

# Univariate Time Series Data and Model Card

Generated by [Cardtale](#) - Automated Model and Data Card Generator

This report provides an automated, comprehensive analysis of univariate time series data. Generated by Cardtale, it explores basic aspects and potential challenges in your data to support informed decision-making and modeling choices.

**Generated:** 2025-01-02 17:07

**Series Name:** series\_0

## Table of Contents

- 1 Data Overview**  
Time series fundamental characteristics and statistical properties
- 2 Trend**  
Long-term time series growth and dynamics. Analysis of level stabilization methods.
- 3 Seasonality**  
Analysing recurring patterns in the time series. Assessing the impact of different seasonality modeling strategies
- 4 Variance**  
Exploring the variability of values over time. Assessing the impact of variance stabilization methods
- 5 Change Detection**  
Change detection in the time series distribution

## Data Overview

This section examines the core characteristics and statistical properties of the time series. Understanding these attributes is important for assessing data quality and gaining a preliminary context. We explore the temporal structure, summary statistics, and distribution patterns to create a baseline understanding of your data.

## Time Series Plot

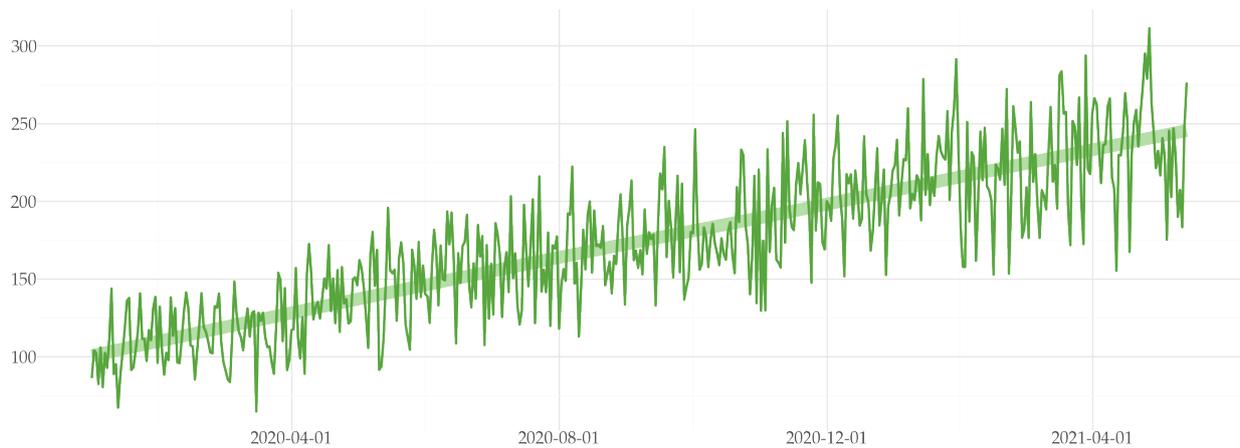


Figure 1: Time series line plot.

- A total of 500 observations spanning from 2020-01-01 to 2021-05-14. These are collected with a daily sampling frequency.
- The data ranges from a minimum of 64.58 to a maximum of 311.28, starting in 86.14789785203725 and ending in 276.4308597125188 during the observed period. The average growth percentage per observation is 2.46% (median equal to -0.41%), with an average value of 173.56. There are no missing values in the time series.

## Trend, Seasonality, and Residuals

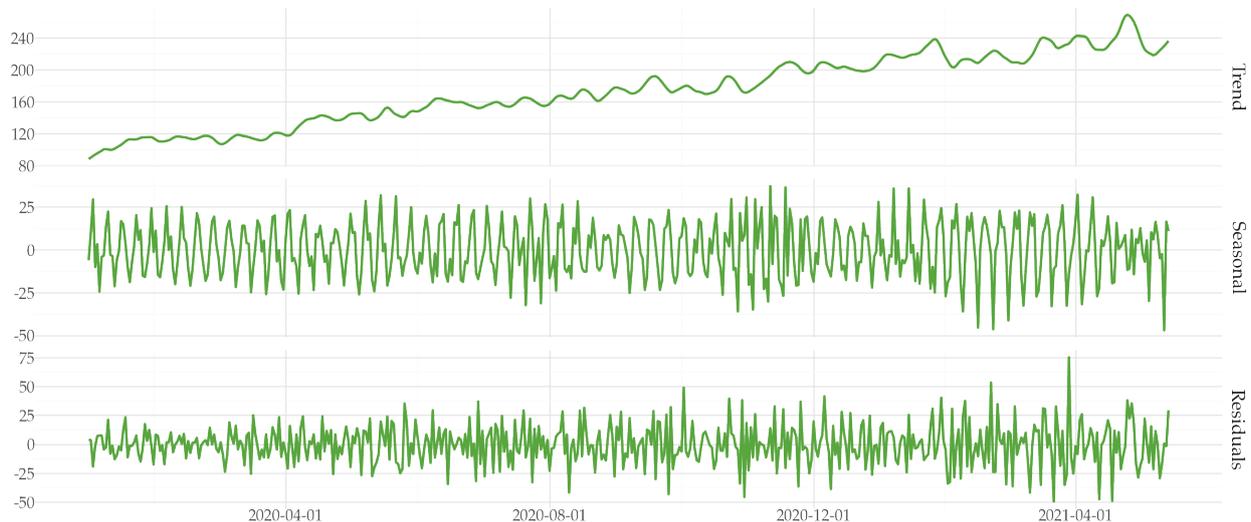


Figure 2: Seasonal, Trend, and Residuals components after decomposition on a daily frequency using the STL (Season-Trend decomposition using LOESS) method.

- The trend strength is 0.86 (ranges from 0 to 1). All hypothesis tests carried out (KPSS, Augmented Dickey-Fuller, and Philips-Perron) indicate that the time series is stationary in trend or level.

- The seasonal strength is 0.5 (ranges from 0 to 1). Tests for weekly seasonal stationarity show mixed results: the OCSB test indicates presence of a seasonal unit root, while the Wang-Smith-Hyndman test suggests stationarity.
- The STL decomposition residuals show balanced behavior: 48.8% of residuals are positive and 51.2% negative. The average magnitude of positive residuals is 13.857 compared to -13.232 for negative residuals. In terms of auto-correlation structure, the residuals show significant temporal dependency in some of the first 7 lags according to the Ljung-Box test. This suggests that the decomposition method is missing some systematic patterns.

## Auto-Correlation

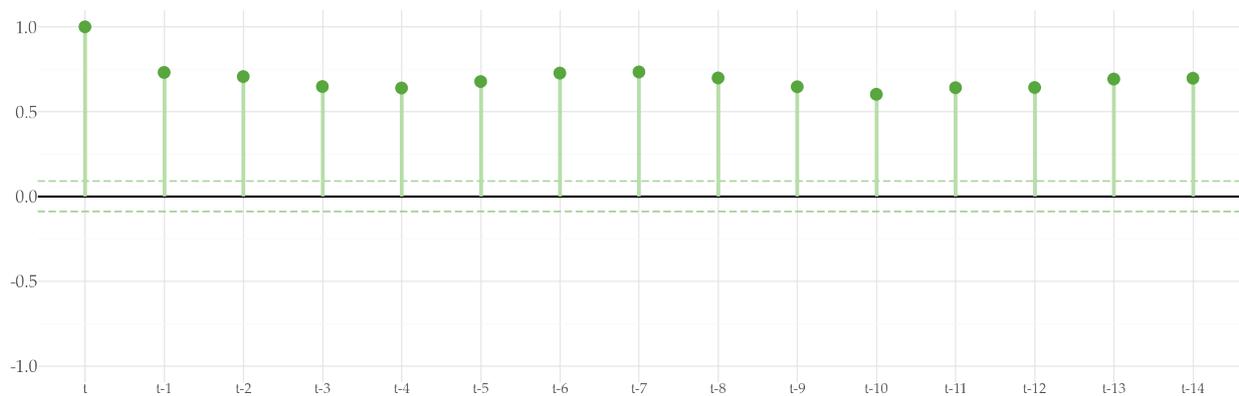


Figure 3: Auto-correlation plot up to 14 lags.

- The following lags show significant autocorrelation: t-1, t-2, t-3, t-4, t-5, t-6, t-7, t-8, t-9, t-10, t-11, t-12, t-13, and t-14. The autocorrelation is positive for all lags with a significant value.
- All lags relative to the seasonal period (t-7 and t-14) show a significant autocorrelation.

## Partial Auto-Correlation

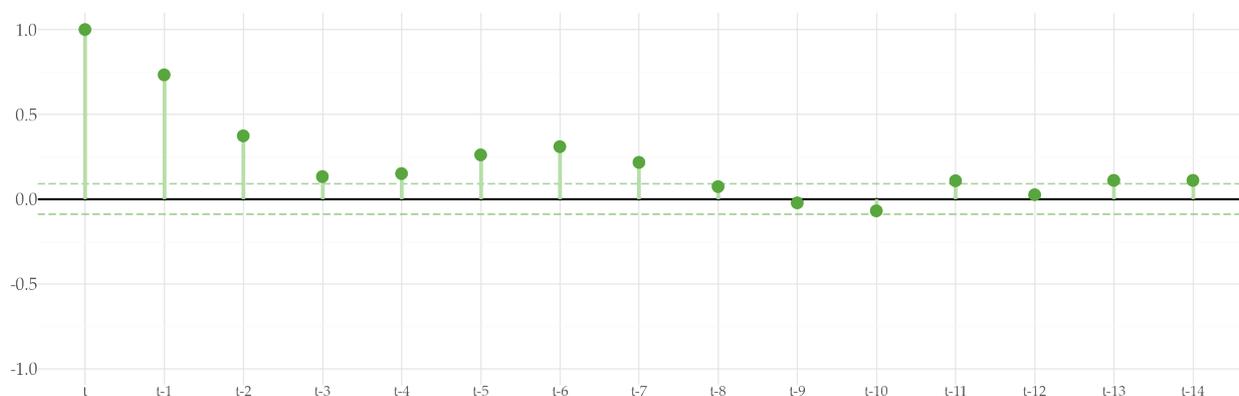


Figure 4: Partial Auto-correlation plot up to 14 lags. At each lag, the partial auto-correlation takes into account the previous correlations.

- The following lags show significant partial autocorrelation: t-1, t-2, t-3, t-4, t-5, t-6, t-7, t-11, t-13, and t-14.
- All lags relative to the seasonal period (t-7 and t-14) show a significant partial autocorrelation.

## Trend

Trend refers to the long-term change in the mean level of a time series. It reflects systematic and gradual changes in the data over time. Understanding the trend is important for identifying long-term growth or decline, structural changes, and making informed modeling decisions. This section examines the characteristics of the trend of the time series.

### Trend Line Plot

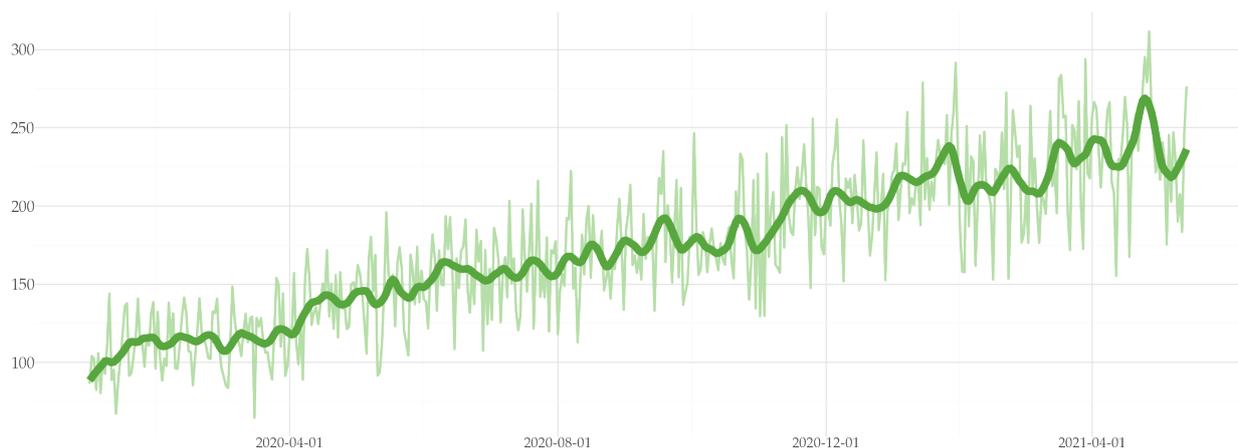


Figure 5: Time series trend plot.

- There is a strong upward trend. The tests KPSS, Augmented Dickey-Fuller, and Philips-Perron did not find evidence for non-stationarity around a deterministic trend.
- The same tests were applied to analyse stationarity around a constant level. The following tests reject this hypothesis: KPSS and Augmented Dickey-Fuller. But, the test Philips-Perron suggest stationarity.
- **Preliminary experiments:** Including a trend explanatory variable which denotes the position (row id) of each observation improves forecasting accuracy. These experiments were conducted using a LightGBM algorithm and evaluated using SMAPE loss function. Using only lag-based features the model achieved a SMAPE of 12.28% on the test set. Including the trend variable improved the SMAPE to 11.02%.

## Long-term Growth

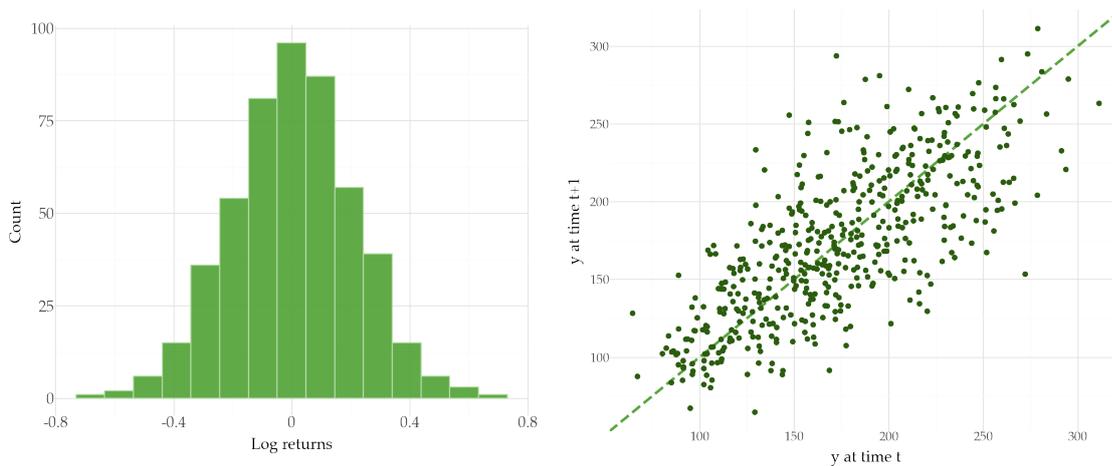


Figure 6: Distribution of log differences (left), and a Lag-plot (right). These plots help to understand how the data changes over consecutive observations. The histogram shows the distribution of these changes using log returns. The lag-plot depicts the randomness in the data. The time series shows greater randomness as the points deviate from the dotted line.

- The time series has an average growth (log returns) of 0.002 (median equal to -0.004). The volatility of the returns in terms of standard deviation is 0.21. The skewness of the log differenced series is equal to -0.03, which is close to zero. This indicates a symmetric distribution, though there is a slight left skewness. The excess kurtosis of the log differenced series is equal to 0.19. This value is similar to that found from data following a Gaussian distribution.
- Concerning the symmetry of returns, 49.3% of the log differences are positive. The average of positive returns is 0.17, while the average of negative returns (50.7% of all returns) is -0.16. Overall, there are 322 return direction changes (64.66% of the data points)
- In the tails, 4.61% of returns fall beyond 2 standard deviations from the mean. The largest positive return is 0.68 on 2020-03-16. Conversely, the largest decline is -0.69 (on 2020-03-15).
- **Preliminary experiments:** Modeling the time series of first differences does not seem to improve forecasting accuracy. Experiments were conducted using a LightGBM algorithm and evaluated using SMAPE loss function. Using the original time series led to a 12.28% SMAPE. The scores using the differenced and log differenced time series are 17.86% and 18.99%, respectively.

## Seasonality

Seasonality represents recurring patterns or cycles that appear at regular intervals in time series data. These are predictable fluctuations that reflect periodic influences such as monthly, quarterly, or yearly cycles. Understanding seasonal patterns is crucial for

forecasting, trend analysis, and identifying anomalies. This section examines the presence, strength, and characteristics of seasonal components in the input time series.

## Seasonal Line Plot (Dayly)



Figure 7: Seasonal plot of daily values grouped by week.

- The seasonal strength is 0.5. This score ranges from 0 to 1 and values above 0.64 are considered significant. The following tests indicate that the time series is non-stationary in seasonality for a weekly period: OCSB. On the other hand, other tests (Wang-Smith-Hyndman) fail to reject the stationary null hypothesis.
- **Preliminary experiments:** Modeling weekly patterns does not improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (12.28% SMAPE):
  - Fourier terms: 12.77% SMAPE
  - Seasonal differencing: 17.6% SMAPE
  - Daily time features: 12.28% SMAPE
- We also analyzed Daily seasonality. All tests indicate seasonal stationarity in these frequencies. Including these seasonal components in the forecasting model led to decreased performance. Statistical analysis of Daily data shows no evidence of systematic differences across Days, with both means (Kruskal-Wallis test) and variances (Levene's test) being statistically similar.

## Seasonal Sub-series Plot (Monthly)



Figure 8: Monthly seasonal sub-series. This plot helps to understand how the data varies within and across monthly groups.

- Tests for monthly seasonal stationarity show mixed results: the OCSB test indicates presence of a seasonal unit root, while the Wang-Smith-Hyndman test suggests stationarity.
- **Preliminary experiments:** There is evidence for a monthly seasonal pattern based on statistical tests. Besides, modeling monthly patterns can improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (12.28% SMAPE):
  - Fourier terms: 12.21% SMAPE
  - Monthly seasonal differencing: 18.75% SMAPE
  - Monthly time features: 11.06% SMAPE

## Mean and Standard Deviation Analysis (Monthly)

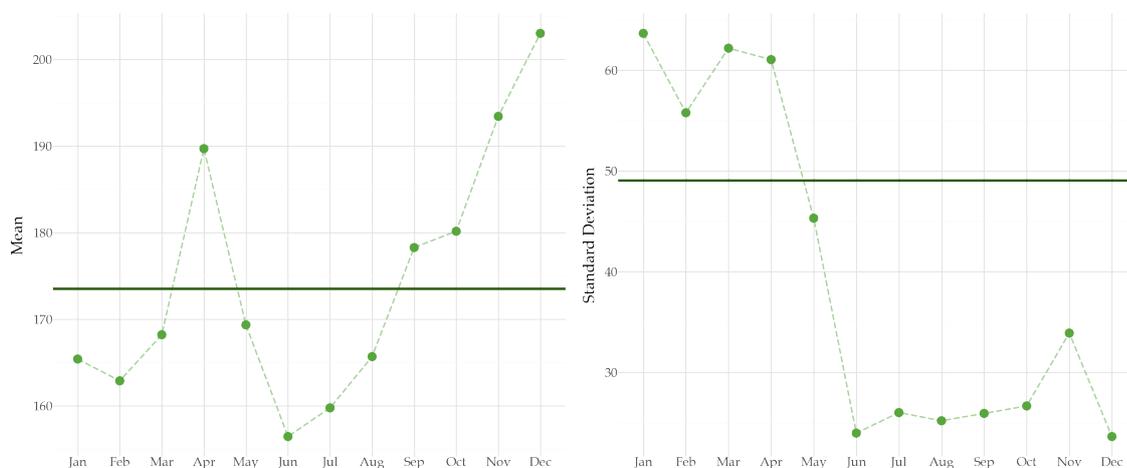


Figure 9: Mean plot (left) and standard deviation plot (right) of monthly values. The horizontal line shows the overall value across groups, while the dots show the value in the corresponding group.

- Analysis of monthly patterns shows significant differences in central tendency (Kruskal-Wallis test) and significant differences in dispersion (Levene's test).

## Variance

Variance measures how data points spread around the average value in your time series. This section examines whether the variability remains stable (homoskedastic) or changes (heteroskedastic) over time. Understanding variance patterns is crucial for selecting appropriate modeling techniques, which can have a significant impact on forecasting accuracy.

## Heteroskedasticity Testing

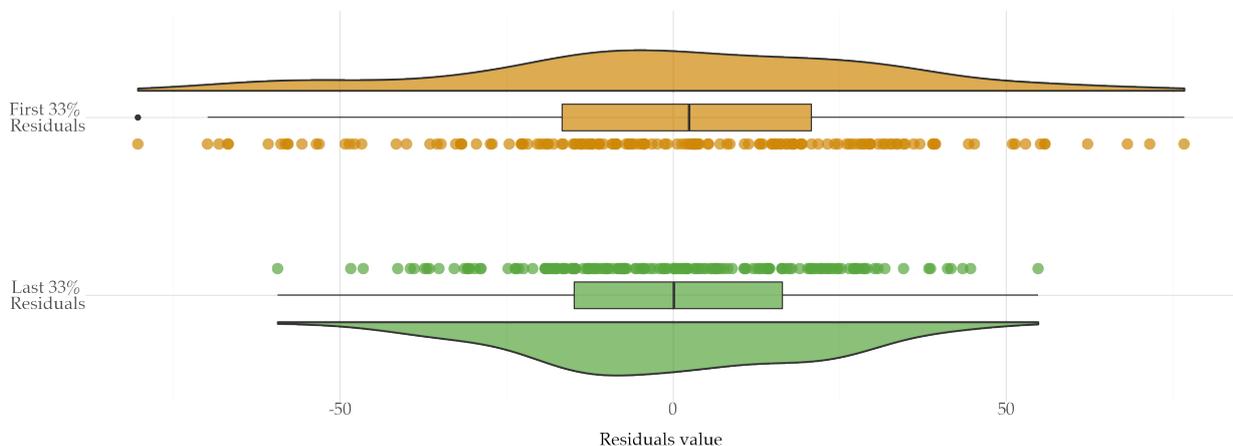


Figure 10: Time series residuals analysis based on a linear trend model. Difference in the distribution of the residuals in the first and last thirds of the series, following a Goldfeld-Quandt partition.

- Statistical evidence was found for the hypothesis that the time series is heteroskedastic, according to the White, Breusch-Pagan, and Goldfeld-Quandt tests. The residuals are based on a linear trend model.
- Variance in seasonal periods according to Levene's test
  - Weekly groups: no differences in variance
  - Monthly groups: differences in variance
- **Preliminary experiments:** Three variance stabilization preprocessing techniques were tested to improve the forecast accuracy of an auto-regressive LightGBM (with 12.28% SMAPE using lag-based features):
  - Log returns: 18.99% SMAPE
  - Log transformation: 12.91% SMAPE
  - Box-Cox transformation: 12.69% SMAPE

# Change Detection

Change points denote significant shifts in the underlying distribution of time series. These structural changes can manifest as sudden shifts in level, trend, variance, or seasonal patterns. Detecting and understanding these points is crucial as they often indicate important events or regime changes that affect modeling decisions. This section identifies potential change points and assesses their impact on the overall analysis strategy.

## Change Points

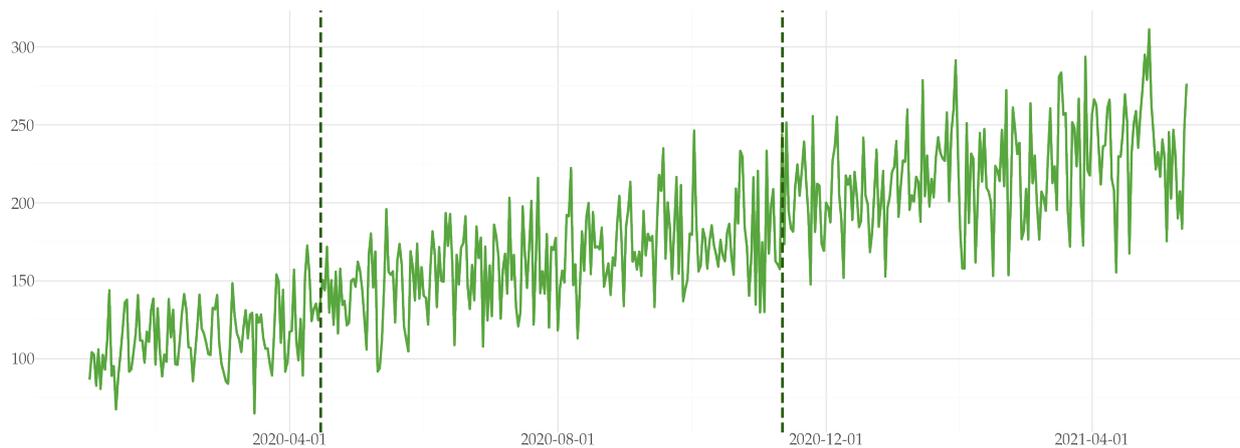


Figure 11: Time series plot with marked change points according to the PELT method.

- There are a total of 2 change points over the time series
- The first change point was found at 2020-04-15 where the time series shows an increasing level.

## Effect on Model Parameters

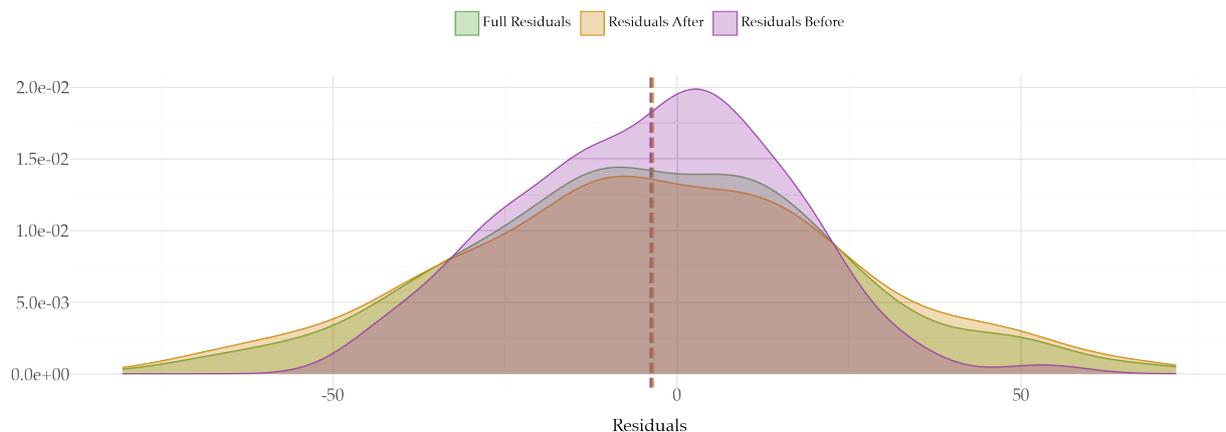


Figure 12: Distribution of the residuals of an ARIMA model before and after the first detected change point. The plot compares three kernel density estimates: residuals from the pre-change model, post-change model,

and full series model. This comparison helps assess whether the structural break affects model adequacy and error distribution properties.

- A Chow test was conducted using an ARIMA(2, 1, 2) model. The test fails to reject the null hypothesis of parameter stability. This suggests that the ARIMA parameters remain stable before and after the first detected change point, suggesting the underlying process structure remained similar despite the level shift.
- **Preliminary experiments:** Adding a step intervention at the change point reduced the model performance. The baseline SMAPE of 12.28% increased to when including the intervention (12.97%).