Causal Inference for Health Data

(STATS C160/C260 — Winter 2026)

Lecture 1: Introduction to Causal Inference

Drago Plečko

College | Physical Sciences
Statistics & Data Science

Causal Inference
FOR HEALTH DATA

Logistics

- When: Tuesday & Thursday, 11:00am
- Where: Mathematical Sciences 5200
- Office hours: Thursday 1-2pm, Mathematical Sciences 8105E (note: 1.15-2.15pm Jan/8)
- Course Website https://dplecko.com/cihd26

  password: causalhealth26
- BruinLearn page also active
- Syllabus published in the website/BruinLearn

Pre-requisites

- You'll need some mathematical background
  - Probabilistic Reasoning (foundations)
  - Statistics (sampling, experimental design)
  - Discrete math (search and graph algorithms)
- It's <u>not</u> a purely theoretical class
  - You need some programming experience (R)
  - Plus: experience with data analysis
  - Plus: experience in health sciences

Grading Scheme

	Undergraduate (C160)	Graduate (C260)
Homeworks [completion based]	10%	10%
Midterm Exam (Feb/12)	50%	—
Midterm Project Report (Feb/12)	—	30%
Final Project Report (Mar/17)	40%	60%

Expectations

From the Class

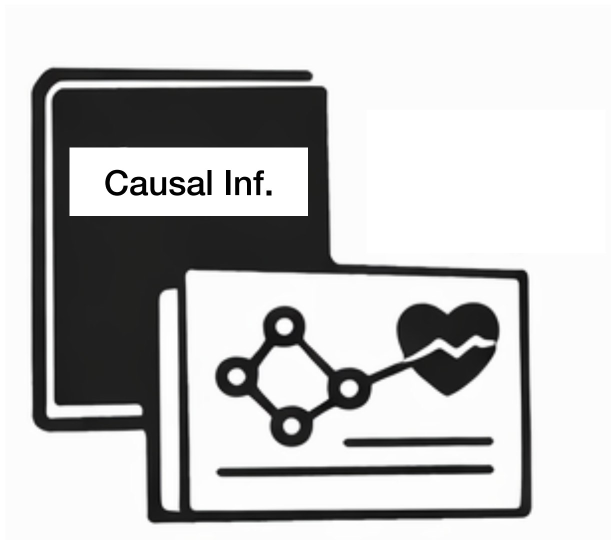

Theory

Methods

Real-World Data

From the Students

Homeworks (w/o AI)

Attend & Participate

Start Project Timely

Data Access for MIMIC-IV

- As said, in the course, we will study interesting real-world causal questions in health data,
- For following the course, getting access to the MIMIC-IV dataset (https://physionet.org/content/mimiciv/2.2/) is strongly recommended,
- For this, some basic training is required (CITI Program Certification),
- After obtaining access to MIMIC-IV, try setting up the ricu R-package (https://github.com/eth-mds/ricu).

Reading Materials

- Additional resources for reading:
  - **The Book of Why**
    J. Pearl, D. Mackenzie
    Basic Books, 2018
  - **Causality: Models, Reasoning, and Inference**
    J. Pearl
    Cambridge Press, 2000
  - **Causal Artificial Intelligence**
    E. Bareinboim
    https://causalai-book.net

Schedule (Tentative)

Week 1. Motivation, Pearl's Causal Hierarchy, SCMs

Weeks 2. d-separation, Layer 1, Identification of Effects

Weeks 3. Truncated Product, Backdoor, Estimation

Weeks 4. Estimation in Practice; Do-calculus

Week 5. Counterfactual Inference

Weeks 6. Heterogenous Effects + Midterm (Feb/12)

Weeks 7. Causal Explanations, Health Equity

Week 8. Sensitivity to Unobserved Confounding

Week 9. Missing Data & Measurement Error

Week 10. Causal Artificial Intelligence in Health Data

Causation in the Sciences

- Standard scientific methodology is built around the idea of combining observations and experiments.

- Scientists formulate hypotheses about unobserved causal mechanisms (e.g., Newtonian Physics).

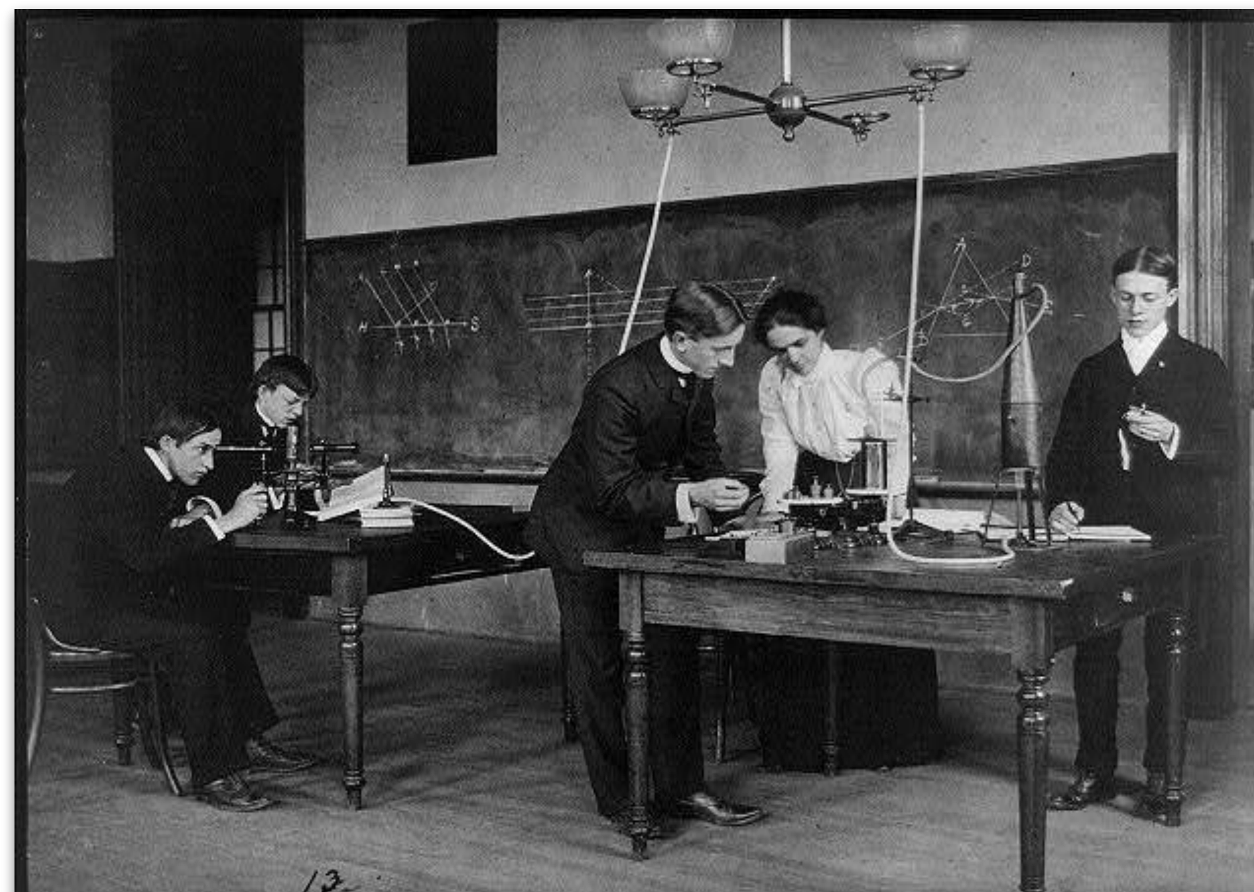

- Submit them to further observation and experimentation, then refine them continuously.

Causation in Medicine: Historical Primer — Scurvy

- in the 16th-18th century, colonial powers increasingly used ships for long voyages,
- soon, sailors began developing specific symptoms, which were shown to be highly lethal (50%-80% mortality) — today known as scurvy,
- the condition was a leading cause of death for sailors, and the cause was unknown,

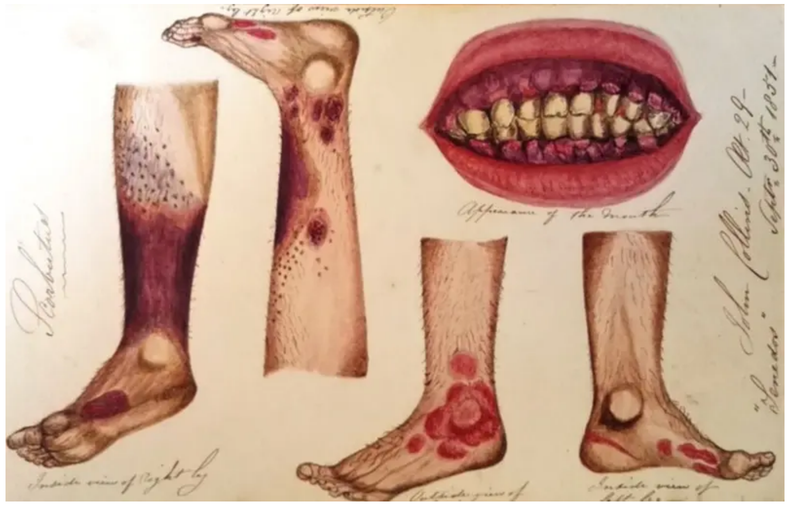

- theories about its cause included “bad air”, dietary problems, and even “laziness”.

Causation in Medicine: Scurvy

- James Lind started serving as Naval Surgeon in 1747, encountering numerous cases of scurvy,

- Lind studied **ship logs**, sailors' medical journals, Admiralty health reports, and previous naval voyages.

In modern terms: health data

- he collected anecdotes, and identified six possible cures: cider, vinegar, seawater, sulphuric acid, nutmeg paste, and citrus fruits.

Causation in Medicine: Scurvy

- Lind then decided to run an experiment: subjecting 12 sailors, in pairs of two, to the 6 possible medicines

In modern terms: clinical trial

- after several days, it was discovered that only citrus fruits resulted in recovery

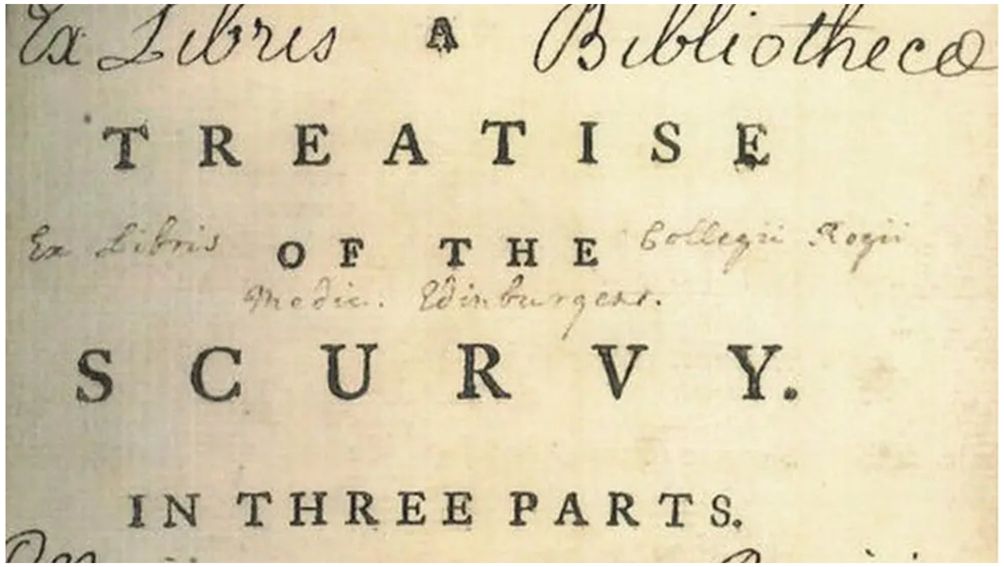
Ex Libris Bibliotheca
TREATISE
Ex Libris OF THE Collegii Regii
Medic. Edinburgens.
SCURVY.
IN THREE PARTS.

- Lind summarized his findings in the "Treatise of the Scurvy"

In modern terms: paper / book

Causal Understanding

- Decompose a phenomenon into *causal components* and describe how they interact.

- Uncover the underlying *data-generating process* (that is, how Nature works) or key aspects of it — how does each arrow above work, mechanistically?
- Infer what *would* or *could* happen under hypothetical or counterfactual situations.
  - What happens if we add seawater?
  - What happens if we add vitamin C?
  - What happens if voyages are made shorter?

+ many others …

The Modern Journey of Causality

Timeline of Modern Causality

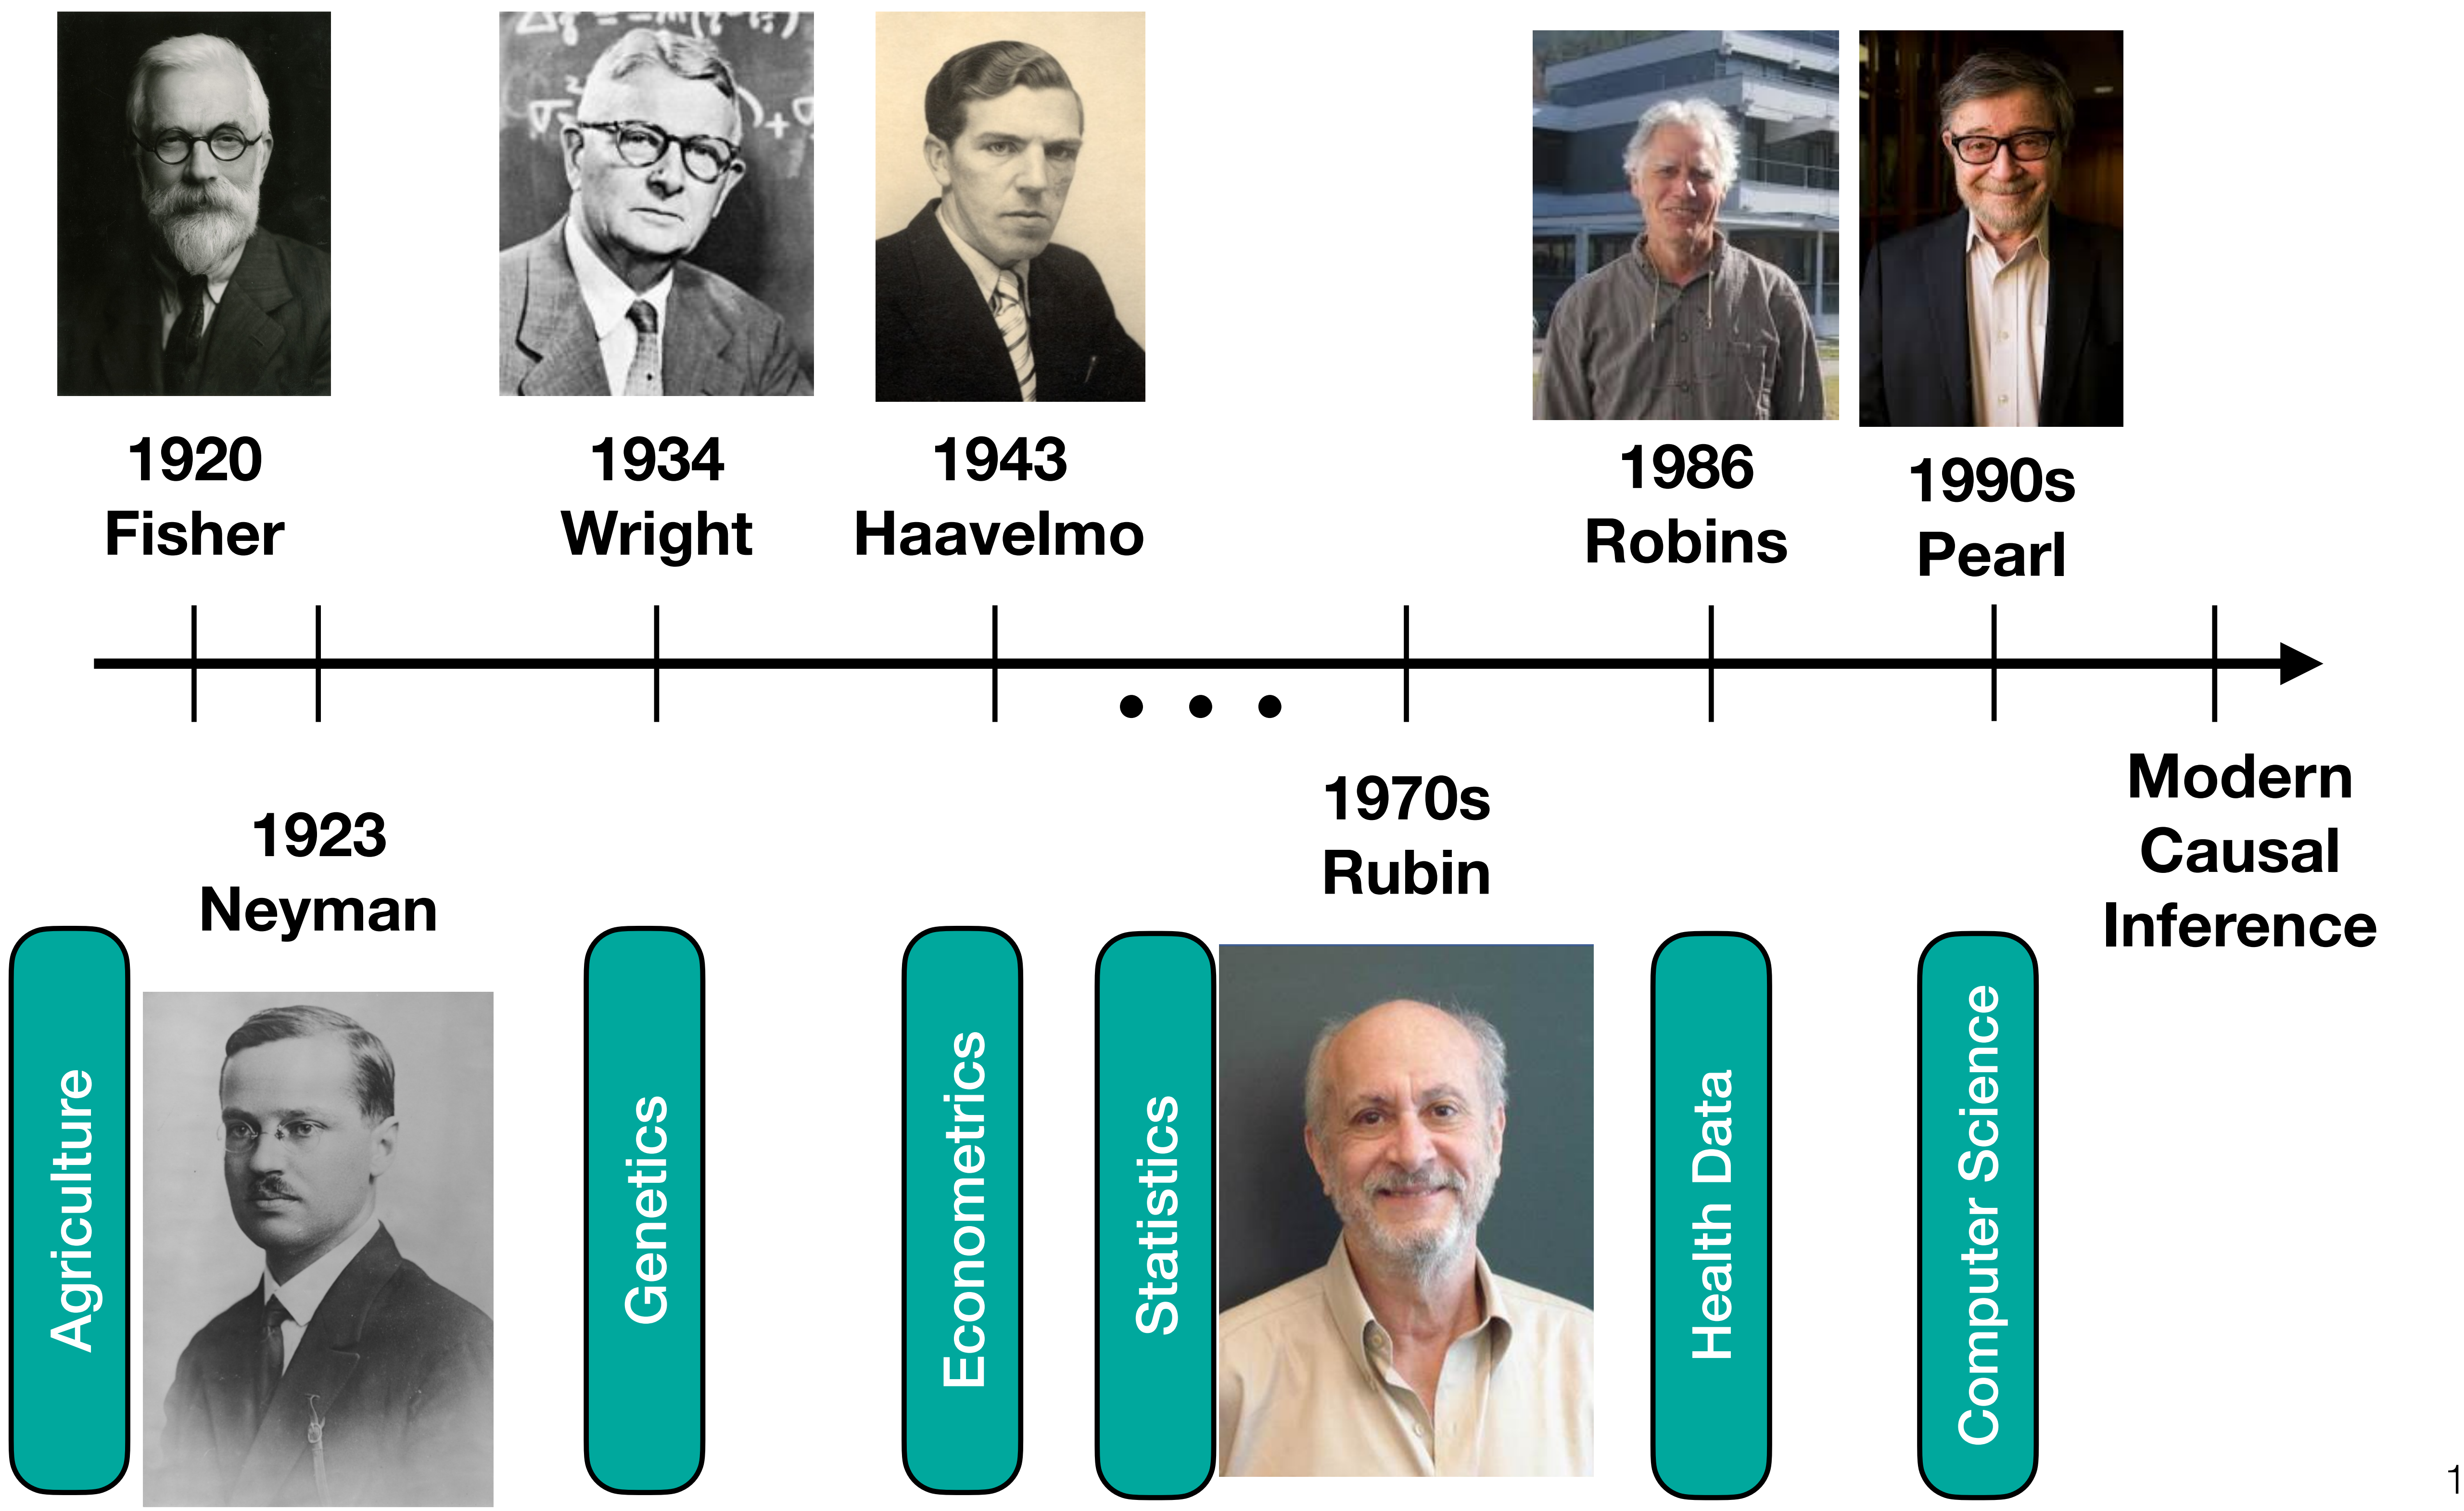

Why Causality Matters for Health Data

Why Causality Matters for Health Data

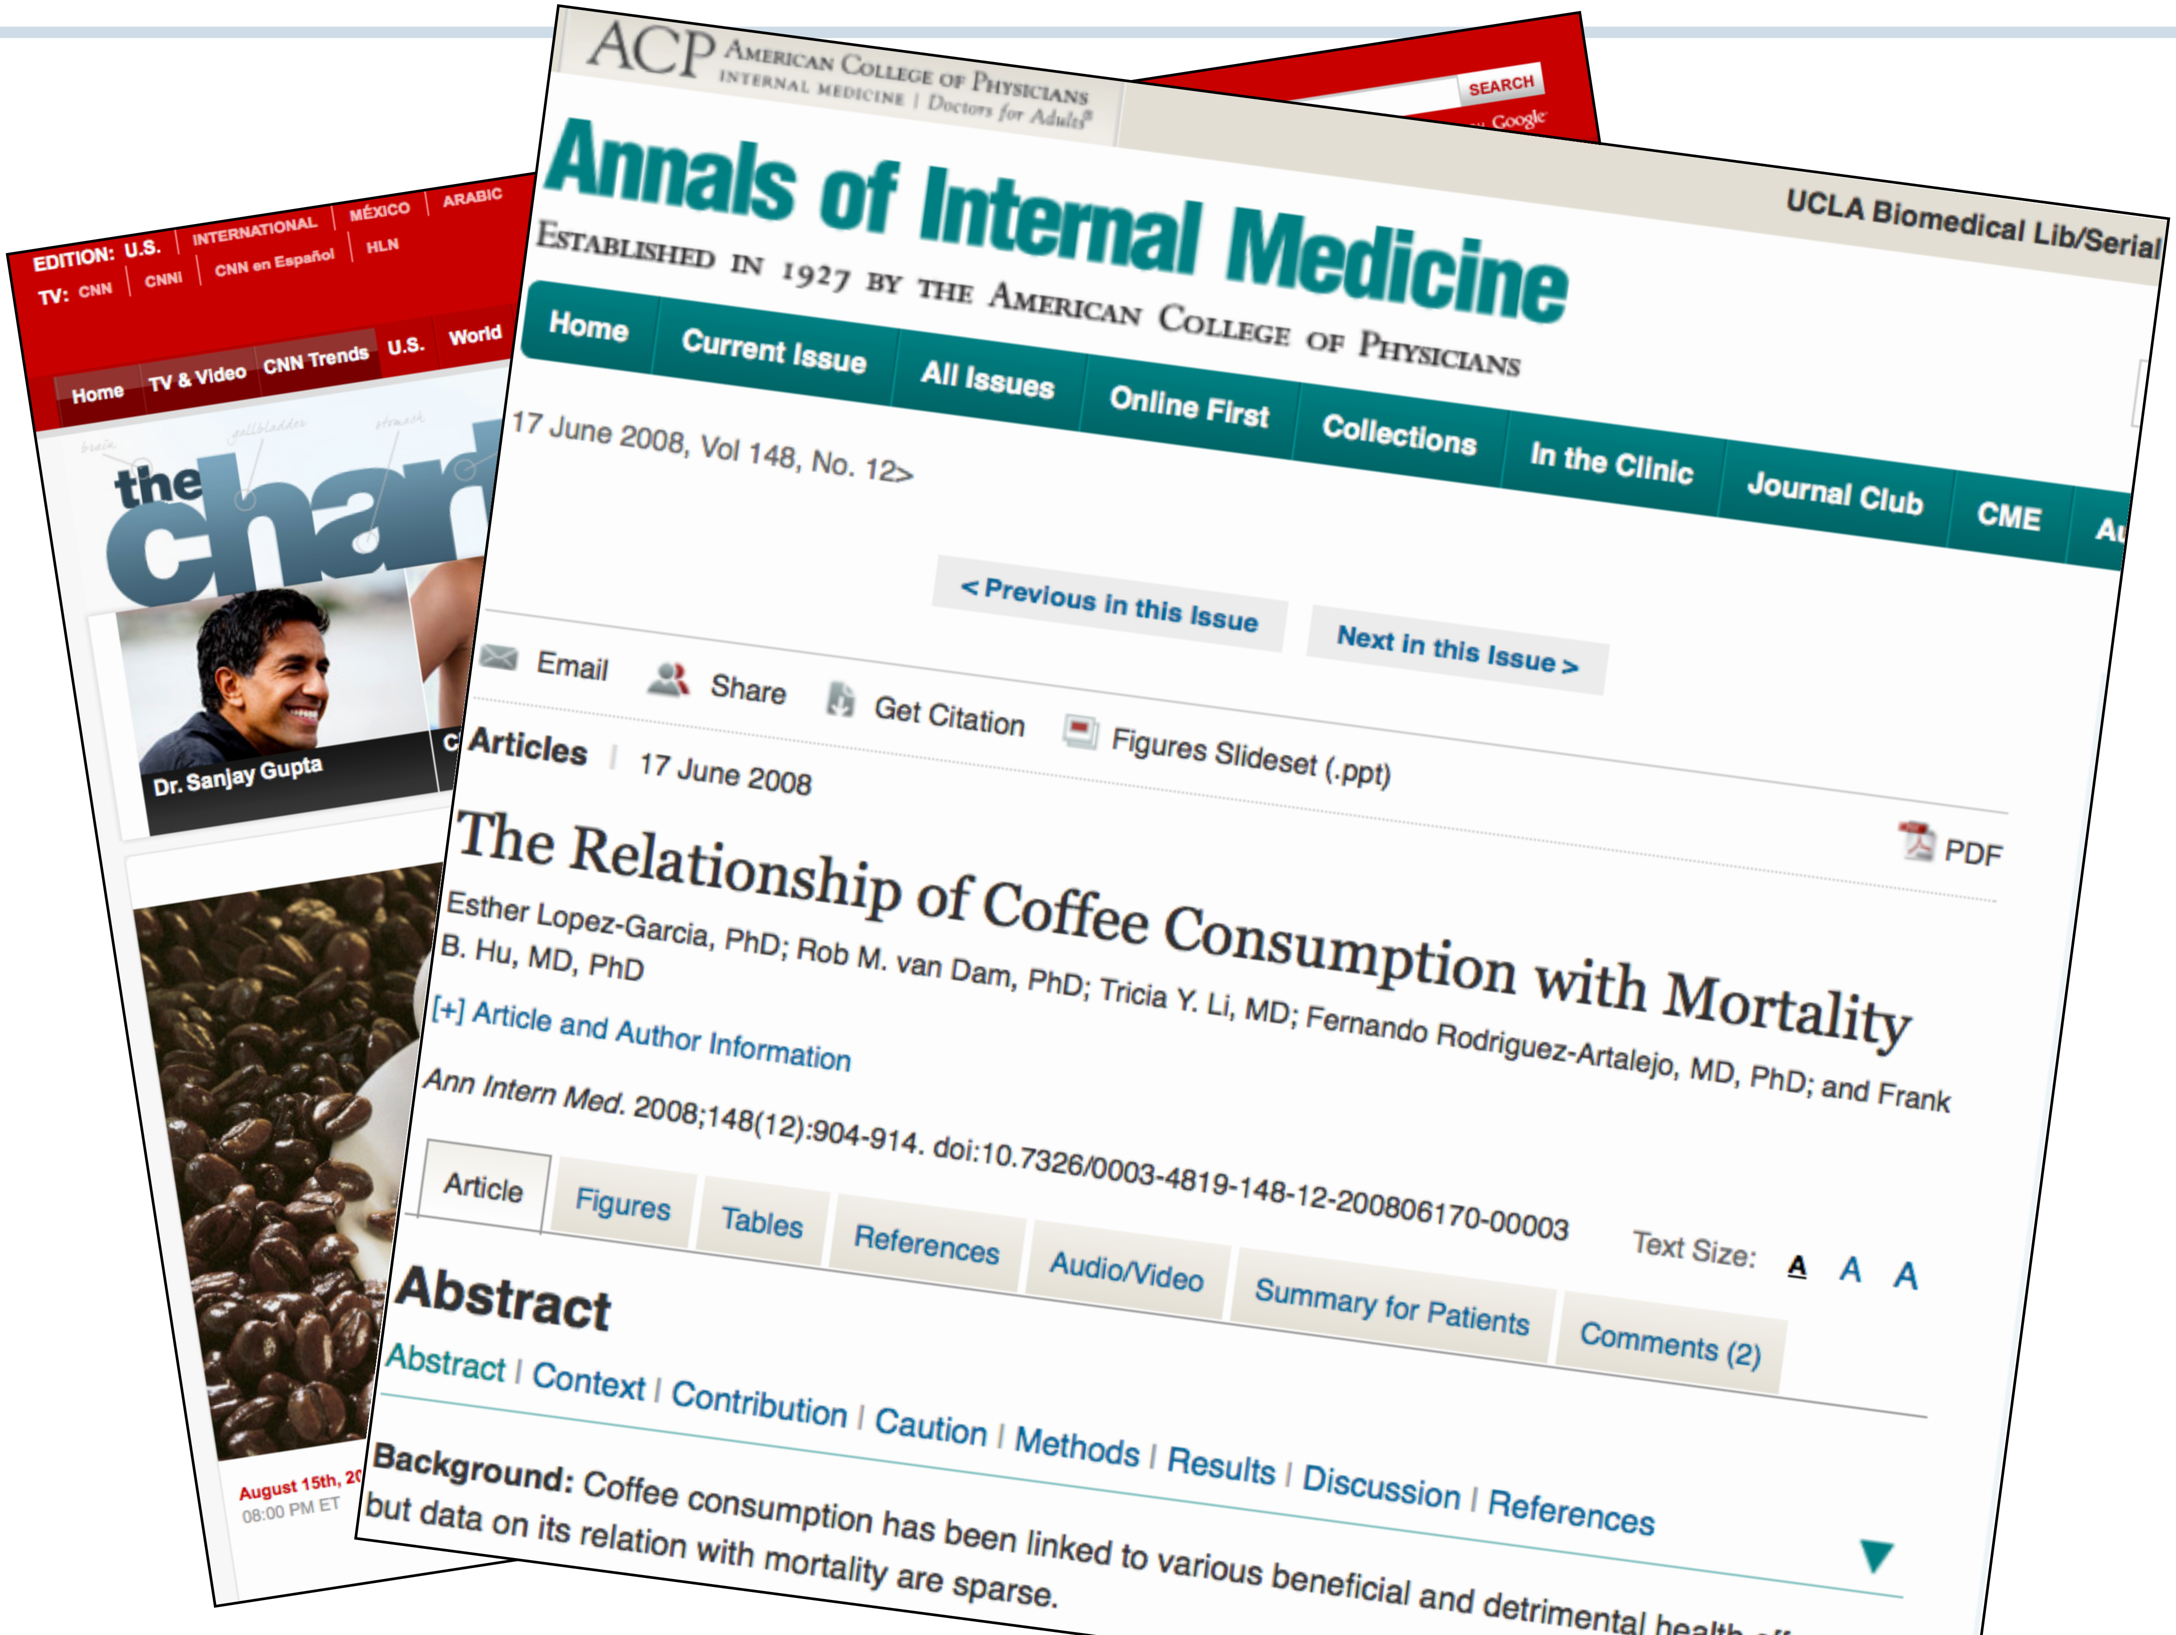

Why Causality Matters for Health Data

N Engl J Med. 1985 Oct 24;313(17):1044-9. doi: 10.1056/NEJM198510243131703.

A prospective study of postmenopausal estrogen therapy and coronary heart disease

M J Stampfer, W C Willett, G A Colditz, B Rosner, F E Speizer, C H Hennekens

PMID: 4047106 DOI: 10.1056/NEJM198510243131703

Abstract

To clarify the possible role of postmenopausal estrogen use in coronary heart disease, we surveyed 121,964 female nurses, aged 30 to 55 years, with mailed questionnaires, beginning in 1976. Information on hormone use and other potential risk factors was updated and the incidence of coronary heart disease was ascertained through additional questionnaires in 1978 and 1980, with a 92.7 per cent follow-up. End points were documented by medical records. During 105,786 person-years of observation among 32,317 postmenopausal women who were initially free of coronary disease, 90 women had either nonfatal myocardial infarctions (65 cases) or fatal coronary heart disease (25 cases). As compared with the risk in women who had never used postmenopausal hormones, the age-adjusted relative risk of coronary disease in those who had ever used them was 0.5 (95 per cent confidence limits, 0.3 and 0.8; P = 0.007), and the risk in current users was 0.3 (95 per cent confidence limits, 0.2 and 0.6; P = 0.001). The relative risks were similar for fatal and nonfatal disease and were unaltered after adjustment for cigarette smoking, hypertension, diabetes, high cholesterol levels, a parental history of myocardial infarction, past use of oral contraceptives, and obesity. These data support the hypothesis that the postmenopausal use of estrogen reduces the risk of severe coronary heart disease.

Why Causality Matters for Health Data

> N Engl J Med. 1985 Oct 24;313(17):1044-9. doi: 10.1056/NEJM198510243131703.

The Daily Telegraph

NEWS BRIEFING

NEWS

NHS 'lacks staffing for seven-day service'

The health service may have too few staff and too little money to deliver government pledges for a "truly seven-day NHS" by 2020, leaked Department of Health documents reveal. Papers drawn up for Jeremy Hunt and other ministers in late July detail a list of risks to the Government's manifesto pledge. Civil servants said the greatest danger was lack of NHS staff, "meaning the full service cannot be delivered". They also suggest staff could provide an obstacle, because "they do not believe in the case for change".
Page 8

NEWS

Mothers earning a third less than men

Mothers end up earning a third less than men as the birth of a child cuts their chances of getting promotions and pay rises, a survey has found. Before having a child the average female worker earns 10 per cent to 15 per cent less per hour than a male employee. After childbirth that

HRT pill triples risk of cancer

Chances of disease found to increase the longer women take hormone treatment

European leaders in step over 'EU army'

estrogen reduces the risk of severe coronary heart disease.

Why Causality Matters for Health Data

N Engl J Med. 1985 Oct 24;313(17):1044-9. doi: 10.1056/NEJM198510243131703.

nature medicine

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature > nature medicine > articles > article

Article | Published: 22 October 2018

The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care

Matthieu Komorowski, Leo A. Celi, Omar Badawi, Anthony C. Gordon & A. Aldo Faisal

longer women take hormone treatment

estrogen reduces the risk of severe coronary heart disease.

Why Causality Matters for Health Data

> N Engl J Med. 1985 Oct 24;313(17):1044-9. doi: 10.1056/NEJM198510243131703.

nature medicine

arXiv > cs > arXiv:1902.03271

Search...
Help | Adv

Computer Science > Artificial Intelligence

[Submitted on 8 Feb 2019]

Does the "Artificial Intelligence Clinician" learn optimal treatment strategies for sepsis in intensive care?

Russell Jeter, Christopher Josef, Supreeth Shashikumar, Shamim Nemati

From 2017 to 2018 the number of scientific publications found via PubMed search using the keyword "Machine Learning" increased by 46% (4,317 to 6,307). The results of studies involving machine learning, artificial intelligence (AI), and big data have captured the attention of healthcare practitioners, healthcare managers, and the public at a time when Western medicine grapples with unmitigated cost increases and public demands for accountability. The complexity involved in healthcare applications of machine learning and the size of the associated data sets has afforded many researchers an uncontested opportunity to satisfy these demands with relatively little oversight. In a recent Nature Medicine article, "The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care," Komorowski and his coauthors propose methods to train an artificial intelligence clinician to treat sepsis patients with vasopressors and IV fluids. In this post, we will closely examine the claims laid out in this paper. In particular, we will study the individual treatment profiles suggested by their AI Clinician to gain insight into how their AI Clinician intends to treat patients on an individual level.

Matthieu Komorowski, Leo A. Celi, Omar Badawi, Anthony C. Gordon & A. Aldo Faisal

longer women take hormone treatment

estrogen reduces the risk of severe coronary heart disease.

Modern language of causality

- Critical observation: causal mechanisms behind the system under study are generally not observable, but they do produce *traces*: **data**.
- David Hume (1748): “Nature has kept us at a great distance from all her secrets, and has afforded only the knowledge of a few superficial qualities of objects; while she conceals from us those powers and principles, on which the influence of these objects entirely depends”.

Causal framework (requirements)

- The observation that "reality" and the data generated by it are fundamentally distinct entities leads to two requirements for any framework for causal inference:

1. The causal mechanisms underlying the phenomenon under investigation should be accounted for —indeed, formalized— in the analysis.
2. This collection of mechanisms (even if primarily unobserved) should be formally tied to its output: the generated phenomena and corresponding datasets.

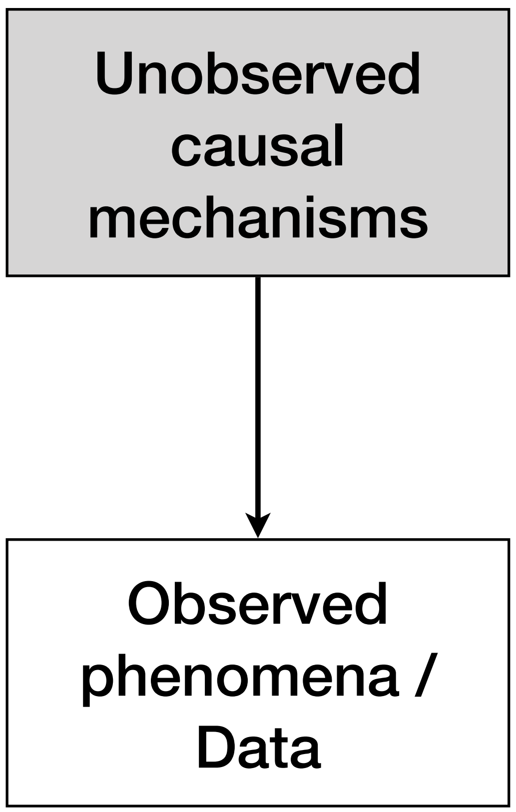

Allegory of the Cave (Plato)

1 Structural Causal Models (intuition)

- We will assume that the underlying reality (or "ground truth"), our putative target, can be represented as a collection of causal mechanisms in the form of a mathematical object called *structural causal model (SCM)*.
- We assume the existence of those mechanisms, independent of our practical capacity to discern their form, nature, and intricate details.

② Pearl's Causal Hierarchy

- Each collection of causal mechanisms (i.e., SCM) induces a causal hierarchy (sometimes called the "ladder of causation").
- This structure is named the Pearl Causal Hierarchy (PCH) after Turing Awardee and UCLA Professor Judea Pearl, who was the first to recognize and discuss it in some depth.
- In this course, we pursue a formal view of the PCH and its implications.

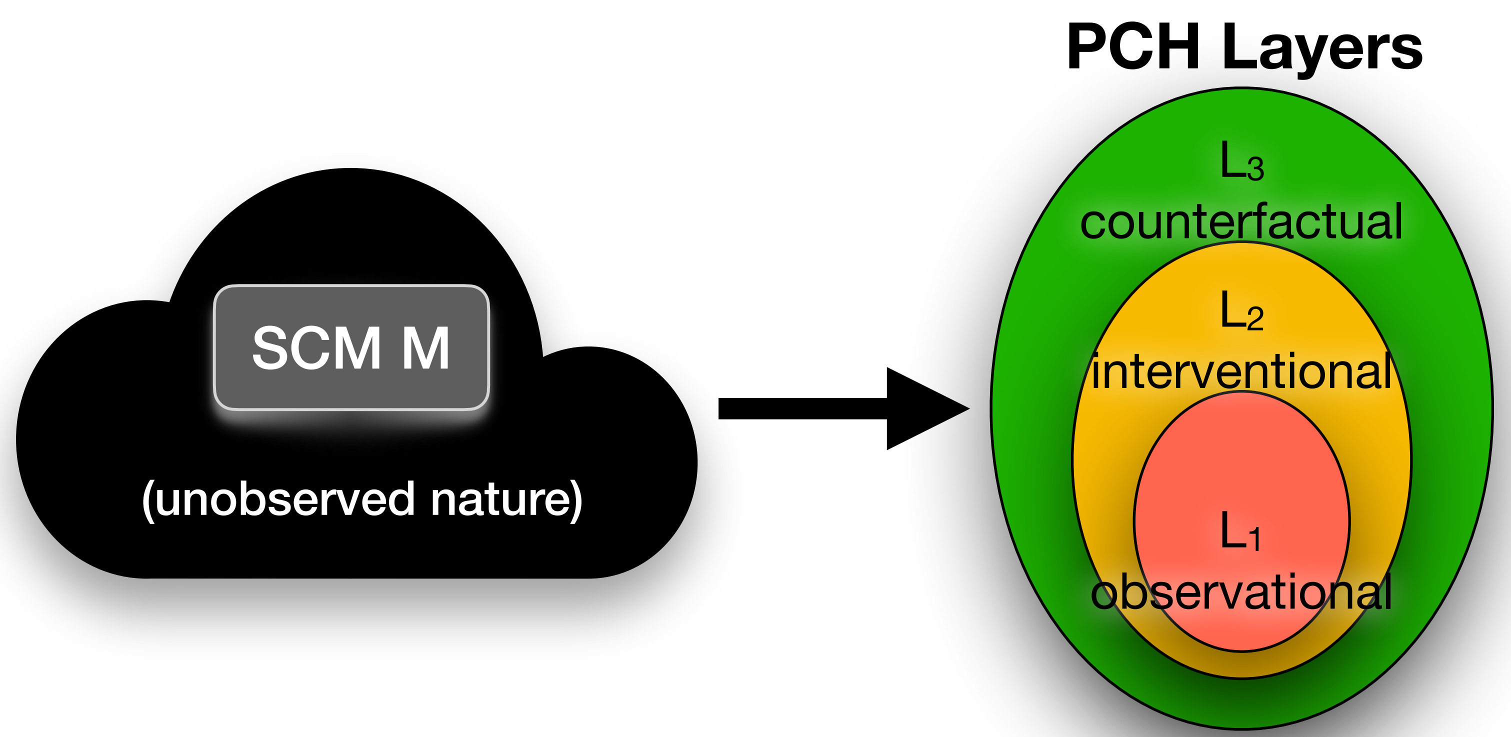

PCH - Comparing to Human Cognitive Capabilities

- **Cognitive capabilities** Pearl's Causal Hierarchy (PCH)
  - human: sensors + agency + retrospection

sensors	agency	retrospection
layer 1	layer 2	layer 3

  - 1. ability to observe the world and infer the 'next' stage
  - 2. ability to intervene in the world and change it
  - 3. ability to reflect about what we did (and didn't do) and infer new behaviors

Big picture - Pearl's Causal Hierarchy

	Level (Symbol)	Typical Activity	Typical Question	Examples
1 👀	Associational $P(y \mid x)$	Seeing (Associations) Predictive ML	What is? How would seeing X change my belief in Y?	Is low blood pressure associated with sepsis?
2 💪	Interventional $P(y \mid do(x), c)$	Doing (Treatment Effects, Interventions) Reinforcement Learning	What if? What if I do X?	Do statins prevent heart attacks?
3 🧠	Counterfactual $P(y_x \mid x', y')$	Imagining, Retrospection (Treatment Harm)	Why? What if I had acted differently?	Had the patient not undergone surgery, would they be alive?

Big picture - Pearl's Causal Hierarchy

Causal Inference =
Cross-layer inferences

	Level (Symbol)	Typical Activity	Typical Question	Examples
1 👀	Associational $P(y \mid x)$	Seeing (Associations)	What is? How would seeing X change my belief in Y?	Is low blood pressure associated with sepsis?
2 💪	Interventional $P(y \mid do(x), c)$	Doing (Treatment Effects, Interventions)	What if? What if I do X?	Do statins prevent heart attacks?
3 🧠	Counterfactual $P(y_x \mid x', y')$	Imagining, Retrospection (Treatment Harm)	Why? What if I had acted differently?	Had the patient not undergone surgery, would they be alive?

Big picture - Pearl’s Causal Hierarchy

Causal Inference = Cross-layer inferences

Task-type 0

Data: Layer 1
Inference: Layer 1

e.g., “*does high cholesterol predict mortality?*”

Level (Symbol)	Typical Activity	Typical Question	Examples
1 Associa… $P(y \mid$ …		What is?	Is low blood pressure associated with sepsis?
2 Interven… $P(y \mid do$ …			Do statins prevent heart attacks?
3 Counterfactual $P(y_x \mid x', y')$	Imagining, Retrospection (Treatment Harm)	Why? What if I had acted differently?	Had the patient not undergone surgery, would they be alive?

Big picture - Pearl's Causal Hierarchy

**Causal Inference =
Cross-layer inferences**

Task-type 1

Data: Layer 1
Inference: Layer 2

e.g., "*will a patient's state improve if we treat them with antibiotics?*"

Level (Symbol) | Typical Activity | Typical Question | Examples

1 Associational $P(y \mid x)$ — Seeing — What is? How would seeing ... — Is low blood pressure associated with sepsis?

2 Interventional $P(y \mid do(x))$ — ... statins prevent heart attacks?

3 Counterfactual $P(y_x \mid x', y')$ — Retrospection (Treatment Harm) — What if I had acted differently? — Had the patient not undergone surgery, would they be alive?

Big picture - Pearl's Causal Hierarchy

Causal Inference = Cross-layer inferences

Level (Symbol)	Typical Activity	Typical Question	Examples
1 Associational $P(y \mid x)$	Seeing (Associations)	How would seeing	Is low blood pressure associated ith sepsis?
2 Interventional			o statins event heart tacks?
3 Counterfactua $P(y_x \mid x', y')$			ad the atient not ndergone rgery, would they be alive?

Task-type 2

Data: Layer 2
Inference: Layer 2

e.g., "*given a randomized controlled trial on drugs A, B, can we say anything about their joint effect?*"

Big picture - Pearl's Causal Hierarchy

Level (Symbol)	Typical Activity	Typical Question	Examples
1 👀 Associational $P(y \mid x)$	Seeing (Associations)	How would seeing X change my belief in Y?	Is low blood pressure associated with sepsis?
2 💪 Interventional $P(y \mid do(x), c)$			
🧠 Counterfactual $P(y_x \mid x', y')$			

Causal Inference =
Cross-layer inferences

Task-type 3

Data: Layer 1, 2
Inference: Layer 3

e.g., "*if we adopted a different treatment policy, which patients would avoid harm?*"

Big Picture

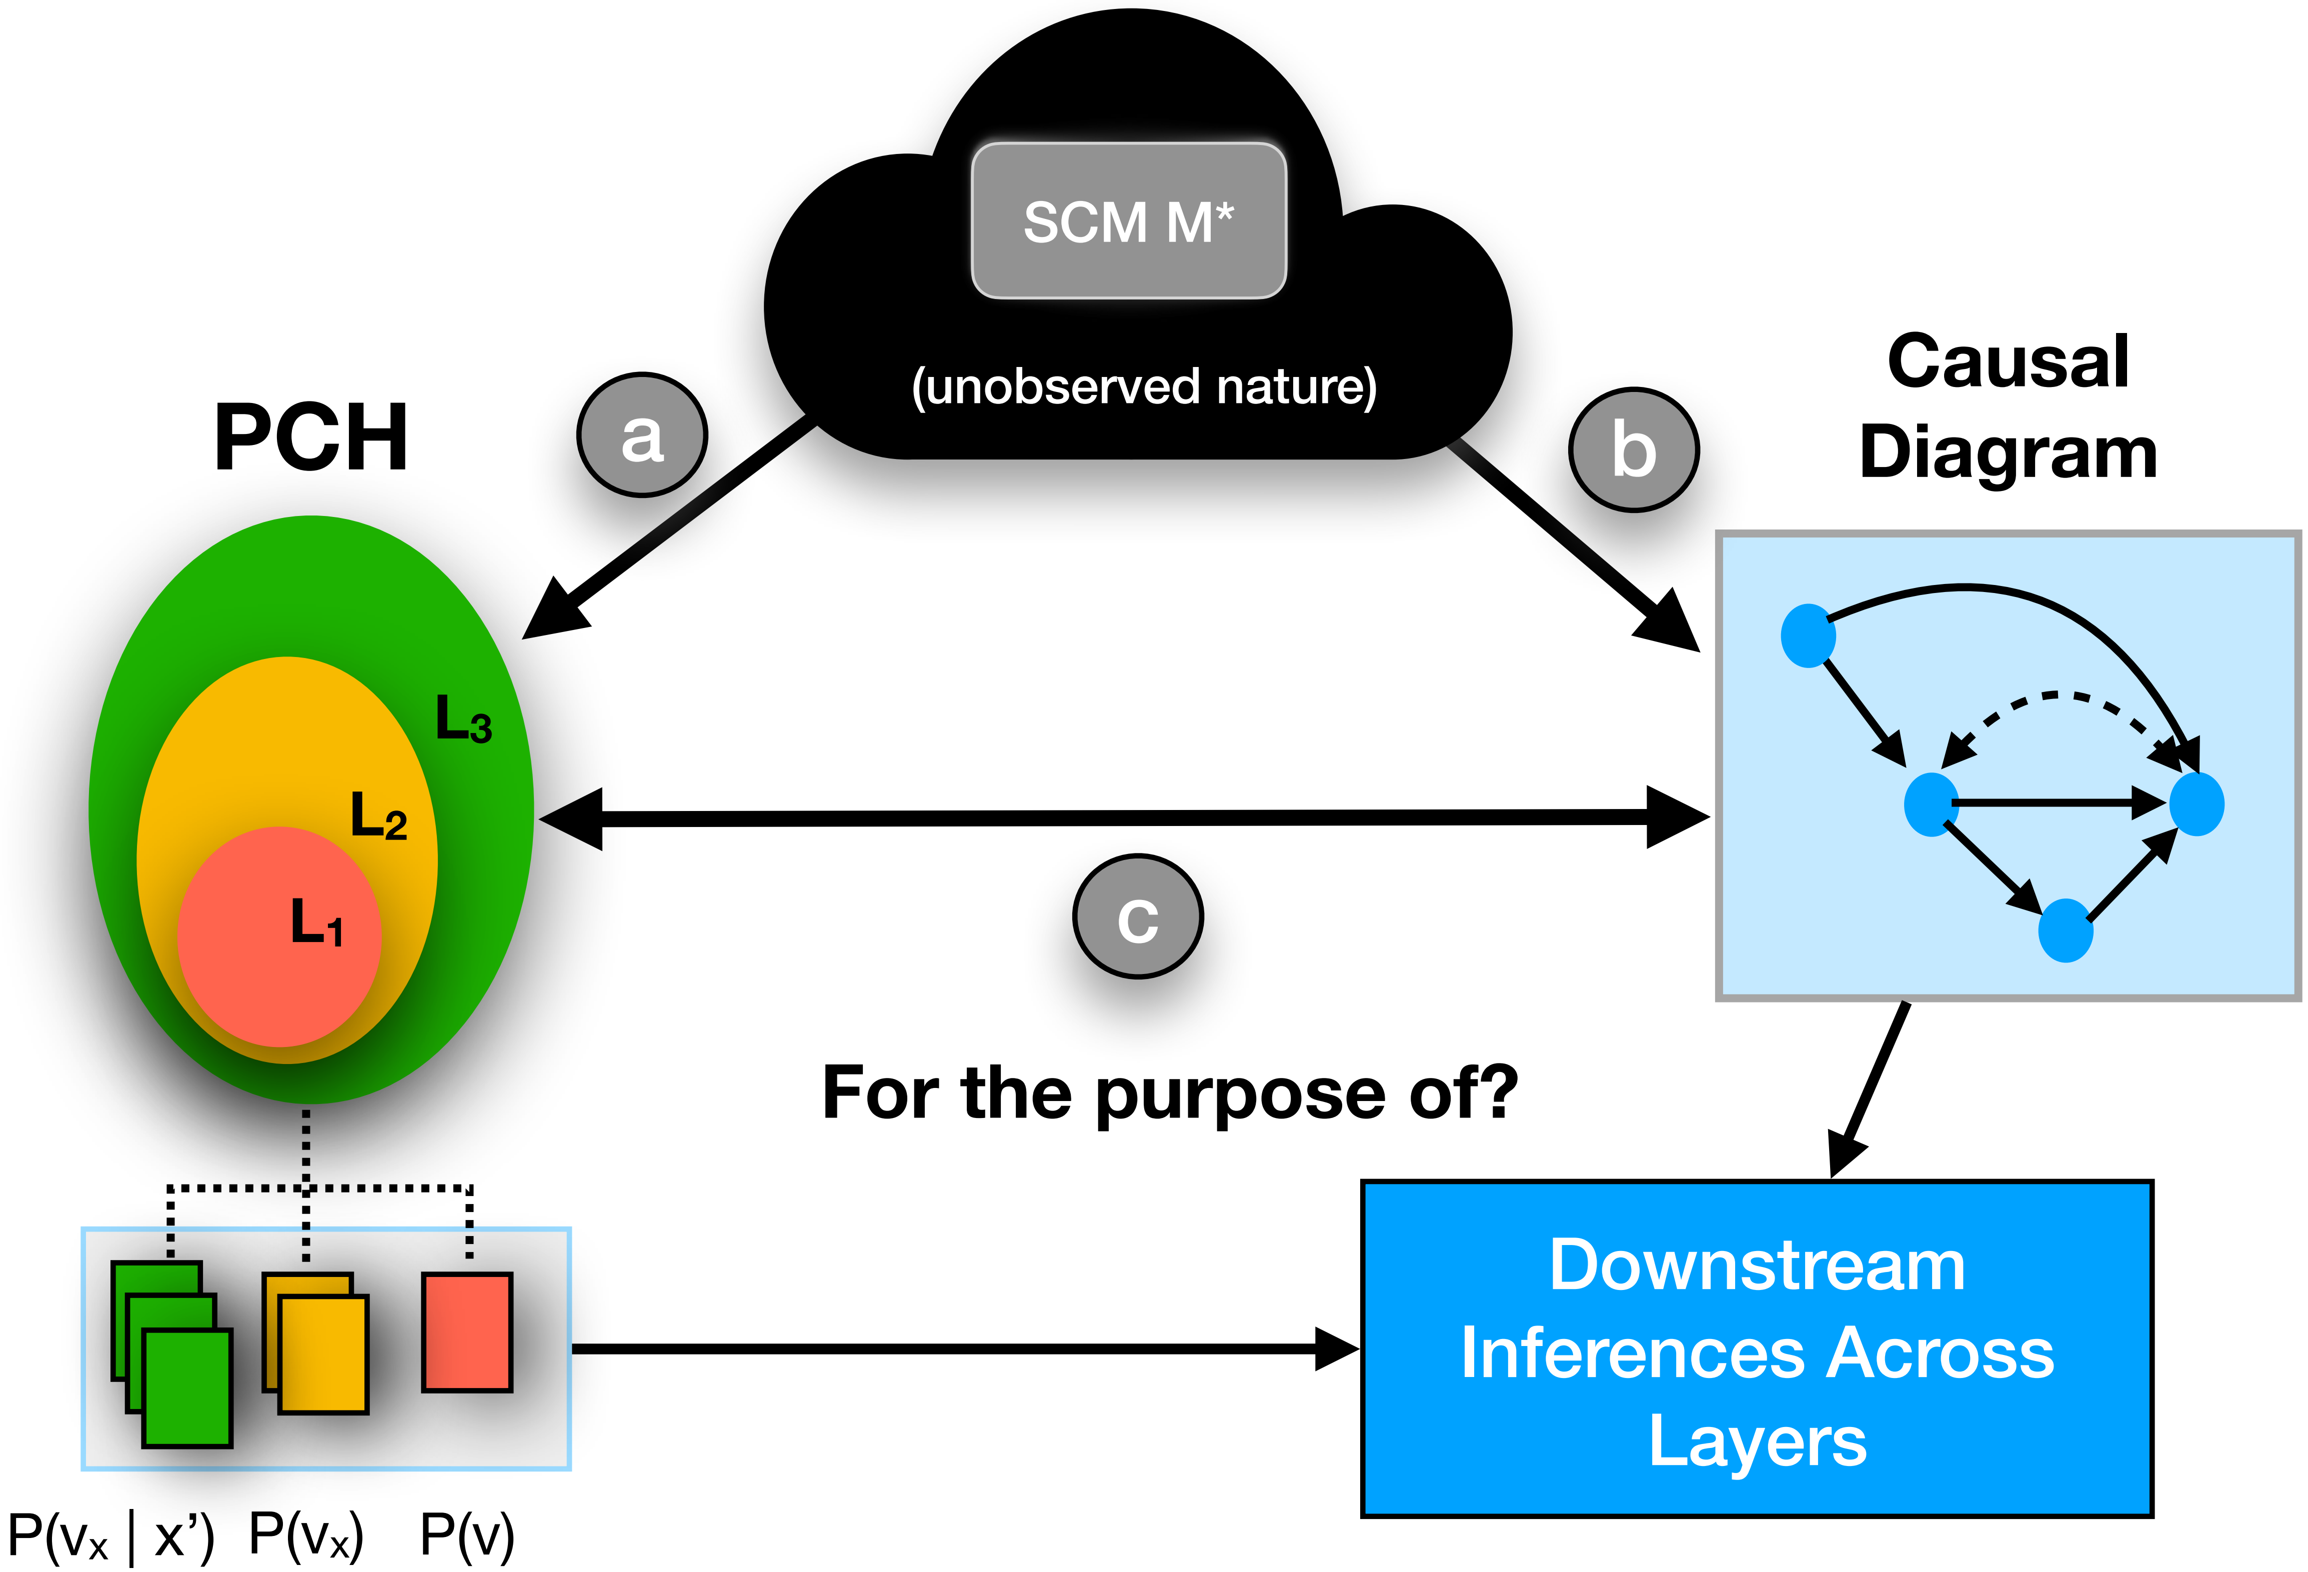
SCM M*
(unobserved nature)
a
b
PCH
Causal Diagram
L3
L2
L1
c
For the purpose of?
Downstream Inferences Across Layers
P(vx | x') P(vx) P(v)

Motivating Examples — Why Causality Matters in Practice

Ex 1. Diabetes Treatment with Semaglutide

(Simpson, 1951; Pearl, 2000)

	HbA1c low (Y)	HbA1c high ($\neg Y$)		Sucess Rate
drug (X)	20	20	40	50%
no-drug ($\neg X$)	16	24	40	40%
	36	44		

Stratified by sex

	HbA1c low (Y)	HbA1c high ($\neg Y$)		Sucess Rate
drug (X)	20	20	40	50%
no-drug ($\neg X$)	16	24	40	40%
	36	44		

male ($\neg F$)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Success Rate
drug (X)	18	12	30	60%
no-drug ($\neg X$)	7	3	10	70%
	25	15	40	

female (F)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Success Rate
drug (X)	2	8	10	20%
no-drug ($\neg X$)	9	21	30	30%
	11	29	40	

Which table should the physician consult?

	HbA1c low (Y)	HbA1c high ($\neg Y$)		Sucess Rate
drug (X)	20	20	40	50%
no-drug ($\neg X$)	16	24	40	40%
	36	44		

$P(Y \mid F, X) < P(Y \mid F, \neg X)$

$P(Y \mid \neg F, X) < P(Y \mid \neg F, \neg X)$

but

$P(Y \mid X) > P(Y \mid \neg X)$!

male ($\neg F$)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Success Rate
drug (X)	18	12	30	60%
no-drug ($\neg X$)	7	3	10	70%
	25	15	40	

female (F)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Success Rate
drug (X)	2	8	10	20%
no-drug ($\neg X$)	9	21	30	30%
	11	29	40	

Identical tables, different variable: Which table should the physician consult now?

	HbA1c low (Y)	HbA1c high ($\neg Y$)		Sucess Rate
drug (X)	20	20	40	50%
no-drug ($\neg X$)	16	24	40	40%
	36	44		

$$P(Y \mid F, X) < P(Y \mid F, \neg X)$$

$$P(Y \mid \neg F, X) < P(Y \mid \neg F, \neg X)$$

but

$P(Y \mid X) > P(Y \mid \neg X)$!

weight loss ($\neg F$)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Success Rate
drug (X)	18	12	30	60%
no-drug ($\neg X$)	7	3	10	70%
	25	15	40	

no weight loss (F)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Success Rate
drug (X)	2	8	10	20%
no-drug ($\neg X$)	9	21	30	30%
	11	29	40	

Would the result be different with a different Story?

($\neg F$)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Sucess Rate
drug (X)	18	12	30	60%
no-drug ($\neg X$)	7	3	10	70%
	25	15	40	

(F)	HbA1c low (Y)	HbA1c high ($\neg Y$)		Sucess Rate
drug (X)	2	8	10	20%
no-drug ($\neg X$)	9	21	30	30%
	11	29	40	

Sex-Stratified

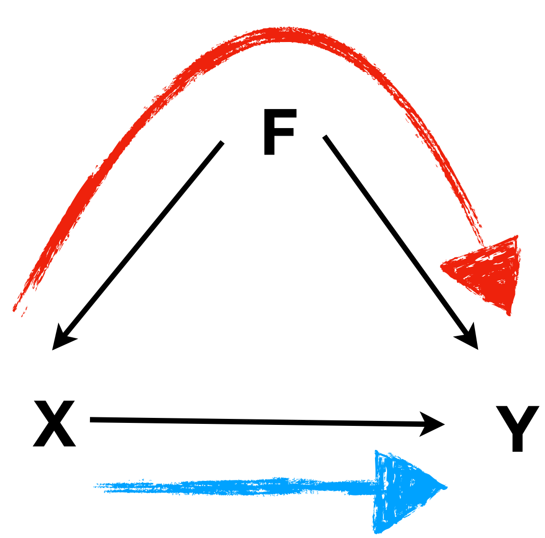

Weight Loss-Stratified

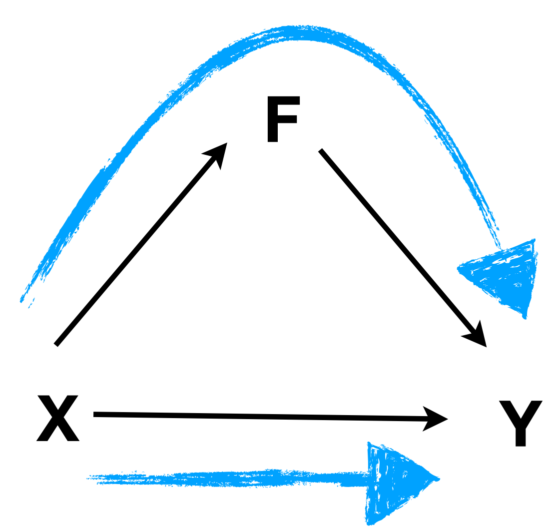

Big Picture

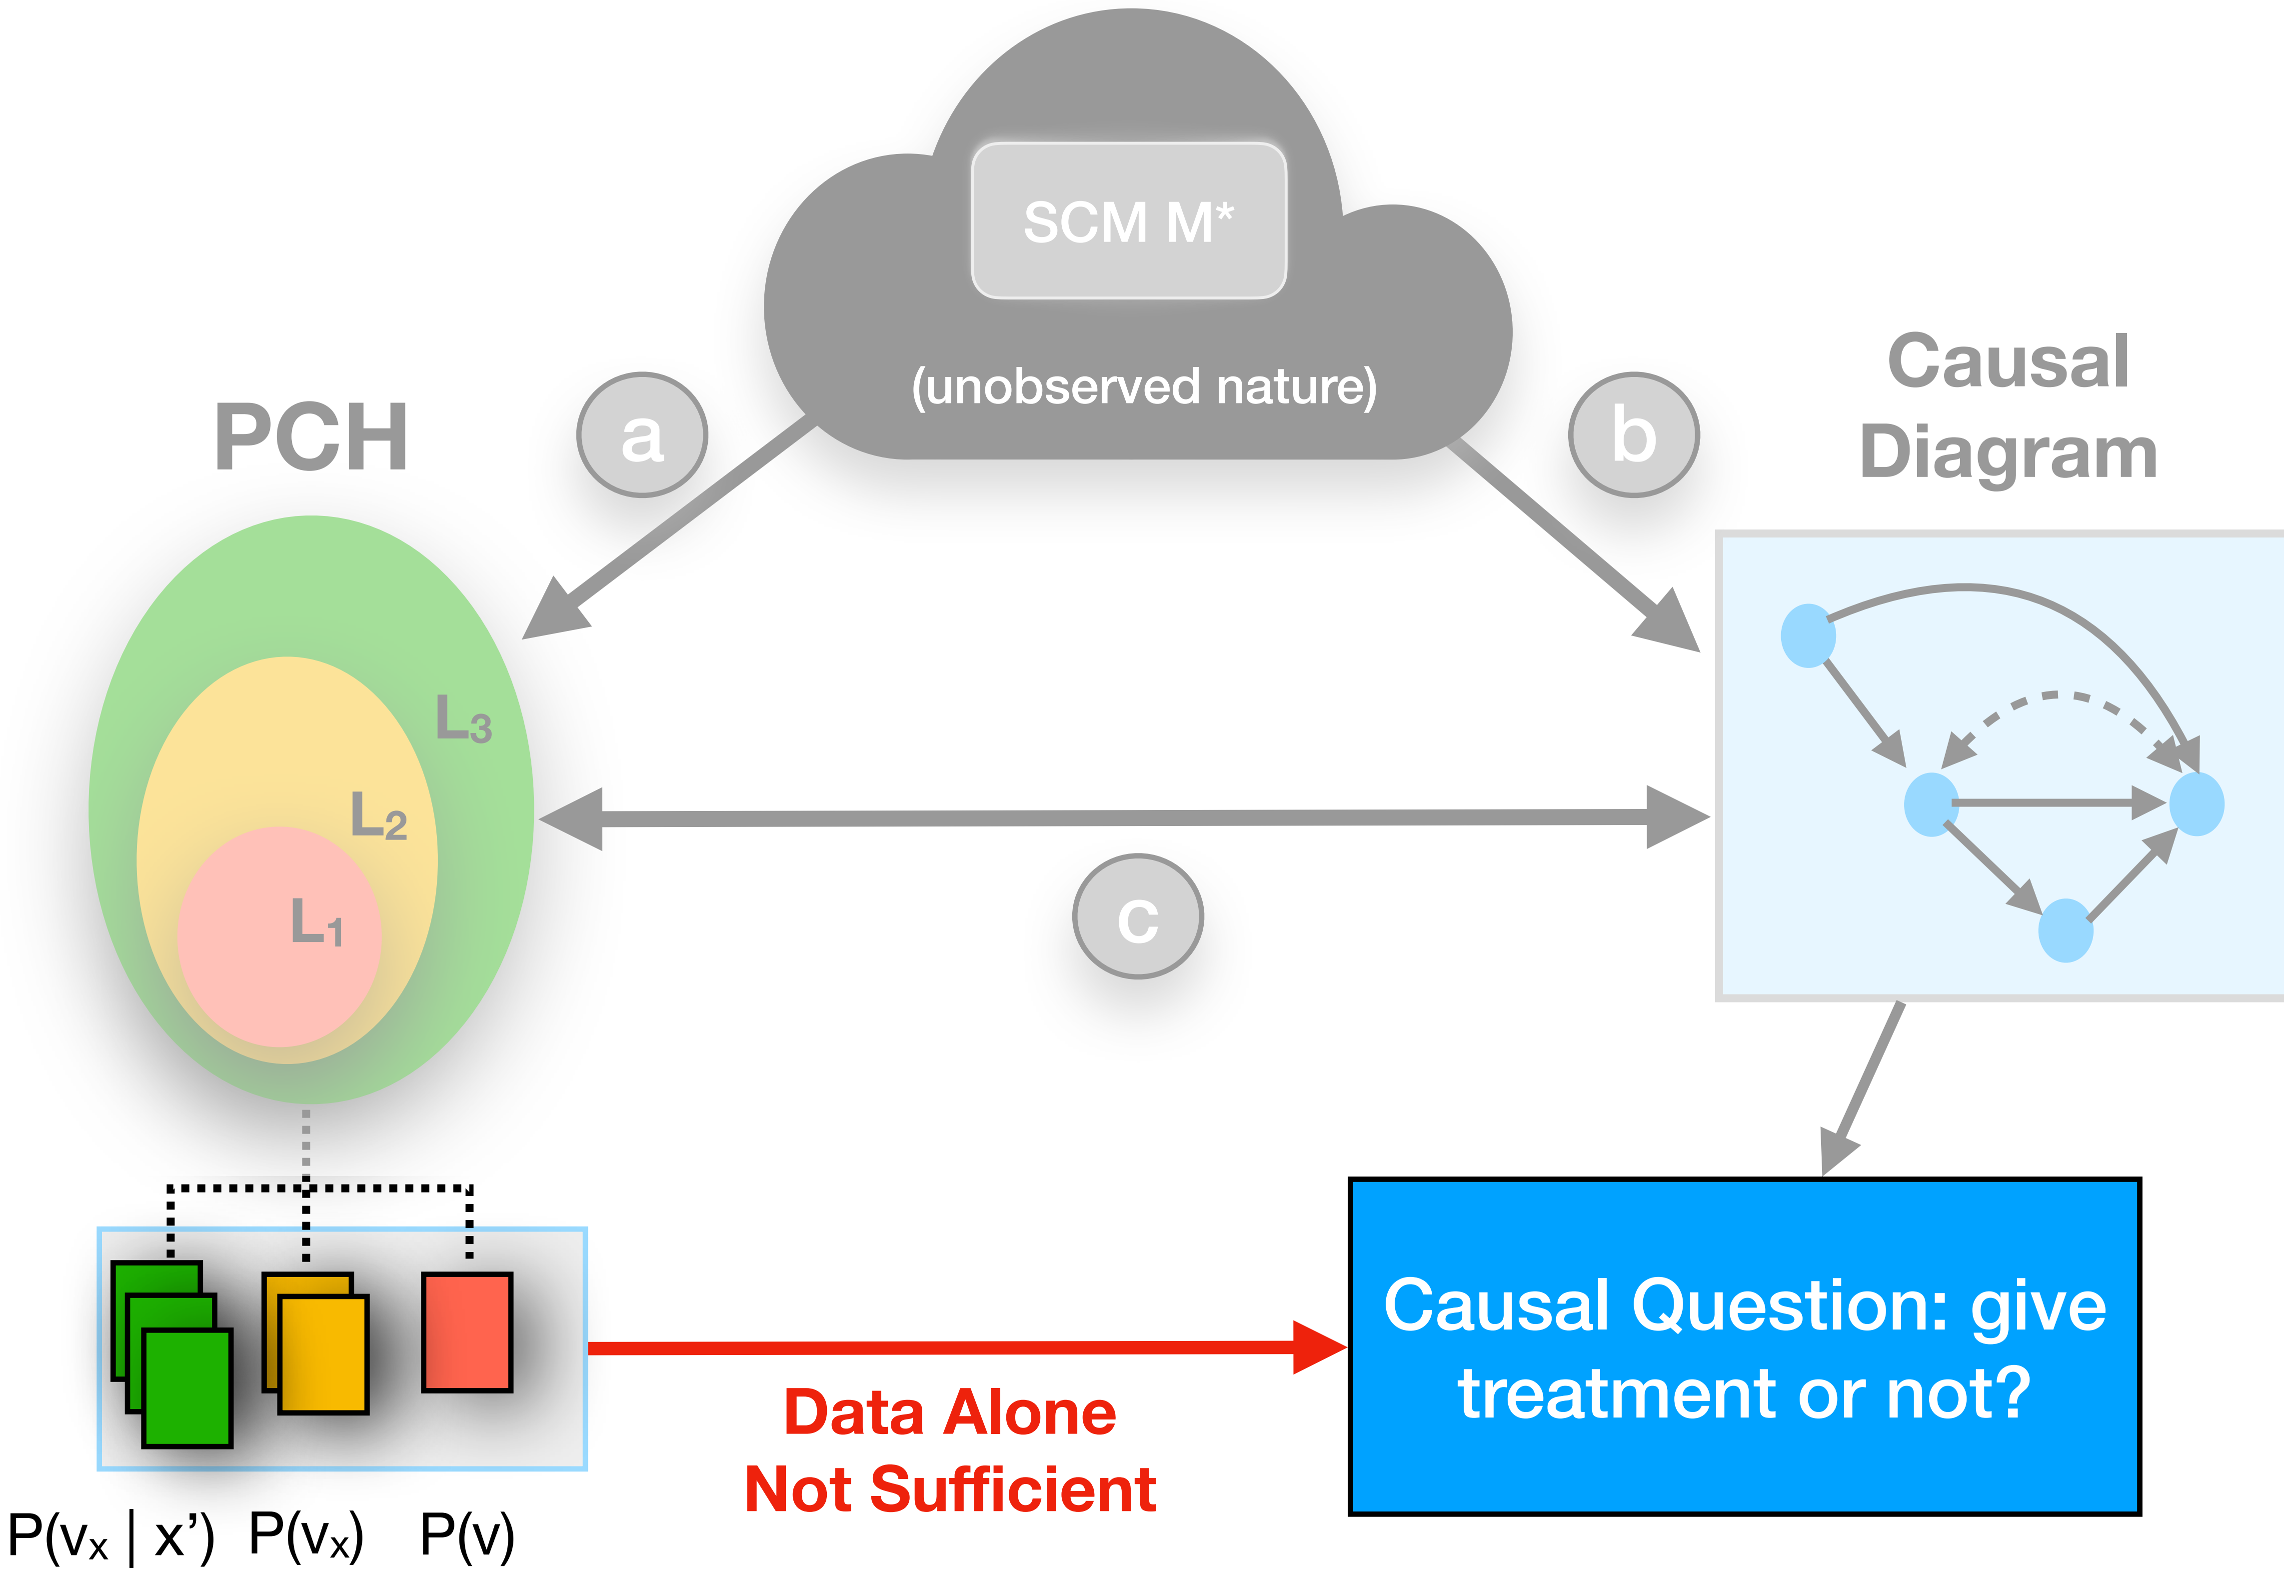

SCM M*
(unobserved nature)
PCH
a
b
c
Causal
Diagram
L3
L2
L1
P(vx | x') P(vx) P(v)
Data Alone
Not Sufficient
Causal Question: give treatment or not?