Advanced Topics: Time Series Analysis

Basic Forecasting, Trend Analysis, and Seasonality

Sarwan Ali

Department of Computer Science
Georgia State University

Understanding Time Series Analysis

Today's Learning Journey

What is Time Series Data?

Definition

A **time series** is a sequence of data points indexed in temporal order, typically collected at successive, equally spaced points in time.

Examples:

- Stock prices over time
- Weather measurements
- Sales data by month
- Website traffic counts
- Sensor readings

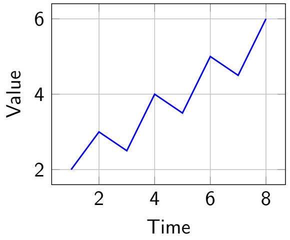

Time Series vs Cross-Sectional Data

Time Series Data

- Observations over time
- Temporal dependencies
- Order matters
- Autocorrelation present

Example: Daily temperature readings

$$T_1, T_2, T_3, \ldots, T_n \tag{1}$$

Cross-Sectional Data

- Observations at one point in time
- Independent observations
- Order doesn't matter
- No temporal structure

Example: Heights of students in a class

$$H_1, H_2, H_3, \ldots, H_n \tag{2}$$

Decomposition of Time Series

A time series can be decomposed into several components:

Additive Model

$$Y_t = T_t + S_t + C_t + I_t \tag{3}$$

Multiplicative Model

$$Y_t = T_t \times S_t \times C_t \times I_t \tag{4}$$

Where:

- T_t = Trend component
- S_t = Seasonal component
- C_t = Cyclical component
- I_t = Irregular (random) component

Trend Component

Trend

The long-term movement or direction in the data over time.

Types of Trends:

- Upward trend: Increasing over time
- Downward trend: Decreasing over time
- No trend: Stationary around a mean

Mathematical representation:

$$T_t = \alpha + \beta t + \gamma t^2 + \dots \tag{5}$$

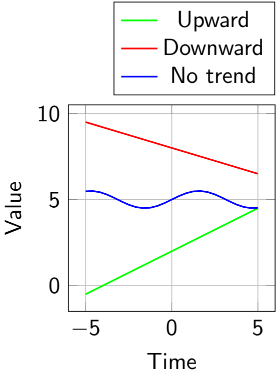

Seasonality

Regular, predictable patterns that repeat over fixed periods (e.g., daily, weekly, monthly, yearly).

Examples:

- Ice cream sales peak in summer
- Electricity usage patterns
- Holiday shopping spikes
- Weekly website traffic patterns

Mathematical representation:

$$S_t = \sum_{i=1}^{s} \gamma_i D_{i,t} \tag{6}$$

where $D_{i,t}$ are seasonal dummy variables.

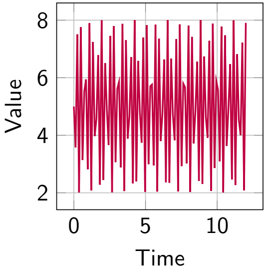

Seasonal pattern repeating every 12 time units

Cyclical vs Seasonal Components

Seasonal

- Fixed period (e.g., 12 months)
- Predictable timing
- Same pattern each cycle
- Caused by calendar events

Example: Retail sales during Christmas season

Cyclical

- Variable period (2-10+ years)
- Unpredictable timing
- Varying amplitude
- Economic/business cycles

Example: Economic recession cycles

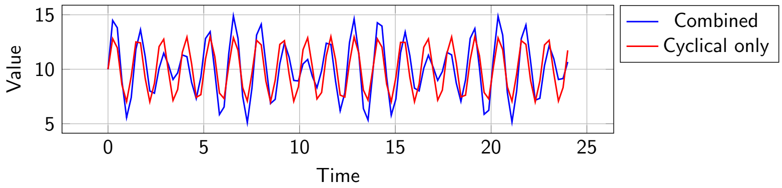

Methods for Trend Analysis

1. Visual Inspection

Plot the data and observe the general direction over time.

2. Moving Averages

$MA_t = \frac{1}{k}\sum_{i=0}^{k-1} Y_{t-i} \rightarrow$ Smooths out short-term fluctuations to reveal underlying trend.

3. Linear Regression

$Y_t = \alpha + \beta t + \epsilon_t \rightarrow$ Fits a straight line through the data points.

4. Polynomial Regression

$$Y_t = \alpha + \beta_1 t + \beta_2 t^2 + \ldots + \beta_k t^k + \epsilon_t \tag{7}$$

Captures non-linear trends.

Moving Averages - Example

Simple Moving Average (SMA):

$$SMA_t = \frac{Y_t + Y_{t-1} + \ldots + Y_{t-n+1}}{n} \qquad (8)$$

Weighted Moving Average (WMA):

$$WMA_t = \frac{\sum_{i=0}^{n-1} w_i Y_{t-i}}{\sum_{i=0}^{n-1} w_i} \qquad (9)$$

Exponential Moving Average (EMA):

$$EMA_t = \alpha Y_t + (1-\alpha) EMA_{t-1} \qquad (10)$$

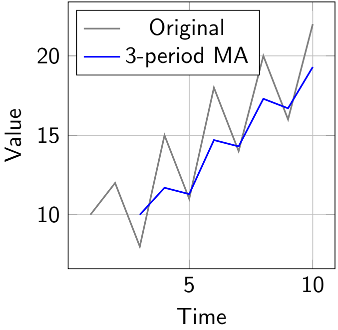

Trend Detection Methods

1. Mann-Kendall Test

Non-parametric test for monotonic trend detection.

$$S = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} \text{sign}(Y_j - Y_i) \tag{11}$$

2. Linear Regression Significance

Test if the slope coefficient β is significantly different from zero.

$$H_0 : \beta = 0 \quad \text{vs} \quad H_1 : \beta \neq 0 \tag{12}$$

3. Augmented Dickey-Fuller Test

Tests for unit roots (non-stationarity): $\Delta Y_t = \alpha + \beta t + \gamma Y_{t-1} + \sum_{i=1}^{p} \delta_i \Delta Y_{t-i} + \epsilon_t$

Detecting Seasonality

1. Seasonal Plots

Plot data by season/period to visualize recurring patterns.

2. Autocorrelation Function (ACF)

Measures correlation between observations at different lags.

$$\rho_k = \frac{\sum_{t=k+1}^{n}(Y_t - \bar{Y})(Y_{t-k} - \bar{Y})}{\sum_{t=1}^{n}(Y_t - \bar{Y})^2} \tag{13}$$

3. Periodogram

Identifies dominant frequencies in the data: $I(\omega) = \frac{1}{n}\left|\sum_{t=1}^{n} Y_t e^{-i\omega t}\right|^2$

4. Seasonal Decomposition

Separates seasonal component from trend and noise.

Seasonal Decomposition Methods

1. Classical Decomposition

Step 1: Estimate trend using moving averages
Step 2: Remove trend to get detrended series
Step 3: Estimate seasonal component
Step 4: Calculate residuals

2. X-13ARIMA-SEATS

Advanced method used by statistical agencies for seasonal adjustment.

3. STL Decomposition

Seasonal and Trend decomposition using Loess smoothing.

- Handles any type of seasonality
- Robust to outliers
- Allows seasonal component to change over time

Handling Seasonality

1. Seasonal Differencing

$\nabla_s Y_t = Y_t - Y_{t-s}$, where s is the seasonal period.

2. Seasonal Dummy Variables

$Y_t = \alpha + \sum_{i=1}^{s-1} \beta_i D_{i,t} + \epsilon_t$, where $D_{i,t}$ are seasonal dummies.

3. Fourier Terms

$$Y_t = \alpha + \sum_{k=1}^{K} \left[\beta_k \cos\left(\frac{2\pi kt}{s}\right) + \gamma_k \sin\left(\frac{2\pi kt}{s}\right) \right] + \epsilon_t \tag{14}$$

Example: Monthly Data

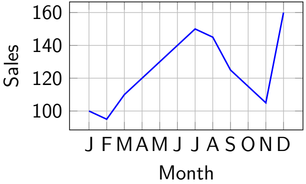

Clear seasonal pattern with December peak

Forecasting Overview

Definition

Forecasting is the process of predicting future values based on historical data and identified patterns.

Key Principles:

- Use all available information
- Forecasts are uncertain
- Accuracy decreases with horizon
- Simple methods often work well

Forecast Horizon:

- Short-term: 1-30 days
- Medium-term: 1-12 months
- Long-term: 1+ years

Notation

Let $\hat{Y}_{t+h|t}$ denote the forecast of Y_{t+h} made at time t, where h is the forecast horizon.

Simple Forecasting Methods

1. Naive Method

$\hat{Y}_{t+h|t} = Y_t, \rightarrow$ Use the last observed value as the forecast.

2. Seasonal Naive Method

$\hat{Y}_{t+h|t} = Y_{t+h-s}, \rightarrow$ Use the value from the same season in the previous year.

3. Simple Mean Method

$\hat{Y}_{t+h|t} = \frac{1}{t}\sum_{i=1}^{t} Y_i, \rightarrow$ Use the average of all historical values.

4. Drift Method

$\hat{Y}_{t+h|t} = Y_t + h \cdot \frac{Y_t - Y_1}{t-1}, \rightarrow$ Extrapolate the trend from first to last observation.

Exponential Smoothing

Simple Exponential Smoothing

$$\hat{Y}_{t+1|t} = \alpha Y_t + (1-\alpha)\hat{Y}_{t|t-1} \quad (15)$$

$$= \alpha Y_t + \alpha(1-\alpha)Y_{t-1} + \alpha(1-\alpha)^2 Y_{t-2} + \dots \quad (16)$$

where $0 < \alpha < 1$ is the smoothing parameter.

Properties:

- Weighted average of past observations
- Weights decrease exponentially
- Recent observations get higher weights
- Only one parameter to estimate

Choosing α:

- High α (≈ 0.8): Fast adaptation
- Low α (≈ 0.2): Smooth forecasts

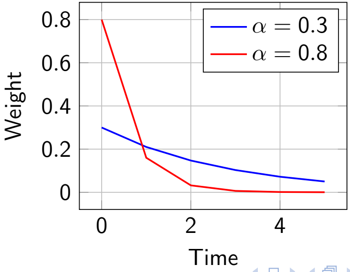

Holt's Linear Trend Method

Holt's Method (Double Exponential Smoothing)

Extends simple exponential smoothing to handle trend.

$$\ell_t = \alpha Y_t + (1-\alpha)(\ell_{t-1} + b_{t-1}) \quad (17)$$

$$b_t = \beta(\ell_t - \ell_{t-1}) + (1-\beta) b_{t-1} \quad (18)$$

$$\hat{Y}_{t+h|t} = \ell_t + h \cdot b_t \quad (19)$$

where:

- ℓ_t is the level (smoothed value)
- b_t is the trend (slope)
- α is the level smoothing parameter
- β is the trend smoothing parameter

Key Advantage: Can forecast multiple steps ahead with linear trend extrapolation.

Holt-Winters Seasonal Method

Holt-Winters Method (Triple Exponential Smoothing), Handles level, trend, and seasonality.

Additive Seasonality:

$$\ell_t = \alpha(Y_t - s_{t-m}) + (1-\alpha)(\ell_{t-1} + b_{t-1}) \tag{20}$$

$$b_t = \beta(\ell_t - \ell_{t-1}) + (1-\beta)b_{t-1} \tag{21}$$

$$s_t = \gamma(Y_t - \ell_{t-1} - b_{t-1}) + (1-\gamma)s_{t-m} \tag{22}$$

$$\hat{Y}_{t+h|t} = \ell_t + h \cdot b_t + s_{t+h-m} \tag{23}$$

Multiplicative Seasonality:

$$\ell_t = \alpha \frac{Y_t}{s_{t-m}} + (1-\alpha)(\ell_{t-1} + b_{t-1}) \tag{24}$$

$$s_t = \gamma \frac{Y_t}{\ell_{t-1} + b_{t-1}} + (1-\gamma)s_{t-m} \tag{25}$$

$$\hat{Y}_{t+h|t} = (\ell_t + h \cdot b_t) \cdot s_{t+h-m} \tag{26}$$

Forecast Accuracy Measures

Let $e_t = Y_t - \hat{Y}_t$ be the forecast error.

Scale-Dependent Measures

$\text{MAE} = \frac{1}{n}\sum_{t=1}^{n}|e_t|$, $\text{RMSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}e_t^2}$

Percentage Measures

$\text{MAPE} = \frac{100}{n}\sum_{t=1}^{n}\left|\frac{e_t}{Y_t}\right|$, $\text{sMAPE} = \frac{100}{n}\sum_{t=1}^{n}\frac{|e_t|}{(|Y_t|+|\hat{Y}_t|)/2}$

Scale-Free Measures

$\text{MASE} = \frac{\text{MAE}}{\frac{1}{n-1}\sum_{t=2}^{n}|Y_t - Y_{t-1}|}$

Cross-Validation for Time Series

Time Series Cross-Validation

Unlike regular cross-validation, maintains temporal order.

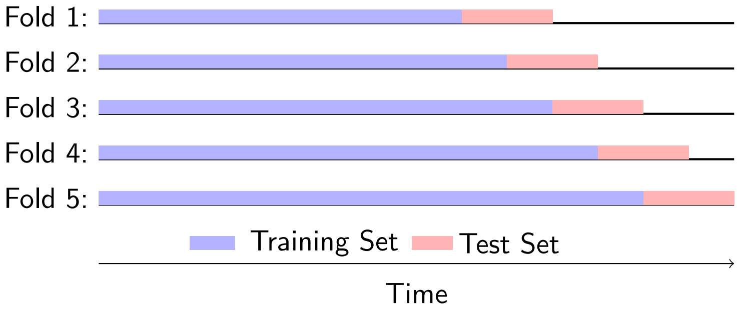

Key Points:

- Always use past data to predict future
- Test set size can be fixed or growing
- Provides more robust performance estimates

Residual Analysis

Good forecasts should have residuals that are:

- Uncorrelated (no pattern)
- Zero mean
- Constant variance (homoscedastic)
- Normally distributed

Ljung-Box Test

Tests for autocorrelation in residuals:

$$Q_{LB} = n(n+2)\sum_{k=1}^{h}\frac{\hat{\rho}_k^2}{n-k} \quad (27)$$

where $\hat{\rho}_k$ is the sample autocorrelation at lag k.

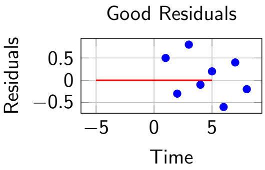

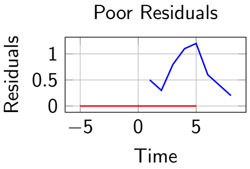

Implementation Steps

1. Data Preparation

- Check for missing values and outliers, Ensure regular time intervals
- Handle irregular observations, Transform data if necessary (log, square root)

2. Exploratory Data Analysis

- Plot the time series, Identify trend, seasonality, and cycles
- Check for structural breaks, Analyze autocorrelation patterns

3. Model Selection and Fitting

- Choose appropriate forecasting method, Estimate model parameters
- Check model assumptions, Validate on holdout data

Python Implementation Example

```python
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
from statsmodels.tsa.holtwinters import ExponentialSmoothing
from statsmodels.tsa.seasonal import seasonal_decompose
from sklearn.metrics import mean_absolute_error, mean_squared_error
# Load and prepare data
df = pd.read_csv('sales_data.csv')
df['date'] = pd.to_datetime(df['date'])
df.set_index('date', inplace=True)
ts = df['sales']
fig, axes = plt.subplots(2, 2, figsize=(12, 8)) # Exploratory analysis
ts.plot(ax=axes[0,0], title='Original Time Series')
# Seasonal decomposition
decomposition = seasonal_decompose(ts, model='additive', period=12)
decomposition.trend.plot(ax=axes[0,1], title='Trend')
decomposition.seasonal.plot(ax=axes[1,0], title='Seasonal')
decomposition.resid.plot(ax=axes[1,1], title='Residuals')
# Holt-Winters forecasting
model = ExponentialSmoothing(ts, trend='add', seasonal='add', seasonal_periods=12)
fitted_model = model.fit()
forecast = fitted_model.forecast(steps=12) # Generate forecasts
forecast_ci = fitted_model.get_prediction(start=-24).conf_int()
train_size = int(len(ts) * 0.8) # Evaluate accuracy
train, test = ts[:train_size], ts[train_size:]
model_eval = ExponentialSmoothing(train, trend='add', seasonal='add', seasonal_periods=12).fit()
predictions = model_eval.forecast(steps=len(test))
mae = mean_absolute_error(test, predictions)
rmse = np.sqrt(mean_squared_error(test, predictions))
print(f"MAE: {mae:.2f}, RMSE: {rmse:.2f}")
```

Common Pitfalls and Best Practices

Common Pitfalls

- Ignoring data quality issues
- Overfitting to historical data
- Not accounting for structural breaks
- Using inappropriate accuracy measures
- Forecasting too far into the future
- Not updating models regularly

Best Practices

- Start with simple methods
- Use multiple forecasting methods
- Combine forecasts when possible
- Regularly monitor and update models
- Provide prediction intervals
- Document assumptions and limitations

Model Selection Guidelines

- **No trend, no seasonality**: Simple exponential smoothing
- **Trend, no seasonality**: Holt's method
- **Trend and seasonality**: Holt-Winters method
- **Complex patterns**: Consider ARIMA or machine learning methods

Beyond Basic Methods

ARIMA Models

- Autoregressive Integrated Moving Average
- Handles non-stationary data
- Box-Jenkins methodology
- Seasonal ARIMA (SARIMA)

Machine Learning

- Neural networks (LSTM, GRU)
- Support Vector Regression
- Random Forest for time series
- Deep learning approaches

State Space Models

- Kalman filtering
- Dynamic linear models
- Structural time series models
- Unobserved components

Multivariate Methods

- Vector Autoregression (VAR)
- Cointegration analysis
- Dynamic factor models
- Cross-correlation analysis

Real-World Applications

Business Applications:

- Demand forecasting
- Inventory management
- Financial planning
- Revenue prediction
- Customer behavior analysis

Finance:

- Stock price prediction
- Risk management
- Portfolio optimization
- Volatility modeling
- Algorithmic trading

Science & Engineering:

- Climate modeling
- Quality control
- Sensor data analysis
- Energy consumption
- Medical monitoring

Technology:

- Web traffic analysis
- System performance monitoring
- User engagement metrics
- A/B testing analysis
- Capacity planning

Key Takeaways

1. Understanding Time Series Components

- Trend: Long-term direction
- Seasonality: Regular, predictable patterns
- Cyclical: Irregular, longer-term fluctuations
- Irregular: Random variation

2. Forecasting Methods

- Simple methods: Naive, seasonal naive, drift
- Exponential smoothing: Simple, Holt's, Holt-Winters
- Choose method based on data characteristics

3. Evaluation and Validation

- Use appropriate accuracy measures, Apply time series cross-validation
- Analyze residuals for model adequacy, Consider forecast intervals and uncertainty

Next Steps

Recommended Learning Path:

1. Master basic decomposition methods
2. Practice with real datasets
3. Learn ARIMA modeling
4. Explore multivariate methods
5. Study advanced ML techniques

Useful Resources:

- "Forecasting: Principles and Practice" by Hyndman & Athanasopoulos
- Python: statsmodels, scikit-learn
- R: forecast, tseries packages

Practice Exercises:

- Download economic data (FRED)
- Analyze retail sales data
- Forecast stock prices
- Model weather patterns
- Predict website traffic

Key Skills to Develop:

- Data visualization
- Statistical testing
- Model selection
- Performance evaluation
- Business interpretation

Questions & Discussion

sali85@student.gsu.edu

Thank you for your attention!