STR-NBEATS: A Novel Hybrid Framework Integrating Time Series Decomposition with Deep Learning for Enhanced Temperature Forecasting

Mohammed Elseidi
Assistant Professor of Statistics
College of Arts and Science, Umm Al Quwain University, Umm Al Quwain, UAE
drmohammed.elseidi@uaqu.ac.ae

Faculty of Commerce, Mansoura University, Mansoura, Egypt
dr_mohamedelsiedy@mans.edu.eg

Abstract

Accurate temperature forecasting plays a pivotal role in environmental management, agriculture, and energy planning, where reliable predictions underpin informed decision-making and resource allocation. However, effectively modeling the intricate seasonal patterns, long-term trends, and nonlinear fluctuations inherent in temperature data remains a persistent challenge.

To address these complexities, we introduce STR-NBEATS, a hybrid framework that integrates Seasonal-Trend Decomposition using Regression (STR) with the Neural Basis Expansion Analysis (N-BEATS) deep learning architecture. STR decomposes temperature time series into interpretable trend and seasonal components, along with a remainder term. This enables a tailored forecasting strategy wherein predictable cyclical behavior is handled through a simple seasonal naive approach, and more complex trend and residual dynamics are captured by the highly flexible N-BEATS network.

We rigorously benchmark our approach against robust and well-established forecasting methods capable of modeling seasonality, including an STL-based hybrid model combined with Exponential Smoothing (STL-ETS) and an automated seasonal ARIMA model. Empirical results on real-world datasets demonstrate that STR-NBEATS consistently outperforms these strong benchmarks, achieving lower error metrics and delivering forecasts that are both more accurate and interpretable. By enhancing the fidelity of temperature predictions and providing clearer insights into underlying climatic patterns, STR-NBEATS offers a valuable tool for stakeholders seeking to navigate the challenges of a changing environment with greater confidence and precision.

Keywords: Temperature forecasting, Seasonal-Trend decomposition, N-BEATS, Seasonal ARIMA, STL-ETS, Time series analysis, Environmental science

1 Introduction

Accurate temperature forecasting is crucial for a wide range of socio-economic and environmental applications, including agriculture, environmental management, and energy planning. Precise predictions enable stakeholders to make informed decisions, allocate resources efficiently, and implement effective adaptation measures in the face of climate variability. As weather patterns become increasingly uncertain due to climate change, the demand for robust forecasting frameworks that offer both accuracy and interpretability has grown more urgent.

In agriculture, reliable temperature forecasts guide farmers in determining optimal planting times, irrigation schedules, and harvesting operations, thereby minimizing the risk of crop failure and enhancing yield quality [1, 2]. Such foresight supports long-term adaptation strategies, including breeding climate-resilient crop varieties and employing precision agriculture techniques to safeguard food security amid evolving climate conditions [3]. In environmental management, anticipating extremes like droughts, heatwaves, and floods enables timely intervention and resource allocation to mitigate damage and support sustainable conservation efforts [4, 5]. Likewise, in the energy sector, temperature-driven forecasts underpin efficient energy generation, distribution, and consumption, informing the integration of renewable resources and enhancing grid stability [6, 7].

Despite these wide-ranging benefits, temperature forecasting remains challenging due to the inherently complex patterns present in climatic data, including multiple seasonal cycles, evolving trends, and stochastic fluctuations. Traditional statistical and numerical methods have shown limitations when grappling with these complexities. Numerical Weather Prediction (NWP) models, while accurate in the short term, are computationally intensive and sensitive to initial conditions, restricting their long-term reliability [8, 9, 10]. ARIMA-based methods and exponential smoothing approaches, though conceptually simpler, often struggle to capture nonlinearities, multiple seasonalities, and shifting trends, leading to suboptimal forecast accuracy [11, 12, 13, 14, 15, 16].

Advances in machine learning (ML) and deep learning (DL) have sparked a paradigm shift in forecasting. Deep neural networks can model complex, nonlinear relationships and subtle patterns that elude traditional methods. Techniques such as LSTM networks, hybrid LSTM-CNN architectures, and attention-based mechanisms have demonstrated remarkable improvements in predictive accuracy [17, 18, 19, 20, 21, 22, 23]. However, these models are often considered "black boxes" and may be prone to overfitting, prompting ongoing research into more interpretable and robust approaches.

A promising strategy to enhance both accuracy and interpretability involves integrating decomposition techniques with advanced ML/DL models. Decomposition methods like Empirical Mode Decomposition (EMD), Seasonal-Trend Decomposition using LOESS (STL), and Singular Spectrum Analysis (SSA) separate time series into seasonal, trend, and residual components, enabling targeted modeling of each element. Such hybrid frameworks have led to improved forecast performance and clarity, as demonstrated by EEMD-LSTM, CEEMDAN-based models, and STL-based neural network hybrids [24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34].

Among these methodologies, Seasonal-Trend Decomposition using Regression (STR) [35] is a powerful recent development. STR employs a regression-based framework to model trend and multiple seasonal patterns with remarkable flexibility. Unlike traditional methods that assume fixed, unchanging seasonality, STR can represent complex and evolving seasonal topologies, handle time-varying and seasonal covariates, and ensure that extracted components remain smooth, interpretable, and stable over time. This opens up new avenues for handling intricate and dynamic climatic phenomena.

Despite STR's strengths, it has rarely been integrated with cutting-edge deep learning frameworks. Recent advances in deep learning have given rise to models that can flexibly and accurately capture the nonlinear, multi-timescale dynamics of complex data. One such model, Neural Basis Expansion Analysis for Time Series (N-BEATS) [36], has emerged as a state-of-the-art forecasting architecture due to its ability to leverage hierarchical basis expansions for improved accuracy and interpretability. Originally introduced to compete in global forecasting challenges, N-BEATS and its variants have demonstrated exceptional performance in diverse applications, including financial markets, energy load forecasting, quality of service in cellular networks, and the control of automated guided vehicles [37, 38, 39, 40, 41, 42, 43, 44, 45]. These studies consistently report that N-BEATS outperforms traditional statistical methods like ARIMA and ETS, as well as other deep learning architectures (e.g., LSTMs and CNNs), both in terms of forecast accuracy and robustness.

In this context, the integration of STR—offering a sophisticated, regression-based decomposition of trend and multiple seasonal components—with an N-BEATS model capable of extracting complex residual dynamics presents a compelling opportunity. STR effectively isolates seasonal and trend variations, yielding interpretable components that can be more accurately modeled once stripped of extraneous patterns. N-BEATS, in turn, excels at modeling and forecasting the decomposed signals, capturing nonlinearities and subtle features that may elude simpler methods. By uniting STR's interpretability and decomposition power with N-BEATS's state-of-the-art forecasting ability, the proposed STR-NBEATS framework aims to deliver accurate, stable, and transparent forecasts.

Contributions of This Study: This paper introduces a novel hybrid framework, STR-NBEATS, that seamlessly integrates Seasonal-Trend Decomposition using Regression with the N-BEATS forecasting architecture. The key contributions are:

1. **Innovative Integration:** We bridge the interpretability and structural clarity of STR with the high-accuracy predictive capabilities of N-BEATS, creating a unified model that leverages both approaches' strengths.

2. **Enhanced Forecasting Performance:** The STR-NBEATS framework yields forecasts that improve upon state-of-the-art benchmarks. By decomposing the time series, we enable N-BEATS to focus on the residual complexity and deliver sharper, more accurate predictions.

3. **Robustness and Interpretability:** Decomposition via STR ensures that seasonal and trend patterns are clearly identified and isolated, making forecasts more interpretable and enabling stakeholders to discern underlying climatic dynamics with greater confidence.

4. **Benchmarking and Validation:** We rigorously compare our framework against established seasonal ARIMA and STL-ETS models, demonstrating consistent improvements across accuracy metrics and offering valuable insights into the benefits of combining advanced decomposition with cutting-edge deep learning.

The remainder of this paper is organized as follows. Section 2 details the STR methodology, the N-BEATS model, and the integration process. Section 3 presents the experimental setup and dataset. In Section 4, we discuss results and compare performance against benchmark models. Finally, Section 5 concludes the paper and outlines avenues for future research.

2 Methodology

In this section, we present the integrated STR-NBEATS forecasting framework. Our approach begins with the Seasonal-Trend Decomposition using Regression (STR), which decomposes the time series into trend, seasonal, and remainder components. Next, we employ a simple seasonal naive (sNaive) model to forecast the seasonal components. We then apply N-BEATS, a state-of-the-art neural network model, to the non-seasonal (trend + remainder) component. Finally, we combine these forecasts to obtain the final prediction.

2.1 Notation

Let $\{y_t\}_{t=1}^{T}$ denote a univariate time series of length T. We aim to produce forecasts $\hat{y}_{t+h}$ for horizons $h = 1, 2, \ldots, H$. We consider I distinct seasonal periods $\{\omega_i\}_{i=1}^{I}$, such as daily or annual cycles, which may be non-integer and heterogeneous.

2.2 Seasonal-Trend Decomposition using Regression (STR)

The STR model [35] decomposes the time series into trend, seasonal, and remainder (noise) components, potentially including covariates. Formally, STR represents the observed series as:

$$y_t = T_t + \sum_{i=1}^{I} S_t^{(i)} + \sum_{p=1}^{P} \phi_{p,t} z_{p,t} + R_t, \tag{1}$$

where:

- T_t is the smooth trend component,
- $S_t^{(i)}$ is the i-th seasonal component with period ω_i,
- $\phi_{p,t} z_{p,t}$ represents the effect of external covariates,
- R_t is the remainder (noise) component.

STR enforces smoothness and identifiability through penalization of differences. For example, the trend is often encouraged to be smooth by penalizing its second difference:

$$\Delta^2 T_t = T_t - 2T_{t-1} + T_{t-2} \sim \mathcal{N}(0, \sigma_T^2), \tag{2}$$

ensuring that abrupt changes in T_t are unlikely.

After fitting the STR model, we obtain estimates $\hat{T}_t$, $\hat{S}_t^{(i)}$, and $\hat{R}_t$. Thus,

$$y_t \approx \hat{T}_t + \sum_{i=1}^{I} \hat{S}_t^{(i)} + \hat{R}_t. \tag{3}$$

2.3 Seasonal Forecasting with sNaive

To forecast the seasonal components, we apply a seasonal naive method. The sNaive forecast for each seasonal component $S_t^{(i)}$ over the horizon h is:

$$\hat{S}_{t+h}^{(i)} = \hat{S}_{t+h-\omega_i}^{(i)}, \tag{4}$$

assuming that seasonal patterns repeat. If multiple seasonal components exist, we sum their forecasts:

$$\hat{S}_{t+h} = \sum_{i=1}^{I} \hat{S}_{t+h}^{(i)}. \tag{5}$$

2.4 N-BEATS Modeling of Trend+Remainder

Next, we combine the estimated trend and remainder to form a non-seasonal series:

$$T\hat{R}_t = \hat{T}_t + \hat{R}_t. \tag{6}$$

Before training N-BEATS, we normalize $T\hat{R}_t$:

$$T\tilde{R}_t = \frac{T\hat{R}_t - \mu_{T\hat{R}}}{\sigma_{T\hat{R}}}, \tag{7}$$

where $\mu_{T\hat{R}}$ and $\sigma_{T\hat{R}}$ are the mean and standard deviation of the training segment of $T\hat{R}_t$.

2.4.1 N-BEATS Architecture

N-BEATS (Neural Basis Expansion Analysis for Time Series) [36] is a deep neural network designed for interpretable and accurate forecasting. It uses stacks of fully-connected layers organized into blocks. Each block:

- Produces backcast and forecast outputs,
- Expands the time series into a basis using learned coefficients.

A block computes expansions of the form:

$$\hat{x}_t^{(l)} = \sum_i \theta_{l,i} g_i(t), \tag{8}$$

where $g_i(t)$ are basis functions (polynomial, Fourier, or learned), and $\theta_{l,i}$ are the corresponding coefficients.

N-BEATS employs residual stacking, where each block refines what previous blocks did not explain:

$$x_{l+1} = x_l - \hat{x}_l^{(l)}, \quad y = \sum_{l=1}^{L} \hat{y}^{(l)}. \tag{9}$$

After training N-BEATS on $\{T\tilde{R}_t\}$, it provides forecasts $T\hat{\tilde{R}}_{t+h}$. We invert the normalization:

$$T\hat{R}_{t+h} = T\hat{\tilde{R}}_{t+h}\sigma_{T\hat{R}} + \mu_{T\hat{R}}. \tag{10}$$

2.5 Final Forecast Combination

Finally, we combine the non-seasonal forecast from N-BEATS and the seasonal forecasts from sNaive:

$$\hat{y}_{t+h} = \hat{TR}_{t+h} + \hat{S}_{t+h}. \tag{11}$$

This fusion preserves the seasonal structure from STR and the complex, non-seasonal patterns captured by N-BEATS.

2.6 Algorithmic Description

Algorithm 1 STR-NBEATS Forecasting Algorithm

Require: Time series $\{y_t\}_{t=1}^{T}$, seasonal periods $\{\omega_i\}_{i=1}^{I}$, forecast horizon H.

STR Decomposition: Decompose y_t into $\hat{T}_t, \{\hat{S}_t^{(i)}\}, \hat{R}_t$ using STR.
Combine Trend and Remainder: $\hat{TR}_t = \hat{T}_t + \hat{R}_t$.
Forecast Seasonal Components (sNaive):
for $h = 1$ to H **do**
  for $i = 1$ to I **do**

$$\hat{S}_{t+h}^{(i)} = \hat{S}_{t+h-\omega_i}^{(i)}$$

  end for

$$\hat{S}_{t+h} = \sum_{i=1}^{I} \hat{S}_{t+h}^{(i)}$$

end for
Normalization: Compute $\mu_{\hat{TR}}$ and $\sigma_{\hat{TR}}$, then

$$\tilde{TR}_t = \frac{\hat{TR}_t - \mu_{\hat{TR}}}{\sigma_{\hat{TR}}}.$$

Train N-BEATS on $\tilde{TR}_t$ and obtain forecasts $\hat{\tilde{TR}}_{t+h}$.
Inverse Normalization:

$$\hat{TR}_{t+h} = \hat{\tilde{TR}}_{t+h}\sigma_{\hat{TR}} + \mu_{\hat{TR}}.$$

Combine Final Forecasts:

$$\hat{y}_{t+h} = \hat{TR}_{t+h} + \hat{S}_{t+h}.$$

return $\{\hat{y}_{t+h}\}_{h=1}^{H}$.

The proposed STR-NBEATS methodology leverages STR decomposition to separate the time series into interpretable components and sNaive forecasts to handle seasonal patterns. N-BEATS then models the remaining non-seasonal complexity. By blending these forecasts, we preserve interpretable structures while benefiting from the flexibility and accuracy of a powerful neural forecasting model.

2.7 Benchmark Models

To rigorously evaluate the performance of the proposed STR-NBEATS framework, we compare it against two well-established benchmarking models: a seasonal ARIMA model capable of capturing seasonal patterns, and an STL-based decomposition model combined with Exponential Smoothing (ETS) [46]. Both of these benchmarks are widely recognized in the forecasting literature and serve as strong baselines.

2.7.1 Seasonal ARIMA Model

The Autoregressive Integrated Moving Average (ARIMA) model is a commonly used statistical approach for time series forecasting. ARIMA(p, d, q) incorporates three types of parameters:

- p: The order of the autoregressive (AR) part.
- d: The degree of differencing required to achieve stationarity.
- q: The order of the moving average (MA) part.

When seasonality is present, a seasonal extension, ARIMA$(p, d, q)(P, D, Q)_m$, can be employed, where:

- P, D, Q: The seasonal AR, differencing, and MA orders respectively.
- m: The number of periods in each seasonal cycle (e.g., $m = 365$ for daily data with annual seasonality).

The general form of a seasonal ARIMA model can be expressed as:

$$\Phi_P(B^m)\phi_p(B)\nabla^d\nabla_m^D y_t = \Theta_Q(B^m)\theta_q(B)\varepsilon_t, \tag{12}$$

where:

- y_t is the time series value at time t,
- B is the backshift operator ($By_t = y_{t-1}$),
- $\nabla = (1 - B)$ and $\nabla_m = (1 - B^m)$ are the non-seasonal and seasonal differencing operators,
- $\phi_p(B)$ and $\Phi_P(B^m)$ represent the non-seasonal and seasonal AR components,
- $\theta_q(B)$ and $\Theta_Q(B^m)$ represent the non-seasonal and seasonal MA components,
- ε_t is white noise with zero mean and constant variance.

An automated seasonal ARIMA approach uses model selection criteria (such as AIC, AICc, or BIC) to choose optimal parameters (p, d, q) and (P, D, Q), providing a robust baseline that can adapt to different seasonal structures in the data.

2.7.2 STL-ETS Model

The STL-ETS model first applies Seasonal-Trend decomposition using LOESS (STL) to the original time series before employing an Exponential Smoothing (ETS) model for forecasting the decomposed components.

The STL decomposition breaks the original time series y_t into three additive components:

$$y_t = T_t + S_t + R_t, \tag{13}$$

where:

- T_t is the trend component,
- S_t is the seasonal component,
- R_t is the remainder (or residual) component.

Following the decomposition, the remainder series or seasonally adjusted series $(y_t - S_t)$ is modeled using an ETS framework. The ETS family includes a range of models allowing for different configurations of Error (E), Trend (T), and Seasonal (S) components. Common forms include:

- Additive or multiplicative error terms,
- Linear or damped trends,

- Additive, multiplicative, or no seasonal patterns.

In its most general form, an ETS model can be represented as $\mathrm{ETS}(\mathrm{E}, \mathrm{T}, \mathrm{S})$, where $\mathrm{E} \in \{A, M\}$, $\mathrm{T} \in \{N, A, Ad, M, Md\}$, and $\mathrm{S} \in \{N, A, M\}$, indicating the chosen combination of error, trend, and seasonality types. An automated model selection procedure is often employed to choose the appropriate configuration that best fits the data.

By first isolating and stabilizing seasonal and trend variations through STL, the subsequent ETS modeling step can focus on forecasting a more homogeneous remainder series. This hybrid STL-ETS approach effectively handles complex seasonal patterns and trends, making it a strong and flexible competitor against more advanced forecasting models.

Both the automated seasonal ARIMA and the STL-ETS models are strong benchmarks. The seasonal ARIMA method adeptly handles seasonal patterns and can be tuned automatically, while the STL-ETS method leverages powerful decomposition to disentangle underlying structures before forecasting. These benchmarks provide a robust baseline against which the performance of the proposed STR-NBEATS framework can be fairly assessed.

3 Data Description

The dataset comprises daily temperature observations recorded in Umm Al Quwain, UAE, obtained from the National Center of Meteorology. It spans from January 1, 2020, to March 31, 2024, providing a comprehensive basis for training and testing the forecasting model.

For the purpose of this study, the dataset was divided into:

- **Training Set**: All observations from January 1, 2020, to March 31, 2023, used to train the forecasting model.
- **Test Set**: The last 365 days, from April 1, 2023, to March 31, 2024, designated for evaluating the model's forecasting performance.

To address the inherent complexities of temperature data, including annual and semi-annual seasonal patterns, the series was modeled as a multi-seasonal time series with seasonal periods of 365 and 182 days. This preprocessing ensures the data is well-prepared for the STR-NBEATS hybrid model.

4 Analysis and Results

This section provides a thorough exploratory analysis of the daily temperature dataset alongside a detailed decomposition and forecasting study. The primary objective is to uncover intrinsic structural patterns, distributional characteristics, and temporal dependencies that can guide the selection and calibration of suitable forecasting methodologies. By leveraging multiple visual representations, summary statistics, and advanced decomposition techniques, we establish a comprehensive understanding of the underlying climatic processes.

Figure 1 presents four key visual diagnostics arranged in a single frame: the daily temperature time series, the autocorrelation function (ACF), the partial autocorrelation function (PACF), and the histogram of daily temperatures. The time series plot reveals a pronounced annual cycle, evidenced by regular peaks and troughs that correspond to seasonal climatic transitions. This repetitive pattern suggests that models incorporating explicit seasonal terms will likely prove effective. The ACF shows significant correlations at lags consistent with the yearly cycle, further confirming the presence of strong seasonality. Moreover, the appearance of additional, albeit weaker, peaks may reflect semi-annual or other sub-annual climatic influences. The PACF declines steeply after the first lag, indicating that much of the short-term memory in the series is captured by immediate past values and the strong seasonal structure. The histogram of daily temperatures, while centered around moderately high values, displays hints of multimodality, reflecting temperature regimes associated with different times of the year. This distributional complexity suggests that standard Gaussian assumptions may be inappropriate, and more flexible approaches could be warranted for both inference and prediction.

Figure 2 complements this initial exploration with a boxplot, Q-Q plot, seasonal subseries plot, and lag plot of the daily temperatures. The boxplot reveals a relatively stable median and only modest dispersion, illustrating a climate that, while seasonally dynamic, does not exhibit frequent extreme outliers. The Q-Q plot confirms deviations from normality, particularly in the tails, underscoring the need for robust or nonparametric methods that can handle such departures. The seasonal subseries plot arranges the data by seasonal cycle, making the annual pattern unmistakably clear and affirming that these temperature fluctuations recur with notable regularity year after year. Such consistency across multiple annual cycles underscores the suitability of seasonal-trend decomposition frameworks. The lag plot, which displays current temperatures against their one-day lagged values, shows a strong linear relationship, confirming that immediate past observations are highly predictive and reinforcing the fundamental premise of using time series models capable of leveraging this strong temporal dependence.

Table 1 offers descriptive statistics that further characterize the data. The mean and median align closely, indicating near-symmetry in the central tendency. Minimal skewness and moderate kurtosis levels suggest no extreme deviations from a roughly symmetric distribution, although the Q-Q plot's nonlinearity emphasizes that subtle distributional complexities remain. The moderate standard deviation points to a reasonable spread, while the strong and persistent seasonal signals suggest that much of the observed variability is structured rather than random. Together, these findings highlight a dataset that is both stable and periodically dynamic, posing distinct challenges and opportunities for modeling.

Figure 3 introduces the Seasonal-Trend decomposition using Regression (STR). By decomposing the data into a smoothly varying trend, multiple seasonal components, and a residual term, STR elucidates the fundamental building blocks of the time series. The original data plot in the top-left panel of the figure exhibits the well-defined annual cycles. Below it, the trend component shows a gradual, long-term progression, possibly reflecting broader climatic shifts or subtle environmental changes over the observed period. The seasonal components isolate both a dominant 365-day cycle and a 182.5-day sub-annual pattern, each capturing distinct cyclical behaviors. The former aligns neatly with known annual climatology, while the latter may stem from regional weather patterns or secondary climate drivers. The residual component appears substantially more stationary after extracting trend and seasonal factors, confirming that the STR approach effectively removes structured variability and leaves behind a noise-like remainder suitable for further modeling steps.

Building upon the decomposition, the forecasting phase aims to leverage both classical and modern methodologies to predict future temperatures. Figure 4 displays the sNaive forecasts for the 182-day and 365-day seasonal components. These forecasts project historical seasonal patterns into the future without imposing additional complexity, thus serving as stable benchmarks. The retained seasonal structures ensure that expected seasonal peaks and troughs are faithfully reproduced beyond the observed data. When combined with the non-seasonal Trend+Random component forecasts produced by a state-of-the-art deep learning model, N-BEATS, the forecasting framework becomes more adaptable and precise. Figure 5 shows the comparison between the historical Trend+Random values and the N-BEATS predictions. The N-BEATS model adeptly captures gradual shifts in the baseline temperature level as well as irregular fluctuations not directly tied to the seasonal cycles. This flexible modeling of residual dynamics yields a more accurate and nuanced understanding of the underlying climatic processes.

Finally, Figure 6 integrates these components into a cohesive predictive solution. By summing the sNaive-based seasonal forecasts and the N-BEATS Trend+Random forecast, the final predictions align closely with the out-of-sample test data. This close alignment attests to the robustness of the decomposition-based approach and the potency of combining classical seasonal modeling techniques with advanced machine learning methods. The resulting forecasts are both interpretable—given the transparent decomposition into trend, seasonal, and residual elements—and empirically effective.

In sum, the exploratory analysis and decomposition confirm that the temperature data exhibit strong and persistent seasonalities, stable yet non-Gaussian distributions, and meaningful autocorrelation structures. Employing the STR decomposition clarifies the drivers of variability, while the hybrid forecasting strategy that merges sNaive seasonal forecasts with N-BEATS-based residual modeling demonstrates the advantage of leveraging complementary methods. These results establish a framework well-suited for handling complex environmental time series and lay the groundwork for further refinements, including the incorporation of additional covariates or advanced hierarchical modeling approaches.

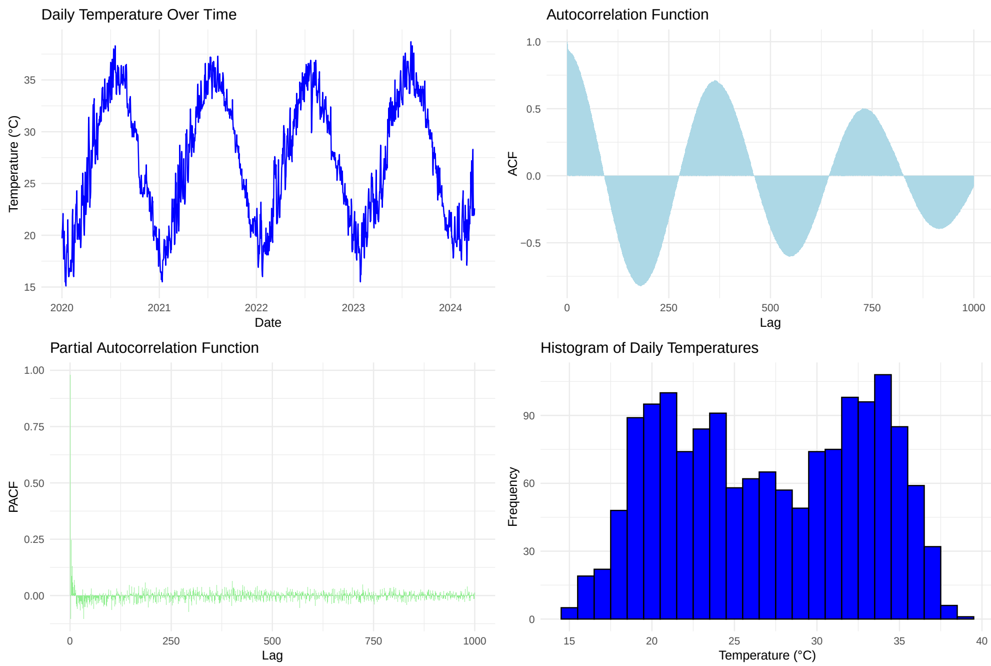

Figure 1: Daily temperature over time (top-left), autocorrelation function (top-right), histogram of daily temperatures (bottom-right), and partial autocorrelation function (bottom-left). The figure demonstrates strong annual seasonality, significant autocorrelation at seasonal lags, and a distribution that may deviate from strict normality.

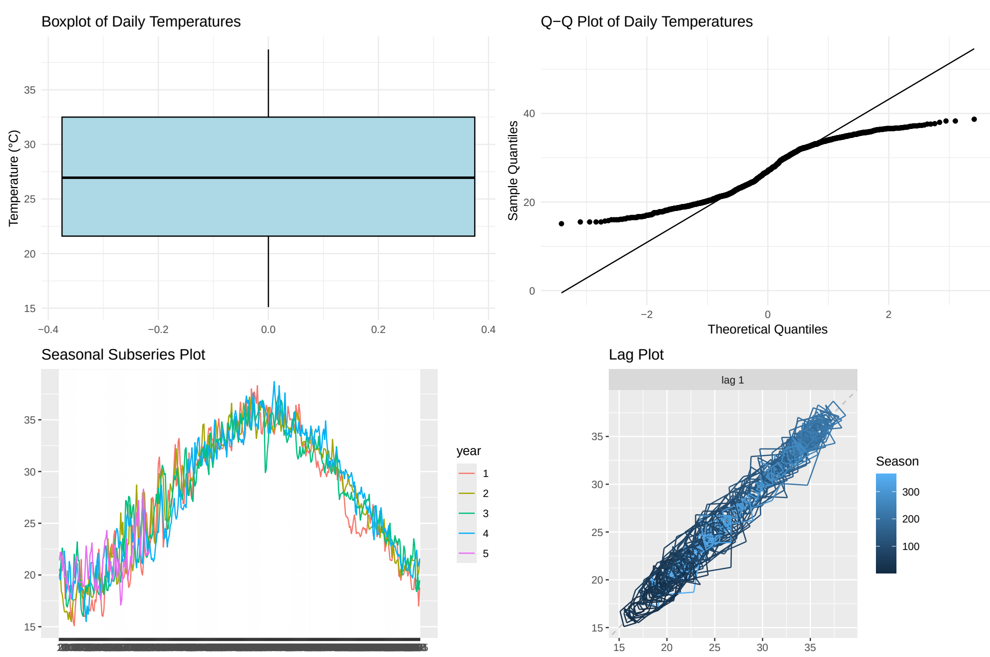

Figure 2: Boxplot, Q-Q plot, seasonal subseries plot, and lag plot of daily temperatures. The boxplot suggests a stable median and limited outliers, the Q-Q plot indicates departures from normality, the seasonal subseries plot confirms recurring annual patterns, and the lag plot shows strong linear autocorrelation.

Statistic	Value
Minimum Temperature	15.10°C
1st Quartile	21.60°C
Median	26.95°C
Mean	27.08°C
3rd Quartile	32.50°C
Maximum Temperature	38.70°C
Standard Deviation	5.94°C
Skewness	-0.03
Kurtosis	-1.29

Table 1: Descriptive statistics of daily temperatures. The distribution is roughly symmetric with moderate variability, and the climate appears stable without frequent extremes.

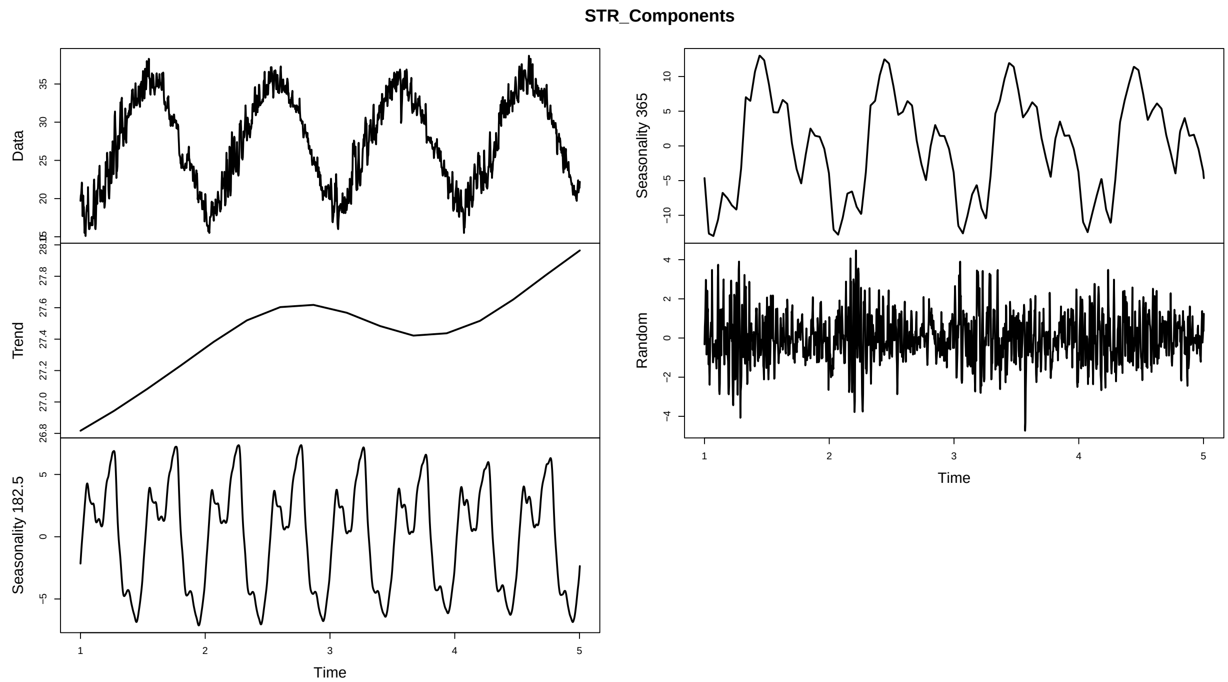

Figure 3: STR decomposition of the temperature time series. The top-left panel shows the raw data with annual cycles. Below it, the trend component reveals gradual long-term changes. To the right, the 365-day seasonal pattern dominates, while a 182.5-day seasonal component captures intra-annual fluctuations. The random residual is substantially reduced and appears stationary, confirming the efficacy of the decomposition.

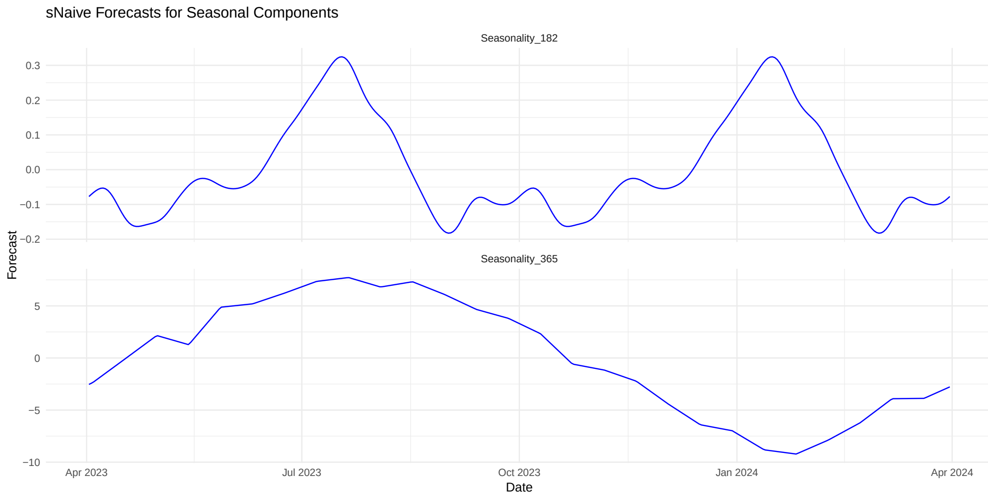

Figure 4: sNaive forecasts for the 182-day (top) and 365-day (bottom) seasonal components. These projections preserve historical seasonal patterns and provide stable baselines for integrating more sophisticated forecasting models.

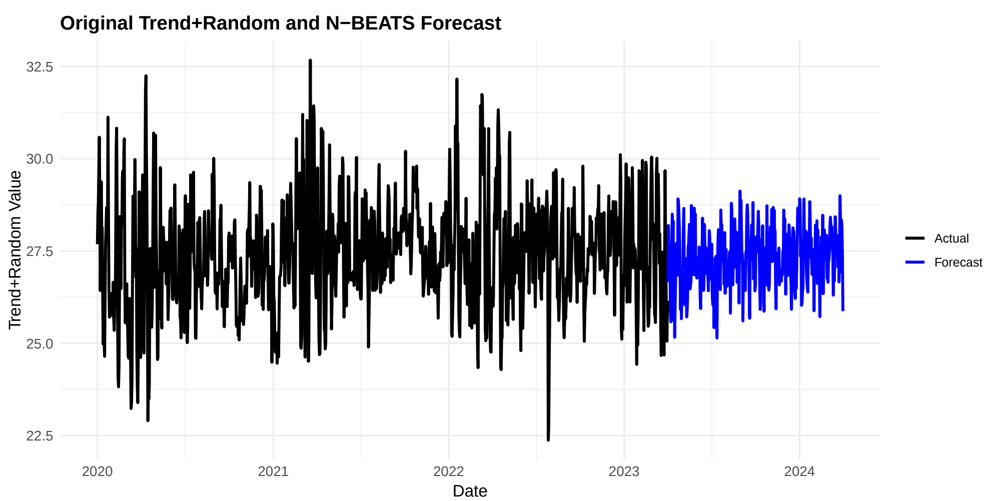

Figure 5: Original Trend+Random component (black) and N-BEATS forecasts (blue). The N-BEATS model accommodates subtle non-seasonal variations and captures evolving baseline conditions, enhancing the predictive performance of the overall model.

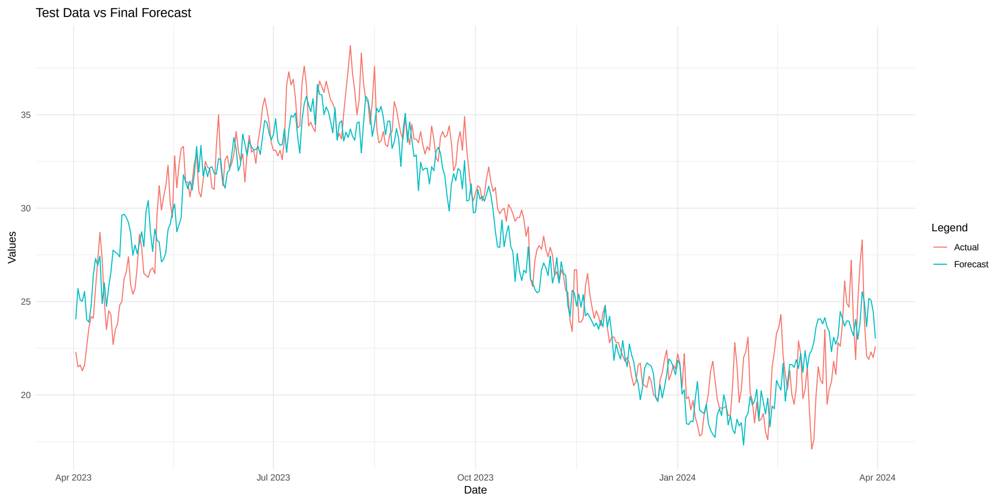

Figure 6: Comparison of the final integrated forecast (cyan) with the actual test data (red). By combining sNaive forecasts of the seasonal components with the N-BEATS-based Trend+Random prediction, the model aligns closely with observed values, validating the effectiveness of the chosen approach.

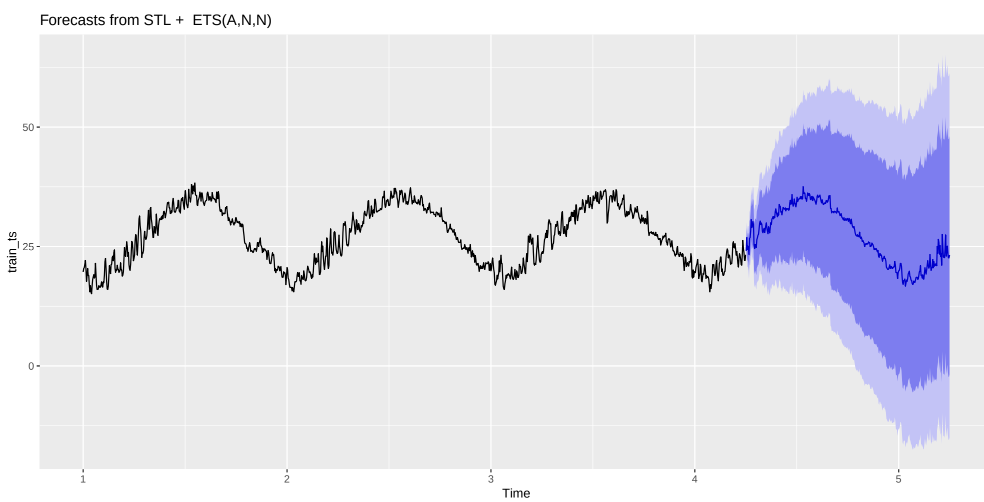

Figure 7: Forecast results from the STL + ETS(A,N,N) model. The black line represents historical temperature data, while the dark blue line and shaded region represent the point forecasts and their associated prediction intervals, respectively. Note the model's ability to capture seasonal patterns to some extent, but also observe the relatively wide prediction intervals, indicating increased uncertainty as the forecast horizon extends.

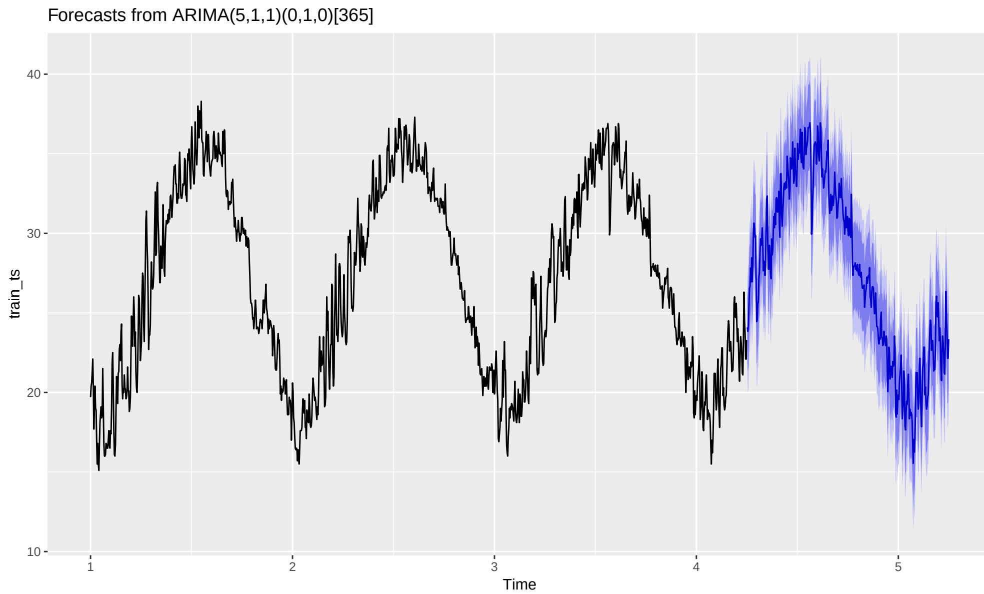

Figure 8: Forecast results from the ARIMA(5,1,1)(0,1,0)[365] model. The historical time series is shown in black, with the forecasts and corresponding intervals represented in blue. While ARIMA-based methods can model seasonal patterns through appropriate differencing and autoregressive terms, their prediction intervals may still widen considerably over longer horizons, reflecting uncertainty and the difficulty in capturing complex, nonlinear relationships.

4.1 Model Fit and Benchmark Discussion

Both the STL-ETS and ARIMA models serve as valuable benchmarks for evaluating our proposed STR-NBEATS framework. The STL-ETS approach applies Seasonal-Trend decomposition using LOESS (STL) to isolate underlying patterns, followed by an ETS(A,N,N) model for forecasting. According to the model summary:

```
Forecast method: STL + ETS(A,N,N)
Smoothing parameters:
    alpha = 0.796
Initial states:
    l = 27.1896
sigma:  1.2788
AIC=8990.861, AICc=8990.881, BIC=9006.098
```

This suggests a relatively simple exponential smoothing model without trend or multiplicative components. The use of STL decomposition provides a clearer separation of seasonal patterns, but as illustrated in Figure 7, the resulting forecasts still exhibit broad uncertainty bands, reflecting the limitations of a purely statistical approach when confronted with complex, high-frequency data and nonlinear dynamics.

On the other hand, the ARIMA(5,1,1)(0,1,0)[365] model incorporates seasonal differencing and multiple autoregressive terms to handle temporal dependencies and seasonal effects:

```
ARIMA(5,1,1)(0,1,0)[365]

Coefficients:
         ar1      ar2     ar3      ar4     ar5      ma1
      0.9226  -0.4145  0.1391  -0.0996  0.0803  -0.9896
s.e.  0.0364   0.0473  0.0495   0.0472  0.0364   0.0112
```

```
sigma^2 = 2.019
log likelihood = -1455.17
AIC=2924.34, AICc=2924.47, BIC=2957.31
```

Despite its sophistication, the ARIMA model's forecasts (Figure 8) still widen significantly as the lead time increases. While ARIMA models can capture certain linear and seasonal structures, they often struggle with more intricate nonlinear behaviors and multiple seasonalities. The observed expansion of prediction intervals over time underscores the challenges associated with relying solely on traditional statistical models for temperature forecasting.

In summary, these benchmark graphical results highlight the inherent complexities and uncertainties associated with temperature forecasting. While STL-ETS and ARIMA methods can provide baseline predictions and capture simpler patterns, they often fail to fully represent the underlying dynamics or achieve tight confidence intervals. These insights underscore the need for more advanced hybrid modeling approaches—such as the proposed STR-NBEATS framework—that combine the interpretability of decomposition techniques with the predictive power and flexibility of deep learning architectures.

5 Results and Discussion

In order to evaluate the effectiveness of the proposed STR-NBEATS framework, we conducted comparative analyses against two benchmark models: (1) a seasonal ARIMA model and (2) an STL-based hybrid model combined with Exponential Smoothing (ETS) forecasting. The test sets from our evaluation were identical across all models, allowing for a consistent assessment of forecasting performance. Table 2 summarizes the forecasting accuracy metrics for STR-NBEATS, the automatic seasonal ARIMA model, and the STL-ETS model.

Table 2: Accuracy metrics for the proposed STR-NBEATS model compared to the seasonal ARIMA model with covariates and the STL-ETS model. Metrics include Mean Error (ME), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Percentage Error (MPE), and Mean Absolute Percentage Error (MAPE).

Model	ME	RMSE	MAE	MPE	MAPE
STR-NBEATS	0.2783	1.8781	1.4685	0.4831	5.6116
ARIMA (Seasonal)	0.0816	2.1830	1.6802	-0.2865	6.4599
STL-ETS	0.3378	1.9459	1.5247	1.0161	5.8563

The STR-NBEATS model demonstrates superior performance across multiple metrics compared to the benchmark models. In particular, the RMSE of STR-NBEATS (1.8781) is notably lower than that of the ARIMA model (2.1830) and the STL-ETS model (1.9459), indicating that the proposed approach consistently reduces the magnitude of forecast errors. Similarly, the MAE for STR-NBEATS (1.4685) is lower than both the ARIMA (1.6802) and STL-ETS (1.5247) models, underscoring its ability to provide more precise point forecasts.

The percentage-based error metrics also highlight the effectiveness of STR-NBEATS. The MAPE for STR-NBEATS (5.6116) outperforms both the ARIMA (6.4599) and STL-ETS (5.8563) approaches, reflecting a reduction in relative error and an improvement in forecast reliability. While the ME and MPE values are relatively small for all models, the STR-NBEATS model maintains a competitive edge, indicating a lower systematic bias and more balanced predictive capabilities.

To provide a visual comparison of the models' forecasting performance, Figure 9 plots the actual temperature values alongside the forecasts generated by STR-NBEATS, ARIMA, and STL-ETS over the test period. The figure clearly illustrates that the STR-NBEATS forecasts closely align with the observed values, demonstrating improved accuracy and reduced forecasting errors compared to the benchmark models. The alignment between STR-NBEATS and the observed data highlights the model's ability to effectively capture the underlying seasonal and trend structures while mitigating residual noise.

These findings align with the core objectives of the STR-NBEATS framework. By integrating Seasonal-Trend Decomposition using Regression (STR) with a powerful deep learning architecture (N-BEATS), the model effectively captures and leverages underlying seasonal and trend patterns in temperature data. This synergy leads to improved performance against both traditional statistical models, such as seasonal ARIMA, and hybrid decomposition approaches like STL-ETS. The results illustrate that the proposed STR-NBEATS framework is not only capable of enhancing predictive accuracy but also of delivering more interpretable and robust forecasts, thereby supporting more informed decision-making and management strategies in climatology, agriculture, energy, and related fields.

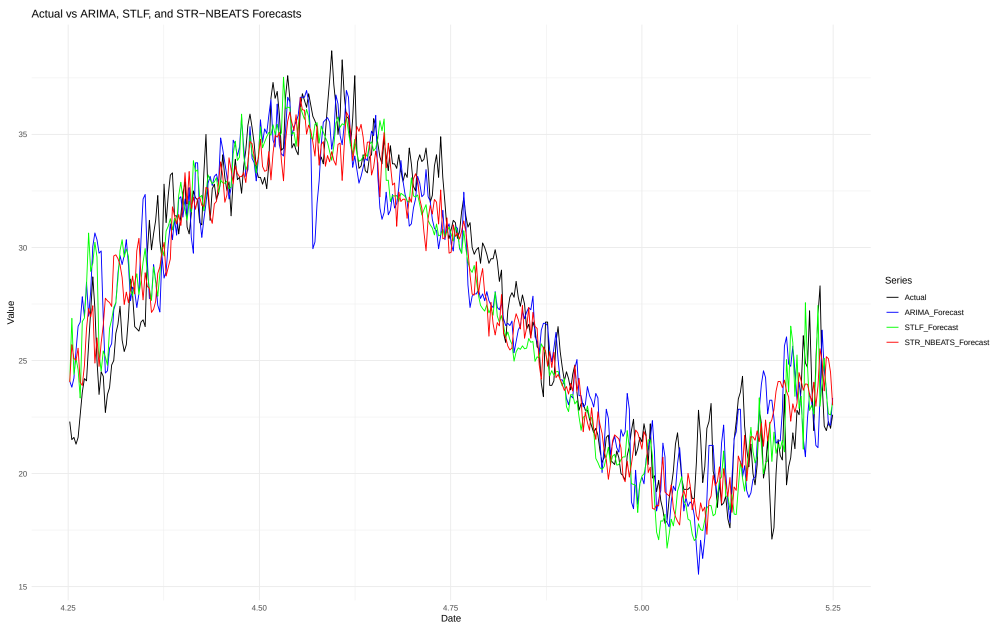

Figure 9: Comparison of actual temperature values (black) with ARIMA (blue), STL-ETS (green), and STR-NBEATS (red) forecasts over the test period. The STR-NBEATS forecasts closely align with the observed values, demonstrating improved accuracy and reduced forecasting errors compared to benchmark models.

6 Conclusion

In this paper, we introduced the STR-NBEATS hybrid framework, integrating Seasonal-Trend Decomposition using Regression with the Neural Basis Expansion Analysis architecture for enhanced temperature forecasting. Our results demonstrate that STR-NBEATS consistently outperforms traditional statistical models and hybrid decomposition approaches, as evidenced by lower RMSE, MAE, and MAPE values. By effectively capturing complex seasonal patterns and long-term trends, while addressing the limitations of classical and purely data-driven methods, the STR-NBEATS framework provides a more accurate, interpretable, and robust solution for environmental and climate-related forecasting tasks.

The improved performance of STR-NBEATS in handling non-stationary temperature data and its ability to deliver meaningful decomposed components underscores its potential for a wide range of applications, including climate risk assessment, agricultural planning, and renewable energy management. Future work may explore the integration of additional external covariates, more sophisticated decomposition techniques, and domain-specific model adaptations to further enhance forecasting capabilities and generalize to other types of time series data.

References

[1] Balyan, S., et al., "Seeding a Sustainable Future: Navigating the Digital Horizon of Smart Agriculture," *Sustainability (Switzerland)*, vol. 16, no. 2, art. no. 475, 2024. DOI: 10.3390/su16020475.

[2] Wang, H., et al., "Application of the Informer Long-Term Sequence Model in Agricultural Temperature Prediction," *Proceedings - 2023 China Automation Congress*, pp. 7122-7127, 2023. DOI: 10.1109/CAC59555.2023.10451131.

[3] Falloon, P., Betts, R., "Climate impacts on European agriculture and water management in the context of adaptation and mitigation," *Science of the Total Environment*, vol. 408, no. 23, pp. 5667-5687, 2010. DOI: 10.1016/j.scitotenv.2009.05.002.

[4] Ukhurebor, K. E., et al., "Climate condition monitoring and automated systems," *AI, Edge and IoT-based Smart Agriculture*, pp. 437-447, 2021. DOI: 10.1016/B978-0-12-823694-9.00031-1.

[5] Sivakumar, M. V. K., "Climate prediction and agriculture: Current status and future challenges," *Climate Research*, vol. 33, no. 1, pp. 3-17, 2006. DOI: 10.3354/cr033003.

[6] Chen, L., et al., "Artificial intelligence-based solutions for climate change: a review," *Environmental Chemistry Letters*, vol. 21, no. 5, pp. 2525-2557, 2023. DOI: 10.1007/s10311-023-01617-y.

[7] Singh, S., Goyal, M. K., "Enhancing climate resilience in businesses: The role of artificial intelligence," *Journal of Cleaner Production*, vol. 418, art. no. 138228, 2023. DOI: 10.1016/j.jclepro.2023.138228.

[8] Woodhouse, N., Musilek, P., "Computational acquisition of meteorological data for applications in electric power systems," *2020 IEEE Electric Power and Energy Conference, EPEC 2020*, 2020. DOI: 10.1109/EPEC48502.2020.9319917.

[9] Abdallah, W., et al., "A vector autoregressive methodology for short-term weather forecasting: tests for Lebanon," *SN Applied Sciences*, vol. 2, no. 9, 2020. DOI: 10.1007/s42452-020-03292-y.

[10] Niu, D., et al., "A Machine-Learning Approach Combining Wavelet Packet Denoising with Catboost for Weather Forecasting," *Atmosphere*, vol. 12, no. 12, 2021. DOI: 10.3390/atmos12121618.

[11] Elmaleh, M. S., "ARIMA Forecasting: Variables without a Cause," *Journal of Business Valuation and Economic Loss Analysis*, vol. 12, no. 1, 2017. DOI: 10.1515/jbvela-2016-0009.

[12] Lah, M. S. C., et al., "Fuzzy-Autoregressive Integrated Moving Average (F-ARIMA) Model to Improve Temperature Forecast," *Lecture Notes in Networks and Systems*, vol. 457, 2022. DOI: 10.1007/978-3-031-00828-3$_5$.

[13] Katušić, D., et al., "A comparison of data-driven methods in prediction of weather patterns in central Croatia," *Earth Science Informatics*, vol. 15, no. 2, 2022. DOI: 10.1007/s12145-022-00792-w.

[14] Zakari, R. Y., et al., "Time series forecasting techniques for climate trend prediction," *Reshaping Environmental Science Through Machine Learning and IoT*, 2024. DOI: 10.4018/979-8-3693-2351-9.ch0014.

[15] Djouama, I., et al., "ML-Based Weather Forecasting Models: A Comparative Study," *Lecture Notes in Networks and Systems*, vol. 1170, 2024. DOI: 10.1007/978-3-031-73344-4$_3$3.

[16] Elseidi, M., "Forecasting temperature data with complex seasonality using time series methods," *Modeling Earth Systems and Environment*, vol. 9, no. 2, 2023. DOI: 10.1007/s40808-022-01632-y.

[17] Yadav, H., Thakkar, A., "NOA-LSTM: An efficient LSTM cell architecture for time series forecasting," *Expert Systems with Applications*, vol. 238, 2024. DOI: 10.1016/j.eswa.2023.122333.

[18] Vo, T., "An Integrated Dual Attention with Convolutional LSTM for Short-Term Temperature Forecasting," *Cybernetics and Systems*, vol. 55, no. 2, 2024. DOI: 10.1080/01969722.2022.2122010.

[19] Khan, M. I., Maity, R., “Hybrid deep learning approach for multi-step-ahead prediction for daily maximum temperature and heatwaves,” *Theoretical and Applied Climatology*, vol. 149, no. 3-4, 2022. DOI: 10.1007/s00704-022-04103-7.

[20] Castro, R., et al., “STConvS2S: Spatiotemporal Convolutional Sequence to Sequence Network for weather forecasting,” *Neurocomputing*, vol. 426, 2021. DOI: 10.1016/j.neucom.2020.09.060.

[21] De Bézenac, E., Pajot, A., Gallinari, P., “Deep learning for physical processes: incorporating prior scientific knowledge,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2019, 2019. DOI: 10.1088/1742-5468/ab3195.

[22] Roy, D. S., “Forecasting the Air Temperature at a Weather Station Using Deep Neural Networks,” *Procedia Computer Science*, vol. 178, 2020. DOI: 10.1016/j.procs.2020.11.005.

[23] Chandrasekaran, R., Paramasivan, S. K., “A State-of-the-Art Review of Time Series Forecasting Using Deep Learning Approaches,” *International Journal on Recent and Innovation Trends in Computing and Communication*, vol. 10, no. 12, 2022. DOI: 10.17762/ijritcc.v10i12.5890.

[24] Zhang, X., Zhang, Q., Zhang, G., et al., “A novel hybrid data-driven model for daily land surface temperature forecasting using long short-term memory neural network based on ensemble empirical mode decomposition,” *International Journal of Environmental Research and Public Health*, vol. 15, no. 5, 2018. DOI: 10.3390/ijerph15051032.

[25] Prasad, R., Deo, R. C., Li, Y., Maraseni, T., “Soil moisture forecasting by a hybrid machine learning technique: ELM integrated with ensemble empirical mode decomposition,” *Geoderma*, vol. 330, pp. 136–161, 2018. DOI: 10.1016/j.geoderma.2018.05.035.

[26] Apaydin, H., Sattari, M. T., Falsafian, K., Prasad, R., “Artificial intelligence modelling integrated with Singular Spectral analysis and Seasonal-Trend decomposition using Loess approaches for streamflow predictions,” *Journal of Hydrology*, vol. 600, 2021. DOI: 10.1016/j.jhydrol.2021.126506.

[27] Zhang, M., Han, G., Wu, X., et al., “SST Forecast Skills Based on Hybrid Deep Learning Models: With Applications to the South China Sea,” *Remote Sensing*, vol. 16, no. 6, 2024. DOI: 10.3390/rs16061034.

[28] Araújo, R. A., Mattos Neto, P. S. G., Nedjah, N., Soares, S. C. B., “An error correction system for sea surface temperature prediction,” *Neural Computing and Applications*, vol. 35, no. 16, pp. 11681–11699, 2023. DOI: 10.1007/s00521-023-08311-8.

[29] Liu, X., Zhang, Q., “Combining Seasonal and Trend Decomposition Using LOESS With a Gated Recurrent Unit for Climate Time Series Forecasting,” *IEEE Access*, vol. 12, pp. 85275–85290, 2024. DOI: 10.1109/ACCESS.2024.3415349.

[30] Zhang, K., Huo, X., Shao, K., “Temperature Time Series Prediction Model Based on Time Series Decomposition and Bi-LSTM Network,” *Mathematics*, vol. 11, no. 9, 2023. DOI: 10.3390/math11092060.

[31] Ji, R., Shi, S., Liu, Z., Wu, Z., “Decomposition-Based Multi-Step Forecasting Model for the Environmental Variables of Rabbit Houses,” *Animals*, vol. 13, no. 3, 2023. DOI: 10.3390/ani13030546.

[32] Elseidi, Mohammed. ”MSTL-NNAR: a new hybrid model of machine learning and time series decomposition for wind speed forecasting.” *Stochastic Environmental Research and Risk Assessment*, Springer, 2024, pp. 1–20.

[33] Elseidi, Mohammed. ”A hybrid Facebook Prophet-ARIMA framework for forecasting high-frequency temperature data.” *Modeling Earth Systems and Environment*, vol. 10, no. 2, Springer, 2024, pp. 1855–1867.

[34] Caporin, Massimiliano, and Mohammed Elseidi. ”Quantile regression-based seasonal adjustment.” *International Journal of Computational Economics and Econometrics*, vol. 13, no. 3, Inderscience Publishers (IEL), 2023, pp. 270–304.

[35] Dokumentov, A., & Hyndman, R. J. (2022). STR: Seasonal-trend decomposition using regression. *INFORMS Journal on Data Science*, *1*(1), 50–62. INFORMS.

[36] Oreshkin, Boris N., Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. "N-BEATS: Neural basis expansion analysis for interpretable time series forecasting." *arXiv preprint arXiv:1905.10437*, 2019.

[37] Van Belle, J., Crevits, R., & Verbeke, W. (2023). Improving forecast stability using deep learning. *International Journal of Forecasting*, *39*(3), 1333–1350. https://doi.org/10.1016/j.ijforecast.2022.06.007.

[38] Karamchandani, A., Mozo, A., Vakaruk, S., Gómez-Canaval, S., Sierra-García, J. E., & Pastor, A. (2023). Using N-BEATS ensembles to predict automated guided vehicle deviation. *Applied Intelligence*, *53*(21), 26139–26204. https://doi.org/10.1007/s10489-023-04820-0.

[39] Sbrana, A., Debiaso Rossi, A. L., & Coelho Naldi, M. (2020). N-BEATS-RNN: Deep learning for time series forecasting. *Proceedings of the 19th IEEE International Conference on Machine Learning and Applications (ICMLA 2020)*, 765–768. https://doi.org/10.1109/ICMLA51294.2020.00125.

[40] Alali, Z. H., & Horne, R. N. (2023). A comparative study of deep learning models and traditional methods in forecasting oil production in the Volve field. *Proceedings of the SPE Annual Technical Conference and Exhibition*, October 2023. https://doi.org/10.2118/214881-MS.

[41] Lee, J., Jeon, M., Kang, M., & Park, K. (2024). Feature-aligned N-BEATS with Sinkhorn divergence. *Proceedings of the 12th International Conference on Learning Representations (ICLR 2024)*. Available at: https://www.scopus.com/inward/record.uri?eid=2-s2.0-85200581724.

[42] Aiwansedo, K., Bosche, J., Badreddine, W., Kermia, M. H., & Djadane, O. (2024). CNN-N-BEATS: Novel hybrid model for time-series forecasting. *Communications in Computer and Information Science*, *2171 CCIS*, 38–57. https://doi.org/10.1007/978-3-031-66694-0_3.

[43] Karanikola, A., Liapis, C. M., & Kotsiantis, S. (2022). A comparison of contemporary methods on univariate time series forecasting. *Learning and Analytics in Intelligent Systems*, *23*, 143–168. https://doi.org/10.1007/978-3-030-76794-5_8.

[44] Makika, H., Romano, J. M. T., & Ballini, R. (2023). Deep learning techniques for time series forecasting: An application for exchange rates. *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, September 2023. https://doi.org/10.1109/MLSP55844.2023.10285930.

[45] Arnab, A. A., Shuvro, A. A., Ma, K., & Leung, H. (2023). A deep learning approach for a QoS prediction system in cellular networks. *Proceedings of the 2023 IEEE World Forum on Internet of Things (WF-IoT 2023)*. https://doi.org/10.1109/WF-IoT58464.2023.10539507.

[46] Hyndman, R. J. *Forecasting: Principles and Practice*. OTexts, 2018.