Time Series Analysis and Forecasting

Guide to AI and Data Courses · Part IV: Data

Trend and season | Stationarity | ARIMA and smoothing | Deep forecasting | Evaluation | Neighboring courses

Sercan KÜLCÜ · Giresun University · Computer Engineering

© Sercan Külcü. All rights reserved. · Interactive HTML version of these slides: https://sercankulcu.github.io/files/ai_courses/slides/Chapter_29_Time_Series_Analysis.html

INTRODUCTION

What Question Does It Ask?

> **How do we find structure in measurements over time and forecast the future?**

Time series analysis studies measurements whose order matters: yesterday's value affects today's. The course teaches how to split a series into trend, seasonality and noise, how to model **autocorrelation**, and how to forecast the future together with its uncertainty.

At a glance

APPROACH time-dependent statistical models	**INPUT** timestamped measurements	**OUTPUT** forecast, trend, seasonality
SUCCESS MEASURE forecast error (MAE, RMSE)	**MATH LOAD** medium-high	**DATA NEEDS** long history

A Short History

1927
Yule explains sunspot numbers with an autoregressive (AR) model.

1960
The Kalman filter is published: it estimates a hidden state step by step from noisy measurements.

1970
Box and Jenkins's book establishes ARIMA models and how to build them.

1982
Engle proposes the ARCH model, in which volatility changes over time.

1997
Hochreiter and Schmidhuber introduce the LSTM; long dependencies become learnable.

2017
Facebook open-sources Prophet, which separates seasonality and holidays.

THE COURSE

Core Concepts

Trend and seasonality
Long-run direction and regularly repeating cycles.

Stationarity
Mean and variance do not change over time.

Autocorrelation (ACF)
Correlation of a series with its own past.

Differencing
Removing the trend to make a series stationary.

ARIMA
Autoregression, differencing and moving average together.

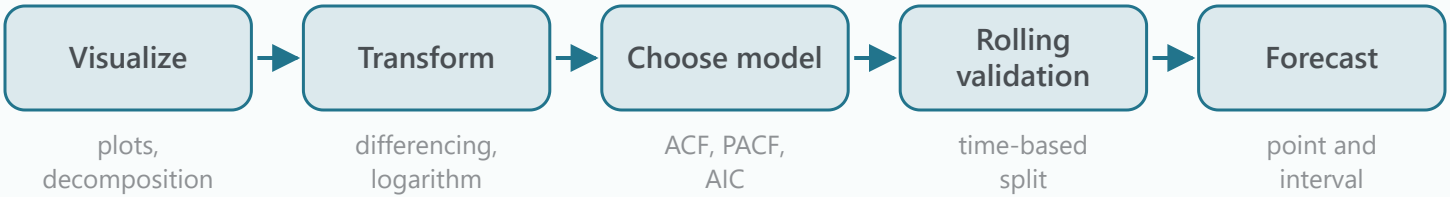
Exponential smoothing
A forecast that weights the recent past more.

Prediction interval
A forecast with its uncertainty, not a single number.

Time-based split
The test set always comes after training; no shuffling.

THE COURSE

A Typical Workflow



Example: A power company plots hourly consumption, separates the daily and weekly cycles, builds a model that also takes temperature as input, backtests it on past weeks, and gives the next day's load with a 95% prediction interval.

THE COURSE

What Does the Course Cover?

1. Components of a series: trend, seasonality, noise
2. Visualization and STL decomposition
3. Stationarity, differencing and unit root tests
4. Autocorrelation and partial autocorrelation
5. Exponential smoothing and Holt-Winters
6. AR, MA and ARMA models
7. ARIMA and seasonal ARIMA (SARIMA)
8. Regression with external regressors, dynamic models
9. Forecast evaluation and rolling-origin validation
10. Volatility models: ARCH and GARCH
11. State space models and the Kalman filter
12. Forecasting with machine learning and deep nets

Skills you will gain

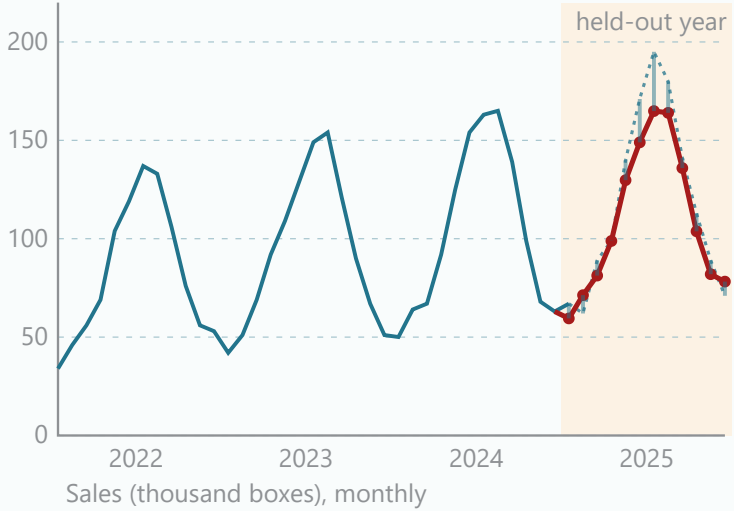
date-indexed data in pandas | fitting ARIMA with statsmodels | Prophet and sktime | designing rolling-origin validation | producing and reading prediction intervals

A typical semester; the order and coverage vary between universities.

TRY IT

Planning Next Summer's Ice Cream

An ice cream maker plans the coming months' production from three years of monthly sales. Pick a method and a horizon; the last year's real sales are held out and the forecast is scored on them.



Naive | Moving avg. | Seasonal naive | **Holt-Winters**

Forecast horizon: **12** months

Holt-Winters, 12-month horizon:
MAE = **11.0** thousand boxes (9.3% of mean sales)

Naive	55.3
Moving average (3 mo)	46.4
Seasonal naive	14.3
Holt-Winters	**11.0**

Naive and moving average forecasts cannot see the season: they draw a flat line through summer. Seasonal methods carry last year's shape forward; Holt-Winters also adds the growth trend.

CONTEXT

Prerequisites

Courses worth knowing before this one, and why they are needed.

Course	Why it is needed
Probability and Statistics	stationarity, autocorrelation, ARIMA
Calculus	series, derivatives and integrals: the language of models and filters
Programming	reading files, handling dates and plotting in Python

CONTEXT

Are You Ready for This Course?

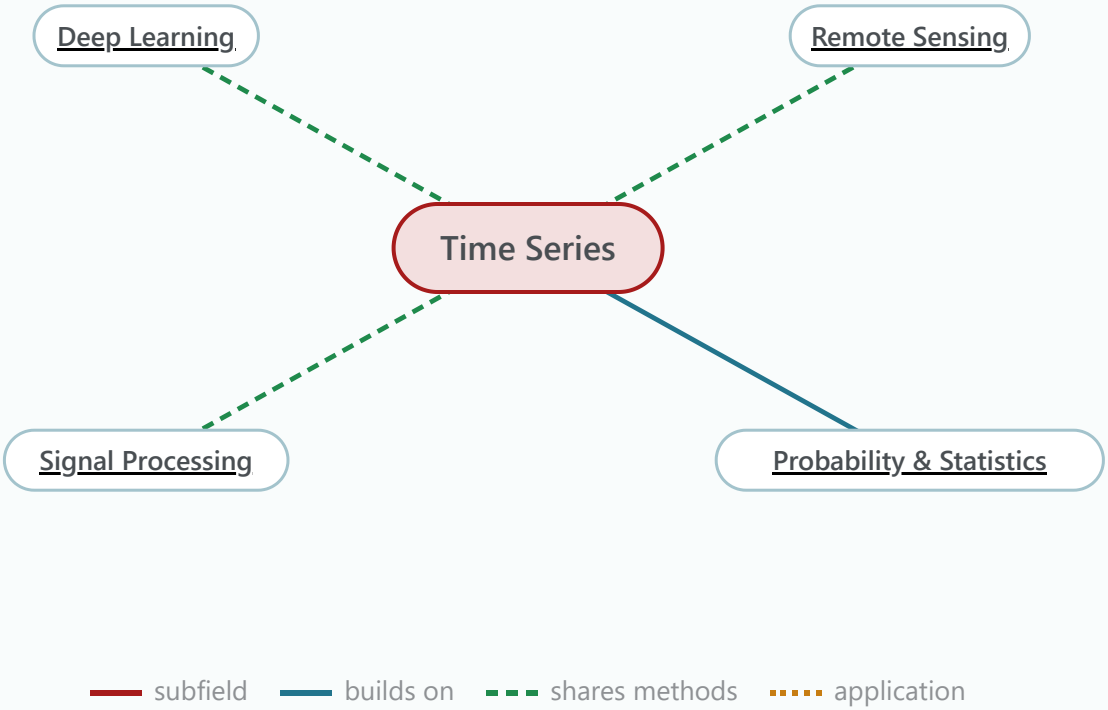
Ask yourself; for every "no", the right-hand column shows where to start.

Ask yourself	If not, start with
Can you explain mean, variance and the correlation coefficient?	Probability and Statistics
Can you fit a linear regression and interpret its coefficients?	Machine Learning
Can you compute the sum of a geometric series?	Sequences and series from calculus
Can you read a CSV file with a date column in Python and plot it?	An intro programming course, then a short pandas tutorial

If you can say "yes" to most of these, you are ready; one or two gaps can be closed during the semester.

CONTEXT

Neighbouring Courses



OVERLAP

Overlapping Courses

Relation	Course	Shared ground
builds on	This course builds on Probability and Statistics	stationarity, autocorrelation, ARIMA
shares methods	Shares methods with Deep Learning	recurrent networks, attention-based forecasting
shares methods	Shares methods with Digital Signal Processing	frequency analysis, filtering
shares methods	Shares methods with Remote Sensing	multitemporal image series

Where It Differs (1/2)

Criterion	Time Series	Probability & Statistics	Signal Processing
Core question	How do we find structure in measurements over time and forecast the future?	How is uncertainty measured, and how do we draw reliable conclusions from limited data?	How are sampled signals analyzed, filtered and transformed?
Approach	time-dependent statistical models	probability models and inference	mathematical transforms and filters
Input	timestamped measurements	sample data	sampled signal
Output	forecast, trend, seasonality	estimate, confidence interval, test result	filtered signal, spectrum
Success measure	forecast error (MAE, RMSE)	unbiasedness, power, significance	signal-to-noise ratio, frequency response
Math load	medium-high	the theory itself	high: complex numbers, Fourier
Data needs	long history	a sample	no learning; the signal suffices

- **Probability & Statistics:** Statistics usually assumes independent observations; this course deals with ordered data whose observations depend on each other.
- **Signal Processing:** Signal processing filters a signal and looks at its frequencies; this course uses the same tools to forecast the future.

DIFFERENCES

Where It Differs (2/2)

Criterion	Time Series	Deep Learning	Machine Learning
Core question	How do we find structure in measurements over time and forecast the future?	How do many-layered networks learn useful representations from raw data?	How does a program generalize from examples?
Approach	time-dependent statistical models	multilayer representation learning	statistical learning
Input	timestamped measurements	raw data: images, text, audio	labeled or unlabeled examples
Output	forecast, trend, seasonality	learned representation and prediction	a model that makes predictions
Success measure	forecast error (MAE, RMSE)	accuracy, loss, scaling	error on unseen data
Math load	medium-high	high: derivatives, matrices, probability	high: linear algebra, probability, optimization
Data needs	long history	large; needs GPUs	moderate to large

- **Deep Learning:** Deep learning offers a powerful family of forecasting models; this course also teaches when a simple model is enough.
- **Machine Learning:** ML shuffles examples and splits them at random; here time order is kept and the model is trained only on the past.

PRACTICE

Applications

Demand forecasting
Planning sales, stock and production.

Energy load
Power consumption and renewable output.

Finance
Volatility, risk and return models.

Weather and climate
Temperature, rainfall and CO_2 series.

Anomaly detection
Deviations in server and sensor metrics.

Economics
Forecasting inflation, unemployment and growth.

PRACTICE

Risks and Ethics

Leaking the future

A random split, or scaling with statistics of the whole series, carries future values into training; test scores become fake.

Structural breaks

In the 2020 pandemic, demand and traffic series changed overnight; models that fit the past well were wrong for weeks.

Forecasts without intervals

A single number suggests certainty; stock or energy decisions get made without knowing the uncertainty.

Forecasts that change behaviour

When everyone acts on the same price or demand forecast, the series moves; a forecast can fulfil or defeat itself.

In depth: Explainable AI and AI Ethics.

PRACTICE

Tools and Datasets

Tools and libraries

pandas | statsmodels | Prophet | sktime | Darts | R forecast / fable

Well-known datasets

Airline passengers (Box-Jenkins) | Sunspot numbers | M4 and M5 competitions

UCI Electricity Load Diagrams | Mauna Loa CO_2 measurements | Jena climate data

PRACTICE

Common Misconceptions

Common belief	In fact
Time series data can be tested with a random split.	Future values leak into training; testing always follows time order.
Deep networks beat classic models on every series.	In the M competitions, simple and statistical methods have often been very strong.
A good model gives one correct number.	Forecasts are uncertain; for decisions the **prediction interval** matters at least as much as the point forecast.
A model that worked in the past always works.	Structural breaks (a pandemic, a crisis) change the series; models must be monitored and refreshed.

NEXT STEPS

After This Course

Courses that build on or extend this one

This course is mostly an end point: it builds on the others.

Related careers

Forecasting analyst | Data scientist | Quantitative analyst | Demand planner | Energy data analyst

After this course you can move on to frequency analysis with signal processing, transformer-based forecasting with deep learning, and multi-temporal satellite images with remote sensing. Econometrics and state space models are natural next steps too.

NEXT STEPS

Further Reading

- Hyndman and Athanasopoulos, *Forecasting: Principles and Practice*, 3rd ed., OTexts, 2021.
- Box, Jenkins, Reinsel and Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed., Wiley, 2015.
- Shumway and Stoffer, *Time Series Analysis and Its Applications*, 4th ed., Springer, 2017.
- Brockwell and Davis, *Introduction to Time Series and Forecasting*, 3rd ed., Springer, 2016.

The End

Chapter 29: Time Series Analysis and Forecasting

← Previous chapter · Map of the field · Next chapter →