



TECHNISCHE
UNIVERSITÄT
DARMSTADT



AIML
Lab

Winter Semester 2025/26 Lecture

Causality for AI & ML

“Introduction”

Prof. Dr. Kristian Kersting


Moritz Willig


Today's speaker

Tim Woydt

Florian Busch

Matej Zečević






Section

1

Why Causality



Inference is Hard!



“What is the weight of the packet?”


Inference is Hard!



Causal Inference is Harder!

We need to understand the underlying *structure* of a problem to perform reliable inference.

- How do we discover causal structures from data?
- Can we always identify the true underlying structures?
- How to reason within causal structures.



Growing Usage of AI Tools

AI systems are in many parts of our everyday life.

- Summarizing texts.
- Taking the role of experts.
- Policy making.

Growing Usage of AI Tools



Growing Usage of AI Tools

The image shows a screenshot of an article from ITN (Imaging Technology News). The article is titled "Making the Role of AI in Breast Cancer Explainable" and is dated March 22, 2021. It discusses an analysis system for breast cancer diagnosis. A central heatmap visualization shows the detection of tumor-infiltrating lymphocytes (TILs) using explainable AI. The heatmap displays red and orange spots (representing TILs) against a background of blue and green (representing other tissues and cells). Below the heatmap, a summary table provides key statistics:

Metric	Value
TILs (as a percentage of the area)	62%
Lymphocytes	3172
Tumor Cells	1217
Streamed Area	0.5 mm ²

The article also includes a caption: "Detection of tumor-infiltrating lymphocytes (TILs) using explainable AI. The AI-technique is used to generate a heatmap showing TILs (red) and other tissues and cells (blue and green). Image courtesy of Klauschen/Charité".

Growing Usage of AI Tools

AI chatbots und summarise new
11 February 2025
Imran Rahman-Jones Technology repo

itn
IMAGING TECHNOLOGY NEWS

COVID-19 IMAGING WOMEN'S HEALTH INFORMAT TECHNOLOGY

NEWS | ARTIFICIAL INTELLIGENCE | MARCH 22, 2021

Making the Role of AI in Breast Cancer Ex

Analysis system for the diagnosis of breast cancer

62% TIL	3172 LYMPHOCYTES	1217 TUMOR CELLS	0.5 mm ² STREAMED AREA
------------	---------------------	---------------------	--------------------------------------

Detection of tumor-infiltrating lymphocytes (TILs) using explainable AI. The AI-technique is used to generate a heatmap showing TILs (red) and other tissues and cells (blue and green). Image courtesy of Klauschen/Charité

nature

Explore content Journal information Publish with us Subscribe

nature > outlook > article

A fairer way forward for AI in health care

Without careful implementation, artificial intelligence could widen health-care inequality.

Linda Nordling

The illustration depicts a group of diverse people in a public square or protest area. They are holding various signs, including one that says 'HEALTH EQUALITY' and another that says 'HEALTH = EQUALITY'. In the center, there is a large computer monitor on a stand, and a person is sitting at a desk with a keyboard. The scene suggests a public discussion or protest about the impact of AI on healthcare equity.

Growing Usage of AI Tools

MIT
Technology
Review

Artificial intelligence

What AI still can't do

Artificial intelligence won't be very smart if computers don't grasp cause and effect. That's something even humans have trouble with.

by **Brian Bergstein**

February 19, 2020



☰ Q

AI chatbot
summaries

11 February 2025

Imran Rahman-Jones

nature

scribe


health care

to widen health-care

Why do we care about Causality?

Understanding the structure of the underlying processes is a key component for making robust and reliable decisions.

You observe a correlation between ice cream sales and sunburns:




What is the right causal direction?

Why do we care about Causality?

Understanding the structure of the underlying processes is a key component for making robust and reliable decisions.

You observe a correlation between ice cream sales and sunburns:




*What is the right causal direction?
“Eating ice cream causes sunburns!”*

Why do we care about Causality?

Understanding the structure of the underlying processes is a key component for making robust and reliable decisions.

You observe a correlation between ice cream sales and sunburns:




*What is the right causal direction?
“Getting a sunburn gets you ice cream!”*

Why do we care about Causality?

Understanding the structure of the underlying processes is a key component for making robust and reliable decisions.

You observe a correlation between ice cream sales and sunburns:




*What is the right causal direction?
"Some factor that influences both!"*

Why do we care about Causality?

Understanding the structure of the underlying processes is a key component for making robust and reliable decisions.

You observe a correlation between ice cream sales and sunburns:

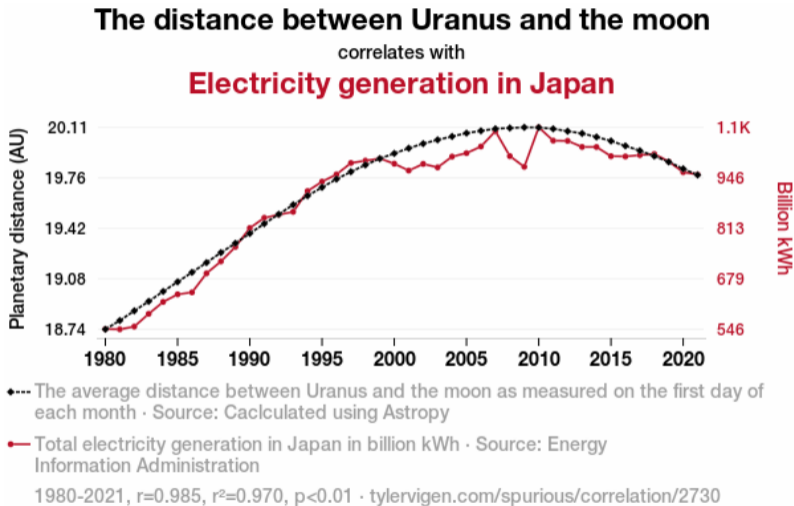


What is the right causal direction?

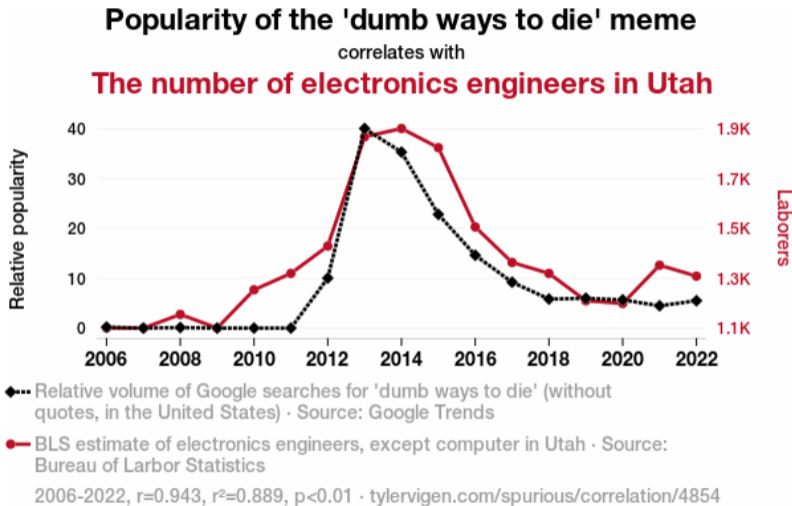
Observing correlations only is *inconclusive!*

Correlational guessing can be **arbitrarily wrong!**

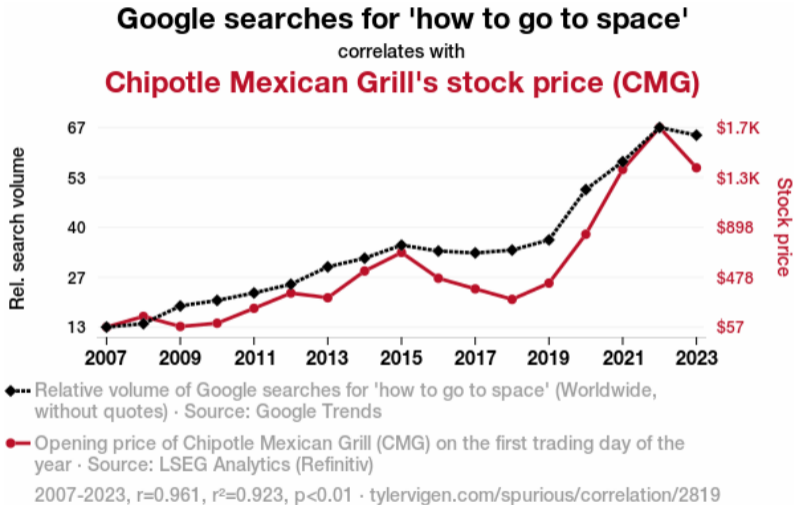
Spurious Correlations?



Spurious Correlations?




Spurious Correlations?



Correlation does not imply causation!



Reichenbach's Common Cause Principle




If two variables X , Y are correlated, there are three possible types of causal structure:

1. X and Y have a common cause Z .
2. X causes Y .
3. Y causes X .

Also: *Conditioning* on a common child of X and Y might correlate variables, leading to collider bias! (More on this in later lectures.)

Inferring Causation From Correlation?

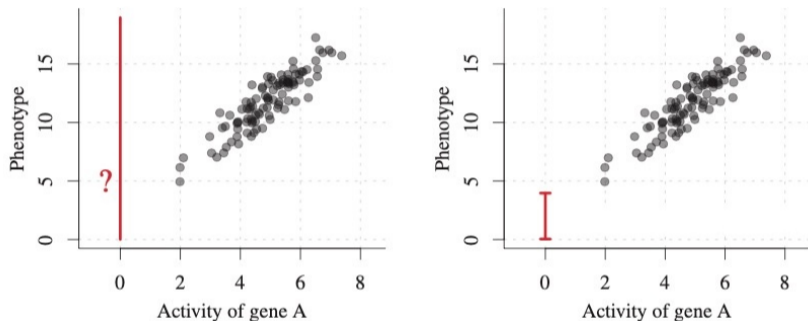


We observe a correlation between the phenotype and the activation of gene A.

Question: If we knock out gene A, will the phenotype, too, go to zero?

Figure: Peters et al., “Elements of Causal Inference”, MIT Press, 2017.

Inferring Causation From Correlation?




We observe a correlation between the phenotype and the activation of gene A.

Question: If we knock out gene A, will the phenotype, too, go to zero?

Answer: Yes! Gene A causes the phenotype.

Figure: Peters et al., "Elements of Causal Inference", MIT Press, 2017.

Inferring Causation From Correlation?




We observe a correlation between the phenotype and the activation of gene B.

Question: If we knock out gene B, will the phenotype, too, go to zero?

Figure: Peters et al., "Elements of Causal Inference", MIT Press, 2017.

Inferring Causation From Correlation?



We observe a correlation between the phenotype and the activation of gene B.

Question: If we knock out gene B, will the phenotype, too, go to zero?

Answer: No! There must be a common cause for gene B and the phenotype.

Figure: Peters et al., “Elements of Causal Inference”, MIT Press, 2017.

Inferring Causation From Correlation?

From a purely correlational perspective, both cases look initially the same.

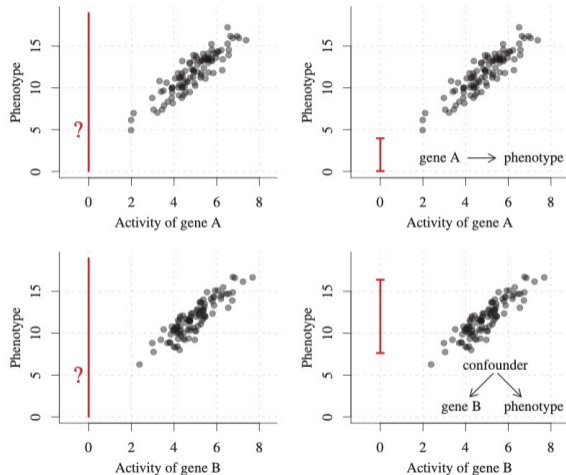


Figure: Peters et al., "Elements of Causal Inference", MIT Press, 2017.

Inferring Causation From Correlation?

From a purely correlational perspective, both cases look initially the same.

Question: Can we tell the outcome, if we come across new data for another gene?

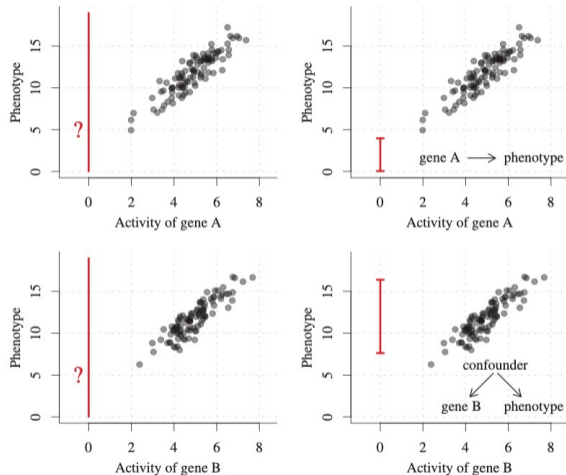


Figure: Peters et al., "Elements of Causal Inference", MIT Press, 2017.

Inferring Causation From Correlation?

From a purely correlational perspective, both cases look initially the same.

Question: Can we tell the outcome, if we come across new data for another gene?

Answer:

Best we can do^[*] is to say:

“We don’t know.”

[*] Without making any further assumptions.

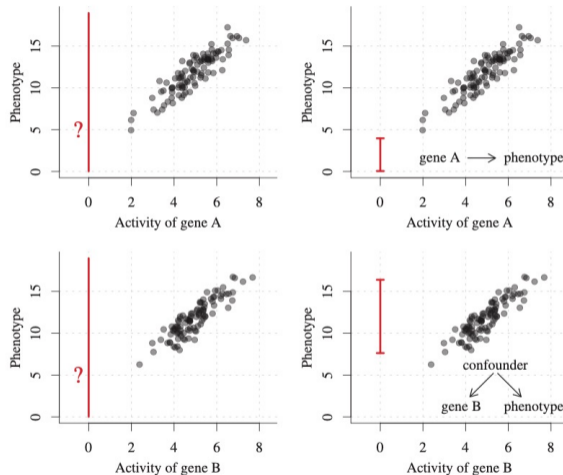


Figure: Peters et al., “Elements of Causal Inference”, MIT Press, 2017.

What can we do about it?

Decompose the problem:

1) Find ways to discover the underlying causal structures.

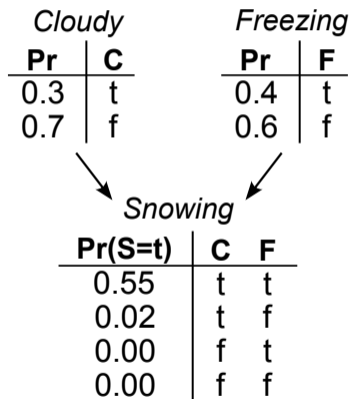
- Model causal relations as *directed acyclic graphs* (DAGs).
- Think about conditions under which we can discover causal directions from data.

2) Estimate causal effects using the causal graph and data.

- Use of the *do-calculus*, a set of rules that determine when causal effects *can be inferred* from data.

Graphical Causal Models

Causal Bayesian Networks



Graphical Causal Models

Causal Bayesian Networks

<i>Cloudy</i>		<i>Freezing</i>	
Pr	C	Pr	F
0.3	t	0.4	t
0.7	f	0.6	f

Snowing

Pr(S=t)	C	F
0.55	t	t
0.02	t	f
0.00	f	t
0.00	f	f

Structural Equation Models

Rain := Bernoulli(0.2)
Sprinkler := Bernoulli(0.5)
Slippery := Rain \vee Sprinkler

Graphical Causal Models

Causal Bayesian Networks

<i>Cloudy</i>		<i>Freezing</i>	
Pr	C	Pr	F
0.3	t	0.4	t
0.7	f	0.6	f

↓ ↓


Snowing

Pr(S=t)	C	F
0.55	t	t
0.02	t	f
0.00	f	t
0.00	f	f

Structural Equation Models

Rain := Bernoulli(0.2)
Sprinkler := Bernoulli(0.5)
Slippery := Rain \vee Sprinkler


Induced Graph:



Causal Discovery

Recall phenotype example? We were jumping right to answering the question.

Before doing inference, we need to **discover the underlying causal model.**



Causal Discovery

Different causal models lead to different outcome predictions.

Causal Discovery

Different causal models lead to different outcome predictions.

Naïve Approach: Iterate over all possible graphs and check for best fit?

#Variables	#DAGS
1	1
2	3
3	25
4	543
5	29281
6	3781503
7	1138779265
8	783702329343
9	1213442454842881
10	4175098976430598143
11	31603459396418917607425
12	521939651343829405020504063
13	18676600744432035186664816926721
14	51439428141044398334941790719839535103
15	237725265553410354992180218286376719253505

Causal Discovery

Different causal models lead to different outcome predictions.

Naïve Approach: Iterate over all possible graphs and check for best fit?

Super exponential growth!

Infeasible even for small problems.


#Variables	#DAGS
1	1
2	3
3	25
4	543
5	29281
6	3781503
7	1138779265
8	783702329343
9	1213442454842881
10	4175098976430598143
11	31603459396418917607425
12	521939651343829405020504063
13	18676600744432035186664816926721
14	51439428141044398334941790719839535103
15	237725265553410354992180218286376719253505

What can we do about it?

- 1) Find ways to discover the underlying causal structures.
 - Model causal relations as *directed acyclic graphs* (DAGs).
 - Think about conditions under which we can discover causal directions from data.
- 2) Estimate causal effects using the causal graph and data.
 - Use of the *do-calculus*, a set of rules that determine when causal effects *can be inferred* from data.


Conditioning \neq Intervening

Conditioning and intervening does not yield the same outcome!




Conditioning \neq Intervening

Conditioning and intervening does not yield the same outcome!




Conditioning \neq Intervening

Conditioning and intervening does not yield the same outcome!



Conditioning \neq Intervening

Conditioning and intervening does not yield the same outcome!



Pearlian Causality

Judea Pearl

ACM Turing Award in 2012 “for fundamental contributions to artificial intelligence through the development of a calculus for probabilistic and causal reasoning.”

Recall Bayesian Networks?



Causal analysis and modeling via **graphical models**.

do-calculus: set of sound and complete rules under which causal queries become identifiable!



Ladder of Causality

The **Ladder of Causality** defines three different levels (or rungs) of causal inference:

1. **Associational**: Observing
2. **Interventional**: Doing
3. **Counterfactual**: Imagining

Levels form a strict hierarchy

Associational \subset **Interventional** \subset **Counterfactual**
with higher levels being strictly more powerful than the lower ones.

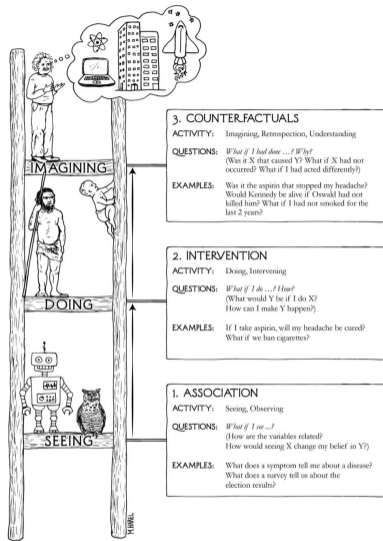



Figure: Pearl and Mackenzie, “The book of why”, Basic books, 2018.

Ladder of Causality - Associational




1. ASSOCIATION

ACTIVITY: Seeing, Observing

QUESTIONS: *What if I see ...?*
(How are the variables related?)
How would seeing X change my belief in Y?)

EXAMPLES: What does a symptom tell me about a disease?
What does a survey tell us about the election results?

Ladder of Causality - Interventional




2. INTERVENTION

ACTIVITY: Doing, Intervening

QUESTIONS: *What if I do ...? How?*
(What would Y be if I do X?
How can I make Y happen?)

EXAMPLES: If I take aspirin, will my headache be cured?
What if we ban cigarettes?

Ladder of Causality - Counterfactual



3. COUNTERFACTUALS

ACTIVITY: Imagining, Retrospection, Understanding

QUESTIONS: *What if I had done ...? Why?*
(Was it X that caused Y? What if X had not occurred? What if I had acted differently?)

EXAMPLES: Was it the aspirin that stopped my headache?
Would Kennedy be alive if Oswald had not killed him? What if I had not smoked for the last 2 years?

Do-calculus

Important result: Even if we know the correct causal graph, **some queries cannot be solved from pure observations!**

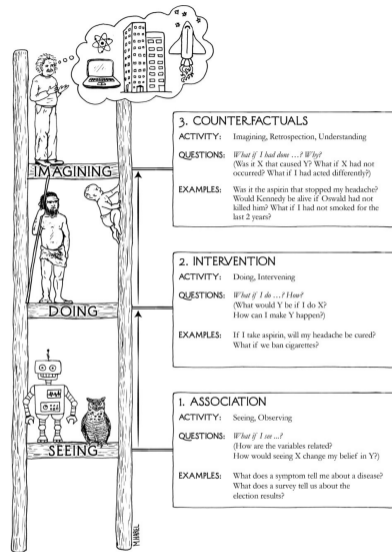


Figure: Pearl and Mackenzie, “The book of why”, Basic books, 2018.

Do-calculus

Important result: Even if we know the correct causal graph, **some queries cannot be solved from pure observations!**

The do-calculus tells us **if** ('completeness') and **how** ('sound') a causal query can be identified from data.

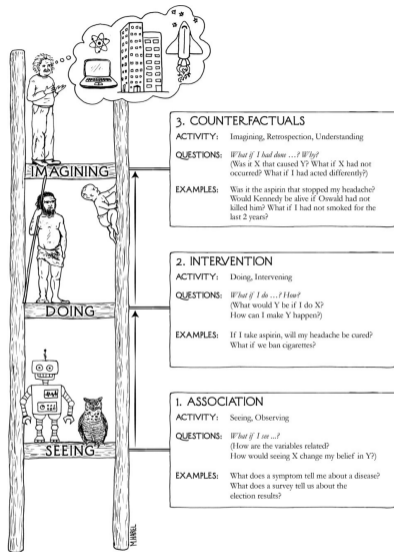


Figure: Pearl and Mackenzie, "The book of why", Basic books, 2018.

Do-calculus

Important result: Even if we know the correct causal graph, **some queries cannot be solved from pure observations!**

The do-calculus tells us **if** ('completeness') and **how** ('sound') a causal query can be identified from data.

Without making any further assumptions:

For some queries we can only find out by *doing* experiments.

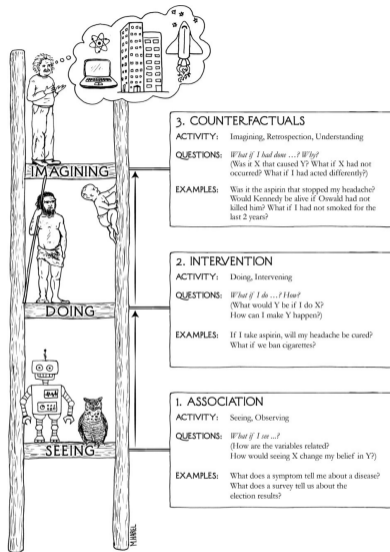


Figure: Pearl and Mackenzie, "The book of why", Basic books, 2018.

Causality in Philosophy


Causality is fundamental to how we make sense of the world. Already the ancient Greek philosophers discussed causal principles.

Universal Causation:

“Everything that becomes or changes must do so owing to some cause; for nothing can come to be without a cause.”

- *Plato in “Timaeus”, c. 360 BC.*


... in plain words: “nothing happens without a cause”.



Causality in Physics

The laws of physics govern and constrain the physical world. Time induces a causal ordering of events.


- Causes precede their effects.



Causality in Physics

The laws of physics govern and constrain the physical world. Time induces a causal ordering of events.


- Causes precede their effects.
- Causes of an event must be in the past light cone of the event.



Causality in Physics


The laws of physics govern and constrain the physical world. Time induces a causal ordering of events.

- Causes precede their effects.
- Causes of an event must be in the past light cone of the event.
- Causes can not have an effect outside their future light cone.



Cognitive Science


Humans use counterfactual simulations to anticipate outcomes and attribute responsibility.



Gerstenberg, "Counterfactual simulation in causal cognition"
Trends in Cognitive Sciences, 2024.

Fairness and Algorithmic Recourse

Causal models can be used to identify and provide algorithmic recourse in unfair or biased situations.




Willig et al., “When Causal Dynamics Matter: Adapting Causal Strategies through Meta-Aware Interventions” NeurIPS, 2025.

Medicine

"Does Smoking cause Cancer?"

The tobacco industry claimed for a long time that genes might be a confounding factor between adopting smoking and cancer.



According to a recent Nationwide survey: **MORE DOCTORS SMOKE CAMELS THAN ANY OTHER CIGARETTE!**

Tobacco of America's leading independent research organizations asked 111,107 doctors name the cigarette they themselves preferred to smoke.

The answers came in by the thousands... from doctors in big cities, in small towns... from general physicians, diagnosticians, surgeons, specialists... doctors in every branch of medicine!

Results? More doctors named Camel as their smoke than any other cigarette!

Yes, a doctor, too, smokes for pleasure. Full flavor, coolness, and mildness mean just as much to him as they do to any other smoker. Smoke Camels and see for yourself.




CAMELS Costlier

Tobacco boredpanda.com

Medicine

"Does Smoking cause Cancer?"

The tobacco industry claimed for a long time that genes might be a confounding factor between adopting smoking and cancer.



Causal methods provide a tool to falsify such claims.




According to a recent Nationwide survey: **MORE DOCTORS SMOKE CAMELS THAN ANY OTHER CIGARETTE!**



...

Create lame jokes by reversing causality?!






Section

2

Why AI and ML?



ML is Omnipresent

ML models have become part of our daily lives: AI assistants, video improvement, recommender systems, ...



The causal graph of some black-box machine, according to Gemini-2.5 Pro.

ML is Omnipresent

ML models have become part of our daily lives: AI assistants, video improvement, recommender systems, ...

Systems do not interact with the world, but are trained on large amounts of data.



The causal graph of some black-box machine, according to Gemini-2.5 Pro.

ML is Omnipresent

ML models have become part of our daily lives: AI assistants, video improvement, recommender systems, ...

Systems do not interact with the world, but are trained on large amounts of data.



The causal graph of some black-box machine, according to Gemini-2.5 Pro.

ML is Omnipresent

ML models have become part of our daily lives: AI assistants, video improvement, recommender systems, ...

Systems do not interact with the world, but are trained on large amounts of data.

“Do those models really *understand* what they are doing?”



The causal graph of some black-box machine, according to Gemini-2.5 Pro.

Do Image Models Capture Causal Relations?



Task:
"Make it a starlit night."

Do Image Models Capture Causal Relations?



StableDiffusion1.5 has learned: Night \rightarrow Lights \checkmark

Do Image Models Capture Causal Relations?




StableDiffusion1.5 has learned: Night \rightarrow Lights \checkmark
Causal outcomes are contained in the training data!

Do Image Models Capture Causal Relations?




Task:
"Turn on the lights."

Do Image Models Capture Causal Relations?



InstructPix2Pix has learned: Light \rightarrow Night ??

Do Image Models Capture Causal Relations?



InstructPix2Pix has learned: Light \rightarrow Night ??
... models also capture anti-causal correlations!

Do Image Models Capture Causal Relations?



Gemini-2.5 has learned: Light \nrightarrow Night \checkmark

Do Image Models Capture Causal Relations?



Gemini-2.5 has learned: Light \nrightarrow Night \checkmark

Under which circumstances can models tell apart causation from correlation?

Limitations of Correlational Learning

Consider the following setup:

Your task is to train a model that predicts B given A and C.



$$A := \mathcal{N}(0, 1.0)$$

$$B := A + \mathcal{N}(0, 3.1)$$

$$C := B + \mathcal{N}(0, 0.23)$$

Limitations of Correlational Learning

Consider the following setup:


Your task is to train a model that predicts B given A and C.



$$A := \mathcal{N}(0, 1.0)$$

$$B := A + \mathcal{N}(0, 3.1)$$

$$C := B + \mathcal{N}(0, 0.23)$$



Limitations of Correlational Learning

Consider the following setup:

Your task is to train a model that predicts B given A and C.



A correlational model will infer B from C because it is more predictive!

```
--- Linear Regression Results ---  
Learned coefficients for A: 0.0419  
                           C: 0.8919  
Intercept: 0.0500  
-----
```

Limitations of Correlational Learning

Consider the following setup:

Your task is to train a model that predicts B given A and C.



A correlational model will infer B from C because it is more predictive!

Limitations of Correlational Learning

Consider the following setup:

Your task is to train a model that predicts B given A and C.



A correlational model will infer B from C because it is more predictive!

What if C changes due to some environmental shift/intervention?

Limitations of Correlational Learning

Consider the following setup:

Your task is to train a model that predicts B given A and C.



A correlational model will infer B from C because it is more predictive!

What if C changes due to some environmental shift/intervention?

The model will still use C to predict B. *It does not generalize!*

Limitations of Correlational Learning

Consider the following setup:

Your task is to train a model that predicts B given A and C.




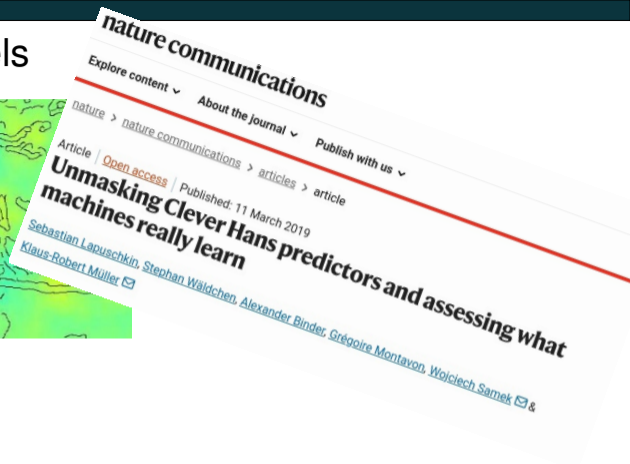

A correlational model will infer B from C because it is more predictive!

What if C changes due to some environmental shift/intervention?


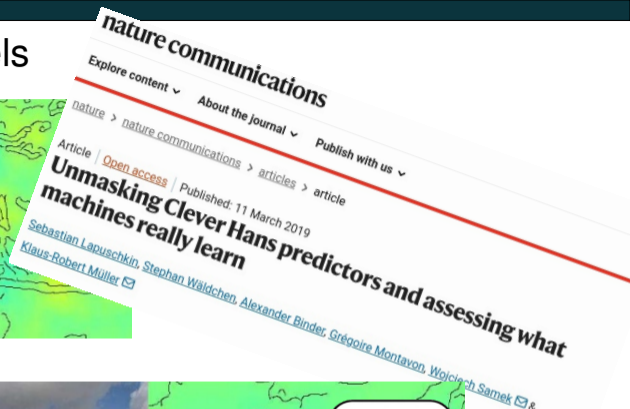

The model will still use C to predict B. *It does not generalize!*

This applies to most associational learners including Neural Networks!


Happens to your Everyday Models



Happens to your Everyday Models




THAT'S A HORSE!



Natural Language Data as an Opportunity

Natural language data contains information *about* causal quantities.
LLMs might be able to infer causal relations from this information.



Zečević* and Willig* et al., "Causal Parrots: Large Language Models May Talk Causality But Are Not Causal" TMLR 2023.

You develop AI models - You are responsible

We will show you how to use data and develop models in a causally principled manner.

- Know the limitations of what can be done with observational data.
- Be aware of the pitfalls of correlational shortcut learning.
- Avoid wasting compute.

You develop AI models - You are responsible

We will show you how to use data and develop models in a causally principled manner.


- Know the limitations of what can be done with observational data.
- Be aware of the pitfalls of correlational shortcut learning.
- Avoid wasting compute.

'Will your healthcare models forbid ice cream sales to prevent sunburns?'



“To Build Truly Intelligent Machines,
Teach Them Cause and Effect”


- Judea Pearl in “The Book of Why”
and in an interview with Quanta magazine (2018)



Section

3

Logistics



Format

3 credit point course.

Lecture every Friday 9:50-11:30, Room S202/C120.

One lecture per week (15 lectures \times 100 min) / **No separate exercise session.**

- 30 minutes exercises and recap (not mandatory / no bonus).
- 70 minutes 'classical' lecture.

Classical lecture part will be recorded. No live-streaming.

Format

3 credit point course.

Lecture every Friday 9:50-11:30, Room S202/C120.

One lecture per week (15 lectures \times 100 min) / **No separate exercise session.**

- 30 minutes exercises and recap (not mandatory / no bonus).
- 70 minutes 'classical' lecture.

Classical lecture part will be recorded. No live-streaming.

Grading: Final written **90 min exam** at the end of the semester.

Likely date: 18. March 2026.

Course will be held in English.

Feel free to ask **questions** in either **English or German.**

Platform

Moodle Course: (Informatik Moodle - please join on your own!)


<https://moodle.informatik.tu-darmstadt.de/course/view.php?id=1851>

Forum for questions.

Upload of weekly exercise sheets. (Starting from the second lecture.)

General questions/thoughts/feedback: Ask in the Moodle forum.


Questions of single concern: Send a mail to moritz.willig@cs.tu-darmstadt.de or just approach us at the lecture.



Section

4

AIML Lab



Who we are

AIML Lab led by Prof. Dr. Kristian Kersting
~40 PhD students, postdocs and researchers on various
topics of AI and ML.

Find us here: <https://www.aiml.informatik.tu-darmstadt.de/>



Causality subgroup (by seniority):

Matej Zečević



Moritz Willig



Florian Busch



Tim Woydt



Nicholas Tagliapietra



People



Applied Causality
Causality & LLM
Tractable Causality



Tractable Causality
Continual Causality




Large Scale Causality
Agentic LLM

Meta-Causality
Causality & LLM



Meta-Causality
Actual Causality






Section

5

On this Course



Suggested Courses

No formal requirements, but...

...good understanding of (Bayesian) Probabilities is helpful.

(Recap in next lecture.)

Helpful courses / concepts to know:

- Probabilistic Graphical Models
- Deep Learning (Training NNs, LLMs, Computer Vision, ...)

What to Expect

- Introduction to Pearl's Causality.
- Insights for training machine learning models that are causally sound.
- General applications for reasoning, fairness, explainability, ...)

Topics I - *Classical Causality*

1. Probabilistic Graphical Models (Probabilities, d-separation, Inference)
2. Structural Causal Models, Causal Hierarchy and do-calculus.
3. Counterfactuals (Implications, Twin Networks, Backtracking)
4. Causal Discovery (Independence- and Score-Based Algorithms)
5. Causal Effect Estimation under Uncertainty
- ⋮

Topics II - *Representation Learning*

⋮

6. Causal Abstractions
 7. Causal Representation Learning (under Interventions)
 8. Causal Representation Learning under Sufficient Variability
 9. Neural Causal Models (Causal VAE, Causal Normalizing Flows, iSPN, ...)
 10. LLM & Causality
- ⋮

Topics III - *General Applications*

⋮

11. Causality & Fairness
12. Explainability & Root-case Analysis
13. Meta-Causality
14. Actual Causality
15. Recap & Questions

Materials

Judea Pearl, Causality, Cambridge University Press, 2009.

Peters et al., Elements of Causal Inference, MIT Press, 2017.



Halpern, Actual Causality, MIT Press, 2016.

Brady Neal's Free Online Course "Introduction to Causal Inference", 2019. <https://www.bradyneal.com/causal-inference-course> Good individual concepts

Jonas Peters Lecture Series "Causality", 2017.

<https://www.youtube.com/watch?v=zvrcyqcN9Wo>



Books are not required for the course, but can be helpful for reading more on particular topics.




Why Graphical Causal Models

“As X-rays are to a surgeon,
graphs are for causation.”

- Judea Pearl in Causality (2009)





See you next week!

