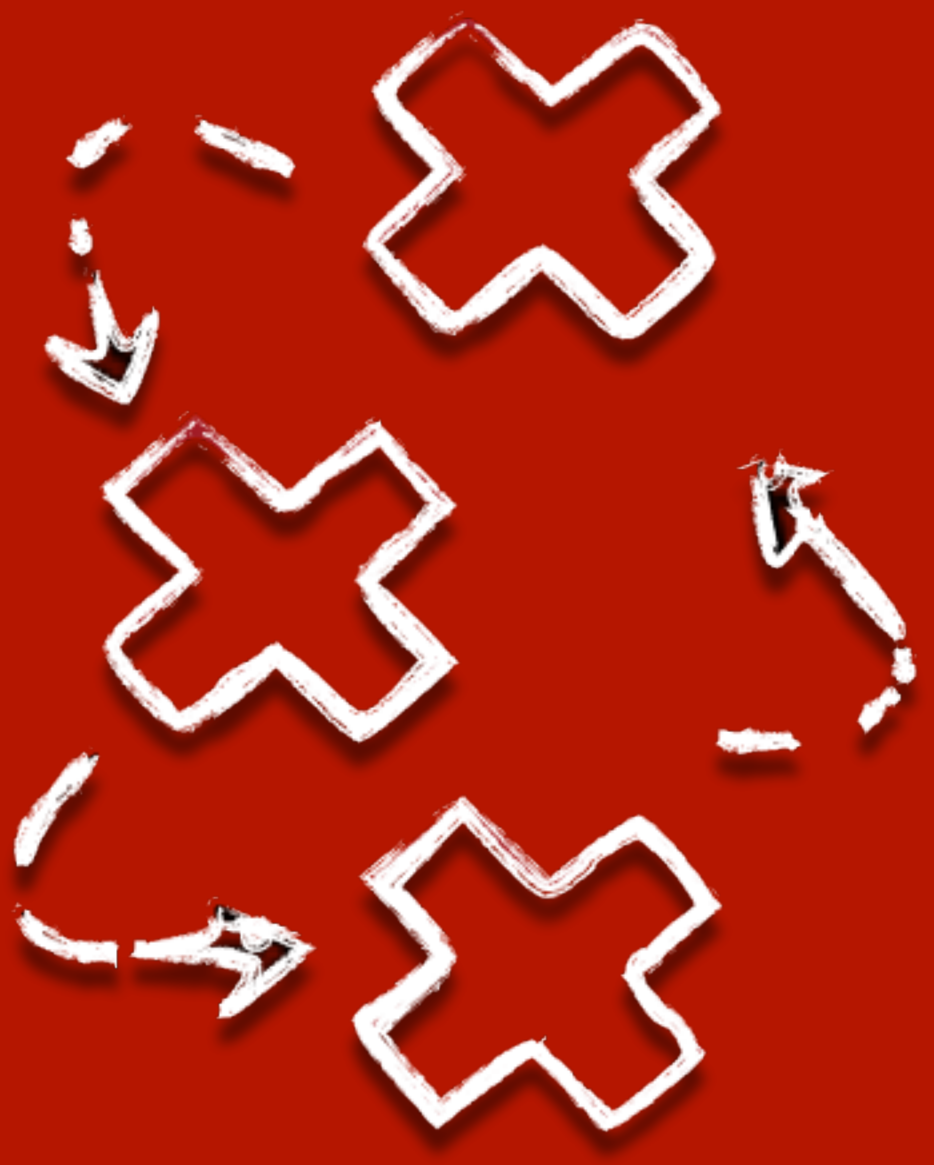

Seminar: Advanced topics in causality and causal ML research

Lecture 1.1: Logistics

Lecturer: Sara Magliacane

UdS - Summer 2026

Lecturer: prof. dr. Sara Magliacane

Currently:

Professor in Causal Machine Learning **(since March)**

Professor at AMLab (part-time from March)

Previously:

MIT-IBM Watson AI Lab — Research Staff Member & Affiliated Professor

IBM Research — Postdoc

Introduction - an overview of my research

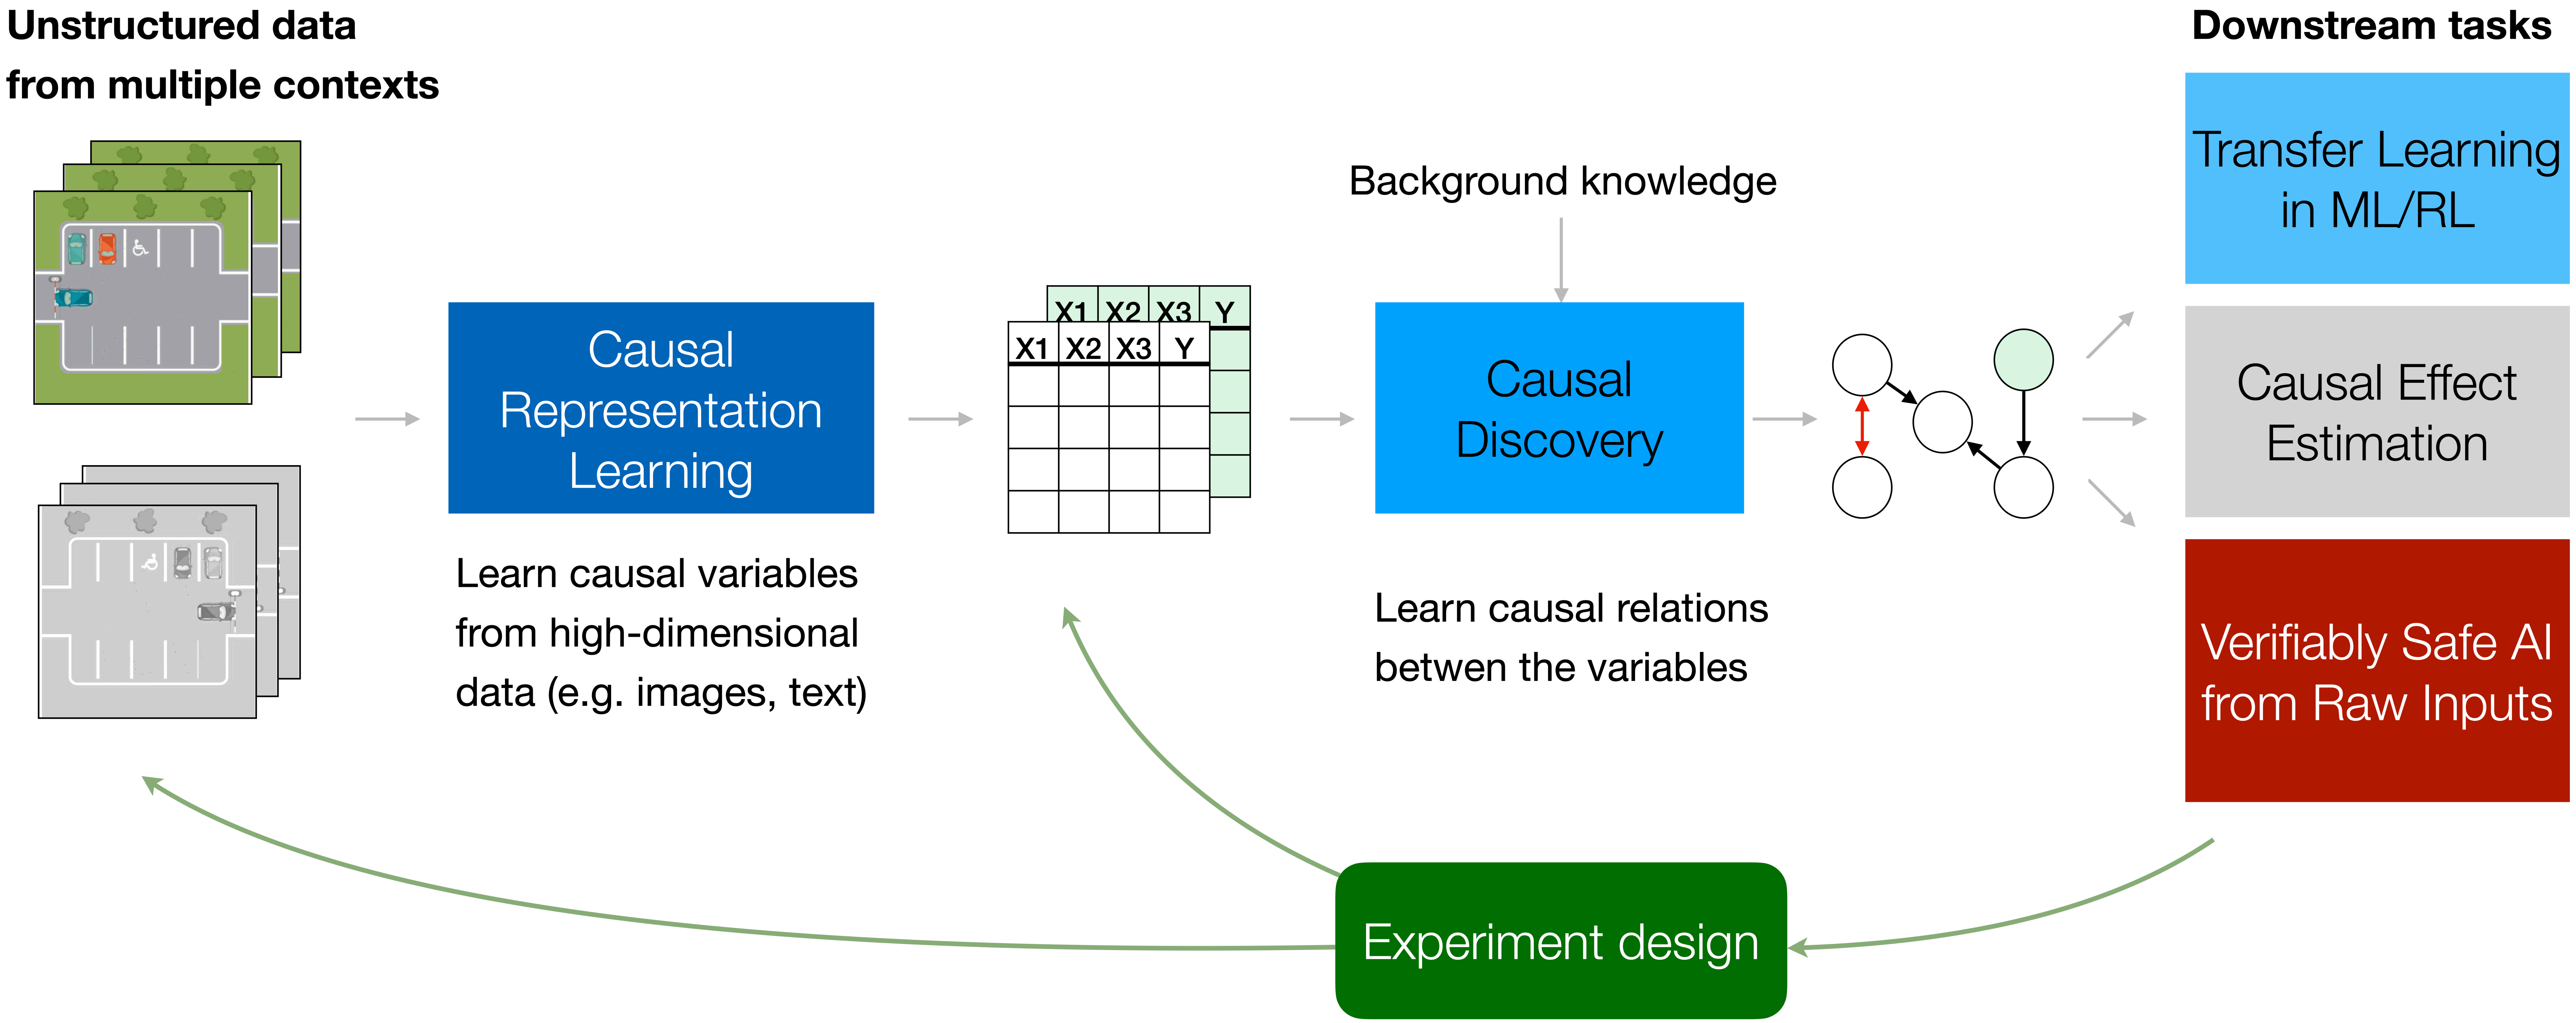

What about you? Round of introductions

- What is your background?
- What is your experience with causality?
- What do you expect from the course?

Tentative timetable

When	Where	Activity
April 16th at 10-12	SR106	Inperson lecture - Course logistics + Intro to causality
April 23rd at 10-12	https://uva-live.zoom.us/j/63564721506	Online lecture - Causal discovery
April 30th at 10-12	https://uva-live.zoom.us/j/63564721506	Online lecture - Causal representation learning
May 7th at 10-12	https://uva-live.zoom.us/j/63564721506	Online lecture - Causality-inspired ML/RL
May 21st at 10-12	SR106	Progress meeting with groups
May 28th at 10-12	SR106	Progress meeting with groups
June 11th at 10-12	SR106	Progress meeting with groups
June 25th at 10-12	SR106	Progress meeting with groups
July 2nd at 10-12	SR106	Progress meeting with groups
July 13-17th (TBD)	SR106	Group Presentations

Background sessions

Choose a paper from the list to reproduce and extend in a group

Projects page on the course CMS

https://cms.sic.saarland/causalml_26/2/List_of_papers_code_for_group_projects

List of papers+code for group projects

Manage Content | + New Page | Edit Page | Delete

General advice

The plan is for each group to choose one paper, read it carefully and reproduce the experiments in the original paper (or at least as much as possible), and propose some small extension. I will be helping with advice and evaluate which extensions are feasible within the course timeframe.

Possible ideas:

- A simple (or not, depending on the ambition) possible extension is to apply the methods to new (ideally real-world) datasets or problem
- Adapting the methods to new settings beyond their current theory or relaxing assumptions at least empirically is always interesting
- Creative combinations of methods are also welcome

Caveats:

- I just moved to UdS so I'm not exactly sure about the computational resources available to you. Once we figure out what are you interested in, I will investigate, but in any case, there are always projects that can be run on laptops.
- I have selected some recent and exciting papers that (afaik) have reasonable or well-documented codebases. In most cases, I know the contributors, so we can ask for help if needed. In principle, you can also propose something else, but this means I might not be able to give you the same level of advice.

Current list of papers (1/2)

Causal discovery

Methods to improve the scalability of causal discovery:

- Sequential Non-Ancestor Pruning - SNAP [Schubert et al. 2025] - project page with paper and code
- Local Optimal Adjustments discovery - LOAD [Schubert et al. 2026] - project page with paper and code

Methods to improve accuracy of causal discovery by integrating background knowledge (e.g. from LLMs):

- Imperfect experts: Guess2Graph [Hiramath et al. 2026] - project page with paper and code
- Expert-in-the-loop [Ankan et al. 2025] - paper, example notebook

Causal representation learning (CRL)

Methods that can disentangle causal variables from realistic images based on interventions or actions:

- Interventional CRL methods: CITRIS [Lippe et al. 2022] and iCITRIS [Lippe et al. 2023] - project page with papers and code, example notebook
- CRL based on actions: BISCUIT [Lippe et al. 2023] - project page with paper, code, pretrained model and datasets

Methods that do not require actions, but instead assume multiple simultaneous views

- Multiview-CRL [Yao et al. 2024] - project page with paper and code

Current list of papers (2/2)

Downstream tasks / Causality-inspired ML/RL

Methods that exploit causality or CRL for guarantees in interpretability or XAI:

- Guarantees for Concept Bottleneck Models: [Fokkema et al. 2025] paper, code
- Interpretability of LLMs through mixtures of causal models: [Pislar et al. 2025] paper, code

Generalization in RL through causal representations:

- AdaRL: What, Where, and How to Adapt in Transfer Reinforcement Learning [Huang et al. 2022] project page with paper and code
- Causality-guided Self-adaptive Representation-based approach - CSR [Yang et al. 2025] project page with paper and code

Using causality to improve bandits:

- (a bit simpler) Structural causal bandits [Lee et al. 2018] project page with code and paper
- (extension of the previous one) Non-stationary causal bandits [Kwon et al. 2025] project page with code and paper

Grading policies

- The grading is based on the group projects and each group member should contribute significantly to the project. The grade will be based on this contribution, in particular:
- **Project presentations (50%)** - each group member is expected to present for 10 minutes and answer questions for 5 minutes
- **Jupyter notebook (50%)** - code + a little bit of theoretical explanation, each group member should mention what part they wrote

Example notebook for style (note that this is only a tutorial so less than what we are expected here)

https://uvadlc-notebooks.readthedocs.io/en/latest/tutorial_notebooks/DL2/Causality_and_CRL/citris-tutorial.html

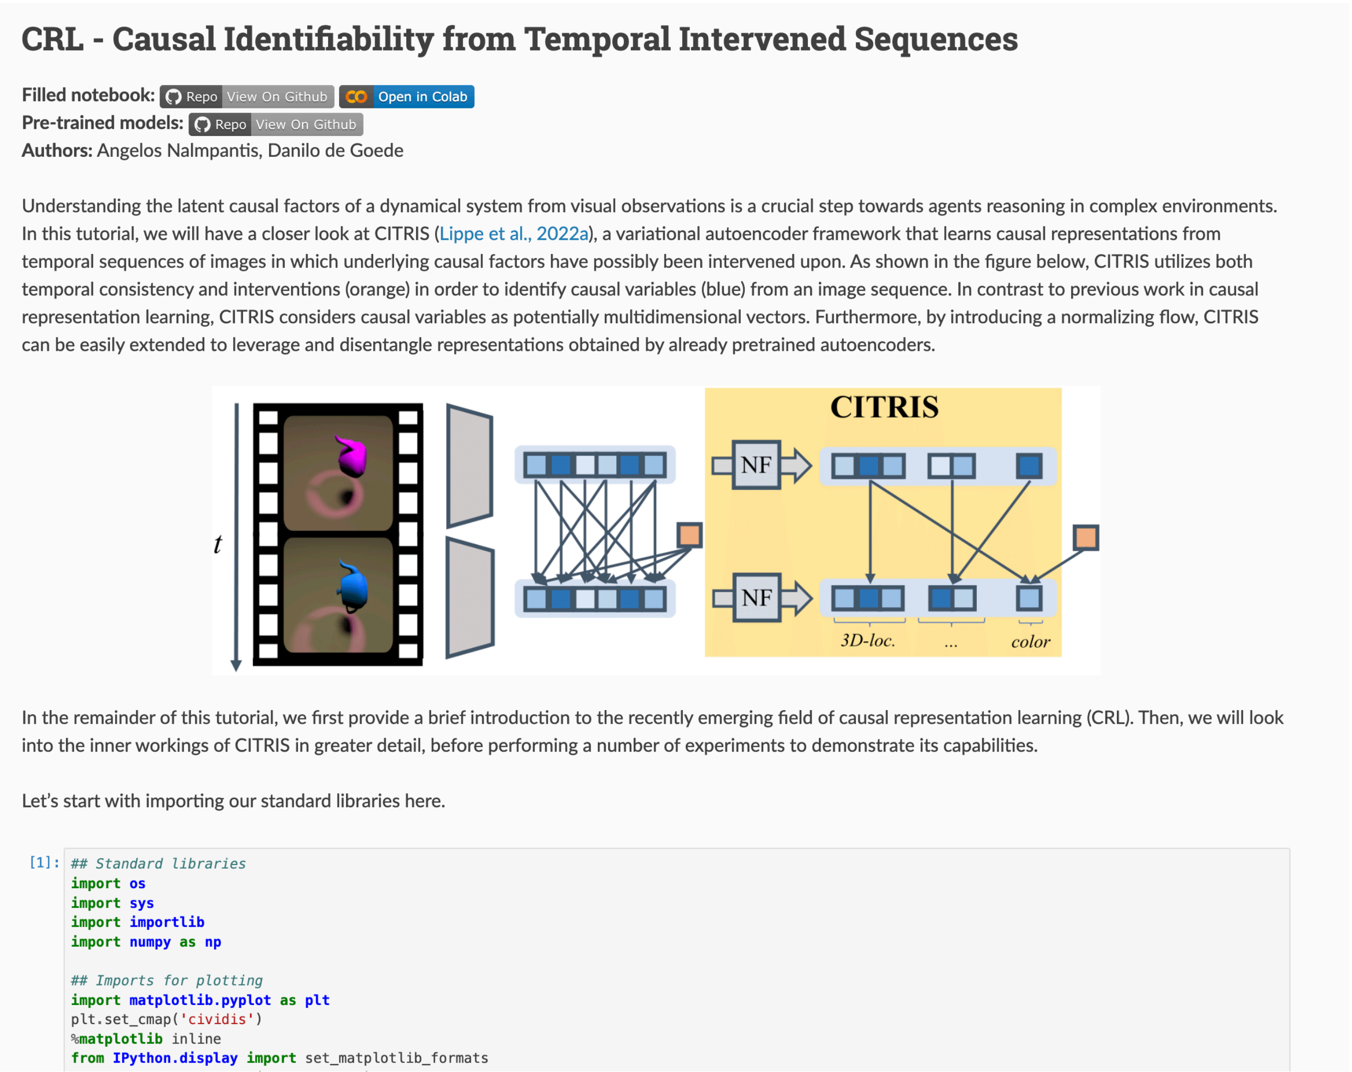

CRL - Causal Identifiability from Temporal Intervened Sequences

Filled notebook: Repo View On Github | Open in Colab
Pre-trained models: Repo View On Github
Authors: Angelos Nalmpantis, Danilo de Goede

Understanding the latent causal factors of a dynamical system from visual observations is a crucial step towards agents reasoning in complex environments. In this tutorial, we will have a closer look at CITRIS (Lippe et al., 2022a), a variational autoencoder framework that learns causal representations from temporal sequences of images in which underlying causal factors have possibly been intervened upon. As shown in the figure below, CITRIS utilizes both temporal consistency and interventions (orange) in order to identify causal variables (blue) from an image sequence. In contrast to previous work in causal representation learning, CITRIS considers causal variables as potentially multidimensional vectors. Furthermore, by introducing a normalizing flow, CITRIS can be easily extended to leverage and disentangle representations obtained by already pretrained autoencoders.

In the remainder of this tutorial, we first provide a brief introduction to the recently emerging field of causal representation learning (CRL). Then, we will look into the inner workings of CITRIS in greater detail, before performing a number of experiments to demonstrate its capabilities.

Let's start with importing our standard libraries here.

```
[1]: ## Standard libraries
import os
import sys
import importlib
import numpy as np

## Imports for plotting
import matplotlib.pyplot as plt
plt.set_cmap('cividis')
%matplotlib inline
from IPython.display import set_matplotlib_formats
```

Background textbooks (all available for free)

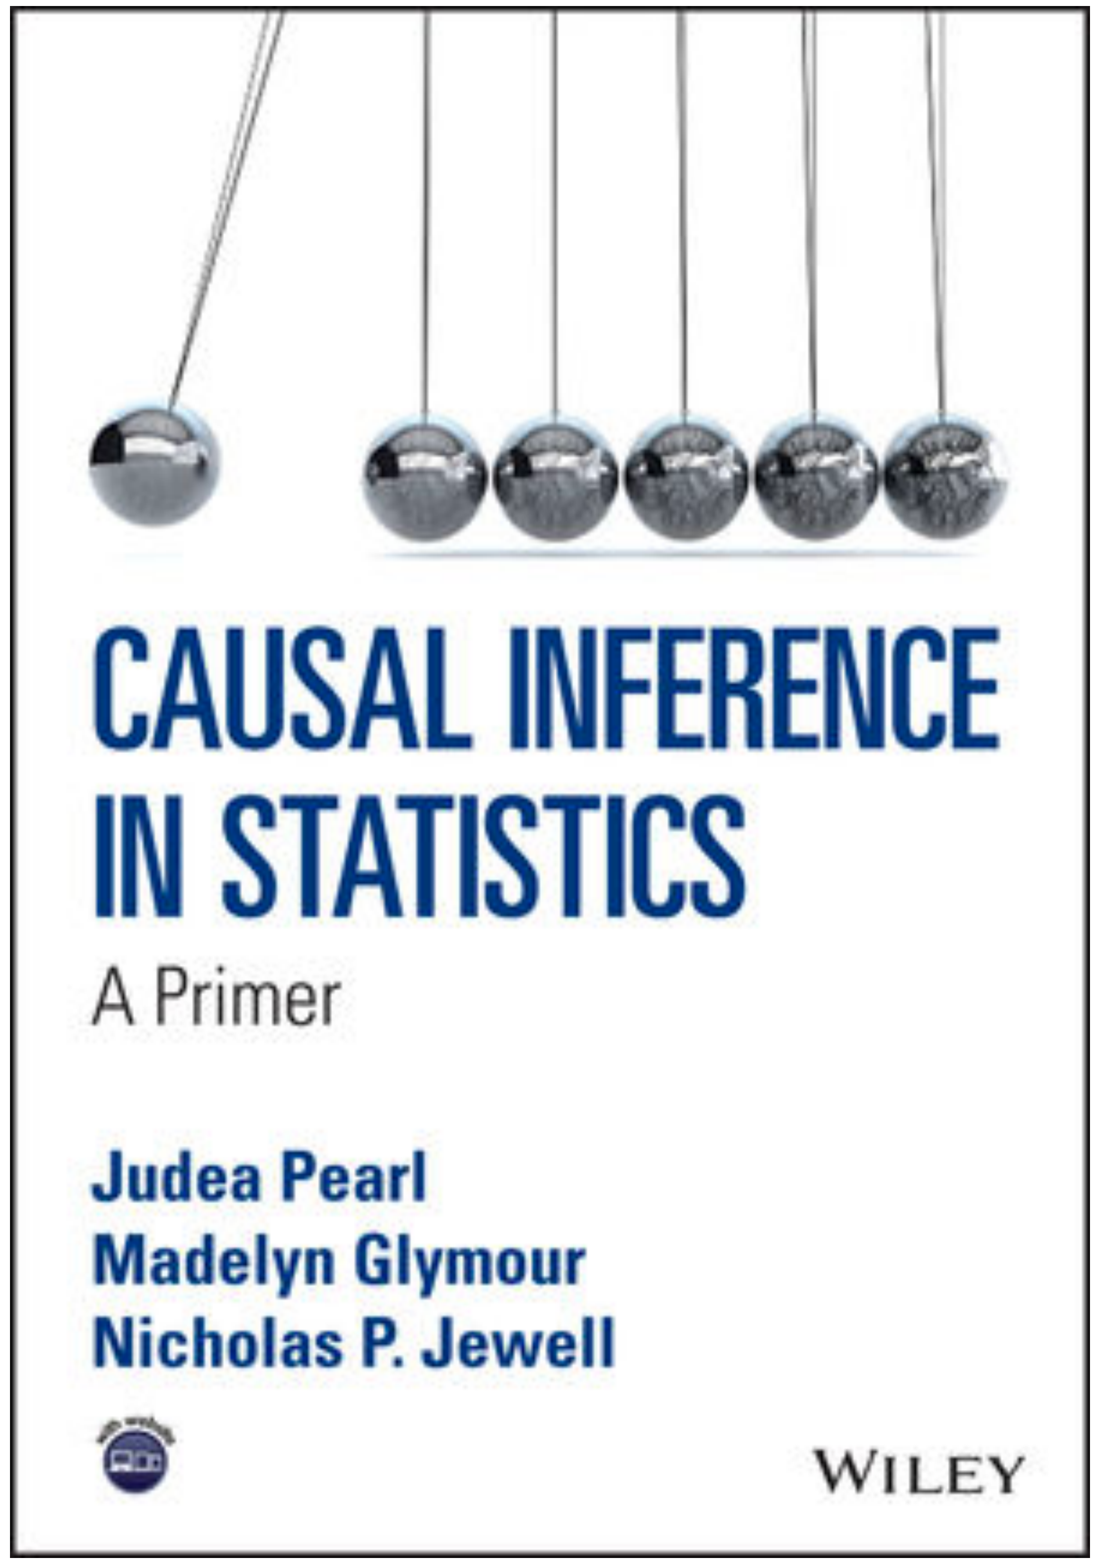

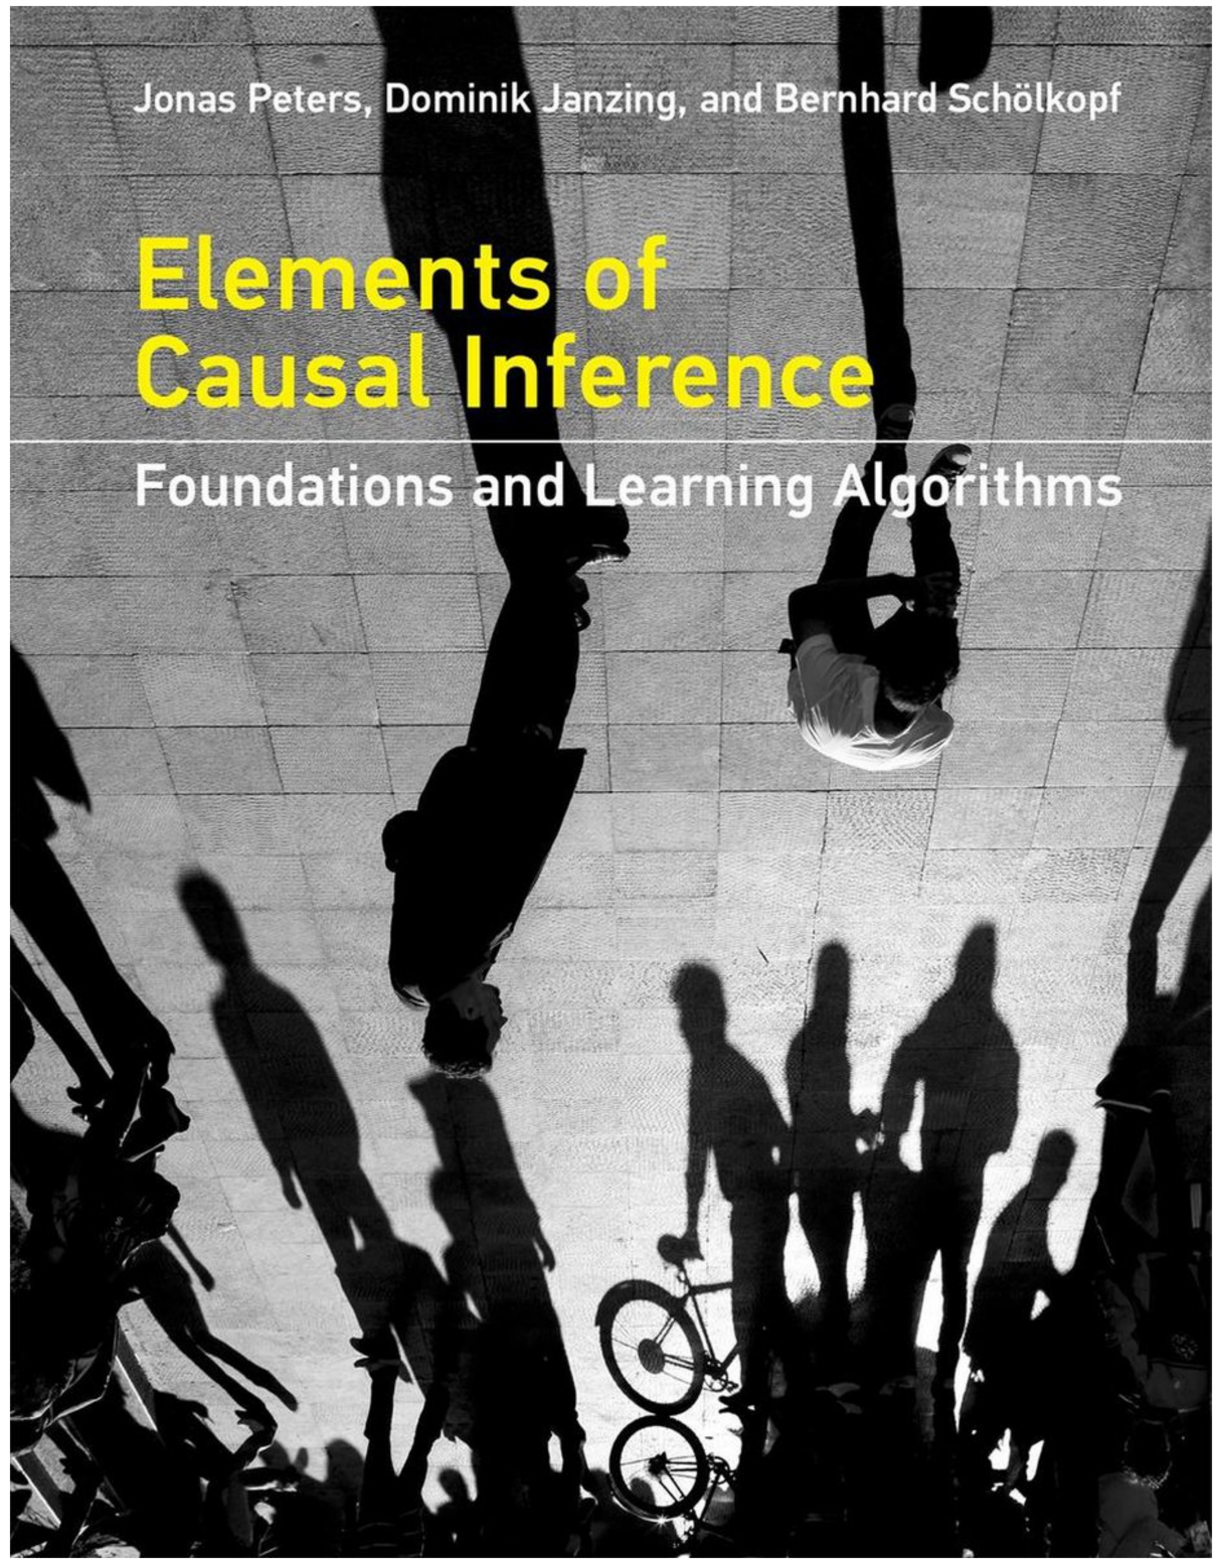

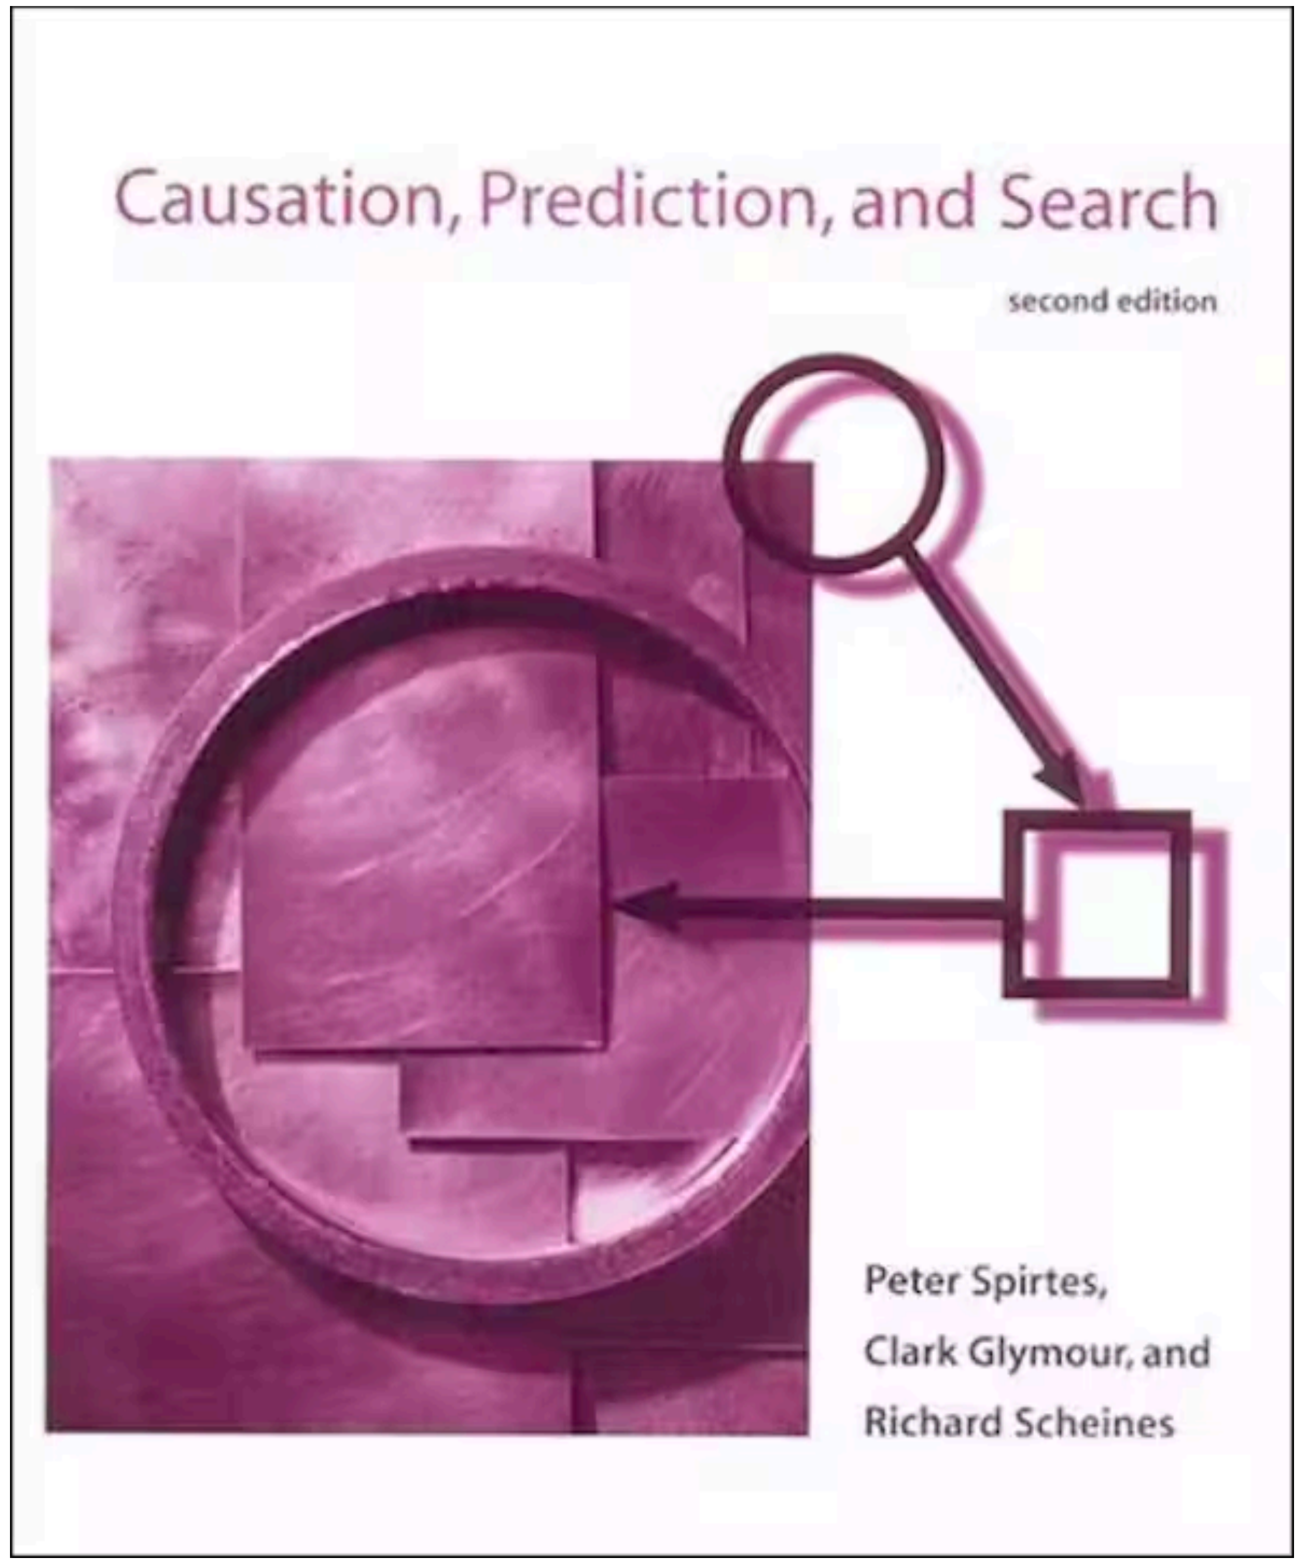

http://bayes.cs.ucla.edu/PRIMER/

http://web.math.ku.dk/~peters/jonas_files/ElementsOfCausalInference.pdf

https://www.researchgate.net/publication/242448131_Causation_Prediction_and_Search

Suggested reading for the first 4 weeks

- Week 1: Intro to causality
  - First two chapters of the Causality Primer by Pearl et al. 2021
- Week 2: Causal discovery
  - Chapter 5 (and possibly 6) of Spirtes, Glymour, Scheines, Causation, Prediction, Search
  - Survey by Glymour et al (2019)
- Week 3: Causal representation learning
  - Towards Causal Representation Learning by Schölkopf et al. 2021
- Week 4: Downstream tasks/causal ML
  - Causal Machine Learning survey by Kaddour et al. 2022 (includes also some of the previous topics)

If you are really curious… UvA course on Causality (24 hours)

1	Introduction
2	Probability recap
3	Graphical models, d-separation
4	Causal graphs, Interventions, SCMs
5	Covariate adjustment: backdoor criterion
6	Covariate Frontdoor criterion, Instrumental variables
7	Counterfactuals, potential outcomes, estimating causal effects 1
8	Estimating causal effects 2 (matching, IPW)
9	Constraint based structure learning
10	Score based structure learning, restricted models
11	Do-calculus, transportability, Joint Causal Inference
12	Causality-inspired ML, recap of the course

Background on causal graphs

We know the causal graph, how do we estimate causal effects?

What happens if the graph is unknown?

Cutting edge research

Videos: https://video.uva.nl/playlist/dedicated/0_ug2t9vyf/

Questions??

- If you have any follow-up question, please use the **Forum on the course page** to reduce duplication instead of using email, I'm still getting used to the UdS email :)

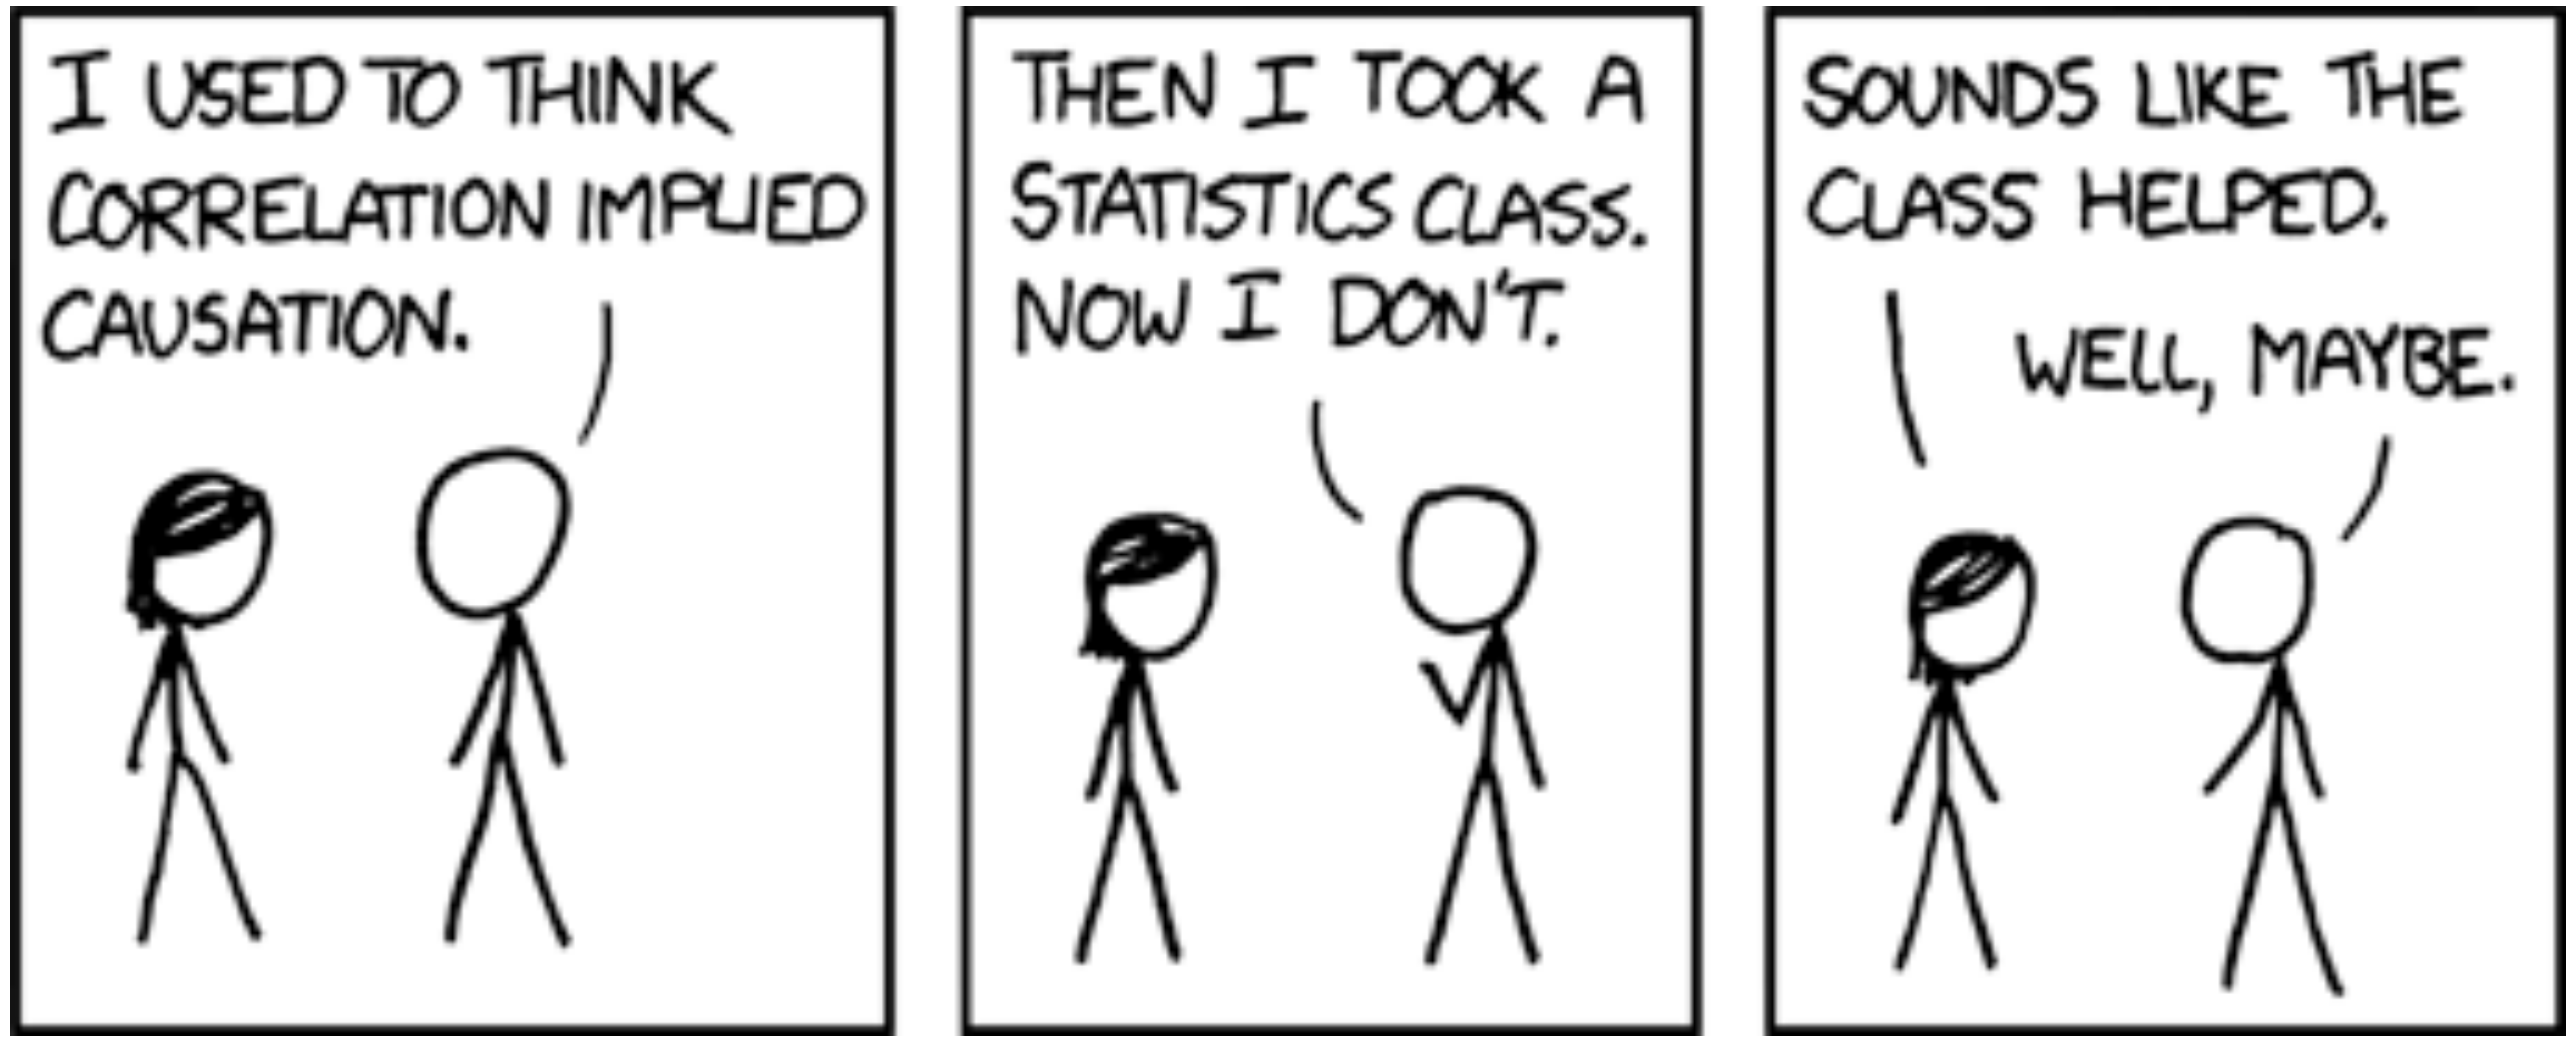

https://xkcd.com/552/