applications can span diverse domains, from financial forecasting to environmental modeling, etc. [2].

Given that a core component of time series analysis builds on forecasting future predictions from past trends/responses, the concept of similarity in time series then arises [3]. Similarity in this context refers to the degree of resemblance or correspondence between different segments of a single time series or between multiple time series [4]. The quantification of similarity enables researchers and practitioners to identify recurring patterns, classify time series into groups with similar characteristics, and detect deviations from expected behavior. In a way, similarity may

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

39 enable the establishment of a notion of "normal" behavior, which can then help identify instances
40 that exhibit low similarity to these reference patterns (i.e., anomalies). Thus, the notion of
41 similarity can be particularly significant in pattern recognition, anomaly detection, and clustering
42 [5].

43 In parallel, this concept is not only about finding sequences that look alike but also aims to
44 understand temporal alignment and variation between data points. For example, two economic
45 time series might exhibit similar patterns of growth and recession cycles but shift in time or at
46 different scales. Accurately measuring such similarities involves techniques that consider
47 alignment and scaling of time data instead of standard distance metrics [6]. This can be apparent
48 in the case of traditional metrics like Euclidean and Manhattan distances, as these quantify
49 similarity by measuring the distances between points in a time series. However, these measures
50 often fall short of capturing the dynamic characteristics of time series data, such as trends and
51 seasonality [7]. The same measures also tend to be sensitive to small fluctuations, which can be
52 misleading in a temporal context where trends and cycles play a significant role [8].

53 As one can see, unlike static data, where similarity can often be measured using straightforward
54 distance metrics, time series data presents additional complexities stemming from the observations'
55 temporal nature. Consequently, specialized similarity measures have been developed to address
56 these challenges, each with its own strengths and limitations. For instance, Dynamic Time
57 Warping (DTW) allows elastic transformations of the time [9]. DTW also allows for non-linear
58 alignment of time series and hence can accommodate differences in speed and duration of patterns.
59 Thus, DTW can align two sequences in a way that minimizes their overall distance. This method
60 is particularly effective in dealing with time series that are similar in shape but vary in speed or
61 timing of events where temporal alignment is fundamental, such as speech recognition and
62 bioinformatics [10,11].

63 Another similarity-based concept is the Longest Common Subsequence (LCSS), which measures
64 the similarity between two sequences by identifying the longest subsequence present in both
65 sequences without altering the order of elements [12]. LCSS can be robust to noise and occlusions
66 and is particularly useful in real-time series applications where missing values may occur. Another
67 approach to quantifying similarity in time series is through the use of feature-based methods and
68 correlation measures. The former methods involve extracting relevant features or summary
69 statistics from the time series and comparing these derived representations rather than the raw data
70 points, including statistical moments, frequency domain characteristics, and model parameters.
71 The latter measures, such as Pearson's or Spearman's correlation coefficients etc., capture the
72 degree of relationship/association between time series. Both techniques have proven useful yet
73 may fail to capture non-linear relationships and are sensitive to outliers and phase differences
74 [13,14].

75 More recently, machine learning (ML) advancements have introduced new models for assessing
76 similarity in time series data. Techniques such as Siamese and triplet networks learn similarity

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

77 metrics directly from data, potentially capturing complex, non-linear relationships that traditional
78 approaches might miss [15]. Furthermore, methods like K-Nearest Neighbors (KNN) for time
79 series rely on identifying similar historical patterns to make predictions about future values. More
80 sophisticated approaches, such as Long Short-Term Memory (LSTM) networks, can now
81 implicitly learn to recognize and utilize similar patterns in their internal representations [16]. A
82 key component that remains opaque in ML-based methods is their blackbox nature and limited
83 interpretability [17].

84 It can then be inferred that the concept of similarity, as well as explainability, can be thought of as
85 elemental to forecasting tasks. Thus, this paper presents the development of the time series variant
86 to the SPINEX (Similarity-based Predictions with Explainable Neighbors Exploration) family. This
87 variant builds upon the concept of similarity and explainability between similarly-identify
88 neighbors and segments. As such, SPINEX hopes to bridge some of the existing challenges. In this
89 study, we examine the algorithm's ability to perform on 49 diverse datasets compared to 18
90 commonly used algorithms.

91 **2.0 Description of the SPINEX for time series and forecasting**

92 *2.1 General description*

93 SPINEX represents a unique approach to time series analysis and prediction. This algorithm
94 integrates multiple techniques to deliver robust, adaptive, and interpretable time series forecasting.
95 For example, at its core, SPINEX employs a multi-method similarity analysis, utilizing various
96 measures such as cosine similarity, Euclidean distance, DTW, Pearson, and Spearman correlation.
97 This ensemble approach enables a comprehensive assessment of segment similarities, capturing
98 diverse aspects of time series behavior. A key feature of SPINEX is its adaptive window sizing
99 mechanism, which adjusts based on data length, variability, and potential seasonality. This
100 adaptability allows the algorithm to handle time series of varying lengths and characteristics
101 effectively. Additionally, SPINEX implements time series cross-validation to provide robust
102 performance estimates and assess model stability across different time periods.

103 SPINEX's multi-level analysis capability provides a hierarchical view of time series patterns to
104 enable robustness to different scales of temporal dependencies. Furthermore, SPINEX incorporates
105 a dynamic thresholding technique for anomaly detection and forecasting validation. This method
106 adjusts the similarity threshold based on the recent performance of the predictions and the
107 distribution of similarity scores, which further enhances the algorithm's flexibility and
108 responsiveness to changing patterns in the data. The same can be crucial for understanding outliers
109 and potential regime changes in the data. Thus, SPINEX effectively identifies the most relevant
110 segments for making predictions to improve the reliability of the forecast. In cases where
111 accessible similarity-based prediction is not feasible, the algorithm switches to a fallback
112 prediction method, which includes trend extraction, multiple seasonality detection, non-linear
113 trend modeling, and anomaly-aware residual prediction with confidence intervals.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

114 As mentioned above, SPINEX focuses on explainability by offering detailed results on the most
115 similar historical segments, their contributions to the prediction, and visualizations of the nearest
116 neighboring similar segments. This enhances understanding of the prediction process and the
117 underlying patterns in the data. Computational efficiency within SPINEX is achieved through the
118 use of techniques like numba for just-in-time compilation and caching mechanisms. For
119 completion, a representative pseudo-code is provided below.

```
120 Input:  
121   data: Time series data  
122   window_size (optional): Length of each segment  
123   forecast_horizon: Number of future steps to predict  
124   similarity_methods: List of similarity methods (e.g., 'cosine', 'euclidean', 'dtw', etc.)  
125   dynamic_window: Flag to enable adaptive window sizing  
126   multi_level: Flag to enable multi-level similarity analysis  
127   dynamic_threshold: Flag to enable dynamic threshold adjustments  
128 Output:  
129   Predicted future values of the time series  
130   Identified anomalies (if applicable)  
131   Explainability insights for predictions  
132 Procedure:  
133 1. Initialization:  
134   - Convert `data` to a numpy array.  
135   - Set `window_size` to default or provided value.  
136   - Set default similarity methods and initialize caches.  
137 2. Dynamic Parameter Adjustment (if `dynamic_window` or `dynamic_threshold` is enabled):  
138   - Calculate volatility or variability in recent data.  
139   - Adjust `window_size` based on variability and predefined bounds.  
140   - Calculate dynamic threshold for similarity scores based on recent errors and scores.  
141 3. Segment Extraction:  
142   - Slide a window of size `window_size` across the data to extract overlapping segments.  
143   - Normalize each segment (mean = 0, std = 1).  
144 4. Similarity Matrix Calculation:  
145   - For each specified similarity method:  
146     - Compute pairwise similarity scores between segments using:  
147       - Cosine similarity  
148       - Euclidean similarity  
149       - DTW (Dynamic Time Warping)  
150       - Other specified methods  
151     - Cache results for reuse.  
152 5. Find Similar Segments:  
153   - Evaluate similarity scores for segments using all methods.  
154   - Combine results across methods to compute an overall similarity score.  
155 6. Prediction:  
156   - Identify top similar segments based on overall similarity score.  
157   - Use weights derived from similarity scores to combine predictions.  
158   - If no valid predictions are possible, use a fallback method (e.g., seasonal decomposition or trend modeling).  
159 7. Anomaly Detection (Optional):
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

160 - Define a threshold for similarity scores (dynamic or static).
161 - Identify segments with scores below the threshold as anomalies.
162 **8. Explainability Analysis (Optional):**
163 - Analyze contributions of features to similarity scores.
164 - Identify top-contributing features for each similar segment.
165 - Compute weighted contributions of segments to predictions.
166 **9. Evaluation:**
167 - Evaluate prediction accuracy using metrics such as MSE, RMSE, MAE, R², etc.
168 - Validate predictions across multiple train-test splits if required.
169 **10. Visualization (Optional):**
170 - Plot predictions alongside actual time series data.
171 - Highlight anomalies or visualize top similar segments and their contributions.
172 **End Algorithm**

173 *2.2 Detailed description*

174 A more detailed description of SPINEX's methods and functions is provided herein. It is worth
175 noting that the presented default settings were arrived at from an empirical analysis of the 49
176 datasets examined in this paper (obtained from synthetic and real datasets).

177 Initialization (`__init__`)

178 The `__init__` method initializes and sets up the operational parameters of SPINEX. The method
179 signature is as follows:

```
180 def __init__(self, data, window_size=None, forecast_horizon=1, similarity_methods=None,  
181             dynamic_window=True, multi_level=True, dynamic_threshold=True):
```

182 More specifically:

- 183 • **data:** The input time series data, converted into a NumPy array for efficient numerical operations.
- 184 • **window_size:** Determines the length of the segments to be compared. If not specified, it is set to the greater
185 of 10 or one-tenth of the data length. This parameter can be dynamically adjusted based on the data's
186 volatility.
- 187 • **forecast_horizon:** Specifies how far into the future predictions are made. By default, it is the smaller of the
188 provided value and one-tenth of the data length.
- 189 • **similarity_methods:** A list of methods used to compute similarity between segments. Defaults to ['cosine',
190 'euclidean', 'dtw'] if not specified.
- 191 • **dynamic_window:** Enables or disables dynamic adjustment of the window size based on data characteristics.
- 192 • **multi_level:** Allows the use of multiple window sizes in the analysis to capture different scales of patterns.
- 193 • **dynamic_threshold:** Enables adaptive thresholding in the similarity calculations to improve forecast
194 reliability.
- 195 • Additionally, the class uses caching mechanisms (`similarity_cache` and `segments_cache`) to store computed
196 results for re-use to optimize performance for large datasets.

197 Method: Similarity Measures

198 This method offers several methods to compute the similarity between time series segments:

- 199 • **Cosine Similarity:**

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 200 ○ Computes the cosine of the angle between two vectors and is defined as the dot product of the
201 vectors divided by the product of their norms. This measure is effective in identifying the
202 similarity in direction regardless of magnitude.
- 203 ○ Equation: $similarity = \frac{X \cdot X^T}{\|X\| \|X^T\|}$
- 204 • **Correlation Similarity:**
- 205 ○ Calculates the Pearson correlation coefficient matrix of the rows of X, providing a measure of
206 linear relationships between segments.
- 207 ○ Equation: $similarity = \text{corrcoef}(X)$
- 208 • **Euclidean Similarity:**
- 209 ○ Uses the Euclidean distance to compute similarity by applying a transformation that inversely
210 relates distance to similarity.
- 211 ○ Equation: $similarity = \frac{1}{1 + \sqrt{cdist(X, X, \text{Euclidean})}^2}$
- 212 • **Spearman Similarity:**
- 213 ○ Calculates the Spearman rank correlation between the columns of X, useful for capturing
214 monotonic relationships between segments that may not necessarily be linear.
- 215 ○ Equation: $similarity = \text{spearmanr}(X^T)[0]$
- 216 • **Dynamic Time Warping (DTW) Similarity:**
- 217 ○ Measures similarity based on the minimal distance that aligns two time series, accounting for
218 shifts and distortions in time. In essence, DTW measures the similarity between two temporal
219 sequences, which may not be of the same length, by aligning their points to minimize the overall
220 distance between them.
- 221 ○ Equation: $similarity = \frac{1}{1 + DTW(X[i], X[j])}$
- 222 ▪ $DTW(x, y) = \min(\text{cost} + \min(\text{dtw}_{matrix}[i-1, j], \text{dtw}_{matrix}[i, j-1], \text{dtw}_{matrix}[i-1, j-1]))$
- 223 ▪
- 224 • **Direction Similarity:**
- 225 ○ Calculates the direction similarity via the direction method to be discussed later on.

226 Method: `adjust_dynamic_parameters`

227 This method adjusts the window size and similarity threshold based on the recent behavior of the
228 time series and the algorithm's performance.

- 229 • **Volatility-based Window Size Adjustment:**
- 230 ○ **Volatility Calculation:** First, this method calculates the volatility of the most recent portion of
231 the data, defined as the standard deviation over the last data segments. The size of this
232 segment is a maximum of 10 or one-tenth of the data length but not exceeding half the length
233 of the data.
- 234 ○ **Window Size Recalculation:** The window size is inversely adjusted based on the calculated
235 volatility to respond to the data's fluctuating nature. If the volatility is low, a larger window size
236 is used to smooth out noise and capture more extended patterns. If the volatility is high, the
237 window size decreases, making the model more responsive to recent changes. A scaling factor
238 controls this adjustment, clipped between 0.1 and 1.0 to prevent extreme values.
- 239 ○ Equation: $window_size =$
- 240 $max(MIN_WINDOW_SIZE, \min(\frac{MAX_WINDOW_SIZE}{scale_factor}, MAX_WINDOW_SIZE))$
- 241 ▪ where $scale_factor = \text{clip}(volatility, 0.1, 1.0)$.
- 242 • **Threshold Adjustment:**

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 243 ○ **Error-based Adjustment:** If recent prediction errors are available, the threshold is adjusted
- 244 based on these errors' mean and standard deviation to accommodate the model's predictive
- 245 accuracy.
- 246 ○ **Similarity Score-based Adjustment:** If recent similarity scores are tracked, the threshold is
- 247 further adjusted to reflect the mean and variability in these scores. This dynamic threshold
- 248 helps maintain the similarity measure's relevance under varying data conditions.
- 249 ○ Equation: $threshold = mean_sim + std_sim + threshold_adjustment$

250 Method: `get_dynamic_threshold`

251 This method computes an adaptive threshold for the similarity scores to decide which time series
252 segments are considered similar enough to be relevant for predictions.

- 253 • **Basic Threshold Calculation:** Calculates a baseline threshold as the sum of the mean and standard deviation
- 254 of the similarity scores. This method aims to keep only the most similar segments, thus ensuring that the
- 255 predictions are based on the most relevant and recent data patterns.
- 256 • **Threshold Adjustment:** If fewer than five segments exceed this baseline threshold, indicating a potential
- 257 over-tightening, the threshold is reduced to the 90th percentile of the scores to include more segments.
- 258 ○ Equation: $adjusted_threshold =$
- 259
$$\begin{cases} percentile(similarities, 90) & \text{if } abs\{s > base_threshold\} < 5 \\ base_threshold, & \text{otherwise} \end{cases}$$
- 260

261 Method: `adjusted_dtw_similarity`

262 This method modifies the DTW similarity measure to be more forgiving by squaring the DTW
263 distance before inversely transforming it into a similarity score. This adjustment makes the
264 similarity measure less sensitive to small variations, emphasizing more significant patterns in the
265 similarity assessment.

266 ○ Equation: $adjusted_scores = \frac{1}{1 + \sqrt{dtw_scores}}$

267 Method: `plot_prediction`

268 This method is designed to visualize the forecasting performance of the SPINEX model by plotting
269 actual data alongside predicted values. This method serves as a tool for assessing the accuracy and
270 relevance of the model's predictions.

271 Method: `extract_segments`

272 This method prepares segments of the time series data for further analysis, such as computing
273 similarities or making predictions such that:

- 274 • **Dynamic Window Size:** If no specific window size is provided, the method calculates an adaptive window
- 275 size using the `adaptive_window_size()` method.
- 276 • **Adjustment for Small Data:** If the total data length is less than the determined window size, the window size
- 277 is adjusted to half the data length to ensure at least some segmentation can be performed.
- 278 • **Segmentation:** Using `np.lib.stride_tricks.sliding_window_view`, the method creates overlapping segments
- 279 of the specified window size from the time series data. This function efficiently generates a new view into
- 280 the data array without copying the data.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 281
- **Normalization:** Each segment is normalized by subtracting its mean and dividing by its standard deviation. This step standardizes the segments, mitigating the effect of different scales or baselines in the data and improving the comparability between segments.
- 282
- 283
- 284
- Equation: $normalized_segments = \frac{segments - segment_means}{segment_stds + 1e-8}$
 - Here, $1e-8$ is added to the standard deviations to prevent division by zero in the case of very uniform segments.
- 285
- 286

287 Method: find_similar_segments

288 This method facilitates the identification of similar segments within the time series data, which is
289 crucial for making accurate predictions.

- **Multi-Level Analysis:** Depending on the multi_level attribute, the method considers multiple window sizes for segmentation. These sizes include a smaller window (half the primary size), the primary window size itself, and a larger window (double the primary size or one-fourth the length of the data, whichever is smaller). This multi-scale approach allows the model to capture similarities at different granularities.
 - **Segment Extraction and Hashing:** For each window size, the method extracts segments and computes a hash to uniquely identify them. This hash is used to cache the segments and avoid redundant calculations.
 - **Similarity Calculation:** For each window size, the method computes similarity matrices using the specified methods (cosine, euclidean, dtw, etc.). If a large number of segments are detected (more than 500), DTW is skipped to avoid performance bottlenecks.
 - **Aggregation of Similarities:** The method averages the similarities across different methods to get a composite similarity measure for each window size. These are then averaged across all window sizes to get the final measure of similarity between segments.
 - **Fallback Method:** If no valid similarities are found (e.g., due to insufficient segments or errors in calculation), a fallback method based on autocorrelation is used.
- 290
- 291
- 292
- 293
- 294
- 295
- 296
- 297
- 298
- 299
- 300
- 301
- 302
- 303

304 Method: fallback_similarity_method

305 This method provides a basic mechanism to calculate similarity based on autocorrelation when
306 other methods fail or are not applicable due to data constraints.

307 Method: analyze_segment_similarity

308 This method quantitatively assesses how similar a particular segment (indexed) is to the most
309 recent segment in the time series.

- **Segment Extraction:** Both the target segment and a reference segment (usually the most recent one) are extracted.
 - **Similarity Calculation:** The method calculates similarity scores using all available similarity methods, providing a detailed breakdown of how each method perceives the similarity.
 - **Feature Contributions:** It calculates the absolute differences between the corresponding features of the two segments to determine which features contribute most to any dissimilarity.
- 310
- 311
- 312
- 313
- 314
- 315

316 Method: get_nearest_neighbors

317 This method identifies the nearest neighbors of the most recent segment based on the computed
318 similarities and can identify tasks for anomaly detection. After calculating similarities for all
319 segments, it sorts these and picks the top k segments most similar to the latest segment, providing
320 their indices and similarity scores.

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

321 Method: detect_seasonality

322 This method is designed to identify seasonality within the time series data, which is crucial for
323 understanding periodic patterns that could influence forecasting and other analytical tasks.

- 324 • **Autocorrelation Calculation:** The method calculates the autocorrelation (ACF) for the data up to a specified
325 lag (*max_lag*). If *max_lag* is not specified, it defaults to half the length of the data.
- 326 • **Peak Detection:** The method then identifies ACF peaks, representing potential seasonal periods. Peaks are
327 detected where the autocorrelation at a given lag is greater than its neighbors, indicating a repeating pattern.
- 328 • **Seasonality Inference:** If any peaks are detected, the first peak is assumed to represent the primary seasonal
329 period, and its lag is returned. An empty list is returned if no peaks are detected, indicating no detectable
330 seasonality.

331 Method: detect_anomalies

332 This method identifies anomalies in the time series data by comparing the similarity of data
333 segments to a dynamically determined threshold.

- 334 • **Segment Extraction and Similarity Calculation:** Segments of the data are extracted, and their similarities are
335 computed.
- 336 • **Threshold Determination:** A threshold is set at a specified percentile (default is the 2nd percentile) of the
337 similarity scores, identifying the least similar segments as potential anomalies.
- 338 • **Anomaly Identification:** Segments whose similarity scores fall below the threshold are marked as anomalies.
- 339 • Equation: $Threshold = percentile(similarities, threshold_percentile)$

340 Method: fallback_prediction

341 This method provides a comprehensive mechanism for generating predictions when standard
342 approaches are not feasible, utilizing multiple time series decomposition and modeling techniques.

- 343 • **Pre-checks:** It first ensures that there is sufficient data for prediction based on the specified number of points
344 required.
- 345 • **Adaptive Window Sizing:** This step dynamically adjusts the window size for trend extraction based on
346 minimizing the mean squared error (MSE) of the trend-subtracted data.
- 347 • **Trend Extraction:** Utilizes a moving average to smooth the data and extract the underlying trend.
- 348 • **Seasonality Detection:** Employs autocorrelation to identify potential seasonality periods and extract these
349 seasonal components.
- 350 • **Residual Calculation:** The residuals (or unexplained components) are analyzed after removing the trend and
351 seasonal components.
- 352 • **Anomaly Detection and Handling:** Anomalies in the residuals are identified and replaced with median values
353 to stabilize the model.
- 354 • **Non-linear Trend Modeling:** Fits a polynomial model to predict the future trend based on past data.
- 355 • **Seasonal Component Prediction:** Projects the identified seasonal patterns into the future.
- 356 • **Residual Prediction:** Uses a weighted average approach to predict future residuals, incorporating confidence
357 intervals to account for uncertainty.
- 358 • **Combination of Components:** The final prediction combines the trend, seasonal, and residual predictions to
359 form a complete forecast.
- 360 • **Trend:** Extracted using a moving average filtered by convolution:
 - 361 ○ Equation: $trend = convolve(data, window)/window_size$
- 362 • **Seasonality:** Identified through peak detection in the autocorrelation function.
- 363 • **Residuals:** Calculated as data – trend

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 364
- **Confidence Intervals for Residuals:** Generated by assuming the residuals follow a normal distribution modulated by an exponential decay in influence.
- 365

366 Method: tune_hyperparameters

367 This method optimizes the hyperparameters of the model, specifically focusing on the detection of
368 seasonalities.

- 369
- **Iterative Testing:** The method iterates over a range of possible numbers of seasonalities to detect (from 1 to 4).
 - **Prediction Generation:** For each candidate setting, it generates predictions using the fallback_prediction method.
 - **Evaluation:** Calculates the MSE for each set of predictions compared to the actual data.
 - **Selection of Optimal Parameter:** Identifies the number of seasonalities that result in the lowest MSE, suggesting the best fit for the data.
- 370
371
372
373
374
375

376 Method: predict

377 This method combines various techniques to generate accurate forecasts based on the similarity of
378 time series segments.

- 379
- **Dynamic Parameter Adjustment:** Initially, dynamic parameters such as window size and thresholds are adjusted based on recent data characteristics.
 - **Similarity Assessment:** It calculates similarities between segments of the time series to identify patterns that can be used for forecasting.
 - **Fallback Prediction:** If no significant similarities are found, it resorts to a fallback prediction method that uses more basic statistical methods.
 - **Threshold Determination:** Determines a dynamic threshold for considering a segment significantly similar to the latest data, adjusting the threshold based on the distribution of similarity scores.
 - **Valid Predictions Identification:** Identifies segments that meet the similarity threshold and ensures that they are within a valid range for making predictions.
 - **Prediction Compilation:** Compiles predictions from multiple segments, weighted by their similarity scores, and adjusts them to align with the most recent actual data point.
 - **Error Handling:** If any step fails, it defaults to the fallback prediction method.
- 380
381
382
383
384
385
386
387
388
389
390
391

392 Method: update_recent_performance

393 This method updates the performance metrics of the model by recording the recent error and
394 similarity scores, which are essential for monitoring and improving the model's accuracy over
395 time.

- 396
- **Dynamic Parameter Adjustment:** Initially, dynamic parameters such as window size and thresholds are adjusted based on recent data characteristics.
 - **Similarity Assessment:** It calculates similarities between segments of the time series to identify patterns that can be used for forecasting.
- 397
398
399

400 Method: validate_prediction

401 This method evaluates the robustness of the model's predictions by using cross-validation,
402 specifically time-series cross-validation, where the order of data points is preserved.

- 403
- **Setup:** Determines the number of splits for cross-validation based on available data, ensuring there are enough points for each training and testing set.
- 404

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

405
406
407
408
409
410
411
412

- **Cross-Validation:**
 - Single Split Handling: If there is insufficient data for multiple splits, perform a single train-test split.
 - Multiple Splits: Uses TimeSeriesSplit from scikit-learn to create training and testing segments. It ensures that predictions are based only on past data, respecting the temporal order.
 - Prediction and Evaluation: For each split, the model predicts future values based on the training set, and the predictions are evaluated using the evaluate_prediction method.
 - Aggregation of Results: The results from each split are aggregated to calculate average performance metrics across all splits.

413 Method: get_explainability_results

414 This method provides insights into why certain predictions were made based on the similarity of
415 time series segments.

416
417
418
419
420
421
422
423

- **Similarity Assessment:** The method first identifies similar segments by calculating and evaluating segment similarities.
- **Threshold Determination:** It dynamically determines a similarity threshold above which segments are considered significantly similar.
- **Top Segments Identification:** Segments surpassing the threshold are marked as key influencers. If no segments exceed the threshold, the top k segments based on similarity scores are selected.
- **Contribution Calculation:** Each top segment calculates how much each segment contributes to the predictions, using weighted averages based on their similarity scores.

424 Method: analyze_and_plot_neighbors

425 This provides a deeper analysis of how and why certain segments are considered similar to the
426 current segment, offering both visual and numerical insights.

427
428
429
430
431
432
433
434
435
436

- **Current and Neighbor Segment Extraction:** Similar to plot_nearest_neighbors, but with added analysis of segment similarities.
- **Similarity Analysis:** For each neighbor, it computes detailed similarity scores using various metrics.
- **Visualization and Reporting:** Each neighbor's segment and its similarity scores are plotted and displayed. This includes a breakdown of the scores for different similarity metrics and the identification of key features contributing to the similarities.
- **Similarity Scores:** Each neighbor's similarity to the current segment is quantified using methods like cosine, euclidean, and DTW similarities.
- **Feature Contributions:** Differences between segments are analyzed to pinpoint which specific elements (data points) contribute most to the observed similarities or discrepancies.

437 Additional functions for optimized clustering:

438 Method: direction_accuracy

439 This method calculates the direction accuracy to compare the directional trends between two time
440 series segments. Given two segments, segment1 and segment2, the following steps compute the
441 direction accuracy:

442
443
444
445

- **Calculate the Differences:** compute the first-order differences of both segments such that:
 - Equation: $Diff\ 1_i = segment\ 1_{i+1} - segment\ 1_i$ (with a similar approach for segment 2)
- **Determine the Direction:** Using the sign function $sign(\cdot)$, the direction of these differences can be calculated.
- **Equation:** $direction\ 1_i = sign(Diff\ 1_i)$ (with a similar approach for segment 2)

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 446 • **Compare Directions:** Compare the directional trends of the two segments by checking if the directions are
447 equal at each time step
- 448 ○ Equation: $match = \begin{cases} 1 & \text{if } direction1_i = direction2_i \\ 0 & \text{otherwise} \end{cases}$

449 **Method: Entropy**

450 The `numba_sample_entropy` method calculates the sample entropy of a sequence x , which is a
451 measure of the complexity or the amount of regularity and unpredictability in time series data.
452 This entropy is useful for determining the complexity of physiological time series signals.

- 453 • **Mathematical Representation:**
- 454 ○ $Sample\ entropy = -\log \frac{A+1e-10}{B+1e-10}$

456 **Method: Hash Array (`hash_array`)**

457 This static method generates a unique hash for a numpy array using MD5 to create keys for caching
458 purposes, allowing efficient retrieval of previously computed results.

459 **Method: `plot_anomalies`**

460 This method visualizes the anomalies detected.

461 **Method: `plot_nearest_neighbors`**

462 This method visualizes the time series segments that are most similar to the most recent segment,
463 facilitating an understanding of the model's decision-making process.

464 **3.0 Description of benchmarking algorithms**

465 We examined SPINEX against 18 commonly used time series forecasting algorithms, namely,
466 ARIMA, SARIMA, ETS, Holt-Winters, Prophet, Theta, Simple Moving Average, VAR, Croston's
467 Method, LSTM, Neural Networks, Gaussian Process Regression, KNN, SVR, Random Forest,
468 XGBoost, Gradient Boosting, CatBoost, and Bagging. As one can see, the first nine algorithms are
469 specifically designed for time series analysis, while the latter group consists of other ML
470 algorithms that can be adapted for time series forecasting with appropriate feature engineering, as
471 seen in [8,18–20]. Each of these algorithms is described in this section, where we showcase a brief
472 historical background and algorithmic logic (with additional details being available in the cited
473 original sources). Table 1 compares these algorithms with respect to their time series forecasting
474 characteristics.

475 *3.1 Algorithms specifically designed for time series analysis*

476 3.1.1 Autoregressive Integrated Moving Average (ARIMA and SARIMA)

477 ARIMA (Autoregressive Integrated Moving Average) and its seasonal variant SARIMA were
478 popularized by Box and Jenkins in the 1970s [21] – however, the concepts of Auto-Regressive and
479 Moving Average models were introduced by Yule in 1926 and by Slutsky in 1937, respectively
480 [22]. The ARIMA algorithm combines these concepts and components with differencing to handle
481 non-stationary data. The ARIMA algorithm is particularly effective for univariate time series

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

482 forecasting, while SARIMA extends this capability to series with seasonal patterns. The models
483 are specified by three main parameters: p (order of the Auto-Regressive term), d (degree of
484 differencing), and q (order of the Moving Average term), and SARIMA adds additional seasonal
485 parameters. These algorithms are widely used due to their flexibility and ability to capture complex
486 temporal dependencies. However, the algorithms assume linear relationships and, hence, may
487 struggle with highly nonlinear patterns. Moreover, the selection of appropriate internal parameters
488 can be challenging and often requires expert/domain knowledge or automated procedures [23].

489 3.1.2 Croston's Method

490 Croston introduced this method in 1972 [24] as a specialized forecasting algorithm designed for
491 intermittent demand patterns. This algorithm separates the time series into two components: the
492 non-zero demand sizes and the intervals between non-zero demands. Each component is then
493 forecasted separately using simple exponential smoothing, and the final forecast is obtained by
494 dividing the demand size forecast by the interval forecast. This method is particularly useful in
495 domains where demand occurs sporadically (such as that commonly seen in inventory
496 management and spare parts forecasting) [25]. Croston's method assumes that the demand sizes
497 and intervals are independent (which often introduces bias as this assumption may not always hold
498 true). Several modifications of Croston's method have been proposed to address this main
499 limitation [26,27].

500 3.1.3 Error, Trend, Seasonality (ETS), and the Holt-Winters Method

501 ETS (Error, Trend, Seasonality) and Holt-Winters methods are exponential smoothing techniques
502 that have evolved since their introduction by Brown and Holt in the 1950s [28,29]. These two
503 methods decompose time series into components (level, trend, and seasonality) and use weighted
504 averages of past observations to forecast future values. ETS provides a framework for selecting
505 the most appropriate model based on the nature of the components (i.e., additive or multiplicative).
506 Holt-Winters [30] is a specific implementation within the ETS family that has been modified to
507 account for time series with both trend and seasonal components. This family of algorithms can
508 be effective in handling a wide range of time series patterns. However, these algorithms may
509 struggle with complex, non-linear relationships and can be sensitive to outliers [31].

510 3.1.4 Long Short-Term Memory (LSTM)

511 The Long Short-Term Memory (LSTM) network was introduced by Hochreiter and Schmidhuber
512 in 1997 [32] as a type of recurrent neural network. This network is designed to capture long-term
513 dependencies in sequential data. LSTMs use a series of gates (input, forget, and output gates) to
514 control the flow of information through the network, allowing them to selectively remember or
515 forget information over long sequences. This architecture makes LSTMs particularly well-suited
516 for time series forecasting, especially when dealing with complex, non-linear patterns and long-
517 term dependencies. LSTMs can handle multivariate time series and can learn from historical data.
518 However, they often require substantial training data to perform well, can be computationally
519 intensive, and may be prone to overfitting if not properly regularized. Moreover, LSTM is a
520 blackbox algorithm and can be challenging to interpret [33].

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

521 3.1.5 Prophet

522 Prophet, developed by Meta (Facebook formally) in 2017 [34]. This algorithm offers a procedure
523 for forecasting time series data based on an additive model that decomposes the time series into
524 trend, seasonality, and holiday components. Prophet is designed to handle daily observations with
525 at least one year of historical data and can accommodate missing values and outliers. The algorithm
526 automatically detects changepoints in the trend and allows for user-specified changepoints. A key
527 advantage of this algorithm is its ability to handle multiple seasonalities and incorporate domain
528 knowledge through easily interpretable parameters. Prophet is particularly effective for forecasting
529 tasks with strong seasonal effects (as well as those with several seasons of historical data). Yet,
530 this algorithm may struggle with short-term forecasts or datasets with limited historical data.
531 Additionally, while it is designed to be robust, it may not always capture complex, non-linear
532 patterns effectively [35].

533 3.1.6 Simple Moving Average (SMA)

534 The Simple Moving Average (SMA) is a basic and widely used time series forecasting method.
535 The origin of SMA can be traced back to the early days of technical and inventory analysis [36].
536 This method calculates the arithmetic mean of a set of values over a specific number of time
537 periods and is often used to smooth out short-term fluctuations and highlight longer-term trends
538 or cycles, and can be effective for short-term forecasting in stable time series with minimal trend
539 or seasonality. However, SMA has several limitations. This method can produce lags behind the
540 most recent data points and may miss sudden changes or turning points. SMA also gives equal
541 weight to all observations within the moving window, which may not be ideal if more recent
542 observations are believed to be more relevant. Despite such limitations, SMA remains a useful tool
543 for forecasting methods [37].

544 3.1.7 Theta Method

545 The Theta algorithm was proposed by Assimakopoulos and Nikolopoulos in 2000 [38]. This
546 algorithm decomposes the time series into two "theta lines." The first line represents the long-term
547 trend, and the other captures short-term behavior. These lines are then extrapolated separately and
548 combined to produce the final forecast. The Theta method is praised for its simplicity and
549 effectiveness, especially for seasonal time series. The Theta algorithm often performs well without
550 requiring extensive parameter tuning, making it accessible for practitioners. However, the method
551 assumes that the time series can be well-represented via decomposition into two lines (which may
552 not always hold true for complex, non-linear time series). Moreover, it may struggle with abrupt
553 changes or structural breaks in the data [39].

554 3.1.8 Vector Autoregression (VAR)

555 Vector Autoregression (VAR), introduced by Sims in 1980 [40], is a multivariate forecasting
556 technique that extends the univariate autoregressive model to capture the linear interdependencies
557 among multiple time series. In VAR, each variable is a linear function of past lags of itself and
558 past lags of the other variables. This methodology makes VAR particularly useful for
559 understanding the relationships between multiple related time series and generating forecasts for

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

560 these interactions. VAR models are widely used in econometrics and financial time series
561 forecasting as they can model feedback effects and provide insights into the dynamics between
562 variables through tools like impulse response functions [41]. However, VAR models assume linear
563 relationships between variables and can become overparameterized when dealing with many
564 variables or long lag structures. This could potentially lead to poor forecasts [42].

565 *3.2 ML algorithms adapted for time series forecasting*

566 3.2.1 Gaussian Process Regression (GPR)

567 Gaussian Process Regression (GPR) is a non-parametric probabilistic approach to regression and
568 time series forecasting that is rooted in Bayesian statistics. This algorithm was formalized for ML
569 applications by Rasmussen and Williams [43]. The method models the target variable as a
570 Gaussian process, assuming that any finite collection of data points has a multivariate Gaussian
571 distribution. GPR is particularly valuable in time series forecasting for its ability to provide
572 uncertainty estimates along with predictions [44]. It can capture complex, non-linear relationships
573 in the data and handles missing values naturally [45]. The flexible method can incorporate various
574 trends and seasonal patterns by choosing kernel functions. However, GPR can be computationally
575 intensive for large datasets due to the need to invert large covariance matrices, and its performance
576 depends on the choice of kernel function, which may require domain expertise or extensive
577 selection procedures [46].

578 3.2.2 Gradient Boosting and CatBoost

579 Gradient Boosting stems from a family of ensemble learning techniques, and CatBoost was
580 recently developed by Yandex [47]. These methods work by building a series of weak learners
581 (typically decision trees) sequentially, with each learner trying to correct the errors of its
582 predecessors. In time series contexts, gradient boosting methods can capture complex, non-linear
583 relationships and handle multiple input variables [48]. CatBoost, in particular, is designed to
584 reduce overfitting and handle categorical variables efficiently, which can be beneficial in
585 forecasting scenarios. However, gradient boosting algorithms do not inherently account for the
586 temporal ordering of data, requiring careful feature engineering to incorporate time-based
587 information. As such, they may also struggle with capturing long-term dependencies without
588 extensive lag features [49].

589 3.2.3 K-Nearest Neighbors (KNN)

590 The K-Nearest Neighbors algorithm is often deployed in regression and classification tasks and
591 can be adapted for time series forecasting [50]. The algorithm is non-parametric and can capture
592 non-linear patterns in the data. In the context of time series forecasting, KNN finds historical
593 periods most similar to the current state and uses their subsequent values to make predictions.
594 KNN can be particularly effective when the time series exhibits recurring patterns or when there
595 are strong analogies between past and future behavior. However, this algorithm's performance can
596 degrade with high-dimensional data and long-term forecasts (especially with the lack of strong

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

597 trends). The choice of distance metric and the number of neighbors (k) can significantly impact
598 forecast accuracy [51].

599 3.2.4 Neural Networks

600 Neural networks, encompassing various architectures beyond LSTM, have become increasingly
601 popular for time series forecasting [52]. Neural networks can capture complex, non-linear
602 relationships in time series data and are capable of handling multiple input variables. The
603 flexibility of their design allows practitioners to tailor architectures to specific forecasting
604 problems. However, neural networks are blackboxes that often require large amounts of training
605 data to perform well and can be prone to overfitting if not properly regularized. Additionally, the
606 selection of appropriate network architecture and hyperparameters often requires significant
607 expertise and computational resources [53].

608 3.2.5 Random Forest, Bagging, and XGBoost

609 Random Forest, Bagging, and XGBoost are ensemble learning methods. Bagging, short for
610 Bootstrap Aggregating is a method to reduce variance in predictive models by creating multiple
611 subsets of the original dataset through bootstrap sampling. This method trains a separate model on
612 each subset and aggregates their predictions. Random Forest, introduced by Breiman in 2001 [54],
613 is a specific implementation of bagging that constructs multiple decision trees and merges their
614 predictions to improve accuracy and control overfitting. XGBoost, developed by Chen and
615 Guestrin in 2016 [55], implements gradient boosted decision trees designed for speed and
616 performance. All these algorithms can handle non-linear relationships and are capable of capturing
617 complex patterns in time series data when properly engineered features are provided. Additionally,
618 while they can handle multiple input variables, they may struggle with capturing long-term
619 dependencies without extensive lag features [56,57].

620 3.2.6 Support Vector Regression (SVR)

621 Support Vector Regression is an extension of Support Vector Machines (SVM) that was developed
622 by Vapnik et al. in the 1990s [58]. In time series forecasting, SVR works by mapping the input
623 data into a high-dimensional feature space and finding a hyperplane that best fits the data while
624 maintaining a specified tolerance margin. SVR is capable of capturing non-linear relationships
625 through the use of kernel functions, making it suitable for complex time series patterns. It is
626 particularly effective when dealing with high-dimensional data and can handle multiple input
627 variables. SVR is less prone to overfitting compared to some other ML algorithms due to its
628 structural risk minimization principle [59]. However, the performance of SVR can be sensitive to
629 the tuning of kernels/hyperparameters and necessitates careful feature engineering to incorporate
630 time-based information [60].

631 **Table 1 A comparison among the examined algorithms in this study**

Algorithm	Family	Methodology & Logic	Typical Use Cases	Strengths/Advantages	Weaknesses/Disadvantages
ARIMA	Statistical	Linear, combines differencing with	Time series data without seasonal patterns.	Flexible, good for none/some seasonal data.	Assumes linearity and stationarity, may not be

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

		autoregression and moving average components.			suitable for complex patterns.
SARIMA	Statistical	Extends ARIMA to include seasonal components.	Seasonal time series data.	Handles seasonality, well-understood.	Computationally intensive, linear assumptions can be overfitted with short time series.
ETS	Statistical	Exponential smoothing (decomposing) techniques with error, trend, and seasonal components.	Short-term forecasting, seasonal and non-seasonal data.	Easy to implement, good for data with trends and seasonality.	May overfit on noisy data and can struggle with abrupt changes.
Holt-Winters	Statistical	Triple exponential smoothing for data with trends and seasonality.	Seasonal time series data.	Simple to implement, effective for additive and multiplicative seasonal patterns.	Assumes additive effects, may not handle high-frequency data well and can struggle with irregular time series.
Prophet	Statistical	Additive and decomposable model with trend, seasonality, and holidays.	Daily data with strong multiple seasonality patterns, missing data, and outliers.	Robust to missing data, handles outliers, automatically detects changepoints, and incorporates domain knowledge easily.	Less effective for non-daily data or complex patterns, and may struggle with short-term forecasts.
Theta	Statistical	Decomposes data into two 'theta lines' for different trend assumptions (e.g., long and short-term components).	Time series data with trends.	Simple, effective for data with a trend.	Less effective for seasonal or non-linear data. Offers limited flexibility for complex patterns.
Simple Moving Average	Statistical	Calculates average over a fixed window of past observations (n).	Smoothing noisy data, simple forecasts.	Simple to understand and implement.	Not adaptive, lags in response to real trend changes.
VAR	Statistical	Vector Autoregression, multivariate linear model relating different time series variables.	Multivariate time series data.	Captures relationships between multiple series, good for stationary series.	Requires all series to be stationary, high computational cost.
Croston's Method	Statistical	Separately forecasts non-zero demand sizes and intervals and adjusts for intermittent demand.	Forecasting intermittent demand.	Good for sparse or intermittent data.	May be biased, assumes demand pattern does not change.
LSTM	ML	A type of recurrent neural network that uses gates to control information flow.	Complex patterns, large datasets, non-linear relationships.	Good for capturing long dependencies, non-linear patterns.	Requires large datasets, computationally intensive. Blackbox nature limits interpretability.
Neural Networks	ML	Layers of interconnected neurons learning data features.	Complex nonlinear patterns, high-dimensional data.	Highly flexible, powerful for complex patterns and can handle non-linear relationships.	Requires large data and careful feature engineering, prone to overfitting, black box.
Gaussian Process Regression	ML	Non-parametric kernel-based probabilistic model.	Small to medium datasets, needing uncertainty estimation.	Provides uncertainty measures, flexible.	Computationally expensive, not for large data.
KNN	ML	Predicts based on similar historical patterns by using 'k' nearest points for prediction.	Small datasets, simple non-linear patterns.	Simple, non-parametric, and effective for non-linearities in small datasets.	Not scalable, sensitive to the choice of k and noisy data.
SVR	ML	Fits within a certain threshold and finds the optimal hyperplane in high-dimensional space.	Regression with clear margin of error.	Effective in high-dimensional space, robust to outliers.	Requires good parameter tuning, and can be computationally intensive for very large datasets.
Random Forest	ML	Ensemble of decision trees, averaging to improve prediction.	Various problems.	Robust, handles overfitting well, good for mixed data types,	Requires feature engineering for temporal aspects.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

				and can provide feature importance	
XGBoost	ML	Gradient boosting with decision trees optimized for speed and performance.	Various problems.	Fast, scalable, high performance, handles various data types.	Prone to overfitting if not tuned properly.
Gradient Boosting	ML	Sequential correction of predecessor's errors, using decision trees.	Various problems.	Reduces bias and variance, powerful.	Computationally intensive and prone to overfitting.
CatBoost	ML	Categorical Boosting.	Datasets with many categorical features.	Efficient with categorical data, less prone to overfitting.	Slightly slower compared to other boosting methods.
Bagging	ML	Bootstrap aggregating, reduces variance by averaging a set of parallel estimators.	Reducing variance in noisy data sets.	Reduces overfitting, robust to noisy data.	Can be less effective on biased models, high memory consumption.

632

633 4.0 Description of benchmarking experiments, metrics, and datasets

634 This benchmarking analysis involves a set of 25 synthetic and 25 real timeseries. This analysis
 635 was run and evaluated in a Python 3.10.5 environment using an Intel(R) Core(TM) i7-9700F CPU
 636 @ 3.00GHz and an installed RAM of 32.0GB. All algorithms were run in default settings to allow
 637 fairness and ensure reproducibility, and the performance of each algorithm was evaluated through
 638 several metrics, as discussed below and listed in Table 2. These metrics, along with the selected
 639 sizes of datasets, followed the recommendations of [7,61].

640 We utilize four primary metrics that can be classified under global/general and specific/internal
 641 metrics. The global metrics are suitable for broad comparisons and evaluations across multiple
 642 datasets and models (e.g., Mean Absolute Scaled Error (MASE) and Dynamic Time Warping
 643 (DTW)). On the other hand, internal metrics are used to provide detailed insights into particular
 644 aspects of model performance (such as mean absolute deviation (MAD) and direction accuracy
 645 (DA)). It is worth noting that other metrics (i.e., Root Mean Square Error (RMSE) and the Mean
 646 Absolute Error (MAE)) were not used herein due to their inherent limitations and vulnerabilities
 647 with respect to time series analysis, as pointed out by [62,63].

648 The MASE is a relative measure of forecast accuracy that scales the forecast error by the in-sample
 649 mean absolute error. MASE is scale-independent and can be used to compare forecast accuracy
 650 across different time series [64,65]. DTW measures the similarity between two time series by
 651 finding an optimal alignment between them. Unlike simple distance measures, DTW can handle
 652 time shifts and distortions by allowing flexible matching of time indices [66]. The MAD measures
 653 the average absolute error between the actual and forecasted values to clearly indicate the average
 654 magnitude of forecast errors [67]. The DA measures how well the model predicts the direction of
 655 the time series movement. This metric evaluates whether the forecast correctly predicts the
 656 increase or decrease in the actual values from one time point to the next.

657 Table 2 List of performance metrics.

Type	Metric	Formula
Specific/Internal Metrics	Mean Absolute Deviation (MAD)	$MAD = \frac{1}{n} \sum_{t=1}^n y_t - \hat{y}_t $ <p>Where: • y_t is the actual value at time t.</p>

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

		<ul style="list-style-type: none"> • y_t^{\wedge} is the forecasted value at time t. • n is the total number of observations.
	Direction Accuracy (DA)	$DA = \frac{1}{n-1} \sum_{t=2}^n \mathbb{I}((y_t - y_{t-1})(\hat{y}_t - \hat{y}_{t-1}) > 0)$ <p>Where:</p> <ul style="list-style-type: none"> • y_t is the actual value at time t. • y_t^{\wedge} is the forecasted value at time t. • \mathbb{I} is the indicator function that equals 1 if the condition inside is true and 0 otherwise. • n is the total number of observations.
Global/External Metrics	Mean Absolute Scaled Error (MASE)	$MASE = \frac{\frac{1}{n} \sum_{t=1}^n y_t - \hat{y}_t }{\frac{1}{n} \sum_{t=1}^n y_t - y_{t-1} }$ <p>Where:</p> <ul style="list-style-type: none"> • y_t is the actual value at time t. • y_t^{\wedge} is the forecasted value at time t. • n is the total number of observations.
	Dynamic Time Warping (DTW)	$DTW(A,B) = \min \sqrt{\sum_{i=2}^n (a_i - b_i)^2}$ <p>Where:</p> <ul style="list-style-type: none"> • $A=(a_1, a_2, \dots, a_n)$ and $B = (b_1, b_2, \dots, b_m)$ are two sequences of length n and m respectively. • i' is the optimal alignment index of b corresponding to a_i.

658

659 4.1 Synthetic datasets

660 Twenty five synthetic timeseries of various scenarios were generated and examined by all
661 algorithms (see Fig. 1 and Table 3). These timeseries were generated via the generate_time_series
662 function, which allows researchers to generate controlled datasets that can be used to benchmark
663 and evaluate the performance of time series forecasting models. This particular function accepts a
664 specific mathematical function and the number of data points (n_points) as input parameters. Then,
665 this function generates a sequence of equally spaced time points over a specified range (t_max).
666 The chosen function is applied to these time points to produce the corresponding time series data.
667 Gaussian noise is added to simulate real-world conditions where data often includes random
668 variations. The generate_time_series function starts by creating an array of time points using
669 numpy.linspace, which ensures an even distribution of points between 0 and the specified t_max.
670 This array of time points, t, is then passed to the provided time series function (func), which applies
671 the mathematical transformation and returns the resulting data series.

672 Table 3 Parameters used in the synthetic timeseries

Function	Description	Mathematical Expression	Noise Level	Characteristics
Linear trend	Linear increase with Gaussian noise	$0.5t + \epsilon$	$\sigma=0.1$	Simple trend
Quadratic trend	Quadratic increase with Gaussian noise	$0.05t^2 + \epsilon$	$\sigma=0.1$	Parabolic trend
Exponential growth	Exponential increase with Gaussian noise	$e^{0.1t} + \epsilon$	$\sigma=0.1$	Exponential trend
Sine wave (seasonal)	Periodic sine wave with Gaussian noise	$\sin(2\pi t) + \epsilon$	$\sigma=0.1$	Seasonal, periodic pattern

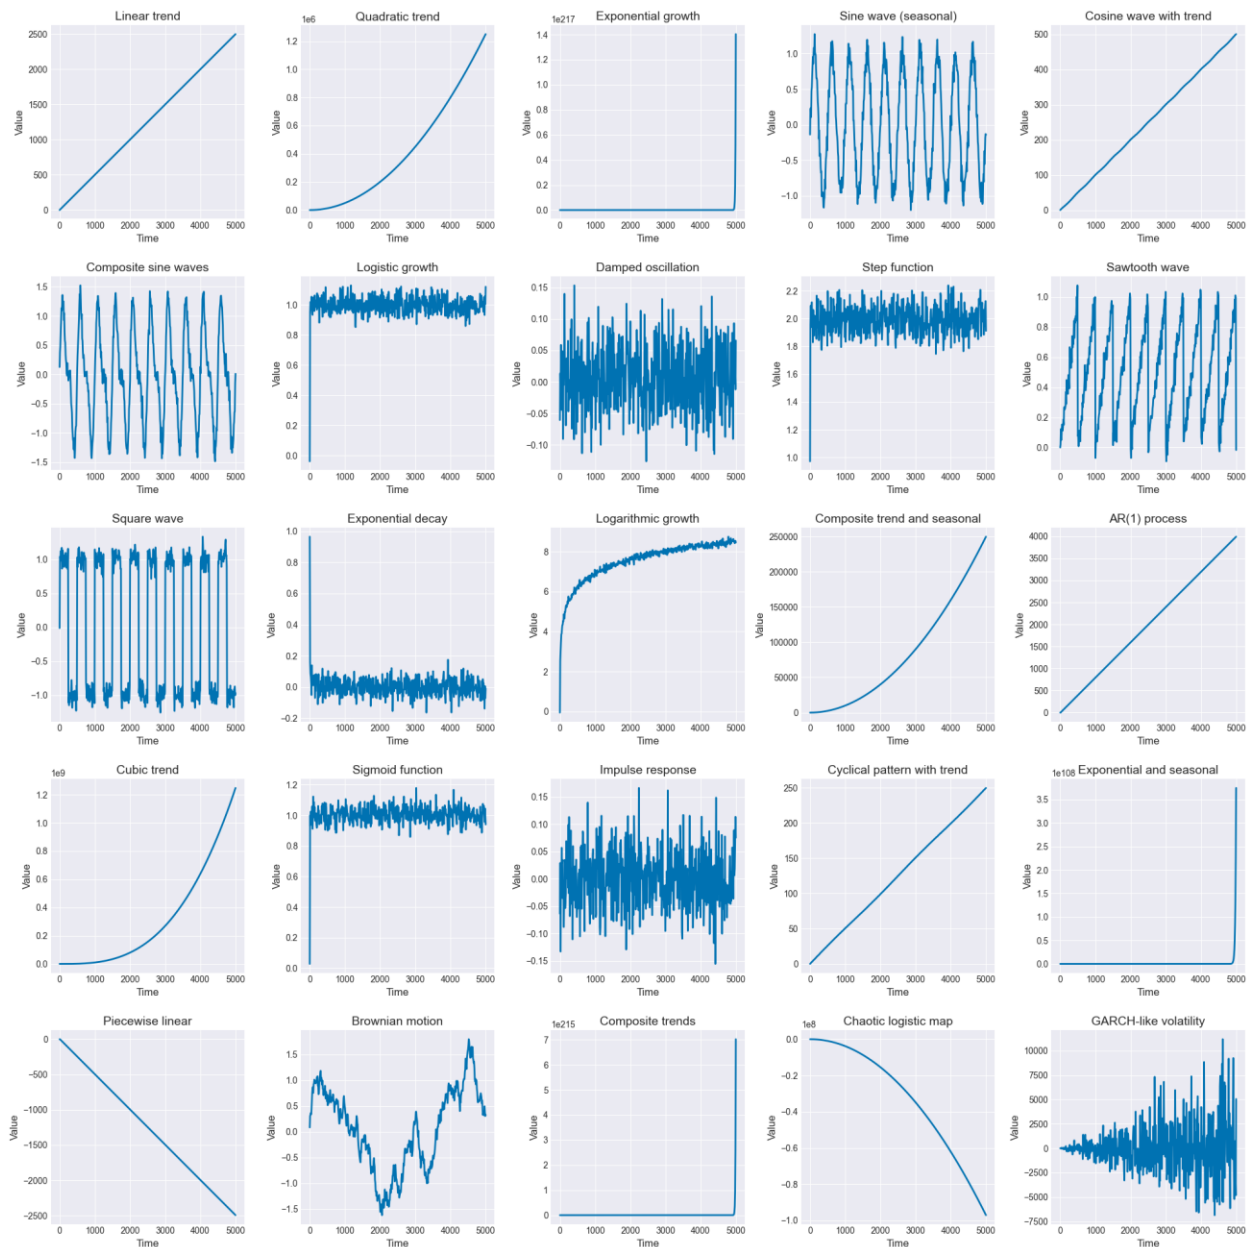
Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

Cosine wave with linear trend	Cosine wave superimposed on a linear trend with noise	$\cos(2\pi t)+0.1t+\epsilon$	$\sigma=0.1$	Trend and seasonality combination
Composite of sine waves	Multiple sine waves combined with Gaussian noise	$\sin(2\pi t)+0.5\sin(4\pi t)+\epsilon$	$\sigma=0.1$	Multiple seasonalities
Logistic growth	Sigmoidal growth with Gaussian noise	$1/1+e^{-t}+5+\epsilon$	$\sigma=0.05$	Non-linear growth
Damped oscillation	Exponentially damped sine wave with noise	$e^{-0.1t}\sin(2\pi t)+\epsilon$	$\sigma=0.05$	Oscillatory effect
Step function	Discrete steps with Gaussian noise	$\text{step}(t)+\epsilon$	$\sigma=0.1$	Abrupt changes
Sawtooth wave	Linear periodic rise with a drop and Gaussian noise	$(t\%1)+\epsilon$	$\sigma=0.05$	Sharp transitions
Square wave	Alternating high and low values with Gaussian noise	$\text{sign}(\sin(2\pi t))+\epsilon$	$\sigma=0.1$	Discrete, binary states
Exponential decay	Exponential decrease with Gaussian noise	$e^{-0.2t}+\epsilon$	$\sigma=0.05$	Decay trend
Logarithmic growth	Logarithmic increase with Gaussian noise	$\log(t+1)+\epsilon$	$\sigma=0.1$	Logarithmic trend
Composite trend, seasonal, and noise	Combination of quadratic trend, sine wave, and noise	$0.01t^2+\sin(2\pi t)+0.5\epsilon$	$\sigma=1$	Complex pattern
Autocorrelated process (AR(1))	Autoregressive process with Gaussian noise	$0.8y_{t-1}+\epsilon$	$\sigma=0.5$	Dependency on previous values
Polynomial trend (cubic)	Cubic polynomial trend with Gaussian noise	$0.01t^3-0.1t^2+0.5t+\epsilon$	$\sigma=0.1$	Higher-order polynomial trend
Sigmoid function	Sigmoidal growth with Gaussian noise	$1/1+e^{-t^5}+\epsilon$	$\sigma=0.05$	Non-linear growth
Impulse response	Exponentially decaying sinusoidal impulse with noise	$e^{-t}\sin(2\pi t)+\epsilon$	$\sigma=0.05$	Impulse-like behavior
Cyclical pattern with trend	Sine wave with linear trend and Gaussian noise	$\sin(2\pi t/5)+0.05t+\epsilon$	$\sigma=0.1$	Cyclic and trending
Composite of exponential growth and seasonal pattern	Exponential growth with superimposed sine wave and noise	$e^{0.05t}+0.5\sin(2\pi t)+\epsilon$	$\sigma=0.1$	Complex trend and seasonality
Piecewise linear function	Linear segments with Gaussian noise	$\text{piecewise}(t)+\epsilon$	$\sigma=0.1$	Segmented linear behavior
Brownian motion (random walk)	Cumulative sum of Gaussian noise	$\sum\epsilon$	$\sigma=0.1$	Stochastic, random walk
Composite of multiple trends	Combination of quadratic, sinusoidal, and exponential trends with noise	$0.01t^2+0.1\sin(2\pi t)+0.05e^{0.1t}+\epsilon$	$\sigma=0.1$	Multiple trend components
Chaotic logistic map	Logistic map function with chaos	$3.9t(1-t)+\epsilon$	$\sigma=0.1$	Chaotic behavior
GARCH-like volatility clustering	Gaussian noise with time-varying volatility	$N[(0, 0.1 + 0.9 \text{abs}(y_{t-1}))]$	$(0, 0.1 + 0.9 \text{abs}(y_{t-1}))$	Volatility clustering

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



674
675

Fig. 1 Visualization of the synthetic datasets

676 The outcome of the benchmarking analysis on the synthetic datasets is listed in Table 4. As one
677 can see, this table showcases three different ranking methods (namely, based on the average
678 ranking, normalized ranking, and wins). All ranking systems used the abovementioned DA, DTW,
679 MASE, and MAD metrics, wherein lower values indicate better performance (rank 1 is best),
680 except for the DA metric, where higher values indicate better performance (rank 1 is best).

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

681 The first ranking system represents averages across all datasets for each algorithm. Each metric
 682 was ranked individually, and the average of these ranks determined the Average Rank for each
 683 algorithm. The Final Rank was then assigned based on the Average Rank values. Then, the second
 684 ranking system involves averaging the original metric values across all datasets for each algorithm.
 685 These averages were then normalized to a 0-1 scale for each metric. The Average Normalized Score
 686 is then computed as the mean of all normalized metric scores for each algorithm, and the Final
 687 Rank is derived from these average normalized scores to allow for a fair comparison across
 688 different metrics and algorithms. Finally, the third ranking system adopted the ranking by Wins
 689 method. In this system, metric values were averaged across all datasets for each algorithm. Each
 690 metric was ranked individually. The rank columns display the rank for each metric, and the
 691 Average Rank reflects the mean of these ranks for each algorithm. The Final Rank was determined
 692 based on the Average Rank values. This ranking method aimed to balance performance across all
 693 metrics, emphasizing how frequently each algorithm performed best in each metric.

694 Table 4 presents a collective view of the algorithmic performance across the different datasets and
 695 systems used. As one can see, SPINEX ranks 5th under the first ranking system and 1st under the
 696 other two ranking systems. Despite not being the top-ranked algorithm in the first ranking system,
 697 SPINEX consistently performed well across different metrics. This performance suggests a well-
 698 rounded response. When comparing SPINEX to other algorithms such as SARIMA, Prophet, Holt-
 699 Winters, and Theta, it is evident that SPINEX stands out regarding consistent performance and
 700 robustness. Similarly, Prophet and Theta showed competitive performance but could not match
 701 SPINEX's consistency across all ranking systems.

702 Table 4 Ranking results on real data

Algorithm	Direction Accuracy	DTW	MASE	MAD	Direction Accuracy (rank)	DTW (rank)	MASE (rank)	MAD (rank)	Average (rank)	Final (rank)
Based on average rankings										
SARIMA	0.578	15.353	2625.313	0.116	3.570	5.830	7.890	9.250	6.640	1
Prophet	0.546	2.058	58.910	0.095	3.870	6.210	7.430	9.030	6.640	2
Holt-Winters	0.572	15.419	46.643	0.291	3.360	6.400	7.980	8.910	6.660	3
Theta	0.550	2.451	82.481	0.100	3.600	7.080	7.520	9.660	6.960	4
SPINEX	0.602	1.956	45.676	0.075	3.230	6.840	10.300	8.950	7.330	5
ARIMA	0.264	2.544	84.369	0.097	7.190	8.800	7.900	6.980	7.720	6
Croston	0.000	2.602	87.962	0.101	10.370	9.130	7.850	5.360	8.180	7
ETS	0.000	2.603	87.962	0.101	10.370	9.360	7.980	5.350	8.260	8
LSTM	0.502	3.756	115.265	0.090	4.390	9.060	9.500	10.140	8.270	9
Random Forest	0.000	2.698	88.205	0.101	10.370	10.120	8.990	5.290	8.690	10
Bagging	0.000	2.698	88.205	0.101	10.370	10.120	8.990	5.290	8.690	10
Gradient Boosting	0.000	2.689	88.144	0.101	10.370	10.210	8.950	5.470	8.750	12
XGBoost	0.000	2.695	88.281	0.101	10.370	10.560	9.310	5.340	8.890	13
SMA	0.000	2.669	88.733	0.101	10.370	10.530	9.640	5.400	8.980	14
KNN	0.000	2.669	88.733	0.101	10.370	10.530	9.590	5.450	8.990	15
CatBoost	0.000	2.666	89.402	0.101	10.370	10.990	9.820	5.330	9.130	16
Neural Network	0.518	6.952	194.656	0.110	3.660	13.880	14.080	11.850	10.870	17
SVR	0.166	4.262	157.608	0.115	7.750	13.230	14.220	11.280	11.620	18
Gaussian Process	0.255	7.354	144.260	0.097	6.840	15.660	16.440	9.300	12.060	19
Based on normalized rankings										
SPINEX	0.602	1.956	45.676	0.075	0.000	0.000	0.000	0.000	0.000	1
Prophet	0.546	2.058	58.910	0.095	0.094	0.008	0.005	0.093	0.050	2

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

Theta	0.550	2.451	82.481	0.100	0.087	0.037	0.014	0.115	0.063	3
LSTM	0.502	3.756	115.265	0.090	0.167	0.134	0.027	0.070	0.100	4
ARIMA	0.264	2.544	84.369	0.097	0.562	0.044	0.015	0.103	0.181	5
Neural Network	0.518	6.952	194.656	0.110	0.139	0.371	0.058	0.164	0.183	6
Gaussian Process	0.255	7.354	144.260	0.097	0.577	0.401	0.038	0.103	0.280	7
SVR	0.166	4.262	157.608	0.115	0.725	0.171	0.043	0.184	0.281	8
Croston	0.000	2.602	87.962	0.101	1.000	0.048	0.016	0.119	0.296	9
ETS	0.000	2.603	87.962	0.101	1.000	0.048	0.016	0.119	0.296	10
KNN	0.000	2.669	88.733	0.101	1.000	0.053	0.017	0.119	0.297	11
SMA	0.000	2.669	88.733	0.101	1.000	0.053	0.017	0.119	0.297	11
CatBoost	0.000	2.666	89.402	0.101	1.000	0.053	0.017	0.119	0.297	13
Gradient Boosting	0.000	2.689	88.144	0.101	1.000	0.054	0.016	0.119	0.297	14
XGBoost	0.000	2.695	88.281	0.101	1.000	0.055	0.017	0.119	0.298	15
Random Forest	0.000	2.698	88.205	0.101	1.000	0.055	0.016	0.119	0.298	16
Bagging	0.000	2.698	88.205	0.101	1.000	0.055	0.016	0.119	0.298	16
Holt-Winters	0.572	15.419	46.643	0.291	0.051	1.000	0.000	1.000	0.513	18
SARIMA	0.578	15.353	2625.313	0.116	0.040	0.995	1.000	0.192	0.557	19
Based on wins										
SPINEX	0.602	1.956	45.676	0.075	1	1	1	1	1	1
Prophet	0.546	2.058	58.910	0.095	5	2	3	3	3.25	2
Theta	0.550	2.451	82.481	0.100	4	3	4	6	4.25	3
ARIMA	0.264	2.544	84.369	0.097	8	4	5	4	5.25	4
Croston	0.000	2.602	87.962	0.101	11	5	6	7	7.25	5
ETS	0.000	2.603	87.962	0.101	11	6	7	7	7.75	6
Gradient Boosting	0.000	2.689	88.144	0.101	11	10	8	7	9	7
SMA	0.000	2.669	88.733	0.101	11	8	12	7	9.5	8
LSTM	0.502	3.756	115.265	0.090	7	14	15	2	9.5	8
KNN	0.000	2.669	88.733	0.101	11	8	12	7	9.5	8
Random Forest	0.000	2.698	88.205	0.101	11	12	9	7	9.75	11
CatBoost	0.000	2.666	89.402	0.101	11	7	14	7	9.75	11
Bagging	0.000	2.698	88.205	0.101	11	12	9	7	9.75	11
XGBoost	0.000	2.695	88.281	0.101	11	11	11	7	10	14
Holt-Winters	0.572	15.419	46.643	0.291	3	19	2	19	10.75	15
Gaussian Process	0.255	7.354	144.260	0.097	9	17	16	5	11.75	16
Neural Network	0.518	6.952	194.656	0.110	6	16	18	16	14	17
SARIMA	0.578	15.353	2625.313	0.116	2	18	19	18	14.25	18
SVR	0.166	4.262	157.608	0.115	10	15	17	17	14.75	19

703

704 Figure 2 shows a more detailed examination of the performance of all algorithm algorithms
 705 evaluated across different settings and metrics. This evaluation was conducted for two different
 706 parameters: maximum time (t_{max}) and number of sequence points (n_{points}).

707 In terms of DA, which measures how well the model predicts the direction of the time series
 708 movement, SPINEX maintained strong performance across both settings. The graphs indicate that
 709 SPINEX's performance remained relatively stable and high compared to other algorithms as the sub-
 710 settings increased. Such stability can be crucial for applications requiring reliable directional
 711 predictions. More specifically, in the t_{max} graph, SPINEX showed a slight improvement with higher
 712 sub-settings, indicating its adaptability to longer forecasting horizons. Similarly, in the n_{points}
 713 graph, SPINEX outperformed most other algorithms, demonstrating its effectiveness in handling
 714 varying data point quantities.

715 For the DTW metric, which measures the alignment between predicted and actual time series,
 716 SPINEX also performed well across different settings. In both graphs, SPINEX maintained lower

This is a preprint draft. The published article can be found at: <https://doi.org/10.1016/j.cie.2024.110812>.

Please cite this paper as:

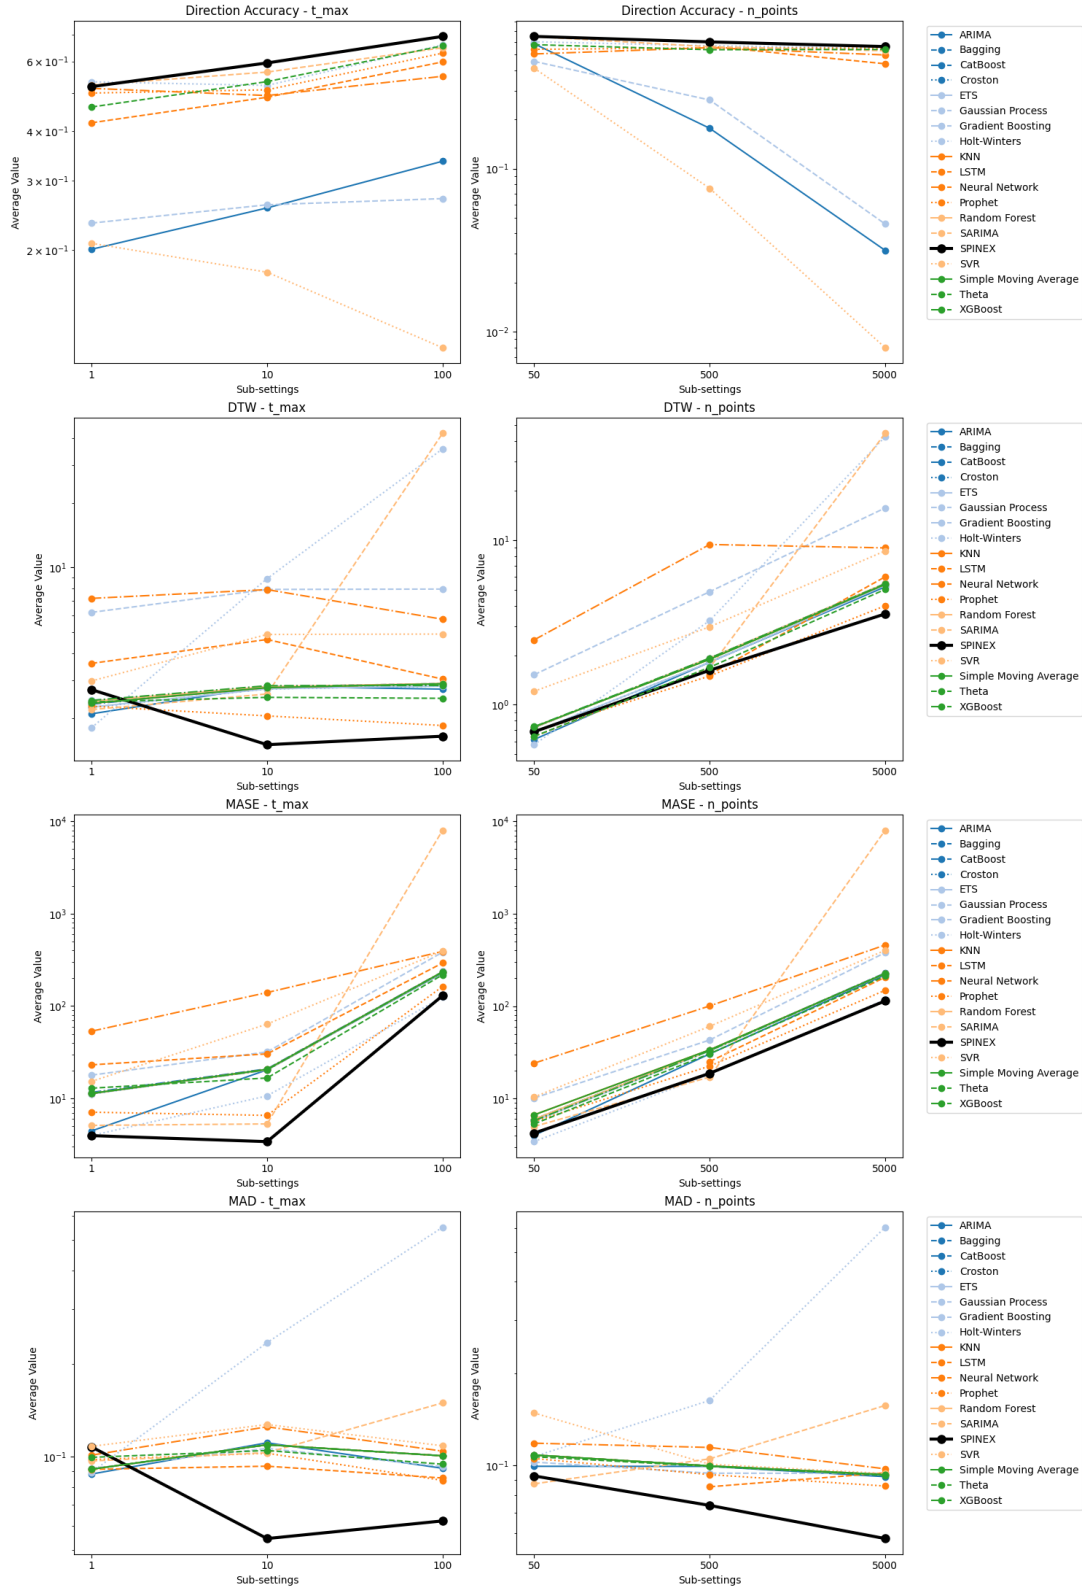
Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

717 DTW values, indicating closer alignment and better predictive accuracy. The t_{\max} graph shows that
718 SPINEX's DTW values remained relatively stable, suggesting its robustness to changes in the
719 forecast length. The n_{points} graph further highlights SPINEX's capability to handle datasets with
720 varying numbers of points without significant loss in accuracy. Furthermore, SPINEX maintained
721 lower MASE and MAD values compared to many other algorithms. This performance indicates
722 that SPINEX can provide accurate forecasts. Figure 3 presents a visual example of two time series
723 as predicted by SPINEX and other algorithms.

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

Algorithm Performance Across Settings



Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

725

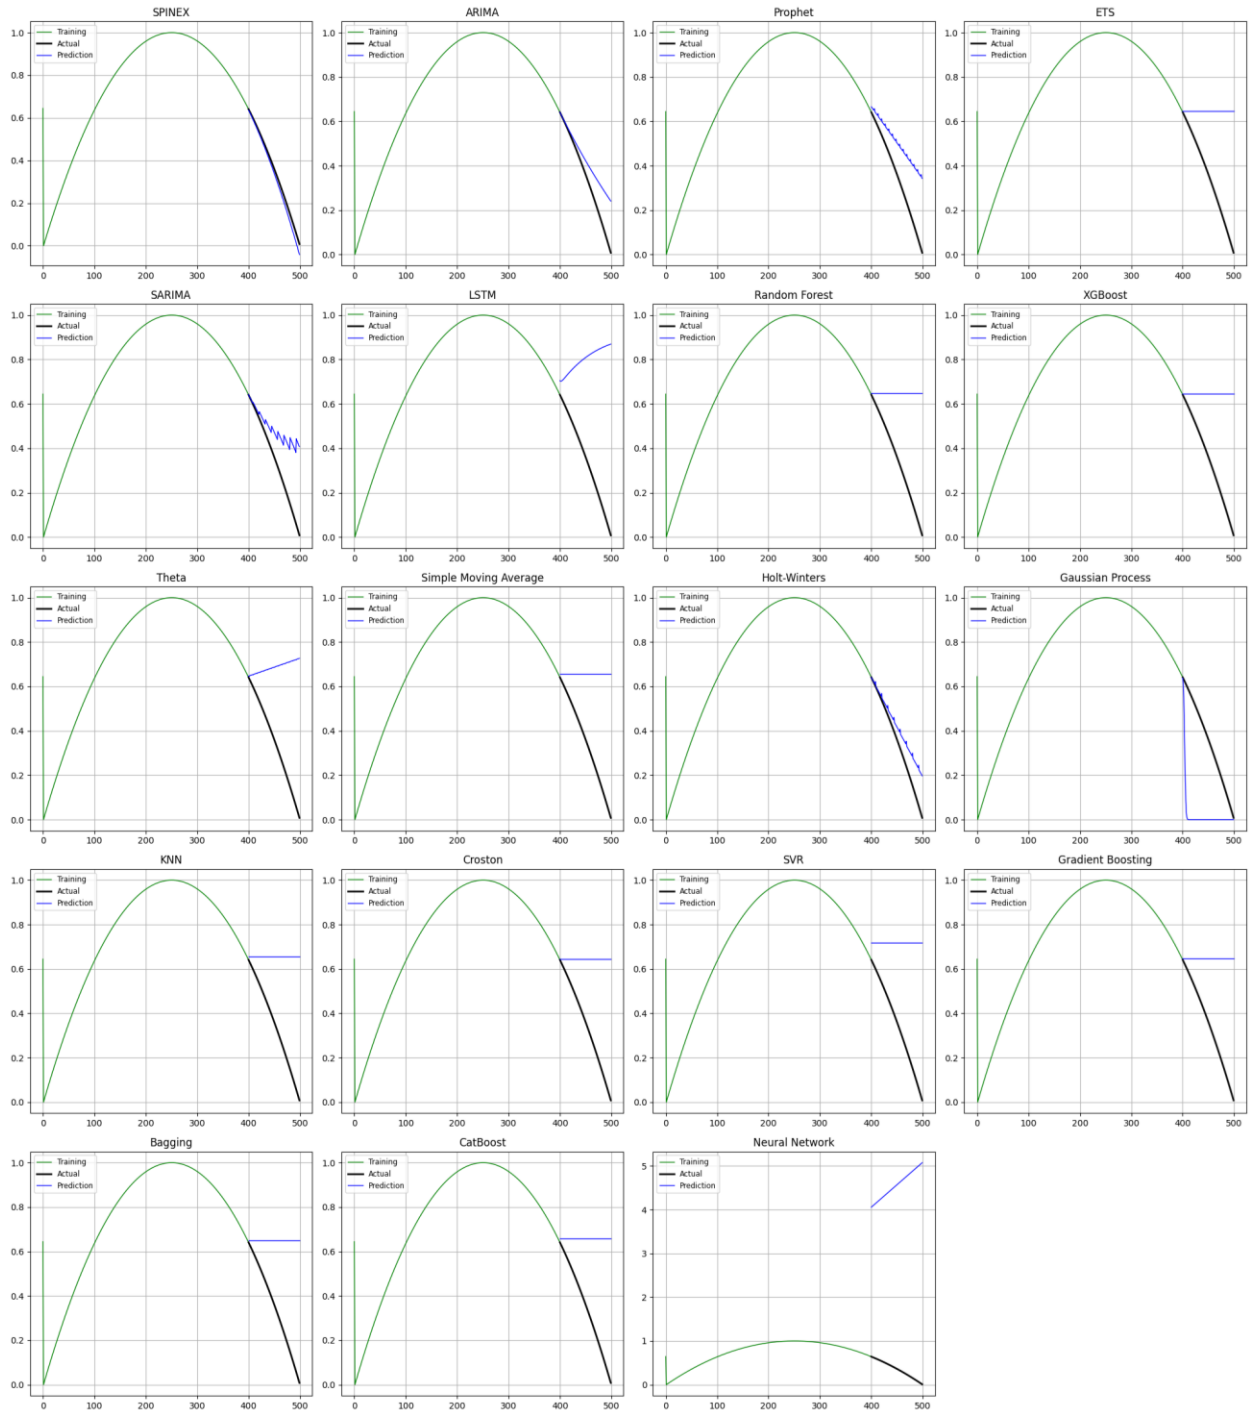
Figure 2 Individual rankings per algorithm for the internal and external metrics.



726

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



727
728
729

Fig. 3 Visualization of forecasting on Dataset no. 2 [$t_{\max} = 10$, $n_{\text{points}} = 50$] (top) and Dataset no. 6 [$t_{\max} = 1$, $n_{\text{points}} = 500$] (bottom)

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

730 **4.2 Real datasets**

731 Twenty four real datasets were used herein to further evaluate the performance of SPINEX against
 732 the other algorithms listed above. These datasets span univariate and multivariate scenarios(see
 733 Table 5 and Fig. 4). Additional details can be found in the cited sources.

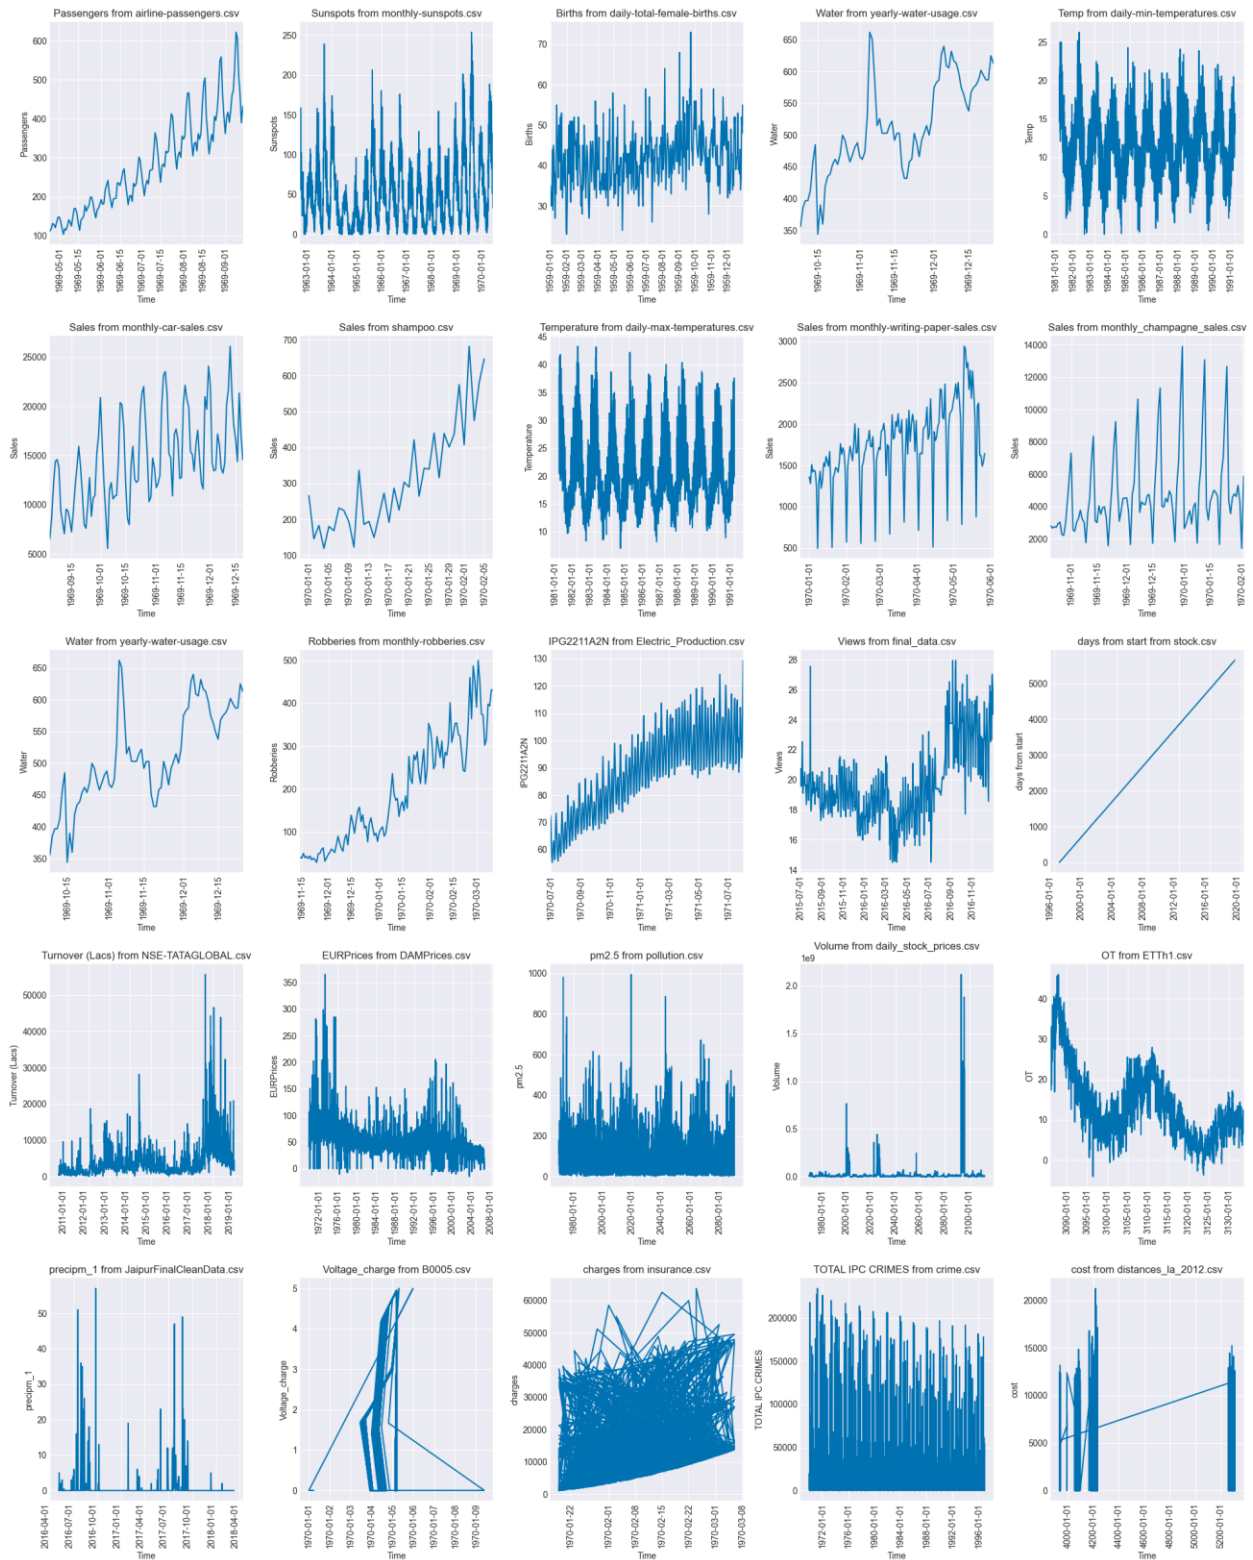
734 **Table 5 Real datasets used in the analysis**

Type	Dataset Name	Samples*	Features	References
Univariate	Airline Passengers	144	2	[68]
	Sunspots	2820	2	[69]
	Daily Female Births	365	2	[70]
	Yearly Water Usage	79	2	[71]
	Daily Minimum Temperatures	3650	2	[72]
	Monthly Car Sales	108	2	[73]
	Shampoo Sales Data	36	2	[74]
	Temperature Data	3650	2	[75]
	Monthly Writing Paper Sales	147	2	[76]
	Monthly Champagne Sales	105	2	[77]
	Monthly Robberies	118	2	[78]
	Electric Production	397	2	[79]
Web Traffic Dataset	550	2	[80]	
Multivariate	Stock and PM2.5 Prediction	5650	10	[81]
	Tata Global Forecasting	2100	8	[82]
	International Airline Passengers	13391	6	[83]
	Pollution Dataset	43824	13	[84]
	Daily Stock Prices	52000	8	[85]
	ETT-small	17420	8	[86]
	Jaipur Final Clean Data	676	40	[87]
	Aprocessed	604802	17	[88]
	Insurance	1338	7	[89]
	Indian Crime Data Analysis Forecasting I	9840	33	[90]
Indian Crime Data Analysis Forecasting II	295374	3	[90]	

735 *Large datasets were stopped at 5000 data points, given the computational resources available during this study.

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

737 Fig. 4 Visualization of the real datasets

738 The benchmarking and ranking analysis results are examined similarly to the case of synthetic data
 739 by using the average, normalized, and win methods. These results are listed in Table 6. This table
 740 shows that SPINEX consistently ranks in the top two positions, along with the Holt-Winters
 741 algorithm.

742 Table 6 Ranking results on real data

Algorithm	Direction Accuracy	DTW	MASE	MAD	Direction Accuracy (rank)	DTW (rank)	MASE (rank)	MAD (rank)	Overall rank	Final (rank)
Based on average rankings										
Holt-Winters	0.549	1.998	3.807	0.060	5.400	9.960	13.960	15.960	11.320	1
SPINEX	0.563	1.635	2.671	0.063	5.320	5.080	17.240	18.520	11.540	2
SARIMA	0.540	2.422	3.778	0.065	5.400	11.160	16.040	15.800	12.100	3
Prophet	0.464	2.100	6.896	0.071	8.440	8.680	13.560	19.640	12.580	4
Theta	0.491	2.299	13.739	0.073	6.760	13.400	18.040	13.160	12.840	5
SMA	0.000	2.472	23.792	0.078	20.040	19.320	14.840	10.600	16.200	6
LSTM	0.402	15.077	48.378	0.075	9.500	17.420	20.670	17.830	16.350	7
KNN	0.000	2.472	23.792	0.078	20.040	19.720	15.160	10.600	16.380	8
XGBoost	0.000	2.594	24.067	0.078	20.040	20.200	16.840	10.280	16.840	9
Gradient Boosting	0.000	2.567	24.043	0.078	20.040	20.120	16.680	10.760	16.900	10
Random Forest	0.000	2.597	23.894	0.078	20.040	20.920	16.680	11.080	17.180	11
Bagging	0.000	2.597	23.894	0.078	20.040	20.920	16.680	11.080	17.180	11
CatBoost	0.000	2.629	24.776	0.078	20.040	21.640	18.200	9.960	17.460	13
Croston	0.000	2.610	23.925	0.078	20.040	21.400	19.400	10.840	17.920	14
ETS	0.000	2.612	23.927	0.078	20.040	21.800	19.480	10.840	18.040	15
ARIMA	0.146	2.536	23.839	0.079	15.400	21.000	20.440	17.640	18.620	16
SVR	0.156	3.348	105.681	0.093	15.080	21.480	23.080	19.880	19.880	17
Neural Network	0.489	5.460	37.621	0.084	6.360	27.320	26.040	20.920	20.160	18
Gaussian Process	0.197	6.240	159.080	0.082	13.080	33.640	32.200	19.560	24.620	19
Based on normalized rankings										
Holt-Winters	0.549	1.998	3.807	0.060	0.024	0.027	0.007	0.000	0.015	1
SPINEX	0.563	1.635	2.671	0.063	0.000	0.000	0.000	0.093	0.023	2
SARIMA	0.540	2.422	3.778	0.065	0.039	0.059	0.007	0.176	0.070	3
Prophet	0.464	2.100	6.896	0.071	0.176	0.035	0.027	0.333	0.143	4
Theta	0.491	2.299	13.739	0.073	0.126	0.049	0.071	0.400	0.161	5
Neural Network	0.489	5.460	37.621	0.084	0.130	0.285	0.223	0.730	0.342	6
ARIMA	0.146	2.536	23.839	0.079	0.740	0.067	0.135	0.580	0.381	7
Simple	0.000	2.472	23.792	0.078	1.000	0.062	0.135	0.546	0.436	8
KNN	0.000	2.472	23.792	0.078	1.000	0.062	0.135	0.546	0.436	9
Gradient Boosting	0.000	2.567	24.043	0.078	1.000	0.069	0.137	0.546	0.438	10
Bagging	0.000	2.597	23.894	0.078	1.000	0.072	0.136	0.546	0.438	11
Random	0.000	2.597	23.894	0.078	1.000	0.072	0.136	0.546	0.438	11
XGBoost	0.000	2.594	24.067	0.078	1.000	0.071	0.137	0.546	0.438	13
Croston	0.000	2.610	23.925	0.078	1.000	0.072	0.136	0.546	0.439	14
ETS	0.000	2.612	23.927	0.078	1.000	0.073	0.136	0.546	0.439	15
CatBoost	0.000	2.629	24.776	0.078	1.000	0.074	0.141	0.546	0.440	16
LSTM	0.402	15.077	48.378	0.075	0.286	1.000	0.292	0.456	0.508	17
SVR	0.156	3.348	105.681	0.093	0.722	0.127	0.659	1.000	0.627	18
Gaussian Process	0.197	6.240	159.080	0.082	0.650	0.343	1.000	0.672	0.666	19
Based on wins										
SPINEX	0.563	1.635	2.671	0.063	1	1	1	2	1.25	1
Holt-Winters	0.549	1.998	3.807	0.060	2	2	3	1	2	2
SARIMA	0.540	2.422	3.778	0.065	3	5	2	3	3.25	3
Prophet	0.464	2.100	6.896	0.071	6	3	4	4	4.25	4
Theta	0.491	2.299	13.739	0.073	4	4	5	5	4.5	5
SMA	0.000	2.472	23.792	0.078	11	6	6	7	7.5	6
KNN	0.000	2.472	23.792	0.078	11	7	7	7	8	7
ARIMA	0.146	2.536	23.839	0.079	10	8	8	16	10.5	8
XGBoost	0.000	2.594	24.067	0.078	11	10	14	7	10.5	8
Croston	0.000	2.610	23.925	0.078	11	13	11	7	10.5	8

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

ETS	0.000	2.612	23.927	0.078	11	14	12	7	11	11
Random Forest	0.000	2.597	23.894	0.078	11	11	9	13	11	11
Bagging	0.000	2.597	23.894	0.078	11	11	9	13	11	11
Gradient Boosting	0.000	2.567	24.043	0.078	11	9	13	13	11.5	14
CatBoost	0.000	2.629	24.776	0.078	11	15	15	7	12	15
LSTM	0.402	15.077	48.378	0.075	7	19	17	6	12.25	16
Neural Network	0.489	5.460	37.621	0.084	5	17	16	18	14	17
Gaussian Process	0.197	6.240	159.080	0.082	8	18	19	17	15.5	18
SVR	0.156	3.348	105.681	0.093	9	16	18	19	15.5	18

743

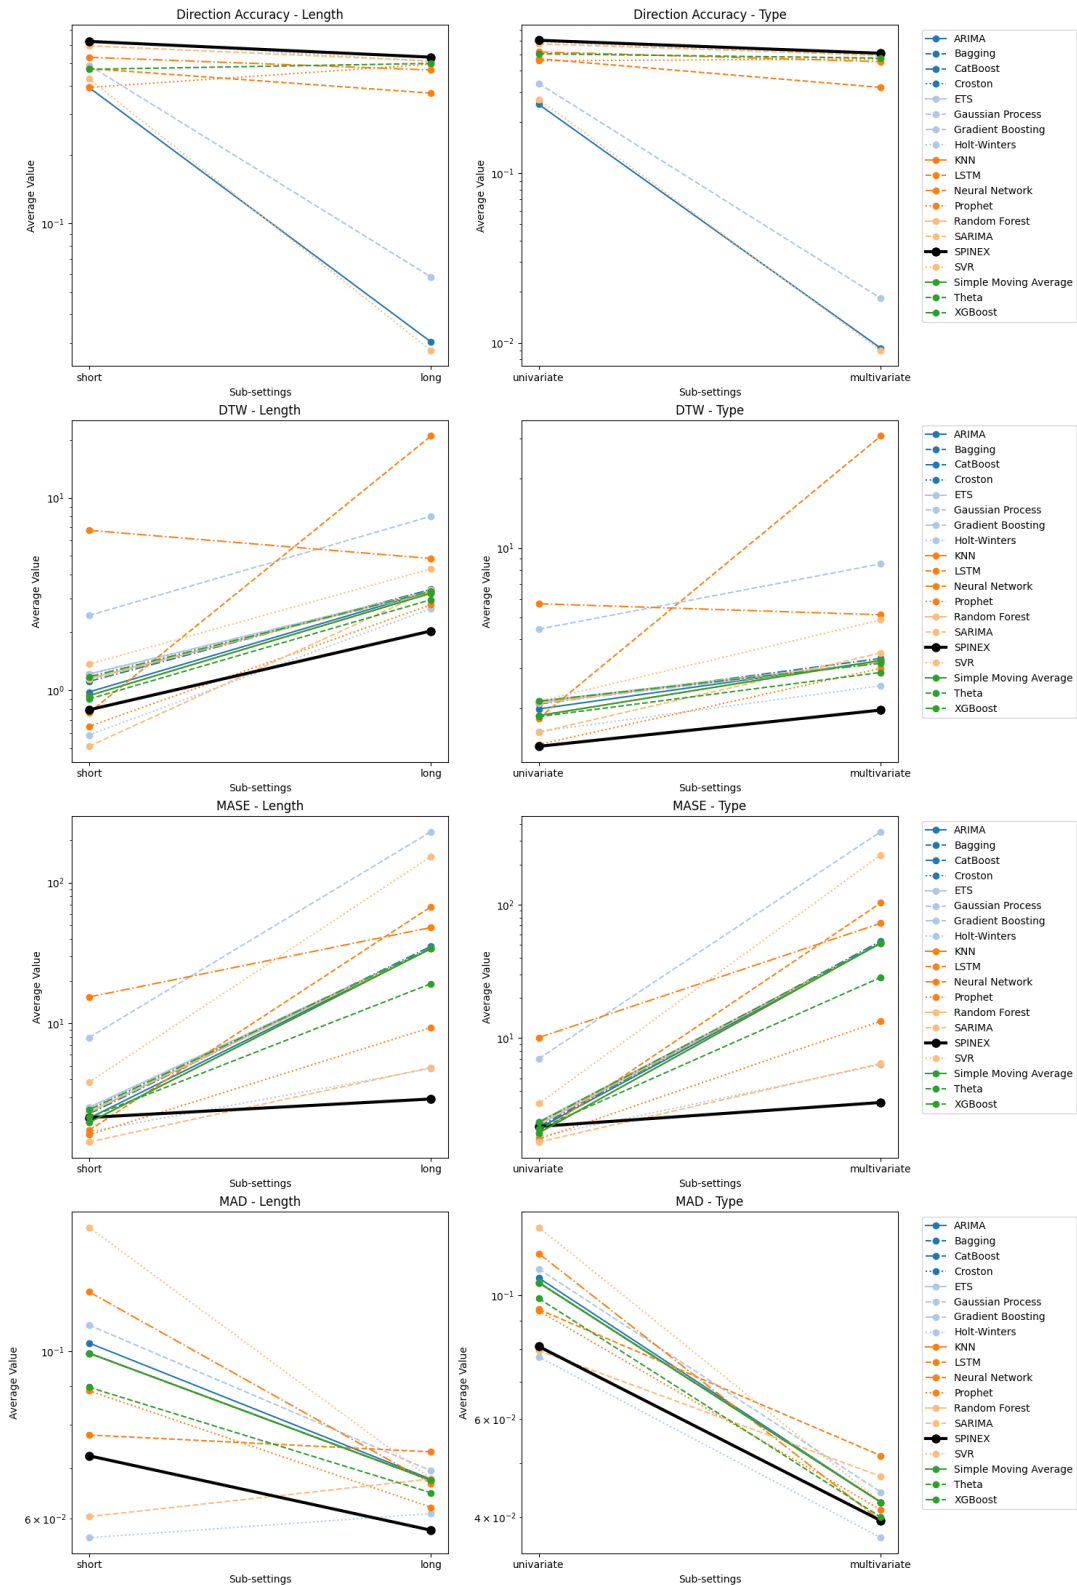
744 Figure 5 illustrates the performance of various algorithms across different settings and metrics.
 745 Each of these metrics is evaluated across two settings: dataset length (short [datasets of less than
 746 200 points] vs. long [datasets of more than 200 points]) and dataset type (univariate vs.
 747 multivariate). Overall, one can see the performance of SPINEX matches well with other algorithms
 748 and, in some cases, outperforms them.

749 For example, in the DA plots, SPINEX demonstrates a notable trend for short and long sequences
 750 and univariate and multivariate data. This suggests that SPINEX is proficient at predicting the
 751 correct direction. In the DTW metric, SPINEX exhibits a lower DTW value for short sequences,
 752 which indicates a higher similarity and better alignment of time series data than other algorithms.
 753 However, as the sequence length extends, SPINEX's DTW value increases, suggesting that its ability
 754 to maintain similarity diminishes slightly with longer sequences. This observation holds for the
 755 uni and multivariate datasets and other algorithms.

756 The MASE metric plots reveal that SPINEX performs consistently well across different lengths and
 757 types, with slightly better performance for short and univariate sequences. This trend continues for
 758 long sequences, where SPINEX remains competitive. It is worth noting that this algorithm maintains
 759 the lowest average MASE for long and multivariate sequences among the other algorithms.
 760 Finally, the MAD metric shows that SPINEX consistently achieves low values across both length
 761 and type dimensions.

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



This is a preprint draft. The published article can be found at: <https://doi.org/10.1016/j.cie.2024.110812>.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

763

Figure 5 Further analysis in terms of dataset length and type

764

Figure 6 presents a sample of a visual representation of two time series as predicted by SPINEX and other algorithms. These two datasets represent those that fall under short and long time series. In both cases, it is clear that the forecasts by SPINEX are in good agreement with the actual series.

765

766


Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



768
769
770

Fig. 6 Visualization of forecasting on Yearly Water Usage dataset (top) and Tata Global Forecasting dataset (bottom)

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

771 **4.3 Dataset examination and Pareto efficiency analysis**

772 Once the above analysis was completed, the same results were examined to identify the most
 773 reoccurring complex datasets that received the poorest performance and the most consistently
 774 ranked datasets with the best performance across all algorithms. The outcome of this analysis is
 775 listed in Table 7. This table shows the top 3 datasets in each category and synthetic and real
 776 datasets. Interestingly, all complex datasets are univariates, while those that fall under the
 777 consistent datasets have a mixture of both types.

778 **Table 7 Dataset examination**

Type	Complex Dataset	Characteristics	Occurrences	Consistent Dataset	Characteristics	Occurrences
Synthetic data	Dataset 27	t _{max} : 100, n _{points} : 5000	17	Dataset 106	t _{max} : 100, n _{points} : 50	9
	Dataset 96	t _{max} : 10, n _{points} : 5000	14	Dataset 115	t _{max} : 100, n _{points} : 50	9
	Dataset 99	t _{max} : 100, n _{points} : 5000	8	Dataset 91	t _{max} : 1, n _{points} : 50	9
Real data	Sunspots	2820/2	15	Yearly Water Usage	79/2	13
	Stock and PM2.5 Prediction	5650/10	14	Tata Global Forecasting	2100/8	10
	Indian Crime II	550/2	8	Jaipur	676/40	7

779

780 A Pareto analysis is performed on synthetic datasets to determine the best-performing time series
 781 algorithms based on the selected evaluation metrics (see Table 8). This analysis employs Pareto
 782 optimality to identify non-dominated solutions (i.e., those representing optimal trade-offs between
 783 different performance metrics: DA, DTW, MASE, and MAD). The concept of Pareto optimality
 784 ensures that the final set of recommended algorithms consists of truly superior options, each
 785 offering a distinct balance of strengths across various performance criteria. The process starts by
 786 normalizing all metrics to a 0-1 scale using min-max normalization. The normalized data is then
 787 grouped by algorithm and dataset to calculate mean values for each metric. Each algorithm further
 788 aggregates these grouped results to evaluate overall performance across all datasets. Then, a
 789 solution is Pareto optimal if no other solution is superior in all metrics simultaneously.

790 **Table 8 Pareto analysis**

Algorithm	Direction Accuracy	DTW	MASE	MAD	Pareto Efficient
SPINEX	0.602	1.956	45.676	0.075	TRUE
Holt-Winters	0.572	15.419	46.643	0.291	TRUE
Theta	0.550	2.451	82.481	0.100	TRUE
LSTM	0.502	3.756	115.265	0.090	TRUE
Prophet	0.546	2.058	58.910	0.095	FALSE
ARIMA	0.264	2.544	84.369	0.097	FALSE
Croston	0.000	2.602	87.962	0.101	FALSE
ETS	0.000	2.603	87.962	0.101	FALSE
Gradient Boosting	0.000	2.689	88.144	0.101	FALSE
Random Forest	0.000	2.698	88.205	0.101	FALSE
Bagging	0.000	2.698	88.205	0.101	FALSE
XGBoost	0.000	2.695	88.281	0.101	FALSE
Simple Moving Average	0.000	2.669	88.733	0.101	FALSE
KNN	0.000	2.669	88.733	0.101	FALSE

Please cite this paper as:


Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

CatBoost	0.000	2.666	89.402	0.101	FALSE
Gaussian Process	0.255	7.354	144.260	0.097	FALSE
SVR	0.166	4.262	157.608	0.115	FALSE
Neural Network	0.518	6.952	194.656	0.110	FALSE
SARIMA	0.578	15.353	2625.313	0.116	FALSE

791

792 4.4 Complexity analysis

793 We evaluate the complexity of the selected algorithms by analyzing the execution time data across
 794 across 50, 500, and 5000 samples (n). Then, we fit three types of models to this data: polynomial
 795 (linear in log-log scale), logarithmic, and exponential. For each algorithm, we compute the
 796 regression parameters and the R² values for these models as a means to measure the goodness of
 797 fit and the model with the highest R² value is considered the best fit, and its corresponding Big O
 798 notation is recorded. This analysis reveals that SPINEX demonstrates logarithmic complexity and
 799 hence indicates that its execution time scales efficiently with the logarithm of the input size,
 800 represented as O(log n), as seen in Fig. 7. It is worth noting that this algorithm was found to have
 801 only logarithmic complexity, while others had polynomial or exponential complexity (see Table
 802 9). However, and from a scalable perspective, this complexity of SPINEX arises from the
 803 integration of multiple similarity measures (as well as the embedded dynamic adjustments) within
 804 the algorithm. The overhead of these settings may pose challenges in practical implementations
 805 when this algorithm is used in large datasets. This can be thought of as one challenge that could
 806 be revisited and overcome in the near future.



807

808

Fig. 7 Outcome of complexity analysis

809 Table 9 Complexity analysis

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

Algorithm	Complexity type	R ²	Big O notation
Prophet	exp	0.84	O(e ⁿ)
ETS	exp	0.94	O(e ⁿ)
XGBoost	exp	0.92	O(e ⁿ)
Theta	exp	0.85	O(e ⁿ)
Simple Moving Average	exp	0.98	O(e ⁿ)
Holt-Winters	exp	0.99	O(e ⁿ)
Croston	exp	0.89	O(e ⁿ)
Gradient Boosting	exp	1.00	O(e ⁿ)
SPINEX	log	0.98	O(log n)
CatBoost	poly	0.62*	O(n ^{0.12})
KNN	poly	0.54*	O(n ^{0.25})
ARIMA	poly	0.66*	O(n ^{0.42})
Bagging	poly	0.83	O(n ^{0.46})
Random Forest	poly	0.84	O(n ^{0.49})
SARIMA	poly	0.93	O(n ^{0.71})
Neural Network	poly	0.85	O(n ^{0.86})
LSTM	poly	1.00	O(n ^{0.93})
SVR	poly	0.72	O(n ^{1.36})
Gaussian Process	poly	0.97	O(n ^{1.72})

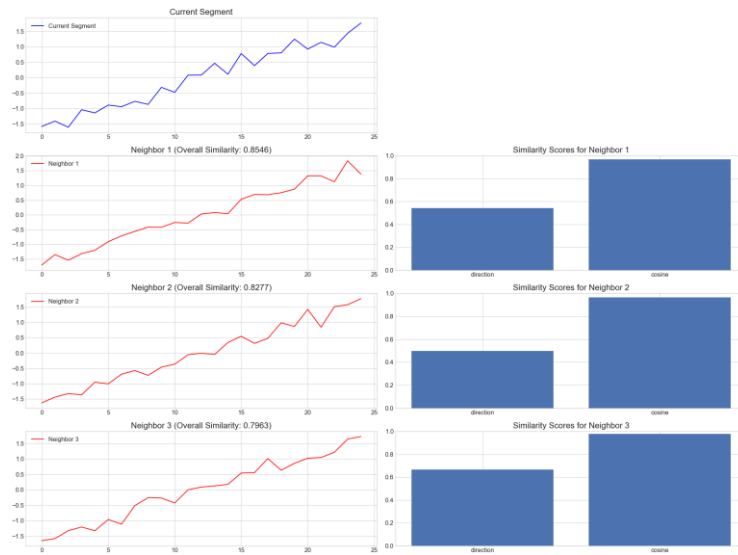
810 *note the low value.

811 4.5 Explainability analysis

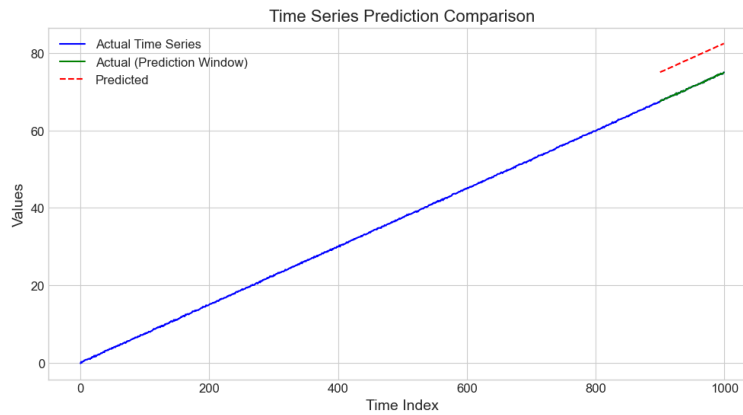
812 To showcase the explainability capabilities of SPINEX, two synthetic datasets (No. 1 and No. 11)
 813 are provided herein, as taken from two different datasets. Figure 8 shows the predicted segment
 814 and three of its neighbors. The same plot also visually represents the neighbors and their overall
 815 similarity as compared to the segment at hand. For example, this figure represents the top three
 816 selected time segments that align the most with the current segment being investigated by SPINEX.
 817 As one can see, the identified segments (i.e., neighbors) align well with the current segment –
 818 which further showcases the applicability of SPINEX. The companion similarity score plot also,
 819 visually, presents the importance of the similarity metrics selected by the user in each case and
 820 notes how each similarity measure relates to the identified segments. Finally, this plot also shows
 821 the individual scores of the similarity measures used to identify them.

Please cite this paper as:

Naser A., Naser M.Z. (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



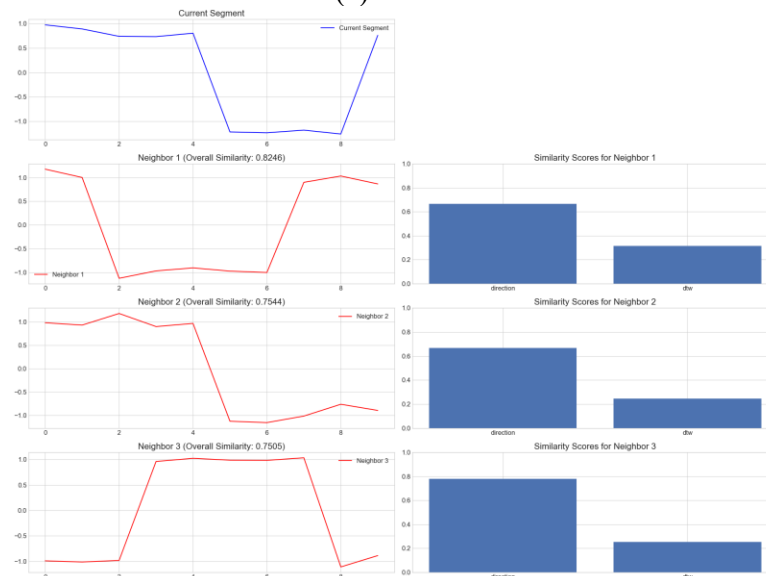
822



823

824

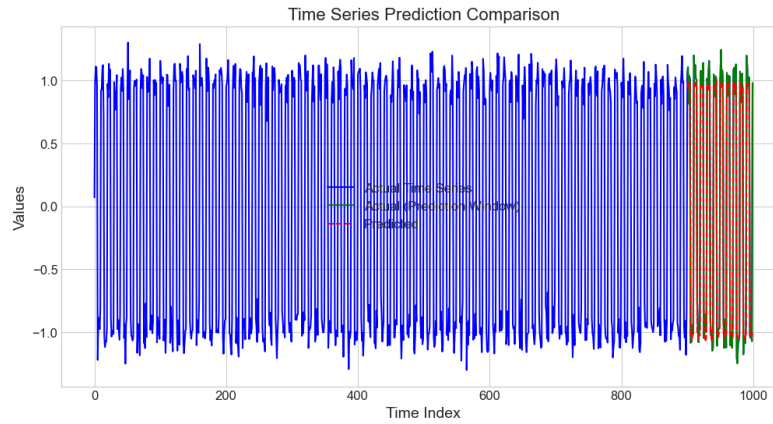
(a) Dataset 1



825

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.



(b) Dataset 11

Fig. 8 Example of explainability capability of SPINEX

826
827
828

5.0 Conclusions

829 This paper introduces a novel member of the SPINEX (Similarity-based Predictions with Explainable
830 Neighbors Exploration) family. This new algorithm enhances time series analysis performance by
831 leveraging the concept of similarity, higher-order temporal interactions across multiple time scales,
832 and explainability. The effectiveness of the proposed SPINEX variant was evaluated through a
833 comprehensive benchmarking study involving 18 time series forecasting algorithms across 49
834 datasets. Our findings from our experiments indicate that SPINEX consistently ranks within the top-
835 5 best-performing algorithms, showcasing its Pareto efficacy in time series forecasting and pattern
836 recognition while maintaining moderate computational complexity on the order of $O(\log n)$.
837 Moreover, the algorithm's explainability features, Pareto efficiency, and medium complexity are
838 demonstrated through detailed visualizations to enhance the prediction and decision-making
839 process.
840

841 Despite the noted positive findings, there are several promising avenues for future research to
842 further enhance SPINEX's capabilities and applicability. To start with, this algorithm can be
843 extended to handle multivariate time series, which could broaden its use cases. Second, while
844 SPINEX dynamically adjusts its internal parameters, further exploration of adaptive mechanisms,
845 such as reinforcement learning or metaheuristics, could dynamically optimize hyperparameters
846 during runtime. It is also worth exploring options to further enhance the algorithmic scalability for
847 large datasets by using sparse similarity matrices or approximate methods for computationally
848 expensive metrics. We are hopeful to be able to improve the proposed algorithm in the near future.
849 In the meantime, we also invite interested readers to spearhead the aforementioned items, as w

Data Availability

850 Some or all data, models, or code that support the findings of this study are available from the
851 corresponding author upon reasonable request.
852

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

853 SPINEX can be accessed from **[to be added]**.

854 **Conflict of Interest**

855 The authors declare no conflict of interest.

856 **References**

- 857 [1] Hamilton, Time series analysis / James D. Hamilton., Time Ser. Anal. (1993).
- 858 [2] J.G. De Gooijer, R.J. Hyndman, 25 years of time series forecasting, Int. J. Forecast. (2006).
859 <https://doi.org/10.1016/j.ijforecast.2006.01.001>.
- 860 [3] M.D. Morse, J.M. Patel, An efficient and accurate method for evaluating time series
861 similarity, in: Proc. ACM SIGMOD Int. Conf. Manag. Data, 2007.
862 <https://doi.org/10.1145/1247480.1247544>.
- 863 [4] S. Aghabozorgi, A. Seyed Shirخورshidi, T. Ying Wah, Time-series clustering - A decade
864 review, Inf. Syst. (2015). <https://doi.org/10.1016/j.is.2015.04.007>.
- 865 [5] S. Schmidl, P. Wenig, T. Papenbrock, Anomaly Detection in Time Series: A
866 Comprehensive Evaluation, in: Proc. VLDB Endow., 2022.
867 <https://doi.org/10.14778/3538598.3538602>.
- 868 [6] S. Lhermitte, J. Verbesselt, W.W. Verstraeten, P. Coppin, A comparison of time series
869 similarity measures for classification and change detection of ecosystem dynamics, Remote
870 Sens. Environ. (2011). <https://doi.org/10.1016/j.rse.2011.06.020>.
- 871 [7] X. Wang, A. Mueen, H. Ding, G. Trajcevski, P. Scheuermann, E. Keogh, Experimental
872 comparison of representation methods and distance measures for time series data, Data Min.
873 Knowl. Discov. (2013). <https://doi.org/10.1007/s10618-012-0250-5>.
- 874 [8] A. Abanda, U. Mori, J.A. Lozano, A review on distance based time series classification,
875 Data Min. Knowl. Discov. (2019). <https://doi.org/10.1007/s10618-018-0596-4>.
- 876 [9] H. Sakoe, S. Chiba, Dynamic Programming Algorithm Optimization for Spoken Word
877 Recognition, IEEE Trans. Acoust. (1978). <https://doi.org/10.1109/TASSP.1978.1163055>.
- 878 [10] Y. Permanasari, E.H. Harahap, E.P. Ali, Speech recognition using Dynamic Time Warping
879 (DTW), in: J. Phys. Conf. Ser., 2019. <https://doi.org/10.1088/1742-6596/1366/1/012091>.
- 880 [11] E. Kostadinova, V. Boeva, L. Boneva, E. Tsiporkova, An integrative DTW-based
881 imputation method for gene expression time series data, in: IS'2012 - 2012 6th IEEE Int.
882 Conf. Intell. Syst. Proc., 2012. <https://doi.org/10.1109/IS.2012.6335145>.
- 883 [12] L. Bergroth, H. Hakonen, T. Raita, A survey of longest common subsequence algorithms,
884 in: Proc. - 7th Int. Symp. String Process. Inf. Retrieval, SPIRE 2000, 2000.
885 <https://doi.org/10.1109/SPIRE.2000.878178>.
- 886 [13] B.D. Fulcher, Feature-Based Time-Series Analysis, in: Featur. Eng. Mach. Learn. Data

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 887 Anal., 2018. <https://doi.org/10.1201/9781315181080-4>.
- 888 [14] K. Mirylenka, M. Dallachiesa, T. Palpanas, Data series similarity using correlation-aware
889 measures, in: ACM Int. Conf. Proceeding Ser., 2017.
890 <https://doi.org/10.1145/3085504.3085515>.
- 891 [15] J. Li, C. Xu, T. Zhang, Similarity Measure of Time Series Based on Siamese and Sequential
892 Neural Networks, in: Chinese Control Conf. CCC, 2020.
893 <https://doi.org/10.23919/CCC50068.2020.9189261>.
- 894 [16] Z. Karevan, J.A.K. Suykens, Transductive LSTM for time-series prediction: An application
895 to weather forecasting, Neural Networks. (2020).
896 <https://doi.org/10.1016/j.neunet.2019.12.030>.
- 897 [17] A. Theissler, F. Spinnato, U. Schlegel, R. Guidotti, Explainable AI for Time Series
898 Classification: A Review, Taxonomy and Research Directions, IEEE Access. (2022).
899 <https://doi.org/10.1109/ACCESS.2022.3207765>.
- 900 [18] A. Nielsen, Practical Time Series: Prediction with Statistics & Machine Learning., 2019.
- 901 [19] A. Babii, E. Ghysels, J. Striaukas, Machine Learning Time Series Regressions With an
902 Application to Nowcasting, J. Bus. Econ. Stat. (2022).
903 <https://doi.org/10.1080/07350015.2021.1899933>.
- 904 [20] Y. Ensafi, S.H. Amin, G. Zhang, B. Shah, Time-series forecasting of seasonal items sales
905 using machine learning – A comparative analysis, Int. J. Inf. Manag. Data Insights. (2022).
906 <https://doi.org/10.1016/j.jjime.2022.100058>.
- 907 [21] G.M. Box, G. E. P., & Jenkins, Time series analysis: forecasting and control. San Francisco,
908 CA: Holden-Day., [University Wisconsin. Madison. WI Univ. Of Lancaster, England].
909 (1976).
- 910 [22] S. Makridakis, M. Hibon, ARMA models and the Box-Jenkins methodology, J. Forecast.
911 (1997). [https://doi.org/10.1002/\(SICI\)1099-131X\(199705\)16:3<147::AID-](https://doi.org/10.1002/(SICI)1099-131X(199705)16:3<147::AID-)
912 [FOR652>3.0.CO;2-X](https://doi.org/10.1002/(SICI)1099-131X(199705)16:3<147::AID-FOR652>3.0.CO;2-X).
- 913 [23] J. Fattah, L. Ezzine, Z. Aman, H. El Moussami, A. Lachhab, Forecasting of demand using
914 ARIMA model, Int. J. Eng. Bus. Manag. (2018).
915 <https://doi.org/10.1177/1847979018808673>.
- 916 [24] J.D. Croston, Forecasting and Stock Control for Intermittent Demands, Oper. Res. Q.
917 (1972). <https://doi.org/10.2307/3007885>.
- 918 [25] A. Segerstedt, E. Levén, A Study of Different Croston-Like Forecasting Methods, Oper.
919 Supply Chain Manag. (2023). <https://doi.org/10.31387/oscm0540400>.
- 920 [26] J.E. Boylan, A.A. Syntetos, The accuracy of a Modified Croston procedure, Int. J. Prod.
921 Econ. (2007). <https://doi.org/10.1016/j.ijpe.2006.10.005>.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 922 [27] R. Teunter, B. Sani, On the bias of Croston's forecasting method, *Eur. J. Oper. Res.* (2009).
923 <https://doi.org/10.1016/j.ejor.2007.12.001>.
- 924 [28] O.L. Davies, R.G. Brown, Statistical Forecasting for Inventory Control., *J. R. Stat. Soc. Ser. A.* (1960). <https://doi.org/10.2307/2342487>.
925
- 926 [29] C.C. Holt, Forecasting seasonals and trends by exponentially weighted moving averages,
927 *Int. J. Forecast.* (2004). <https://doi.org/10.1016/j.ijforecast.2003.09.015>.
- 928 [30] P.R. Winters, Forecasting Sales by Exponentially Weighted Moving Averages Author(s):
929 Peter R. Winters Source, *Manage. Sci.* (1960).
- 930 [31] C. Chatfield, M. Yar, Holt-Winters Forecasting: Some Practical Issues, *Stat.* (1988).
931 <https://doi.org/10.2307/2348687>.
- 932 [32] S. Hochreiter, J. Schmidhuber, Long Short-Term Memory, *Neural Comput.* (1997).
933 <https://doi.org/10.1162/neco.1997.9.8.1735>.
- 934 [33] C. Han, X. Fu, Challenge and Opportunity: Deep Learning-Based Stock Price Prediction by
935 Using Bi-Directional LSTM Model, *Front. Business, Econ. Manag.* (2023).
936 <https://doi.org/10.54097/fbem.v8i2.6616>.
- 937 [34] S.J. Taylor, B. Letham, Forecasting at Scale, *Am. Stat.* (2018).
938 <https://doi.org/10.1080/00031305.2017.1380080>.
- 939 [35] M. Daraghmeh, A. Agarwal, R. Manzano, M. Zaman, Time Series Forecasting using
940 Facebook Prophet for Cloud Resource Management, in: 2021 IEEE Int. Conf. Commun.
941 Work. ICC Work. 2021 - Proc., 2021.
942 <https://doi.org/10.1109/ICCWorkshops50388.2021.9473607>.
- 943 [36] I. Svetunkov, F. Petropoulos, Old dog, new tricks: a modelling view of simple moving
944 averages, *Int. J. Prod. Res.* (2018). <https://doi.org/10.1080/00207543.2017.1380326>.
- 945 [37] C. Chiarella, X. He, C.H. Hommes, A Dynamic Analysis of Moving Average Rules, *SSRN*
946 *Electron. J.* (2011). <https://doi.org/10.2139/ssrn.742386>.
- 947 [38] V. Assimakopoulos, K. Nikolopoulos, The theta model: A decomposition approach to
948 forecasting, *Int. J. Forecast.* (2000). [https://doi.org/10.1016/S0169-2070\(00\)00066-2](https://doi.org/10.1016/S0169-2070(00)00066-2).
- 949 [39] E. Spiliotis, V. Assimakopoulos, S. Makridakis, Generalizing the Theta method for
950 automatic forecasting, *Eur. J. Oper. Res.* (2020). <https://doi.org/10.1016/j.ejor.2020.01.007>.
- 951 [40] J.H. Stock, M.W. Watson, Vector autoregressions, *J. Econ. Perspect.* (2001).
952 <https://doi.org/10.1257/jep.15.4.101>.
- 953 [41] C.J. Lu, T.S. Lee, C.C. Chiu, Financial time series forecasting using independent component
954 analysis and support vector regression, *Decis. Support Syst.* (2009).
955 <https://doi.org/10.1016/j.dss.2009.02.001>.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 956 [42] F. Canova, 53 2 Vector Autoregressive Models: Specification, Estimation, Inference, and
957 Forecasting, *Handb. Appl. Econom.* (1999).
- 958 [43] C.K.I. Williams, C.E. Rasmussen, Gaussian Processes for Regression, *Adv. Neural Inf.*
959 *Process. Syst.* 8 (1995).
- 960 [44] S. Aigrain, D. Foreman-Mackey, Gaussian Process Regression for Astronomical Time
961 Series, *Annu. Rev. Astron. Astrophys.* (2023). [https://doi.org/10.1146/annurev-astro-](https://doi.org/10.1146/annurev-astro-052920-103508)
962 [052920-103508](https://doi.org/10.1146/annurev-astro-052920-103508).
- 963 [45] J.P. Cunningham, Z. Ghahramani, C.E. Rasmussen, Gaussian Processes for time-marked
964 time-series data, in: *J. Mach. Learn. Res.*, 2012.
- 965 [46] L.P. Swiler, M. Gulian, A.L. Frankel, C. Safta, J.D. Jakeman, A SURVEY OF
966 CONSTRAINED GAUSSIAN PROCESS REGRESSION: APPROACHES AND
967 IMPLEMENTATION CHALLENGES, *J. Mach. Learn. Model. Comput.* (2020).
968 <https://doi.org/10.1615/jmachlearnmodelcomput.2020035155>.
- 969 [47] L. Prokhorenkova, G. Gusev, A. Vorobev, A.V. Dorogush, A. Gulin, Catboost: Unbiased
970 boosting with categorical features, in: *Adv. Neural Inf. Process. Syst.*, 2018.
- 971 [48] J.T. Hancock, T.M. Khoshgoftaar, CatBoost for big data: an interdisciplinary review, *J. Big*
972 *Data.* (2020). <https://doi.org/10.1186/s40537-020-00369-8>.
- 973 [49] S.R. Karingula, N. Ramanan, R. Tahmasbi, M. Amjadi, D. Jung, R. Si, C. Thimmisetty,
974 L.F. Polania, M. Sayer, J. Taylor, C.N. Coelho, Boosted Embeddings for Time-Series
975 Forecasting, in: *Lect. Notes Comput. Sci. (Including Subser. Lect. Notes Artif. Intell. Lect.*
976 *Notes Bioinformatics)*, 2022. https://doi.org/10.1007/978-3-030-95470-3_1.
- 977 [50] Y.H. Lee, C.P. Wei, T.H. Cheng, C.T. Yang, Nearest-neighbor-based approach to time-
978 series classification, *Decis. Support Syst.* (2012). <https://doi.org/10.1016/j.dss.2011.12.014>.
- 979 [51] G. Lin, A. Lin, J. Cao, Multidimensional KNN algorithm based on EEMD and complexity
980 measures in financial time series forecasting, *Expert Syst. Appl.* (2021).
981 <https://doi.org/10.1016/j.eswa.2020.114443>.
- 982 [52] G.P. Zhang, D.M. Kline, Quarterly time-series forecasting with neural networks, *IEEE*
983 *Trans. Neural Networks.* (2007). <https://doi.org/10.1109/TNN.2007.896859>.
- 984 [53] S.F. Crone, M. Hibon, K. Nikolopoulos, Advances in forecasting with neural networks?
985 Empirical evidence from the NN3 competition on time series prediction, *Int. J. Forecast.*
986 (2011). <https://doi.org/10.1016/j.ijforecast.2011.04.001>.
- 987 [54] L. Breiman, Random Forests, *Mach. Learn.* 45 (2001) 5–32.
988 <https://doi.org/10.1023/A:1010933404324>.
- 989 [55] T. Chen, C. Guestrin, XGBoost: A scalable tree boosting system, in: *Proc. ACM SIGKDD*
990 *Int. Conf. Knowl. Discov. Data Min.*, 2016. <https://doi.org/10.1145/2939672.2939785>.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 991 [56] J. Luo, Z. Zhang, Y. Fu, F. Rao, Time series prediction of COVID-19 transmission in
992 America using LSTM and XGBoost algorithms, *Results Phys.* (2021).
993 <https://doi.org/10.1016/j.rinp.2021.104462>.
- 994 [57] X. Qiu, L. Zhang, P. Nagaratnam Suganthan, G.A.J. Amaratunga, Oblique random forest
995 ensemble via Least Square Estimation for time series forecasting, *Inf. Sci. (Ny)*. (2017).
996 <https://doi.org/10.1016/j.ins.2017.08.060>.
- 997 [58] C. Cortes, V. Vapnik, Support-Vector Networks, *Mach. Learn.* (1995).
998 <https://doi.org/10.1023/A:1022627411411>.
- 999 [59] J. Shawe-Taylor, P.L. Bartlett, R.C. Williamson, M. Anthony, Structural risk minimization
1000 over data-dependent hierarchies, *IEEE Trans. Inf. Theory.* (1998).
1001 <https://doi.org/10.1109/18.705570>.
- 1002 [60] N. Sapankevych, R. Sankar, Time series prediction using support vector machines: A
1003 survey, *IEEE Comput. Intell. Mag.* (2009). <https://doi.org/10.1109/MCI.2009.932254>.
- 1004 [61] M.Z. Naser, Amir, H. Alavi, A.H. Alavi, Amir, H. Alavi, Error Metrics and Performance
1005 Fitness Indicators for Artificial Intelligence and Machine Learning in Engineering and
1006 Sciences, *Archit. Struct. Constr.* 1 (2021) 1–19.
1007 <https://doi.org/https://doi.org/10.1007/s44150-021-00015-8>.
- 1008 [62] D.S.K. Karunasingha, Root mean square error or mean absolute error? Use their ratio as
1009 well, *Inf. Sci. (Ny)*. (2022). <https://doi.org/10.1016/j.ins.2021.11.036>.
- 1010 [63] C. Chen, J. Twycross, J.M. Garibaldi, A new accuracy measure based on bounded relative
1011 error for time series forecasting, *PLoS One.* (2017).
1012 <https://doi.org/10.1371/journal.pone.0174202>.
- 1013 [64] R.J. Hyndman, A.B. Koehler, Another look at measures of forecast accuracy, *Int. J.*
1014 *Forecast.* (2006). <https://doi.org/10.1016/j.ijforecast.2006.03.001>.
- 1015 [65] P.H. Franses, A note on the Mean Absolute Scaled Error, *Int. J. Forecast.* (2016).
1016 <https://doi.org/10.1016/j.ijforecast.2015.03.008>.
- 1017 [66] R.J. Kate, Using dynamic time warping distances as features for improved time series
1018 classification, *Data Min. Knowl. Discov.* (2016). [https://doi.org/10.1007/s10618-015-0418-](https://doi.org/10.1007/s10618-015-0418-x)
1019 [x](https://doi.org/10.1007/s10618-015-0418-x).
- 1020 [67] U. Khair, H. Fahmi, S. Al Hakim, R. Rahim, Forecasting Error Calculation with Mean
1021 Absolute Deviation and Mean Absolute Percentage Error, in: *J. Phys. Conf. Ser.*, 2017.
1022 <https://doi.org/10.1088/1742-6596/930/1/012002>.
- 1023 [68] J. Brownlee, *Airline Passengers*, (2017).
- 1024 [69] J. Brownlee, *Sunspots*, (2017).
- 1025 [70] J. Brownlee, *Daily Female Births*, (2017).

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

- 1026 [71] J. Brownlee, Yearly Water Usage, (2017).
1027 [72] J. Brownlee, Daily Minimum Temperatures, (2017).
1028 [73] J. Brownlee, Monthly Car Sales, (2017).
1029 [74] J. Brownlee, Shampoo Sales Data, (2017).
1030 [75] J. Brownlee, Temperature Data, (2017).
1031 [76] J. Brownlee, Monthly Writing Paper Sales, (2017).
1032 [77] J. Brownlee, Monthly Champagne Sales, (2017).
1033 [78] J. Brownlee, Monthly Robberies, (2017).
1034 [79] E. Gao, Electric Production, (2018).
1035 [80] S. Subikshaa, Web Traffic Dataset, (2019).
1036 [81] I. Kim, Stock and PM2.5 Prediction, (2020).
1037 [82] M. Jennings, Tata Global Forecasting, (2020).
1038 [83] N. Volfango, International Airline Passengers, (2020).
1039 [84] J. Brownlee, Pollution Dataset, (2017).
1040 [85] ReadyTensor, Daily Stock Prices, (2021).
1041 [86] Z. Haoyi, ETT-small, (2020).
1042 [87] R. Dey, Jaipur Final Clean Data, (2020).
1043 [88] B. Ciranni, Aprocessed, (2021).
1044 [89] K. Krishnan, Insurance, (2019).
1045 [90] S. Chake, Indian Crime Data Analysis Forecasting, (2020).
1046 [91] A.E. Ezugwu, A.M. Ikotun, O.O. Oyelade, L. Abualigah, J.O. Agushaka, C.I. Eke, A.A.
1047 Akinyelu, A comprehensive survey of clustering algorithms: State-of-the-art machine
1048 learning applications, taxonomy, challenges, and future research prospects, *Eng. Appl.*
1049 *Artif. Intell.* (2022). <https://doi.org/10.1016/j.engappai.2022.104743>.
1050 [92] A. Shahraki, A. Taherkordi, O. Haugen, F. Eliassen, A Survey and Future Directions on
1051 Clustering: From WSNs to IoT and Modern Networking Paradigms, *IEEE Trans. Netw.*
1052 *Serv. Manag.* (2021). <https://doi.org/10.1109/TNSM.2020.3035315>.
1053 [93] Y. Perez-Riverol, J.A. Vizcaíno, J. Griss, Future Prospects of Spectral Clustering
1054 Approaches in Proteomics, *Proteomics*. (2018). <https://doi.org/10.1002/pmic.201700454>.
1055 [94] W. Xiao, J. Hu, A Survey of Parallel Clustering Algorithms Based on Spark, *Sci. Program*.

This is a preprint draft. The published article can be found at: <https://doi.org/10.1016/j.cie.2024.110812>.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

1056 (2020). <https://doi.org/10.1155/2020/8884926>.

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

1081

1082

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

1083 **Appendix**

1084 The complete class of SPINEX is shown below. SPINEX can be installed via: **pip install**
1085 **spinex_timeseries**

1086 **@jit(nopython=True)**

1087 **def numba_dtw(x, y):**

1088 n, m = len(x), len(y)

1089 dtw_matrix = np.zeros((n+1, m+1))

1090 for i in range(1, n+1):

1091 for j in range(1, m+1):

1092 cost = abs(x[i-1] - y[j-1])

1093 dtw_matrix[i, j] = cost + min(dtw_matrix[i-1, j], dtw_matrix[i, j-1], dtw_matrix[i-1, j-1])

1094 return dtw_matrix[n, m]

1095

1096 **@jit(nopython=True)**

1097 **def numba_dtw_similarity(X):**

1098 n = X.shape[0]

1099 sim_matrix = np.zeros((n, n))

1100 for i in range(n):

1101 for j in range(i, n):

1102 dist = numba_dtw(X[i], X[j])

1103 sim_matrix[i, j] = sim_matrix[j, i] = 1 / (1 + dist)

1104 return sim_matrix

1105

1106 **@jit(nopython=True)**

1107 **def numba_sample_entropy(x, m=2, r=0.2):**

1108 n = len(x)

1109 B = 0.0

1110 A = 0.0

1111 for i in range(n - m):

1112 for j in range(i + 1, n - m):

1113 matches = 0

1114 for k in range(m):

1115 if abs(x[i+k] - x[j+k]) <= r:

1116 matches += 1

1117 else:

1118 break

1119 if matches == m:

1120 B += 1

1121 if abs(x[i+m] - x[j+m]) <= r:

1122 A += 1

1123 return -np.log((A + 1e-10) / (B + 1e-10))

1124

1125 **def direction_accuracy(segment1, segment2):**

1126 direction1 = np.sign(np.diff(segment1))

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1127     direction2 = np.sign(np.diff(segment2))
1128     return np.mean(direction1 == direction2)
1129
1130 class SPINEX_Timeseries:
1131     def __init__(self, data, window_size=None, forecast_horizon=1, similarity_methods=None,
1132                 dynamic_window=True, multi_level=True, dynamic_threshold=True):
1133         self.data = np.array(data)
1134         if window_size is None:
1135             self.window_size = max(10, len(data) // 10)
1136         else:
1137             self.window_size = min(window_size, len(data) // 2)
1138         self.forecast_horizon = min(forecast_horizon, len(data) // 10)
1139         self.forecast_horizon = forecast_horizon
1140         self.similarity_methods = similarity_methods if similarity_methods else ['cosine', 'euclidean', 'dtw']
1141         self.similarity_cache = {}
1142         self.dynamic_window = dynamic_window
1143         self.multi_level = multi_level
1144         self.dynamic_threshold = dynamic_threshold
1145         self.segments_cache = {}
1146         self.recent_errors = []
1147         self.recent_similarity_scores = []
1148         if self.dynamic_window:
1149             self.window_size = self.adaptive_window_size()
1150
1151     @staticmethod
1152     def hash_array(arr):
1153         return hashlib.md5(arr.data.tobytes()).hexdigest()
1154
1155     @lru_cache(maxsize=128)
1156     def get_similarity_matrix(self, method, segments_hash):
1157         if (segments_hash, method) in self.similarity_cache:
1158             return self.similarity_cache[(segments_hash, method)]
1159         segments = self.segments_cache[segments_hash]
1160         if method == 'cosine':
1161             similarity_matrix = self.cosine_similarity(segments)
1162         elif method == 'correlation':
1163             similarity_matrix = self.correlation_similarity(segments)
1164         elif method == 'euclidean':
1165             similarity_matrix = self.euclidean_similarity(segments)
1166         elif method == 'spearman':
1167             similarity_matrix = self.spearman_similarity(segments)
1168         elif method == 'dtw':
1169             similarity_matrix = numba_dtw_similarity(segments)
1170         elif method == 'direction':
1171             similarity_matrix = self.direction_similarity(segments)
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1172     else:
1173         raise ValueError(f"Invalid similarity method: {method}")
1174     self.similarity_cache[(segments_hash, method)] = similarity_matrix
1175     return similarity_matrix
1176
1177     @staticmethod
1178     def cosine_similarity(X):
1179         norm = np.linalg.norm(X, axis=1)
1180         return np.dot(X, X.T) / np.outer(norm, norm)
1181
1182     @staticmethod
1183     def correlation_similarity(X):
1184         return np.corrcoef(X)
1185
1186     @staticmethod
1187     def euclidean_similarity(X):
1188         sq_dists = cdist(X, X, metric='euclidean')**2
1189         return 1 / (1 + np.sqrt(sq_dists))
1190
1191     @staticmethod
1192     def spearman_similarity(X):
1193         return spearmanr(X.T)[0]
1194
1195     def adjust_dynamic_parameters(self):
1196         MIN_WINDOW_SIZE = 10
1197         MAX_WINDOW_SIZE = len(self.data) // 2
1198         BASELINE_WINDOW_SIZE = max(MIN_WINDOW_SIZE, len(self.data) // 10)
1199         if len(self.data) > BASELINE_WINDOW_SIZE:
1200             volatility = np.std(self.data[-BASELINE_WINDOW_SIZE:])
1201         else:
1202             volatility = np.std(self.data)
1203         scale_factor = np.clip(volatility, 0.1, 1.0) # Limiting scale factor to avoid extreme values
1204         self.window_size = int(MAX_WINDOW_SIZE / scale_factor)
1205         self.window_size = max(MIN_WINDOW_SIZE, min(self.window_size, MAX_WINDOW_SIZE))
1206         if hasattr(self, 'recent_errors'):
1207             recent_error_mean = np.mean(self.recent_errors)
1208             recent_error_std = np.std(self.recent_errors)
1209             threshold_adjustment = recent_error_mean + recent_error_std
1210         else:
1211             threshold_adjustment = 0
1212         if hasattr(self, 'recent_similarity_scores') and self.recent_similarity_scores:
1213             mean_sim = np.mean(self.recent_similarity_scores)
1214             std_sim = np.std(self.recent_similarity_scores)
1215             self.threshold = mean_sim + std_sim + threshold_adjustment
1216         else:
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1217     self.threshold = 0.5 # Default threshold if no recent similarities are recorded
1218     print(f"Adjusted Window Size: {self.window_size}, Threshold: {self.threshold}")
1219
1220     def get_dynamic_threshold(self, similarities):
1221         if self.dynamic_threshold:
1222             mean_sim = np.mean(similarities)
1223             std_sim = np.std(similarities)
1224             base_threshold = mean_sim + std_sim
1225             if len(similarities[similarities > base_threshold]) < 5:
1226                 # If less than 5 indices are above threshold, reduce it to include more indices
1227                 adjusted_threshold = np.percentile(similarities, 90) # Adjusting percentile upward
1228             else:
1229                 adjusted_threshold = base_threshold
1230             print(f"Dynamic Threshold Adjusted: {adjusted_threshold}")
1231             return adjusted_threshold
1232         else:
1233             return np.percentile(similarities, 95)
1234
1235     def adjusted_dtw_similarity(self, X):
1236         dtw_scores = numba_dtw_similarity(X)
1237         adjusted_scores = 1 / (1 + np.sqrt(dtw_scores)) # Squaring DTW scores for more lenience
1238         return adjusted_scores
1239
1240     def plot_prediction(self):
1241         predicted_values = self.predict()
1242         if predicted_values.size > 0:
1243             prediction_start_index = len(self.data) - self.forecast_horizon
1244             plt.figure(figsize=(12, 6))
1245             plt.plot(self.data, label='Actual Time Series', color='blue')
1246             plt.plot(np.arange(prediction_start_index, len(self.data)),
1247                    self.data[prediction_start_index:], label='Actual (Prediction Window)', color='green')
1248             plt.plot(np.arange(prediction_start_index, len(self.data)),
1249                    predicted_values, label='Predicted', color='red', linestyle='--')
1250             plt.title('Time Series Prediction Comparison')
1251             plt.xlabel('Time Index')
1252             plt.ylabel('Values')
1253             plt.legend()
1254             plt.show()
1255         else:
1256             print("No valid predictions could be made.")
1257
1258     @lru_cache(maxsize=32)
1259     def extract_segments(self, window_size=None):
1260         if window_size is None:
1261             window_size = self.adaptive_window_size()
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1262 data_length = len(self.data)
1263 if data_length < window_size:
1264     print(f"Data length ({data_length}) is less than window size ({window_size}). Adjusting window size.")
1265     window_size = data_length // 2 # Use half of data length as window size
1266 n = data_length - window_size + 1
1267 if n <= 1:
1268     return np.array([self.data[-window_size:]])
1269 segments = np.lib.stride_tricks.sliding_window_view(self.data, window_size)
1270 segment_means = np.mean(segments, axis=1)
1271 segment_stds = np.std(segments, axis=1)
1272 normalized_segments = (segments - segment_means[:, np.newaxis]) / (segment_stds[:, np.newaxis] + 1e-8)
1273 return normalized_segments
1274 def find_similar_segments(self):
1275     window_sizes = [self.window_size]
1276     if self.multi_level:
1277         window_sizes = [max(2, self.window_size // 2)] + window_sizes + [min(len(self.data) // 4, self.window_size *
1278 2)]
1279     all_similarities = []
1280     for w_size in window_sizes:
1281         segments = self.extract_segments(w_size)
1282         if len(segments) < 2:
1283             print(f"Not enough segments for window size {w_size}, skipping.")
1284             continue
1285         segments_hash = self.hash_array(segments)
1286         self.segments_cache[segments_hash] = segments
1287         method_similarities = []
1288         for method in self.similarity_methods:
1289             if method == 'dtw' and len(segments) > 500:
1290                 print(f"DTW skipped for large dataset with {len(segments)} segments.")
1291                 continue
1292             try:
1293                 sim_matrix = self.get_similarity_matrix(method, segments_hash)
1294                 if sim_matrix.ndim > 1:
1295                     method_similarities.append(sim_matrix[-1, :-1])
1296             else:
1297                 method_similarities.append(sim_matrix[:,-1])
1298         except Exception as e:
1299             print(f"Error calculating similarity for method {method}: {str(e)}")
1300     if not method_similarities:
1301         print(f"No valid similarity methods for window size {w_size}, skipping.")
1302         continue
1303     min_length = min(len(sim) for sim in method_similarities)
1304     method_similarities = [sim[-min_length:] for sim in method_similarities]
1305     method_similarities_array = np.array(method_similarities)
1306     overall_similarity = np.nanmean(method_similarities_array, axis=0)
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1307     all_similarities.append(overall_similarity)
1308     if not all_similarities:
1309         print("No similarities found for any window size. Using fallback similarity.")
1310         return self.fallback_similarity_method()
1311     min_length = min(len(s) for s in all_similarities)
1312     all_similarities = [s[-min_length:] for s in all_similarities]
1313     all_similarities_array = np.array(all_similarities)
1314     combined_similarities = np.nanmean(all_similarities_array, axis=0)
1315     return combined_similarities
1316
1317     def fallback_similarity_method(self):
1318         # Simple autocorrelation-based similarity
1319         acf = np.correlate(self.data, self.data, mode='full')[len(self.data)-1:]
1320         return acf / acf[0] # Normalize
1321
1322     def analyze_segment_similarity(self, segment_index):
1323         current_segment = self.extract_segments(self.window_size)[-1]
1324         historical_segment = self.extract_segments(self.window_size)[segment_index]
1325         similarity_scores = {}
1326         for method in self.similarity_methods:
1327             if method == 'cosine':
1328                 score = np.dot(current_segment, historical_segment) / (np.linalg.norm(current_segment) *
1329 np.linalg.norm(historical_segment))
1330             elif method == 'euclidean':
1331                 score = 1 / (1 + np.linalg.norm(current_segment - historical_segment))
1332             elif method == 'dtw':
1333                 score = 1 / (1 + numba_dtw(current_segment, historical_segment)) # Use the global function
1334             similarity_scores[method] = score
1335         feature_contributions = np.abs(current_segment - historical_segment)
1336         top_contributing_features = np.argsort(feature_contributions)[::-1][:5]
1337         return {
1338             'similarity_scores': similarity_scores,
1339             'top_contributing_features': top_contributing_features.tolist(),
1340             'feature_contributions': feature_contributions.tolist()
1341         }
1342
1343     def get_nearest_neighbors(self, k=5):
1344         similarities = self.find_similar_segments()
1345         nearest_indices = np.argsort(similarities)[::-1][:k]
1346         return [(idx, similarities[idx]) for idx in nearest_indices]
1347
1348     def dtw_similarity(self, X):
1349         return numba_dtw_similarity(X) # Use the global function
1350
1351     def adaptive_window_size(self):
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1352     data_length = len(self.data)
1353     if data_length < 100:
1354         base_window = max(2, data_length // 20)
1355     elif data_length < 1000:
1356         base_window = max(5, data_length // 40)
1357     else:
1358         base_window = max(25, data_length // 80)
1359     potential_seasons = self.detect_seasonality()
1360     variability = np.std(self.data) / (np.mean(self.data) + 1e-8)
1361     if potential_seasons:
1362         window = min(max(potential_seasons), base_window)
1363     else:
1364         window = int(base_window * (1 + variability))
1365     return max(2, min(window, data_length // 8)) # Ensure window is at most 1/8 of data length
1366
1367     def detect_seasonality(self, max_lag=None):
1368         if max_lag is None:
1369             max_lag = len(self.data) // 2
1370         acf = np.correlate(self.data, self.data, mode='full')[-max_lag:]
1371         peaks = np.where((acf[1:-1] > acf[:-2]) & (acf[1:-1] > acf[2:]))[0] + 1
1372         if len(peaks) > 0:
1373             return [int(peaks[0])] # Return a list with the first peak
1374         return [] # Return an empty list if no peaks found
1375     def detect_anomalies(self, threshold_percentile=2):
1376         segments = self.extract_segments(self.window_size)
1377         similarities = self.find_similar_segments()
1378         threshold = np.percentile(similarities, threshold_percentile)
1379         anomaly_indices = np.where(similarities < threshold)[0]
1380         anomalies = []
1381         for idx in anomaly_indices:
1382             start = idx
1383             end = idx + self.window_size
1384             anomalies.append({
1385                 'start_index': start,
1386                 'end_index': end,
1387                 'segment': self.data[start:end].tolist(),
1388                 'similarity_score': similarities[idx]
1389             })
1390         return anomalies, threshold
1391
1392     def plot_anomalies(self, threshold_percentile=5):
1393         anomalies, threshold = self.detect_anomalies(threshold_percentile)
1394         plt.figure(figsize=(12, 6))
1395         plt.plot(self.data, label='Time Series', color='blue')
1396         for anomaly in anomalies:
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1397     plt.axvspan(anomaly['start_index'], anomaly['end_index'], color='red', alpha=0.3)
1398 plt.title(f'Time Series with Detected Anomalies (Threshold: {threshold:.4f})')
1399 plt.xlabel('Time Index')
1400 plt.ylabel('Values')
1401 plt.legend()
1402 if not anomalies:
1403     plt.text(0.5, 0.5, 'No anomalies detected', horizontalalignment='center',
1404             verticalalignment='center', transform=plt.gca().transAxes)
1405 else:
1406     print(f"Detected {len(anomalies)} anomalies")
1407 plt.show()
1408
1409 similarities = self.find_similar_segments()
1410 print(f"Similarity score range: {similarities.min():.4f} to {similarities.max():.4f}")
1411 print(f"Similarity score mean: {similarities.mean():.4f}")
1412 print(f"Similarity score median: {np.median(similarities):.4f}")
1413 print(f"Anomaly threshold: {threshold:.4f}")
1414
1415 def calculate_mean_squared_error(self, actual, predicted):
1416     return np.mean((actual - predicted) ** 2)
1417
1418 def calculate_basic_similarity(self, actual, predicted):
1419     # Ensuring that neither actual nor predicted are empty to avoid runtime errors
1420     if actual.size == 0 or predicted.size == 0:
1421         return np.nan
1422     correlation = np.corrcoef(actual, predicted)[0, 1]
1423     return correlation
1424
1425 def fallback_prediction(self, num_points):
1426     if len(self.data) < num_points * 2:
1427         raise ValueError("Insufficient data for prediction")
1428     def adaptive_window(data):
1429         def mse(window):
1430             trend = extract_trend(data, int(window))
1431             return np.mean((data[int(window)-1:] - trend)**2)
1432         result = minimize_scalar(mse, bounds=(10, len(data)//2), method='bounded')
1433         return int(result.x)
1434
1435 def extract_trend(data, window_size):
1436     return np.convolve(data, np.ones(window_size), 'valid') / window_size
1437
1438 def detect_seasonalities(data, max_period, num_seasons=2):
1439     correlations = [np.corrcoef(data[:-i], data[i:])[0, 1] for i in range(1, max_period)]
1440     seasons = []
1441     for _ in range(num_seasons):
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1442     if len(correlations) > 0:
1443         season = np.argmax(correlations) + 1
1444         seasons.append(season)
1445         correlations[season-1] = -1 # Remove detected season
1446     return seasons
1447
1448     def model_nonlinear_trend(data, x):
1449         coeffs = np.polyfit(x, data, 3)
1450         return np.poly1d(coeffs)
1451
1452     def detect_anomalies(data, threshold=3):
1453         mean = np.mean(data)
1454         std = np.std(data)
1455         return np.abs(data - mean) > threshold * std
1456     window_size = adaptive_window(self.data)
1457     trend = extract_trend(self.data, window_size)
1458     detrended = self.data[window_size-1:] - trend
1459     seasonality_periods = detect_seasonalities(detrended, num_points)
1460     seasonals = []
1461     for period in seasonality_periods:
1462         seasonal = np.zeros(period)
1463         for i in range(period):
1464             seasonal[i] = np.mean(detrended[i::period])
1465         seasonals.append(seasonal)
1466     combined_seasonal = np.zeros_like(detrended)
1467     for seasonal in seasonals:
1468         combined_seasonal += np.tile(seasonal, len(detrended) // len(seasonal) + 1)[:len(detrended)]
1469     residuals = detrended - combined_seasonal[:len(detrended)]
1470     anomalies = detect_anomalies(residuals)
1471     cleaned_residuals = residuals.copy()
1472     cleaned_residuals[anomalies] = np.median(residuals)
1473     x = np.arange(len(self.data))
1474     trend_model = model_nonlinear_trend(self.data, x)
1475     future_x = np.arange(len(self.data), len(self.data) + self.forecast_horizon)
1476     future_trend = trend_model(future_x)
1477     future_seasonal = np.zeros(self.forecast_horizon)
1478     for seasonal in seasonals:
1479         future_seasonal += np.tile(seasonal, self.forecast_horizon // len(seasonal) + 1)[:self.forecast_horizon]
1480
1481     def predict_residuals_with_ci(residuals, horizon, confidence=0.95):
1482         weights = np.exp(np.linspace(-1, 0, len(residuals)))
1483         weighted_mean = np.sum(residuals * weights) / np.sum(weights)
1484         weighted_std = np.sqrt(np.sum(weights * (residuals - weighted_mean)**2) / np.sum(weights))
1485         predictions = np.random.normal(weighted_mean, weighted_std, (1000, horizon))
1486         mean_prediction = np.mean(predictions, axis=0)
```


Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1487     ci_lower = np.percentile(predictions, (1 - confidence) / 2 * 100, axis=0)
1488     ci_upper = np.percentile(predictions, (1 + confidence) / 2 * 100, axis=0)
1489     return mean_prediction, ci_lower, ci_upper
1490     future_residuals, ci_lower, ci_upper = predict_residuals_with_ci(cleaned_residuals, self.forecast_horizon)
1491     predictions = future_trend + future_seasonal + future_residuals
1492     ci_lower += future_trend + future_seasonal
1493     ci_upper += future_trend + future_seasonal
1494     return predictions, ci_lower, ci_upper
1495
1496 def tune_hyperparameters(self):
1497     # Example: tune the number of seasonalities to detect
1498     best_num_seasons = 1
1499     best_mse = float('inf')
1500     for num_seasons in range(1, 5):
1501         predictions, _, _ = self.fallback_prediction(num_points=20)
1502         mse = np.mean((self.data[-len(predictions):] - predictions)**2)
1503         if mse < best_mse:
1504             best_mse = mse
1505             best_num_seasons = num_seasons
1506     return {'num_seasons': best_num_seasons}
1507
1508 def predict(self):
1509     self.adjust_dynamic_parameters()
1510     try:
1511         similarities = self.find_similar_segments()
1512         if len(similarities) == 0:
1513             print("No similarities found. Using fallback prediction.")
1514             return self.fallback_prediction(self.forecast_horizon)[0]
1515         threshold = self.get_dynamic_threshold(similarities)
1516         valid_indices = []
1517         for percentile in range(95, 70, -5): # Start at 95th percentile, go down to 70th
1518             top_indices = np.where(similarities > np.percentile(similarities, percentile))[0]
1519             valid_indices = top_indices[top_indices + self.window_size + self.forecast_horizon <= len(self.data)]
1520             if len(valid_indices) >= 3:
1521                 break
1522         if len(valid_indices) == 0:
1523             print("No valid indices found. Using fallback prediction.")
1524             return self.fallback_prediction(self.forecast_horizon)[0]
1525         predictions = []
1526         weights = []
1527         for idx in valid_indices:
1528             start = idx + self.window_size
1529             end = start + self.forecast_horizon
1530             if end <= len(self.data):
1531                 segment = self.data[start:end]
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1532         predictions.append(segment)
1533         weights.append(similarities[idx])
1534     if predictions:
1535         min_length = min(len(p) for p in predictions)
1536         predictions = [p[:min_length] for p in predictions]
1537         predictions = np.array(predictions)
1538         weights = np.array(weights)
1539         last_actual = self.data[-1]
1540         for i in range(len(predictions)):
1541             shift = last_actual - predictions[i][0]
1542             predictions[i] += shift
1543         predicted_values = np.average(predictions, axis=0, weights=weights)
1544     else:
1545         print("No valid predictions. Using fallback prediction.")
1546         predicted_values = self.fallback_prediction(self.forecast_horizon)[0]
1547     except Exception as e:
1548         print(f"Error in predict: {str(e)}")
1549         predicted_values = self.fallback_prediction(self.forecast_horizon)[0] # Return only predictions, not CI
1550     if predicted_values.size > 0:
1551         actual_values = self.data[-len(predicted_values):]
1552         prediction_error = self.calculate_mean_squared_error(actual_values, predicted_values)
1553         recent_similarity_score = self.calculate_basic_similarity(actual_values, predicted_values)
1554         self.update_recent_performance(prediction_error, recent_similarity_score)
1555     else:
1556         self.update_recent_performance(np.nan, np.nan)
1557     return predicted_values
1558
1559     def update_recent_performance(self, new_error, new_similarity_score):
1560         self.recent_errors.append(new_error)
1561         self.recent_similarity_scores.append(new_similarity_score)
1562         # Optionally, trim these lists to avoid unlimited growth
1563         self.recent_errors = self.recent_errors[-100:] # Keep the last 100 records
1564         self.recent_similarity_scores = self.recent_similarity_scores[-100:]
1565
1566     def evaluate_prediction(self, actual, predicted):
1567         if len(actual) != len(predicted):
1568             raise ValueError("Actual and predicted arrays must have the same length.")
1569         if len(actual) == 0:
1570             return {metric: np.nan for metric in ['MSE', 'MAE', 'RMSE', 'MAPE', 'SMAPE', 'R-squared', 'Direction Accuracy',
1571 'Theil\'s U']}
1572         mse = np.mean((actual - predicted) ** 2)
1573         mae = np.mean(np.abs(actual - predicted))
1574         rmse = np.sqrt(mse)
1575         mape = np.mean(np.abs((actual - predicted) / (actual + 1e-8))) * 100
1576         smape = np.mean(2 * np.abs(predicted - actual) / (np.abs(actual) + np.abs(predicted) + 1e-8)) * 100
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1577     r2 = r2_score(actual, predicted)
1578     direction_actual = np.sign(np.diff(actual))
1579     direction_pred = np.sign(np.diff(predicted))
1580     direction_accuracy = np.mean(direction_actual == direction_pred) * 100
1581     actual_changes = np.diff(actual)
1582     predicted_changes = np.diff(predicted)
1583     theil_u = np.sqrt(np.sum(predicted_changes**2) / np.sum(actual_changes**2)) if np.sum(actual_changes**2) !=
1584     0 else np.nan
1585     return {
1586         'MSE': mse, 'MAE': mae, 'RMSE': rmse, 'MAPE': mape, 'SMAPE': smape,
1587         'R-squared': r2, 'Direction Accuracy': direction_accuracy, 'Theil\'s U': theil_u
1588     }
1589
1590     def validate_prediction(self, splits=3):
1591         n_samples = len(self.data)
1592         max_splits = (n_samples - self.window_size) // self.forecast_horizon
1593         splits = min(splits, max_splits)
1594         if splits < 2:
1595             print("Warning: Not enough data for multiple splits. Performing single train-test split.")
1596             train_size = int(0.8 * n_samples)
1597             train, test = self.data[:train_size], self.data[train_size:]
1598             self.data = train
1599             self.similarity_cache = {}
1600             predicted = self.predict()
1601             if predicted.size > 0:
1602                 actual = test[:len(predicted)]
1603                 metrics = self.evaluate_prediction(actual, predicted)
1604                 self.data = np.concatenate((train, test)) # Restore original data
1605                 return metrics
1606             else:
1607                 print("Insufficient data to make a prediction.")
1608                 return None
1609         tscv = TimeSeriesSplit(n_splits=splits, test_size=self.forecast_horizon)
1610         errors = []
1611         for train_index, test_index in tscv.split(self.data):
1612             if len(train_index) < self.window_size:
1613                 print(f"Warning: Train set too small for window size. Skipping split.")
1614                 continue
1615             train, test = self.data[train_index], self.data[test_index]
1616             original_data = self.data
1617             self.data = train
1618             self.similarity_cache = {}
1619             predicted = self.predict()
1620             if predicted.size > 0:
1621                 actual = test[:len(predicted)]
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1622         metrics = self.evaluate_prediction(actual, predicted)
1623         errors.append(metrics)
1624     else:
1625         print("Insufficient data to predict for this split.")
1626         self.data = original_data
1627     if errors:
1628         avg_metrics = {metric: np.mean([e[metric] for e in errors if metric in e]) for metric in errors[0]}
1629         return avg_metrics
1630     else:
1631         print("No valid predictions could be made across splits.")
1632         return None
1633
1634     def get_explainability_results(self, top_k=5):
1635         similarities = self.find_similar_segments()
1636         threshold = self.get_dynamic_threshold(similarities)
1637         top_indices = np.where(similarities > threshold)[0]
1638         if len(top_indices) == 0:
1639             top_indices = np.argsort(similarities)[-top_k:]
1640         results = {
1641             'top_similar_segments': top_indices.tolist(),
1642             'similarity_scores': similarities[top_indices].tolist(),
1643             'threshold': threshold,
1644             'segment_contributions': []
1645         }
1646         predictions = []
1647         valid_indices = []
1648         for idx in top_indices:
1649             start = idx + self.window_size
1650             if start + self.forecast_horizon <= len(self.data):
1651                 predictions.append(self.data[start:start + self.forecast_horizon])
1652                 valid_indices.append(idx)
1653         if not predictions:
1654             return results
1655         predictions = np.array(predictions)
1656         weights = similarities[valid_indices]
1657         weighted_predictions = predictions * weights[:, np.newaxis]
1658         for i, (index, score, prediction, contribution) in enumerate(zip(valid_indices, similarities[valid_indices],
1659 predictions, weighted_predictions)):
1660             results['segment_contributions'].append({
1661                 'segment_index': int(index),
1662                 'similarity_score': float(score),
1663                 'prediction': prediction.tolist(),
1664                 'weighted_contribution': contribution.tolist(),
1665                 'contribution_percentage': (contribution / np.sum(weighted_predictions, axis=0) * 100).tolist()
1666             })
```

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1667     return results
1668
1669     def plot_nearest_neighbors(self, k=5):
1670         current_segment = self.extract_segments(self.window_size)[-1]
1671         neighbors = self.get_nearest_neighbors(k)
1672         plt.figure(figsize=(15, 10))
1673         plt.subplot(k+1, 1, 1)
1674         plt.plot(current_segment, color='blue', label='Current Segment')
1675         plt.title('Current Segment')
1676         plt.legend()
1677         for i, (idx, similarity) in enumerate(neighbors, start=2):
1678             neighbor_segment = self.extract_segments(self.window_size)[idx]
1679             plt.subplot(k+1, 1, i)
1680             plt.plot(neighbor_segment, color='red', label=f'Neighbor {i-1}')
1681             plt.title(f'Neighbor {i-1} (Similarity: {similarity:.4f})')
1682             plt.legend()
1683         plt.tight_layout()
1684         plt.show()
1685
1686     def analyze_and_plot_neighbors(self, k=5):
1687         current_segment = self.extract_segments(self.window_size)[-1]
1688         neighbors = self.get_nearest_neighbors(k)
1689         plt.figure(figsize=(20, 5*k))
1690         plt.subplot(k+1, 2, 1)
1691         plt.plot(current_segment, color='blue', label='Current Segment')
1692         plt.title('Current Segment')
1693         plt.legend()
1694         for i, (idx, overall_similarity) in enumerate(neighbors, start=1):
1695             neighbor_segment = self.extract_segments(self.window_size)[idx]
1696             analysis = self.analyze_segment_similarity(idx)
1697             plt.subplot(k+1, 2, 2*i+1)
1698             plt.plot(neighbor_segment, color='red', label=f'Neighbor {i}')
1699             plt.title(f'Neighbor {i} (Overall Similarity: {overall_similarity:.4f})')
1700             plt.legend()
1701             plt.subplot(k+1, 2, 2*i+2)
1702             methods = list(analysis['similarity_scores'].keys())
1703             scores = list(analysis['similarity_scores'].values())
1704             plt.bar(methods, scores)
1705             plt.title(f'Similarity Scores for Neighbor {i}')
1706             plt.ylim(0, 1)
1707             print(f"\nNeighbor {i} Analysis:")
1708             print(f"Overall Similarity: {overall_similarity:.4f}")
1709             print("Similarity Scores:")
1710             for method, score in analysis['similarity_scores'].items():
1711                 print(f" {method}: {score:.4f}")
```

This is a preprint draft. The published article can be found at: <https://doi.org/10.1016/j.cie.2024.110812>.

Please cite this paper as:

Naser A., **Naser M.Z.** (2025). SPINEX-TimeSeries: Similarity-based predictions with explainable neighbors exploration for time series and forecasting problems. *Computers & Industrial Engineering*. <https://doi.org/10.1016/j.cie.2024.110812>.

```
1712     print("Top Contributing Features:", analysis['top_contributing_features'])
1713     plt.tight_layout()
1714     plt.show()
1715
```