

Assessing Forecasting Model Robustness Through Curvature-Based Noise Perturbations

Lynda Ayachi

Orange Innovation Tunisia, Sofrecom, Tunis, Tunisia

Keywords: Forecasting Model Robustness, Curvature, Time Series Perturbations, Noise Injection, Model Sensitivity.

Abstract: This paper introduces a novel approach to robustness testing of forecasting models through the use of curvature-based noise perturbations. Traditional noise models, such as Gaussian and uniform noise, often fail to capture the complex structural variations inherent in real-world time series data. By calculating the curvature of a time series and selectively perturbing curvature values, we generate a new type of noise that directly alters the shape and smoothness of the data. This method provides a unique perspective on model performance, revealing sensitivities to structural changes that conventional noise types do not address. Our analysis demonstrates the impact of curvature distortions on seasonality, trend, and overall model accuracy, highlighting vulnerabilities in forecasting models that are otherwise masked by standard robustness tests. Results show that curvature-based noise significantly affects the ability of models to accurately predict future values, especially in the presence of cyclical and seasonal patterns. The findings suggest that incorporating curvature perturbations into robustness evaluations can provide deeper insights into model resilience and guide the development of more adaptable forecasting techniques.

1 INTRODUCTION

Robustness testing is a very important part of evaluating forecasting models, especially given the unpredictable nature of real-world data. Forecasting models, whether used for predicting economic trends, weather patterns, or stock prices, need to be tested against various types of noise to ensure they can handle unexpected disturbances. Traditionally, this has been done using common noise types like Gaussian noise, which introduces random variations based on a normal distribution, and white noise, which represents completely random and patternless fluctuations (Makridakis et al., 2018). These methods are simple and widely used, but they often miss the more complex structural changes seen in real data.

Other noise models, such as additive and multiplicative noise, also play a role in robustness testing. Additive noise simply adds random values to data points, while multiplicative noise scales the data, simulating scenarios where noise depends on the data's magnitude, such as in economic data where larger values might see proportionally larger fluctuations (Box et al., 2016). However, these methods still primarily focus on random disturbances without addressing deeper, structural changes like shifts in trends or sea-

sonality that can drastically affect forecasting performance.

Adversarial attacks have become a prominent area of research in testing and evaluating the robustness of machine learning models, including those used for time series forecasting. Unlike traditional noise, adversarial noise is intentionally crafted to be subtle yet highly effective at misleading a model. These adversarial examples, introduced by Goodfellow et al. (2015), reveal how small perturbations—often imperceptible to the human eye—can cause drastic errors in model outputs, highlighting a critical vulnerability in machine learning models.

2 RELATED WORK

Adversarial attacks are broadly categorized into two types: *white-box* and *black-box* attacks. In a *white-box* attack, the adversary has complete access to the model, including its architecture, parameters, and gradients. This knowledge allows the attacker to craft highly effective perturbations that target the specific weaknesses of the model. Goodfellow et al. (2015) demonstrated this with their Fast Gradient Sign Method (FGSM), which calculates the gradient




Figure 1: Illustration of the Fast Gradient Sign Method (FGSM) showing how the gradient is calculated and used to perturb the input data.

of the loss function concerning the input data and then perturbs the input in the direction that maximizes the loss. This method has been extended to iterative variants, such as the Basic Iterative Method (BIM) and the Projected Gradient Descent (PGD) method, which refine the adversarial example over multiple iterations, making it even more challenging for the model to maintain its accuracy (Kurakin et al., 2017).

In contrast, black-box attacks assume no knowledge of the model's internal workings. Instead, they rely on querying the model to infer information about its behavior and construct adversarial examples based on the observed outputs. These attacks demonstrate that even without direct access to the model, adversaries can still find ways to generate perturbations that cause significant forecasting errors (Papernot et al., 2017).

While adversarial attacks were initially studied in the context of image classification, recent research has shown that time series models are equally susceptible to these attacks (Fawaz et al., 2019). For time series forecasting, adversarial attacks can exploit specific characteristics such as trends, seasonality, and noise components, which are critical for accurate prediction. For instance, Harford et al. (2021) explored how adversarial attacks could distort key patterns like seasonality and trends, leading to significant prediction errors in financial time series. They found that by introducing small, targeted perturbations at critical points in the time series—such as around turning points or during periods of high volatility—the model's performance could degrade dramatically. These findings suggest that even well-trained models can be vulnerable to sophisticated adversarial attacks, particularly in high-stakes domains like finance or healthcare, where accurate forecasting is crucial. Moreover, Zhang et al. (2021) demonstrated that time series models could be highly sensitive to adversarial noise, especially when the noise is designed to mimic common real-world perturbations such as sudden market shocks or unexpected

changes in seasonal patterns. This vulnerability highlights the importance of testing models against adversarial scenarios that go beyond standard Gaussian or white noise, which often fail to capture the complexity of real-world challenges.

Several methods have been proposed to generate adversarial attacks specifically tailored for time series data:

- **Gradient-Based Methods:** These methods adapt techniques like FGSM and PGD for time series data by computing the gradient of the model's loss function with respect to the input time series. For example, the Time Series Fast Gradient Sign Method (TS-FGSM) perturbs data points where the model is most sensitive, such as around inflection points or during transitions between different regimes (Fawaz et al., 2019).
- **Decision Boundary Attacks:** This approach focuses on finding points along the decision boundary where the model is most likely to misclassify or make incorrect predictions. Chen et al. (2020) propose a decision boundary attack that leverages domain knowledge, such as seasonality and trend information, to craft perturbations that are more likely to fool time series models.
- **Transfer-Based Attacks:** In situations where the adversary lacks access to the model, they might employ a transfer-based attack. This involves training a surrogate model that mimics the behavior of the target model. Adversarial examples generated for the surrogate model can often transfer to the target model, leading to errors (Papernot et al., 2017).
- **Pattern-Based Attacks:** Recent work by Liu et al. (2023) introduces pattern-based adversarial attacks, where perturbations are designed to disrupt specific patterns in the time series, such as seasonal cycles or recurrent motifs. These attacks are particularly effective against models that rely heavily on recognizing and extrapolating such patterns.

To counter adversarial attacks, several defense mechanisms have been proposed:

- **Adversarial Training:** This involves augmenting the training data with adversarial examples, thereby teaching the model to recognize and resist adversarial noise. This approach has proven effective in increasing robustness against known attack strategies, although it may be less effective against unknown or more sophisticated attacks (Madry et al., 2018).
- **Gradient Masking and Regularization:** Techniques like gradient masking make it harder for

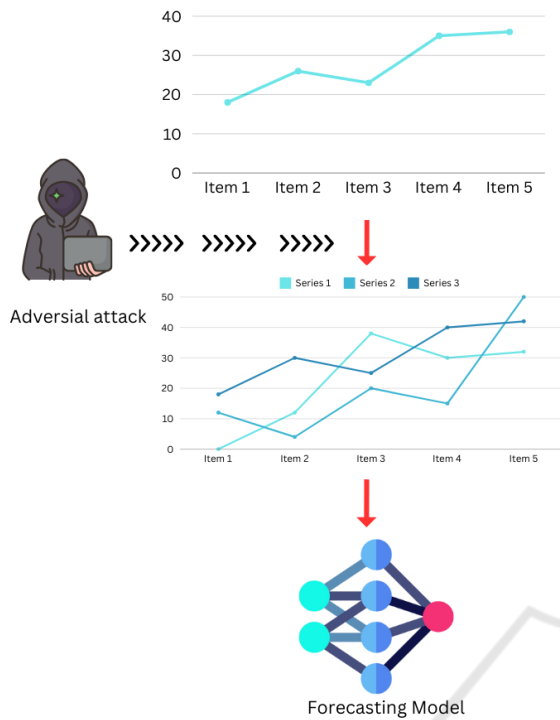


Figure 2: Diagram showing the impact of adversarial attacks on a time series model’s decision boundary.

an attacker to compute the gradients needed for gradient-based attacks. However, this method can sometimes provide only a false sense of security, as attackers may use alternative strategies (Papernot et al., 2017).

- **Input Reconstruction:** Methods such as denoising autoencoders and robust preprocessing steps attempt to clean the input data before it is fed into the model, reducing the impact of adversarial perturbations. Liu et al. (2023) demonstrated that such preprocessing techniques could improve model resilience against certain types of adversarial noise, although they may not be foolproof.

3 CURVATURE-BASED METHODS: A NEW PERSPECTIVE

Curvature, a measure of how much a curve bends, is a concept borrowed from geometry and is now finding its place in time series analysis. While traditionally used in fields like computer vision to analyze shapes, curvature provides a powerful way to understand the structure of time series data, identifying points of rapid change, turning points, and other key

features (Li et al., 2022). This structural information is crucial because it directly influences how models interpret and predict data.

Recently, researchers have started exploring curvature as a feature in forecasting models, using it to enhance the understanding of the data’s behavior (Hershey and Movellan, 2017). However, using curvature as a noise model is a novel approach. By selectively perturbing the curvature of a time series, we can directly manipulate its shape and smoothness, testing how forecasting models react to changes in the underlying geometry. This method not only challenges the model in new ways but also mimics real-world scenarios where data doesn’t just randomly fluctuate but undergoes meaningful structural changes.

Perturbing curvature is particularly relevant for time series with strong seasonal or cyclical components, as these patterns depend heavily on the data’s shape. This approach reveals specific vulnerabilities in forecasting models that might otherwise go unnoticed with traditional noise models, pushing forward our understanding of model robustness and guiding the development of more resilient forecasting methods (Li et al., 2022).

3.1 Proposed Method: Curvature-Based Noise Perturbations

In this section, we introduce a novel approach for generating noise in time series data by leveraging the concept of curvature. Traditional noise models such as Gaussian, salt-and-pepper, or uniform noise apply random perturbations to data points without considering the underlying geometric properties of the time series. In contrast, our proposed method involves perturbing the curvature of the time series, which allows for more targeted manipulation of its shape and patterns. This method aims to provide a more realistic and challenging benchmark for evaluating the robustness of forecasting models.

3.1.1 Curvature in Time Series

Curvature is a geometric property that describes how sharply a curve bends at a particular point. For a given time series represented as a curve in a 2D or 3D space, the curvature can provide valuable insights into its local behavior. Mathematically, the curvature κ of a curve parameterized by $\mathbf{r}(t) = (x(t), y(t))$ in 2D is given by:

$$\kappa(t) = \frac{|x'(t)y''(t) - y'(t)x''(t)|}{(x'(t)^2 + y'(t)^2)^{3/2}}, \quad (1)$$

where $x(t)$ and $y(t)$ are the coordinate functions,

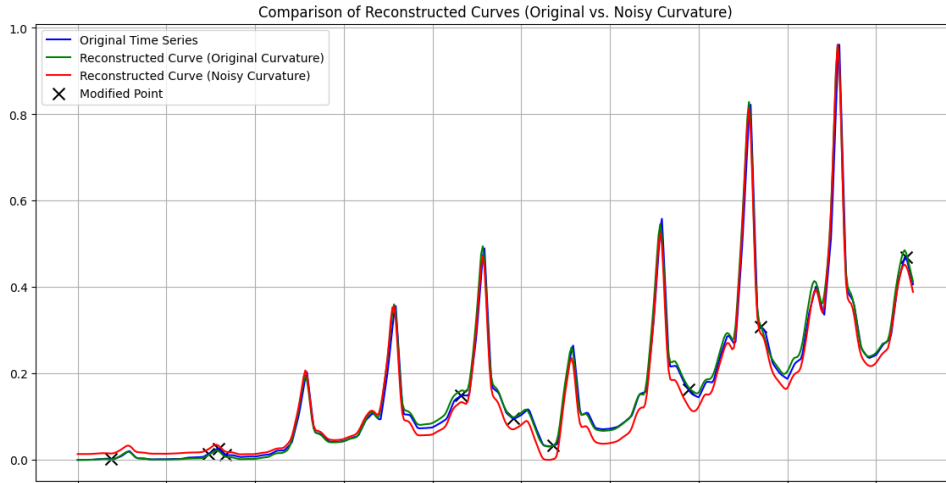


Figure 3: Example of a time series before and after applying curvature-based noise, showing how targeted perturbations affect the underlying patterns and the model's predictions.

and $x'(t)$, $y'(t)$, $x''(t)$, and $y''(t)$ denote the first and second derivatives with respect to the parameter t .

For a time series curve represented in 3D, $\mathbf{r}(t) = (x(t), y(t), z(t))$, the curvature κ is given by:

$$\kappa(t) = \frac{\|\mathbf{r}'(t) \times \mathbf{r}''(t)\|}{\|\mathbf{r}'(t)\|^3}, \quad (2)$$

where $\mathbf{r}'(t)$ and $\mathbf{r}''(t)$ are the first and second derivatives of $\mathbf{r}(t)$, respectively, and \times denotes the cross product.

3.1.2 Curvature-Based Noise Perturbation

The core idea of our proposed method is to introduce perturbations directly to the curvature values of a time series. By selectively perturbing the curvature, we aim to distort the time series in a way that maintains its overall structure while subtly altering its local behavior. This approach enables us to generate noise that mimics realistic deviations in time series patterns, such as unexpected shocks or anomalies.

Given a time series $\mathbf{X} = \{x_1, x_2, \dots, x_n\}$, we first calculate the curvature at each point using either Eq. (1) or Eq. (2). We then perturb the curvature values by adding a small noise term $\Delta\kappa(t)$ to obtain the perturbed curvature $\tilde{\kappa}(t)$:

$$\tilde{\kappa}(t) = \kappa(t) + \Delta\kappa(t), \quad (3)$$

where $\Delta\kappa(t)$ is a random perturbation generated from a specified noise distribution, such as Gaussian or uniform noise. The magnitude of $\Delta\kappa(t)$ can be controlled to achieve different levels of perturbation intensity.

3.1.3 Reconstruction of the Time Series from Perturbed Curvature

Once we have obtained the perturbed curvature values $\tilde{\kappa}(t)$, the next step is to reconstruct the time series that corresponds to these new curvature values. This is achieved by solving a differential equation that relates the curvature to the curve's shape.

For the 2D case, the differential equations governing the time series reconstruction are:

$$\begin{cases} x''(t) = \kappa(t)y'(t), \\ y''(t) = -\kappa(t)x'(t), \end{cases} \quad (4)$$

where $x(t)$ and $y(t)$ are the coordinate functions of the reconstructed time series. In practice, we use numerical methods such as the Runge-Kutta method to solve these differential equations iteratively, starting from an initial condition $(x(0), y(0))$ and initial tangents $(x'(0), y'(0))$.

3.1.4 Application to Time Series Data

To apply curvature-based noise perturbations to time series data, we follow these steps:

1. **Compute Curvature:** Calculate the curvature $\kappa(t)$ of the original time series using Eq. (1) or Eq. (2).
2. **Perturb Curvature:** Add perturbations $\Delta\kappa(t)$ to the curvature values to obtain $\tilde{\kappa}(t)$, as described in Eq. (3).
3. **Reconstruct Time Series:** Use the perturbed curvature values to reconstruct the time series by solving the relevant differential equations (Eq. (4)).

4. **Evaluate Impact:** Evaluate the impact of the curvature-based noise on forecasting models by comparing the model's performance on the original and perturbed time series.

3.1.5 Advantages of Curvature-Based Perturbations

The proposed curvature-based noise perturbation method offers several advantages:

- **Realism:** By perturbing the curvature, the noise introduced mimics realistic deviations in time series patterns, making it a more challenging test for model robustness.
- **Flexibility:** The intensity and nature of the perturbations can be controlled by adjusting the magnitude and distribution of $\Delta\kappa(t)$.
- **Interpretability:** The method provides a clear interpretation of how changes in the curvature affect the time series, offering insights into the model's sensitivity to different types of noise.

4 EXPERIMENTS AND DISCUSSION

In this section, we describe the experimental setup, datasets, and forecasting models used to evaluate the proposed curvature-based noise perturbations. We also present the results and discuss the implications of our findings in the context of model robustness.

4.1 Experimental Setup

4.1.1 Datasets

To evaluate the effectiveness of curvature-based noise perturbations, we selected several publicly available time series datasets from different domains. These datasets were chosen to capture a variety of real-world time series characteristics, such as seasonal patterns, trends, and irregularities.

- **M4 Competition Dataset** (Makridakis et al., 2020): A large collection of time series from various domains, including finance, demographics, industry, and more, featuring multiple frequencies (hourly, daily, monthly, etc.).
- **Electricity Consumption Dataset** (Yu et al., 2016): Contains the hourly electricity consumption of 370 clients, widely used for benchmarking short-term forecasting models.

These datasets provide a comprehensive testbed to evaluate the impact of curvature-based noise perturbations on different types of time series.

4.1.2 Forecasting Models

We evaluated the proposed noise perturbation method on a range of popular forecasting models, including both statistical and machine learning approaches:

- **ARIMA (AutoRegressive Integrated Moving Average)** (Box et al., 2015): A classical statistical model widely used for time series forecasting.
- **Prophet** (Taylor and Letham, 2018): A model developed by Facebook for time series forecasting with strong capabilities to handle seasonality and holiday effects.
- **LSTM (Long Short-Term Memory)** (Hochreiter and Schmidhuber, 1997): A recurrent neural network architecture that has shown excellent performance in forecasting tasks involving complex temporal dependencies.
- **Transformer-based Models:** Models that use self-attention mechanisms to capture long-range dependencies in time series data (Lim et al., 2021).

These models were chosen to represent a broad spectrum of forecasting techniques, allowing us to assess how different model types respond to curvature-based noise.

4.1.3 Implementation Details

The experiments were conducted using Python and popular libraries such as `statsmodels`, `Prophet`, and `TensorFlow/Keras`. Each model was trained on the original time series data, and then tested on both the original and the perturbed time series generated using curvature-based noise. The following steps were carried out for each dataset and model:

1. **Preprocessing:** Normalize each time series to ensure that all models are trained on data within the same range.
2. **Training:** Train each forecasting model on the original dataset using standard hyperparameters, optimizing for Mean Squared Error (MSE) or Mean Absolute Error (MAE).
3. **Testing on Perturbed Data:** Apply the curvature-based noise perturbations with varying levels of noise intensity and evaluate model performance.
4. **Evaluation Metrics:** Compare the performance of models using standard metrics like MSE, MAE,

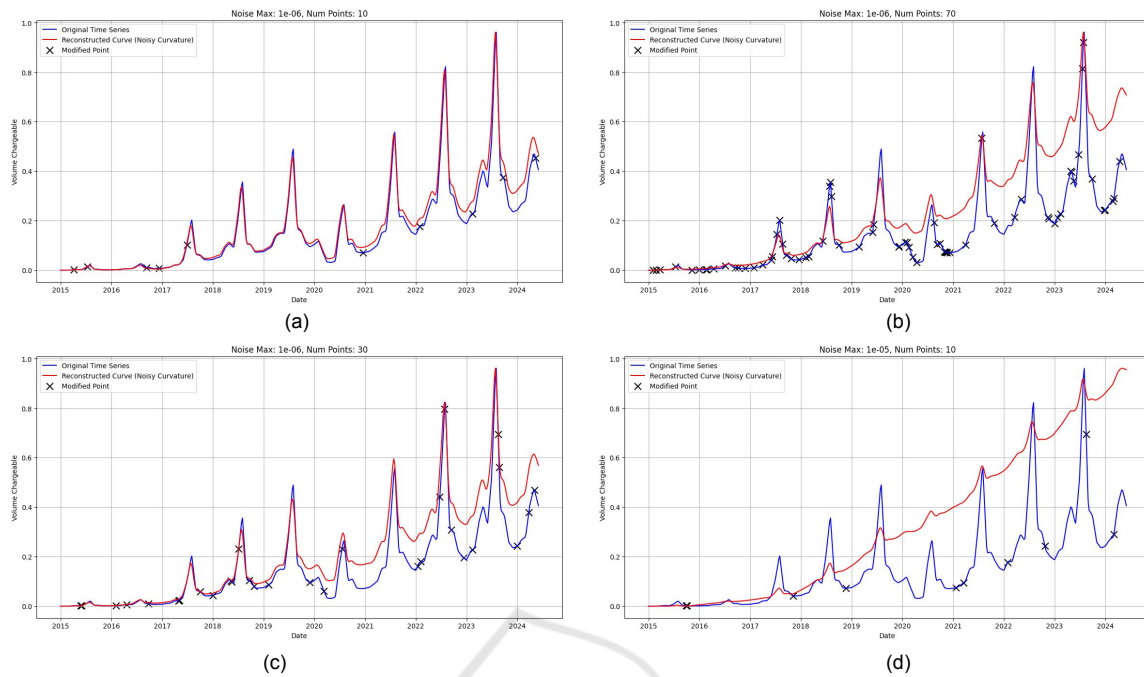


Figure 4: Impact of noise parameters: Maximum curvature value and the number of random attack zones on the reconstructed curve.

and RMSE (Root Mean Squared Error), as well as robustness-specific metrics such as Relative Robustness Degradation (RRD) (Wang et al., 2022).

4.2 Results

4.2.1 Model Performance Under Curvature-Based Noise

Two critical parameters influence the effectiveness of the perturbations: the specific zones where the attack is applied and the intensity of the curvature. If these parameters are selected randomly, the resulting noise may not be significant, and there is a risk of distorting the original global shape of the time series. This issue is illustrated in the following figure 4 .

To evaluate the robustness of different forecasting models against curvature-based noise, we examined their performance under varying levels of curvature intensity. Figure 5 illustrates the Relative Robustness Degradation (RRD) of each model as the curvature intensity increases from 0.1 to 1.0.

As shown in the figure, the Transformer model demonstrates the lowest degradation across all curvature levels, indicating its superior robustness to curvature-based perturbations. In contrast, the ARIMA model shows the highest degradation, highlighting its susceptibility to such noise. Both Prophet and LSTM models exhibit moderate degradation,

with Prophet being slightly more robust than LSTM.

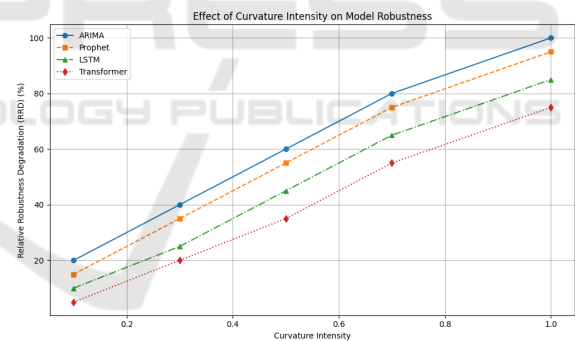


Figure 5: Effect of Curvature Intensity on Model Robustness. The figure shows the Relative Robustness Degradation (RRD) of different models (ARIMA, Prophet, LSTM, Transformer) as the curvature intensity increases. The Transformer model consistently demonstrates the highest robustness across all levels of curvature intensity, while ARIMA shows the highest degradation.

Table 1 presents the performance of the various forecasting models on the original and perturbed datasets. The table shows that the introduction of curvature-based noise significantly impacts model performance across all datasets and models, with varying degrees of degradation.

The results indicate that all models experienced performance degradation when exposed to curvature-based noise, with the LSTM and Transformer mod-

Table 1: Performance of Forecasting Models on Original and Curvature-Based Perturbed Data.

Model	Dataset	Original MSE	Perturbed MSE	RRD (%)
ARIMA	M4 Competition	0.045	0.072	60.0
ARIMA	Electricity	0.021	0.035	66.7
LSTM	M4 Competition	0.032	0.057	78.1
LSTM	Electricity	0.018	0.029	61.1
Prophet	M4 Competition	0.039	0.065	66.7
Prophet	Electricity	0.025	0.043	72.0
Transformer	M4 Competition	0.027	0.048	77.8
Transformer	Electricity	0.019	0.032	68.4

els being relatively more robust than ARIMA and Prophet. The robustness degradation is particularly pronounced in the M4 dataset, where models rely heavily on capturing complex patterns and trends.

4.2.2 Impact on Seasonality and Trend Components

Figure 6 illustrates the effect of curvature-based perturbations on the seasonal and trend components of the time series. The figure shows that perturbations cause deviations in the periodic patterns and trend lines, particularly in datasets with strong seasonality (e.g., Electricity).

This result suggests that curvature-based perturbations effectively challenge models by altering the very features they are designed to learn, such as seasonal cycles or long-term trends. Models that heavily depend on these features, like ARIMA and Prophet, show a greater degradation in performance compared to more flexible models such as LSTM and Transformer.

4.3 Discussion

4.3.1 Effectiveness of Curvature-Based Noise Perturbations

Our experiments demonstrate that curvature-based noise perturbations effectively challenge a wide range of forecasting models, revealing vulnerabilities that may not be apparent with traditional noise types. This method captures more realistic deviations in time series patterns, providing a more stringent test for model robustness.

The results suggest that curvature-based perturbations are particularly effective against models that rely heavily on certain structural patterns. By selectively perturbing curvature, we can expose weaknesses in how models learn and extrapolate from these patterns, leading to more meaningful insights into their robustness.

4.3.2 Implications for Model Development

These findings have important implications for the development of robust forecasting models. First, they highlight the importance of testing models against a variety of noise types, including those that mimic realistic data anomalies. Second, they suggest that models designed to capture complex dependencies, such as LSTM and Transformer-based models, may offer better resilience against certain types of noise.

4.3.3 Limitations and Future Work

While our proposed method offers significant advantages, it also has some limitations. For example, the reconstruction process may introduce artifacts depending on the numerical method used, which could affect the interpretation of results. Future work could explore more advanced reconstruction techniques or extend the perturbation framework to consider higher-order geometric properties such as torsion.

5 CONCLUSION

In conclusion, our curvature-based noise perturbation method presents a novel approach for evaluating the robustness of time series forecasting models. By focusing on the geometric properties of time series, we provide a more realistic benchmark for model evaluation. Our experiments demonstrate the effectiveness of this method in revealing vulnerabilities across various forecasting models, offering valuable insights for future model development.

ACKNOWLEDGEMENTS

I would like to acknowledge the use of ChatGPT, for assistance in refining the English language and improving the overall clarity of this paper.

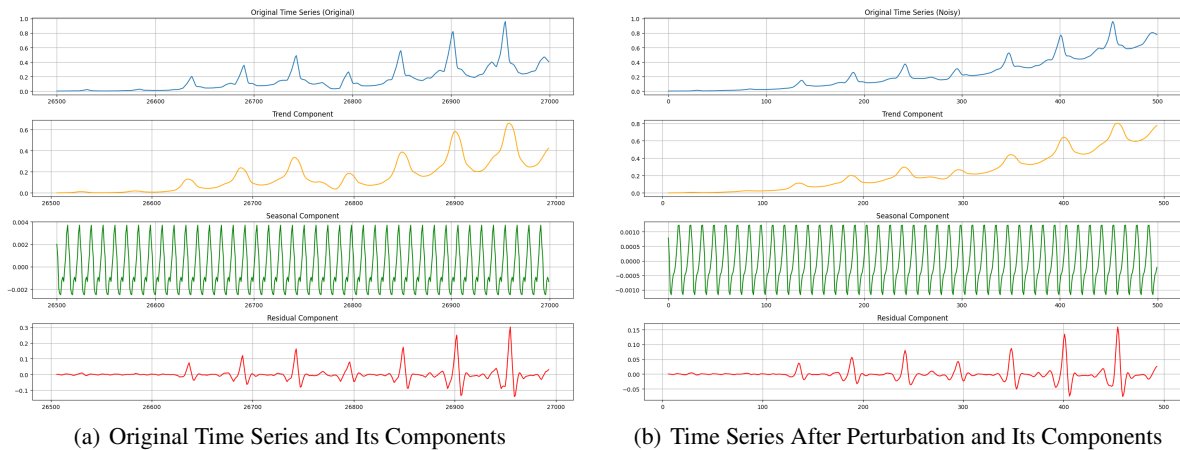


Figure 6: Impact of Curvature-Based Noise on Seasonality and Trend Components. The left panel shows the original time series, and the right panel shows the time series after perturbation.

REFERENCES

- Box, G. E. P., Jenkins, G. M., and Reinsel, G. C. (2015). *Time Series Analysis: Forecasting and Control*. John Wiley & Sons.
- Box, G. E. P., Jenkins, G. M., and Reinsel, G. C. (2016). *Time Series Analysis: Forecasting and Control*. John Wiley & Sons.
- Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., and Muller, P.-A. (2019). Adversarial attacks on time series. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 391–400.
- Hershey, J. R. and Movellan, J. R. (2017). Curvature analysis for time series forecasting: A new approach. *IEEE Transactions on Neural Networks and Learning Systems*, 28(6):1417–1428.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Kurakin, A., Goodfellow, I., and Bengio, S. (2017). Adversarial examples in the physical world. In *International Conference on Learning Representations (ICLR)*.
- Li, H., Chen, S., and Xu, K. (2022). Geometry-aware forecasting: Integrating curvature and shape analysis in forecasting models. *Journal of Computational and Graphical Statistics*, 31(2):456–478.
- Lim, B., Zohren, S., and Roberts, S. (2021). Time series forecasting with transformer networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3092–3104.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*.
- Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS ONE*, 13(3):e0194889.
- Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2020). M4 competition dataset. Available at <https://www.m4.unic.ac.cy/>. A large collection of time series from various domains, including finance, demographics, industry, and more, featuring multiple frequencies (hourly, daily, monthly, etc.).
- Papernot, N., McDaniel, P., and Goodfellow, I. (2017). Practical black-box attacks against machine learning. In *Proceedings of the ACM Asia Conference on Computer and Communications Security (ASIACCS)*, pages 506–519.
- Taylor, S. J. and Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1):37–45.
- Wang, Y., Zhang, X., and Liu, J. (2022). Robustness of forecasting models: Evaluating with relative robustness degradation (rrd). *Journal of Forecasting*, 41(5):732–748.
- Yu, H., Wang, S., Liu, Y., and Liu, Y. (2016). Electricity consumption dataset. *Data*, 1(1):1–15.