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: 2025-02-04 23:32

Series Name: M15

Table of Contents

Other aspects were explored but omitted from the final report:

Trend

Hypothesis testing (KPSS, Augmented Dickey-Fuller) indicates that the time series does not exhibit a significant trend. Taking first differences or feature extraction for trend inclusion did not improve forecasting accuracy in preliminary tests.

Change Detection

No change point was found according to offline change detection methods

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

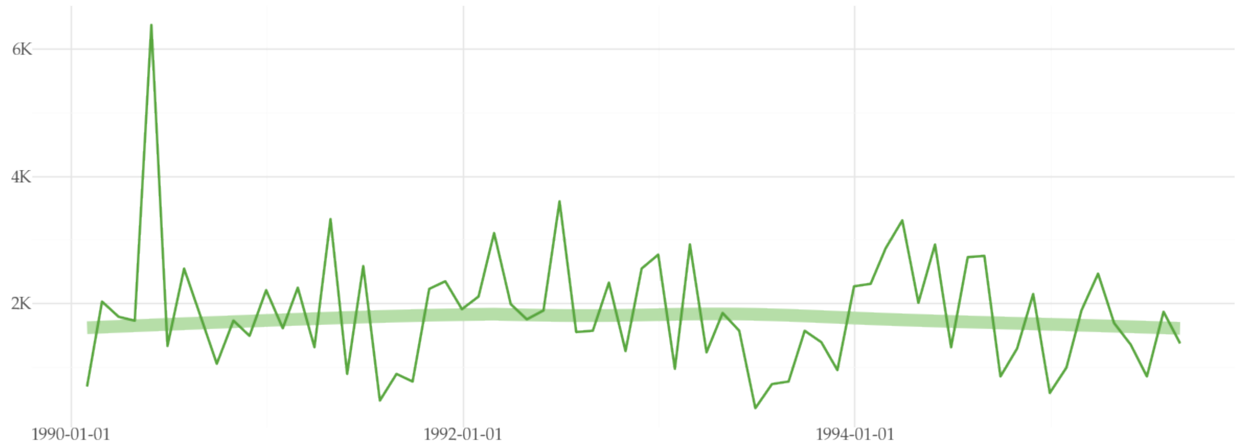

Figure 1: Time series line plot.

- A total of 68 observations spanning from January 1990 to August 1995. These are collected with a monthly sampling frequency.
- The data ranges from a minimum of 340 to a maximum of 6380, starting in 680 and ending in 1360 during the observed period. The average growth percentage per observation is 25.65% (median equal to 5.41%), with an average value of 1843.24. There are no missing values in the time series.

Trend, Seasonality, and Residuals

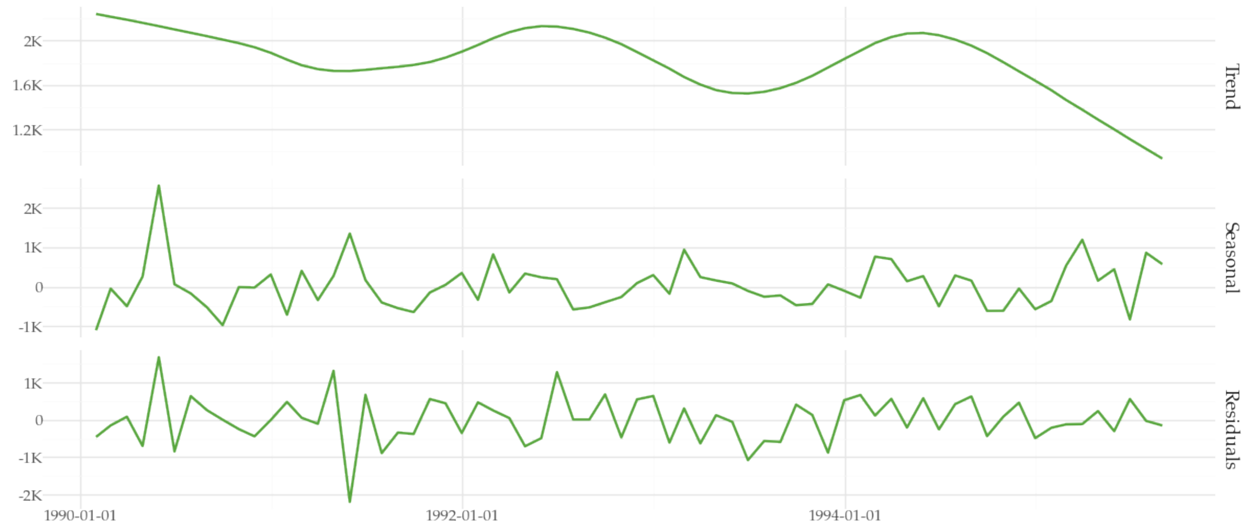

Figure 2: Seasonal, Trend, and Residuals components after decomposition on a monthly frequency using the STL (Season-Trend decomposition using LOESS) method.

- The trend strength is 0.25 (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 0.54 (ranges from 0 to 1). Tests for yearly 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: 51.47% of residuals are positive and 48.53% negative. The average magnitude of positive residuals is 457.463 compared to -467.885 for negative residuals. In terms of auto-correlation structure, the residuals show significant temporal dependency in some of the first 12 lags according to the Ljung-Box test. This suggests that the decomposition method is missing some systematic patterns.

Auto-Correlation

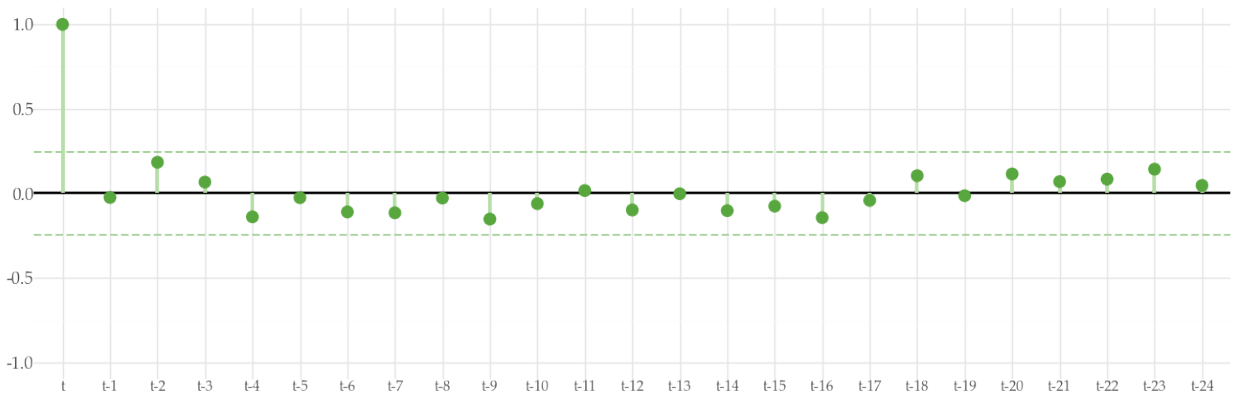

Figure 3: Auto-correlation plot up to 24 lags.

- There are no lags with significant correlation, which suggest that the series may be white noise.
- None of the lags relative to the seasonal period (t-12 and t-24) show any significant autocorrelation.

Partial Auto-Correlation

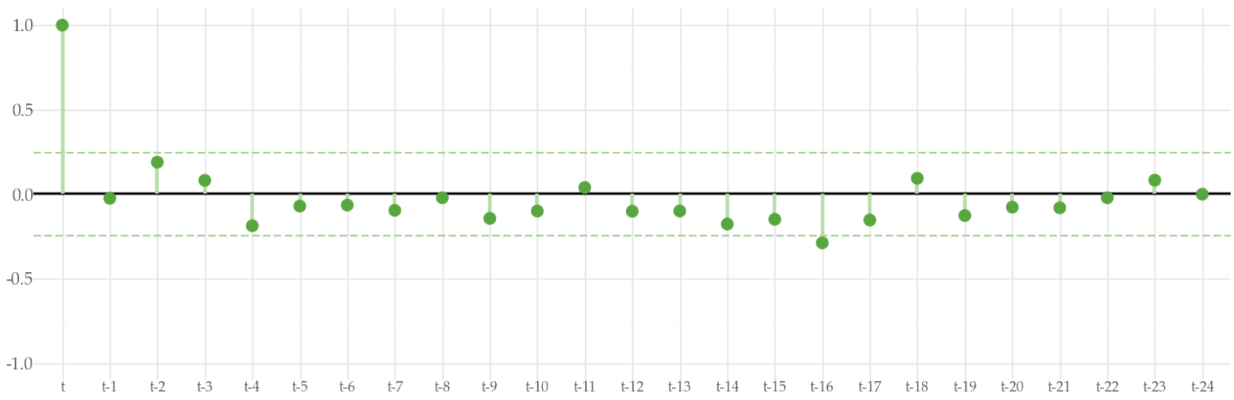

Figure 4: Partial Auto-correlation plot up to 24 lags. At each lag, the partial auto-correlation takes into account the previous correlations.

- The following lags show significant partial autocorrelation: t-16.
- None of the lags relative to the seasonal period (t-12 and t-24) show any significant partial autocorrelation.

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 (Monthly)

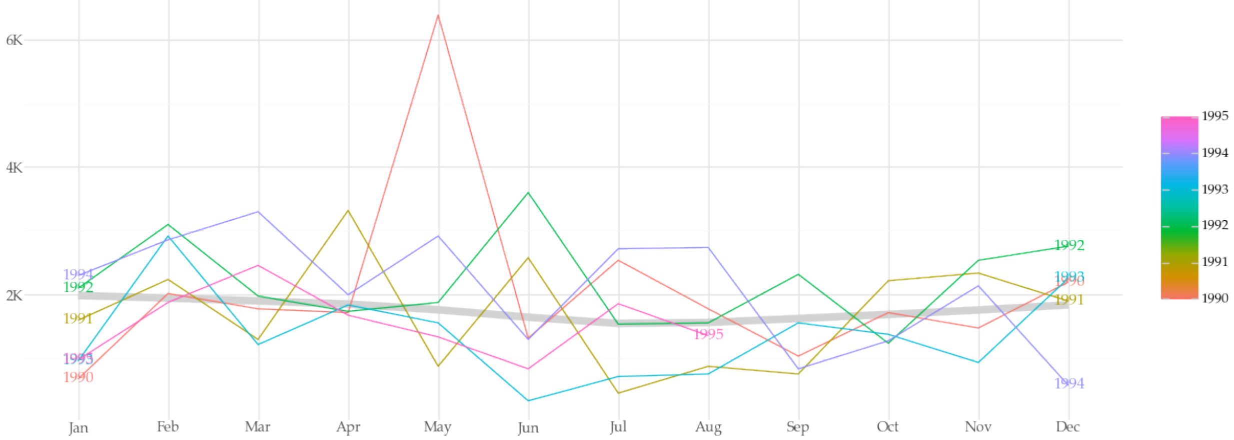

Figure 5: Seasonal plot of monthly values grouped by year.

- The seasonal strength is 0.54. This score ranges from 0 to 1 and values above 0.64 are considered significant. The following tests indicate that the time series is nonstationary in seasonality for a yearly period: OCSB. On the other hand, other tests (Wang-Smith-Hyndman) fail to reject the stationary null hypothesis.

- **Preliminary experiments:** Modeling yearly patterns does not improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (52.86% SMAPE):
    - Fourier terms: 53.71% SMAPE
    - Seasonal differencing: 58.55% SMAPE
    - Monthly time features: 53.87% SMAPE

Seasonal Sub-series Plot (Quarterly)

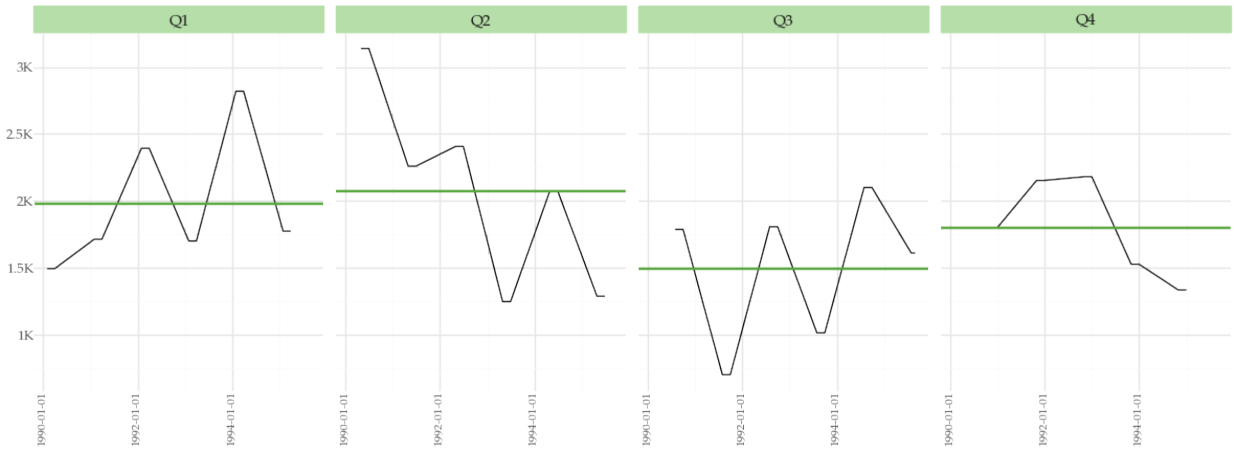

Figure 6: Quarterly seasonal sub-series. This plot helps to understand how the data varies within and across quarterly groups.

- Statistical analysis of quarterly data shows no evidence of systematic differences across quarters, with both means (Kruskal-Wallis test) and variances (Levene's test) being statistically similar.
- Tests for quarterly 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 quarterly seasonal pattern based on statistical tests. Besides, modeling quarterly patterns can improve forecast accuracy. Different approaches were tested relative to a base model using only lag-based features (52.86% SMAPE):
    - Fourier terms: 50.78% SMAPE
    - Quarterly seasonal differencing: 60.59% SMAPE
    - Quarterly time features: 47.83% SMAPE

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

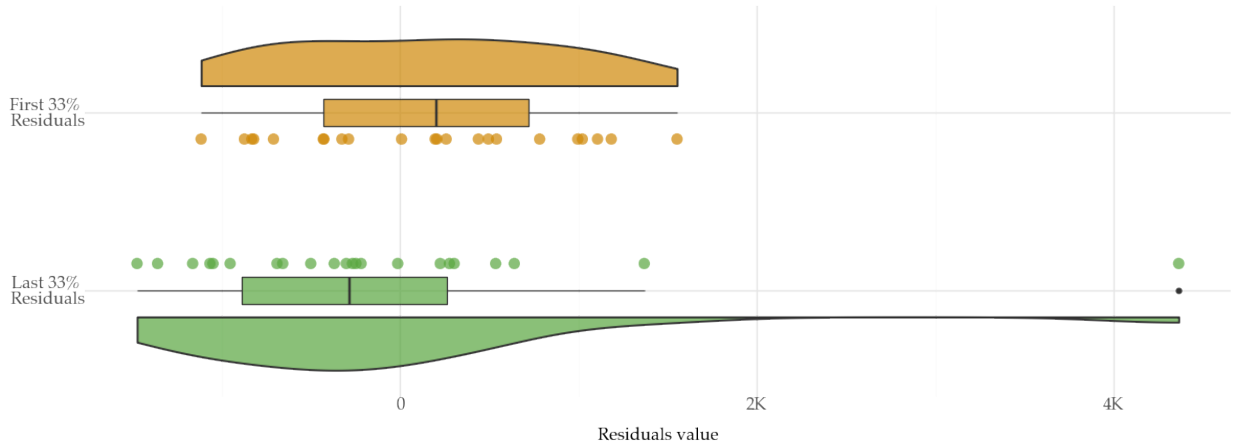

Figure 7: Time series residuals analysis based on a linear trend model. Difference in the distribution of the residuals in the first and last thirds of the series, following a Goldfeld-Quand partition.

- No 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
  - Quarterly groups: no differences in variance
  - Monthly groups: no differences in variance
- **Preliminary experiments:** Three variance stabilization preprocessing techniques were tested to improve the forecast accuracy of an auto-regressive LightGBM (with 52.86% SMAPE using lag-based features):
  - Log returns: 48.53% SMAPE
  - Log transformation: 46.15% SMAPE
  - Box-Cox transformation: 45.26% SMAPE