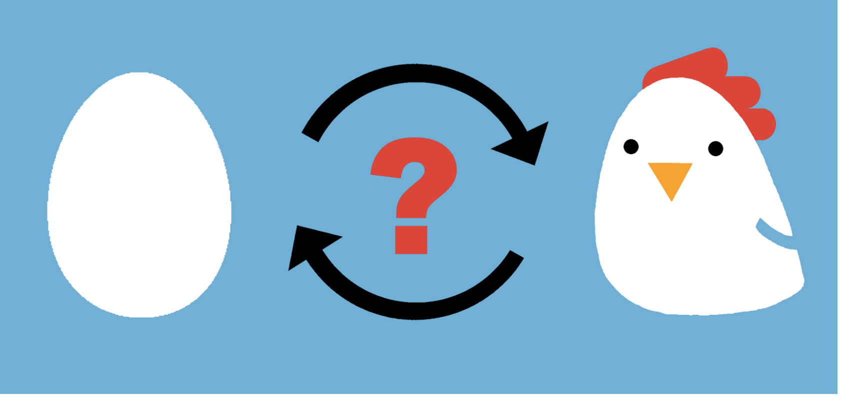

CBMS Conference -- Foundations of Causal Graphical Models and Structure Discovery

Lecture 1

Introduction: A Big Picture of Causality

Instructor: Kun Zhang

Carnegie Mellon University

MOHAMED BIN ZAYED UNIVERSITY OF ARTIFICIAL INTELLIGENCE

Representing Causal Relations with Directed Graphs

- A directed graph represents a causally sufficient causal structure

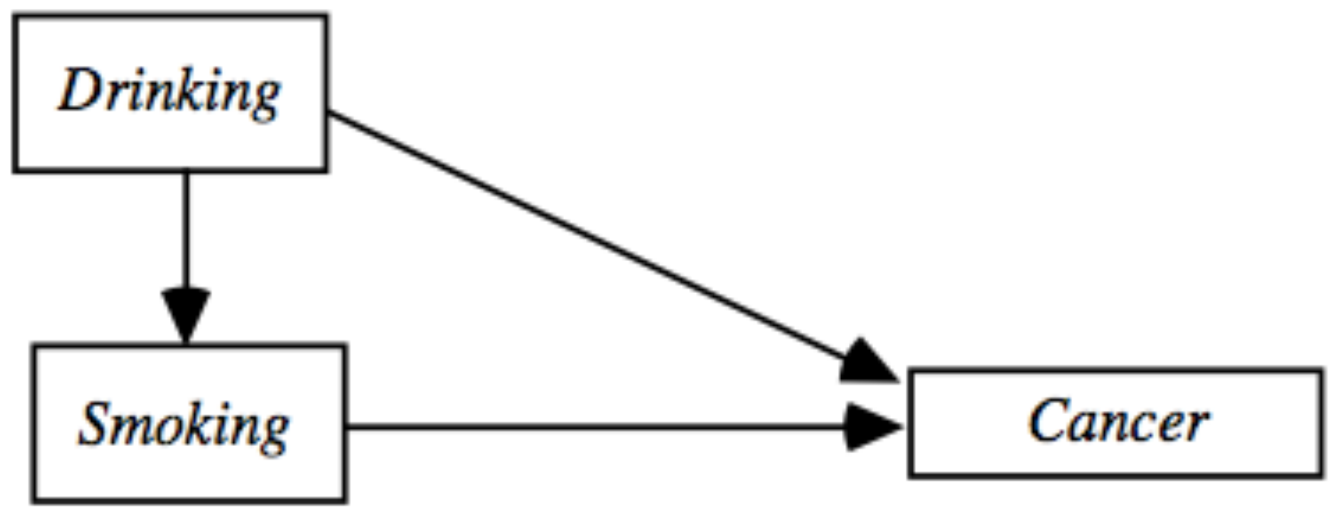

(adapted from "Causation, Prediction, and Search" by SGS, 1995)

- Directed edge from A to B means A is a direct cause of B relative to the given variable set $\boldsymbol{V}$

Causality vs. Dependence

- Causality ➜ dependence ! Dependence ➜ causality

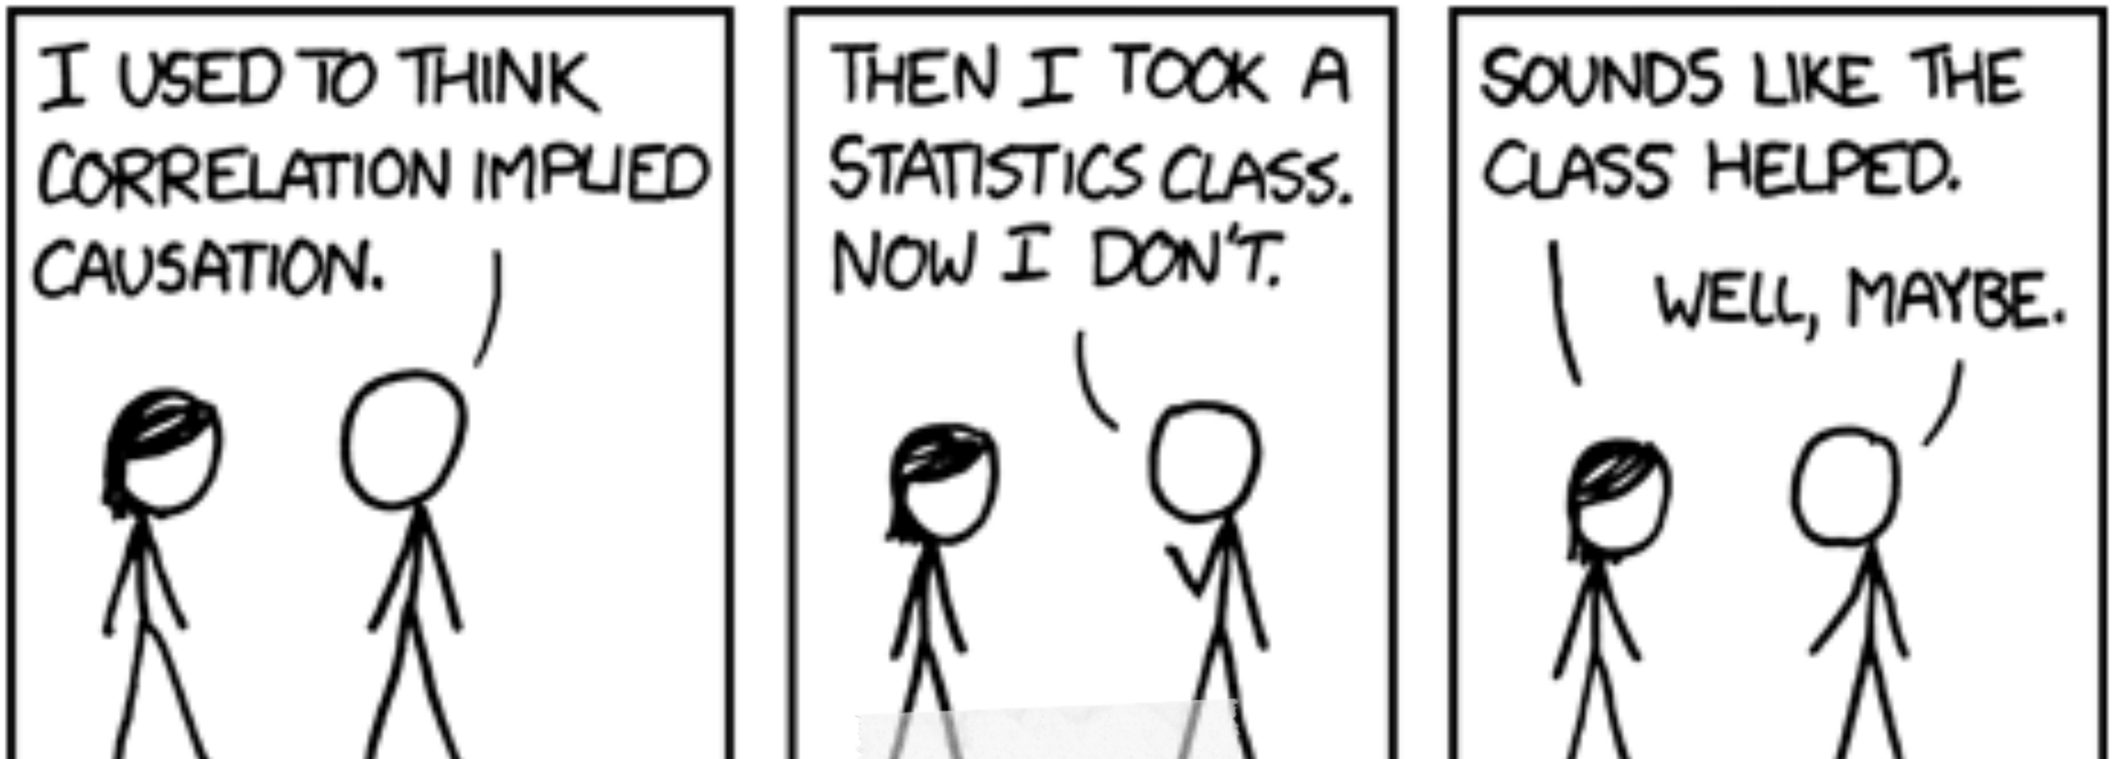

An intervention on X changes only the target variable X, leaving any other variable unchanged, at least for the moment.

X and Y are associated iff

$\exists x_1 \neq x_2 \; P(Y|X=x_1) \neq P(Y|X=x_2)$

X is a cause of Y iff

$\exists x_1 \neq x_2 \; P(Y|do\ X=x_1) \neq P(Y|do\ X=x_2)$

intervention

Classic Ways to Find Causal Information (i.i.d. Case)

- What if X and Y are dependent?
- What if you change X and see Y also changes?
  - A **manipulation/ intervention** directly changes only the target variable X

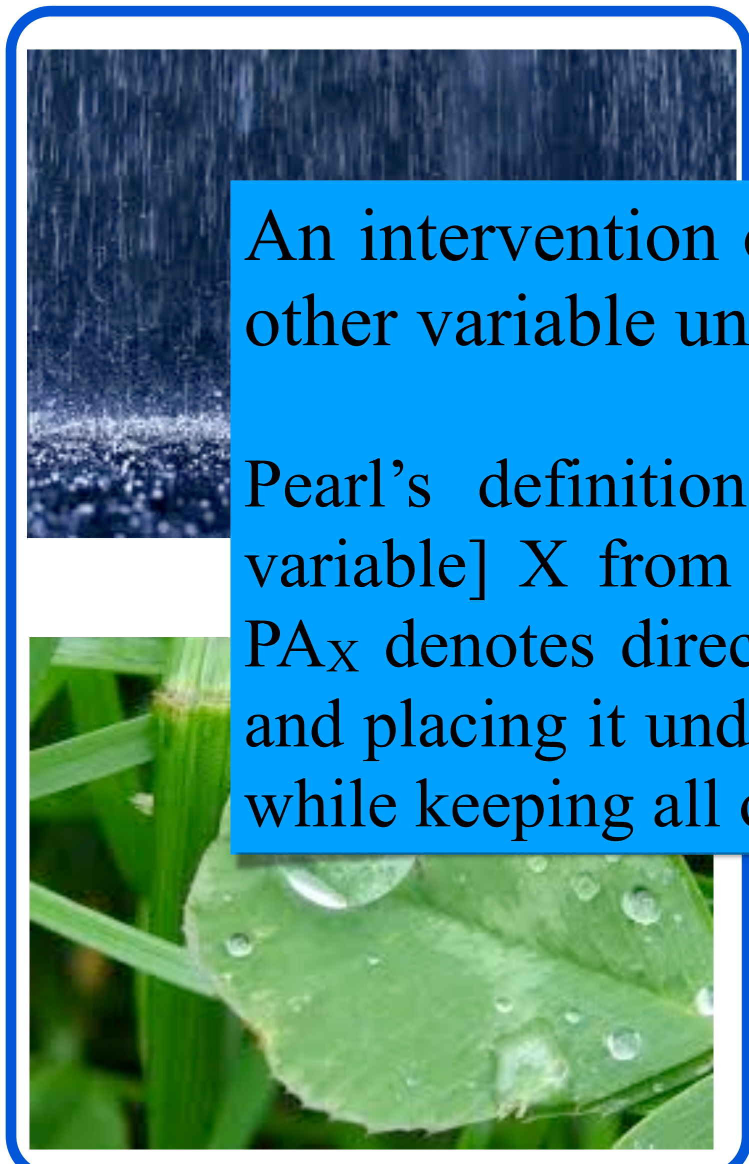

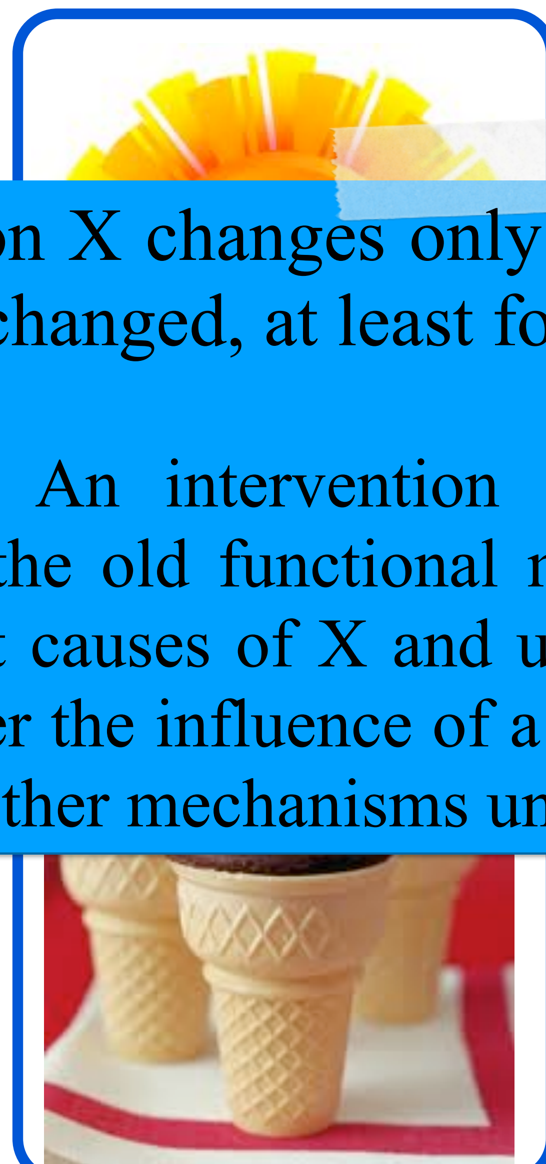

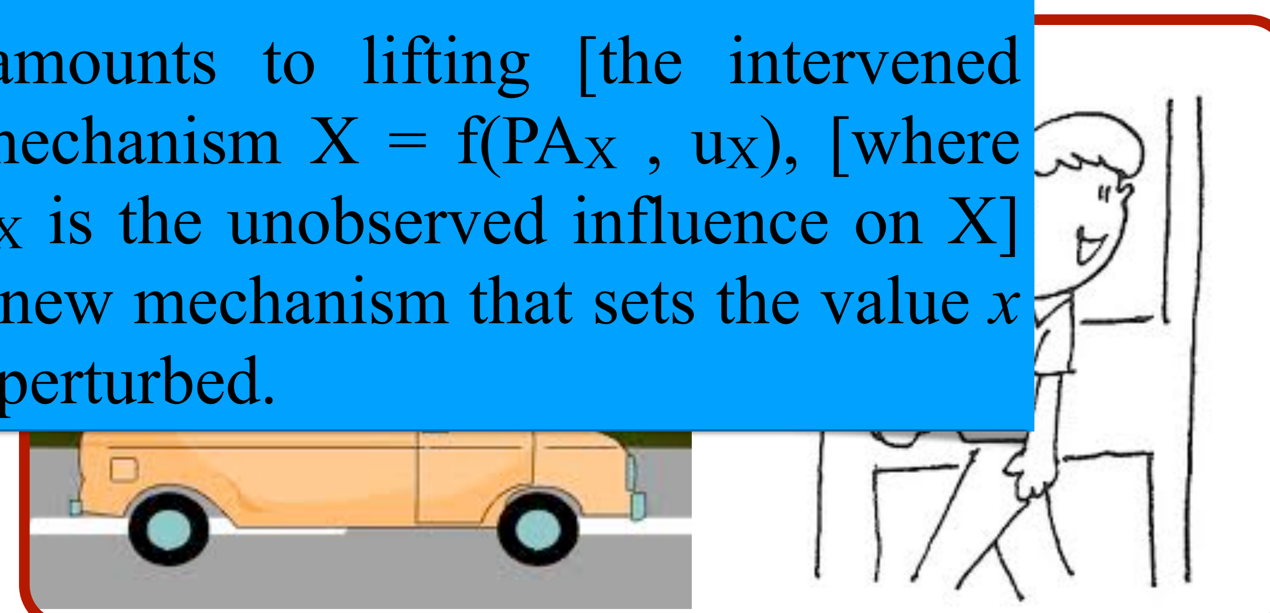

An intervention on X changes only the target variable X, leaving any other variable unchanged, at least for the moment.

Pearl's definition: An intervention amounts to lifting [the intervened variable] X from the old functional mechanism $X = f(PA_X, u_X)$, [where PA_X denotes direct causes of X and u_X is the unobserved influence on X] and placing it under the influence of a new mechanism that sets the value x while keeping all other mechanisms unperturbed.

** Definition of "interventions"*

Introduction

- What is causality?
  - Classic ways to find causal information
- Introduction to ML and AI, and some connections with causality
- Causal thinking

1. Making “changes”
2. Understanding & information fusion
3. Prediction in complex environments
4. Artificial “intelligence”…

- Typical problems in causality research
  - Identification of causal effects
  - Counterfactual reasoning
  - Causal discovery & causal representation learning

Computational Systems Are Reshaping Our World…

What is AI? ML?

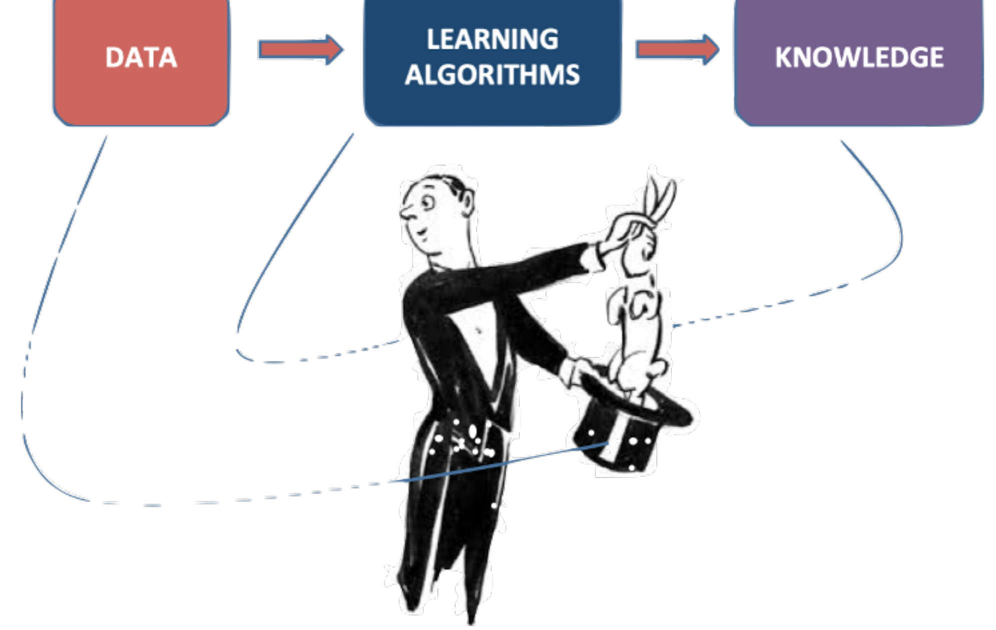

- Artificial intelligence is the capacity of a machine to imitate intelligent human behavior

- Machine learning
  - Arthur Samuel (1959): the field of study that gives computers the ability to learn without being explicitly programmed
  - Tom Mitchell (1998): The goal is identifying the underlying mechanisms and algorithms that <u>allow improving our knowledge with more data</u>

Why Machine Learning?

- Solve problems automatically and efficiently almost everywhere
  - Image classification, cyber fraud detection, healthcare, finance, logistics, entertainment, autonomous driving…

What is AI? ML?

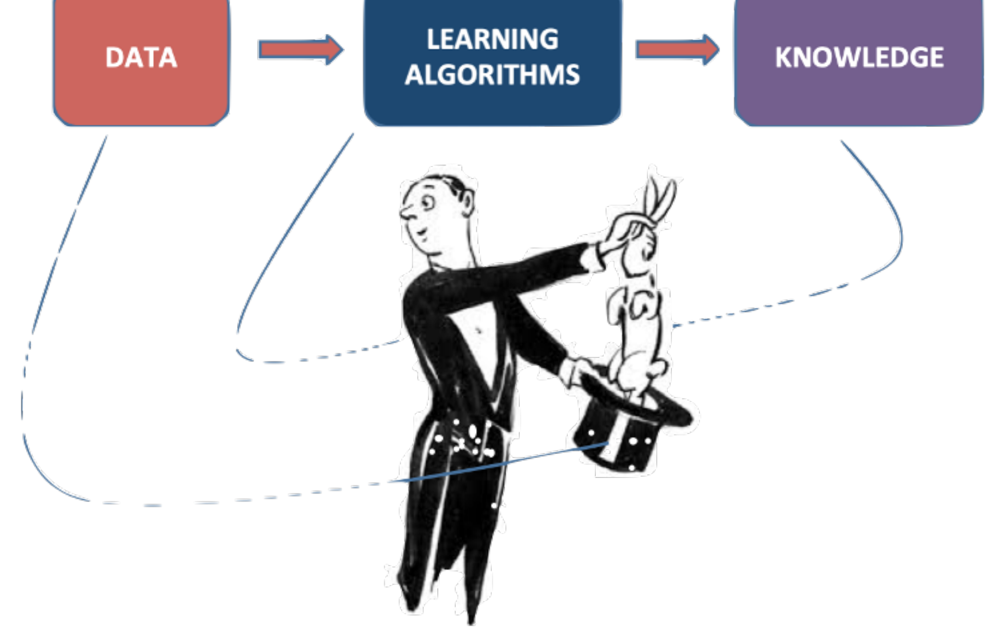
DATA
LEARNING ALGORITHMS
KNOWLEDGE

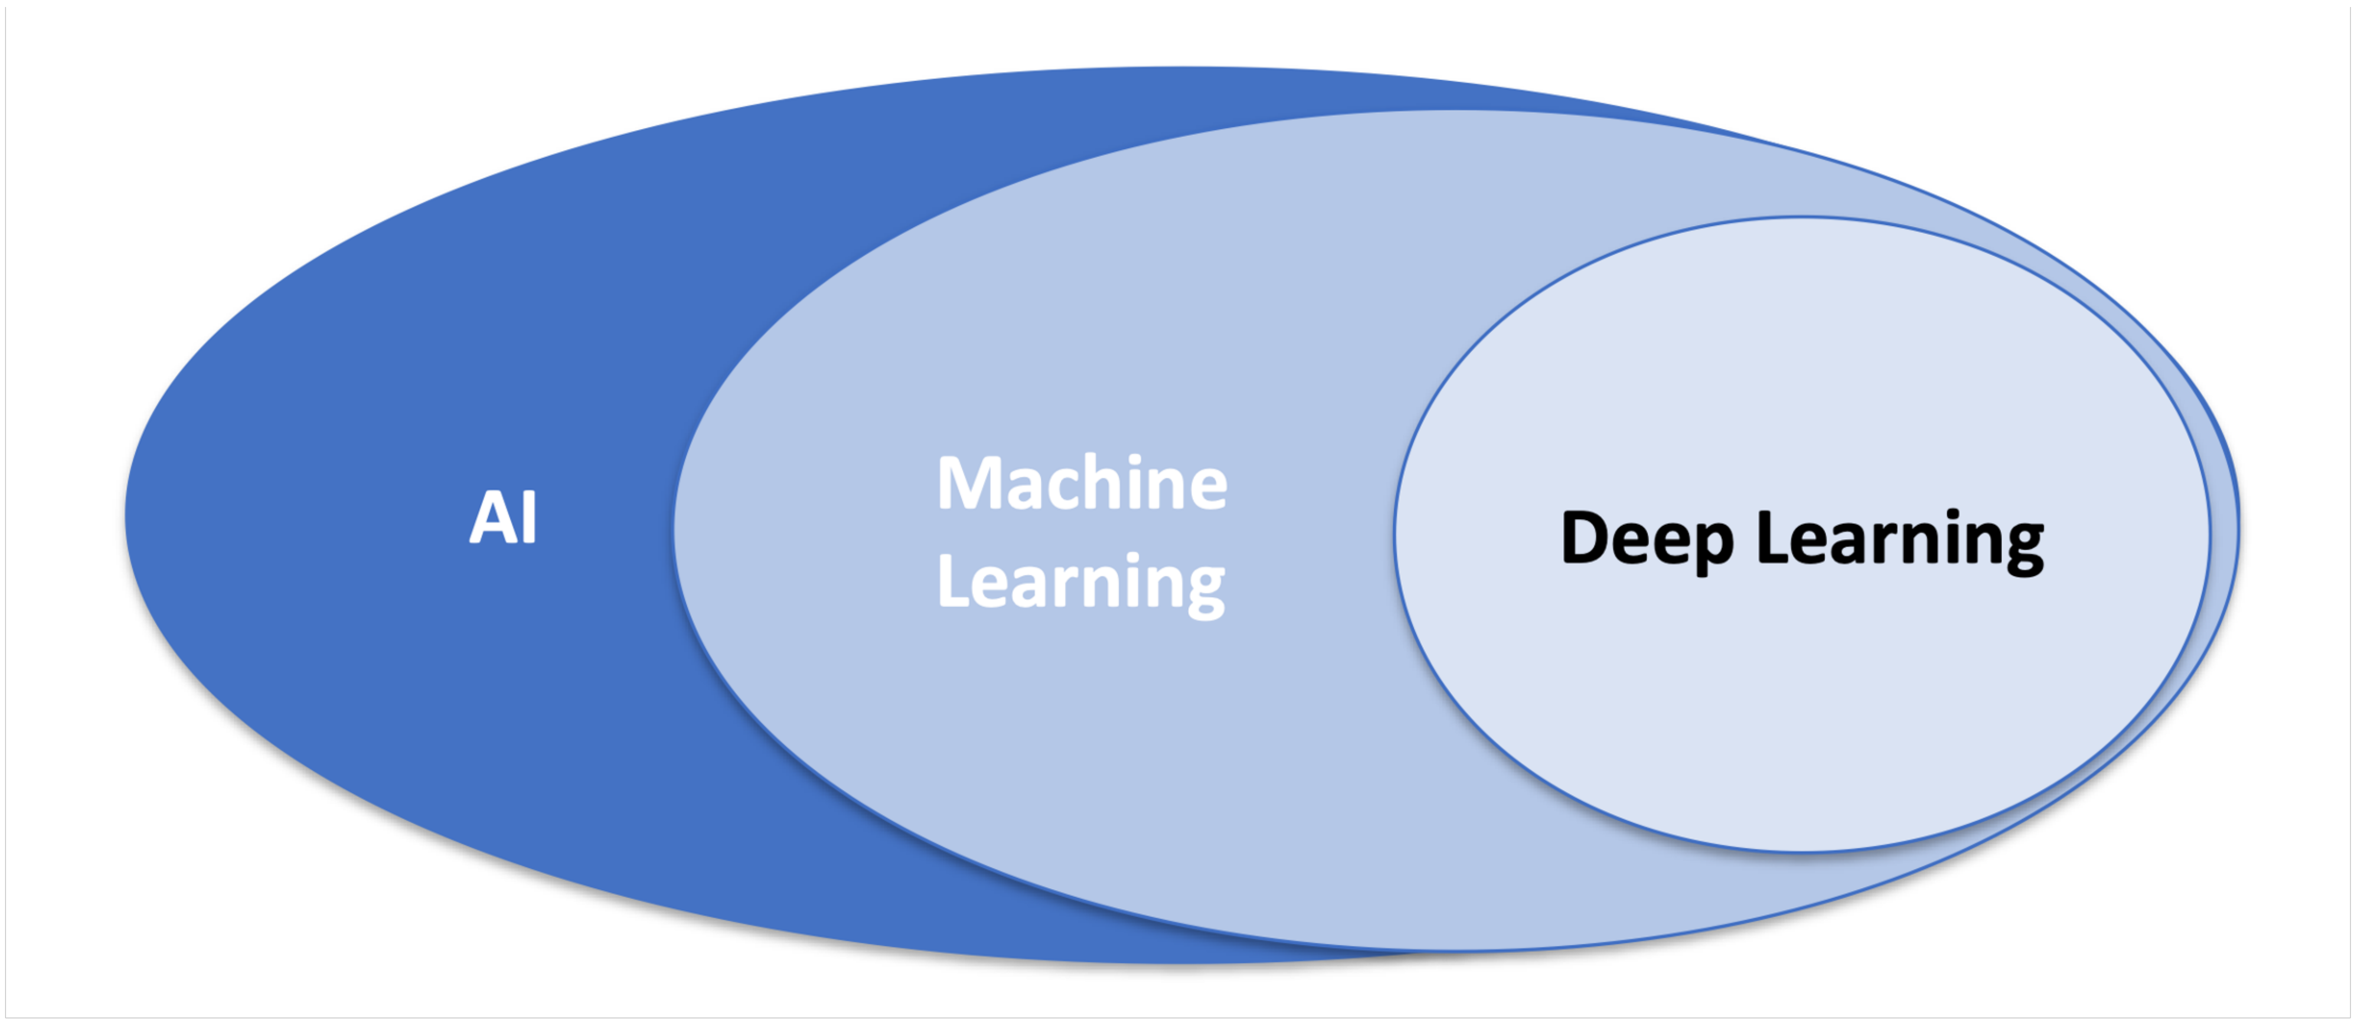
AI
Machine Learning
Deep Learning

Basic Categories of of ML

Supervised
learning

Unsupervised
learning

Reinforcement
learning

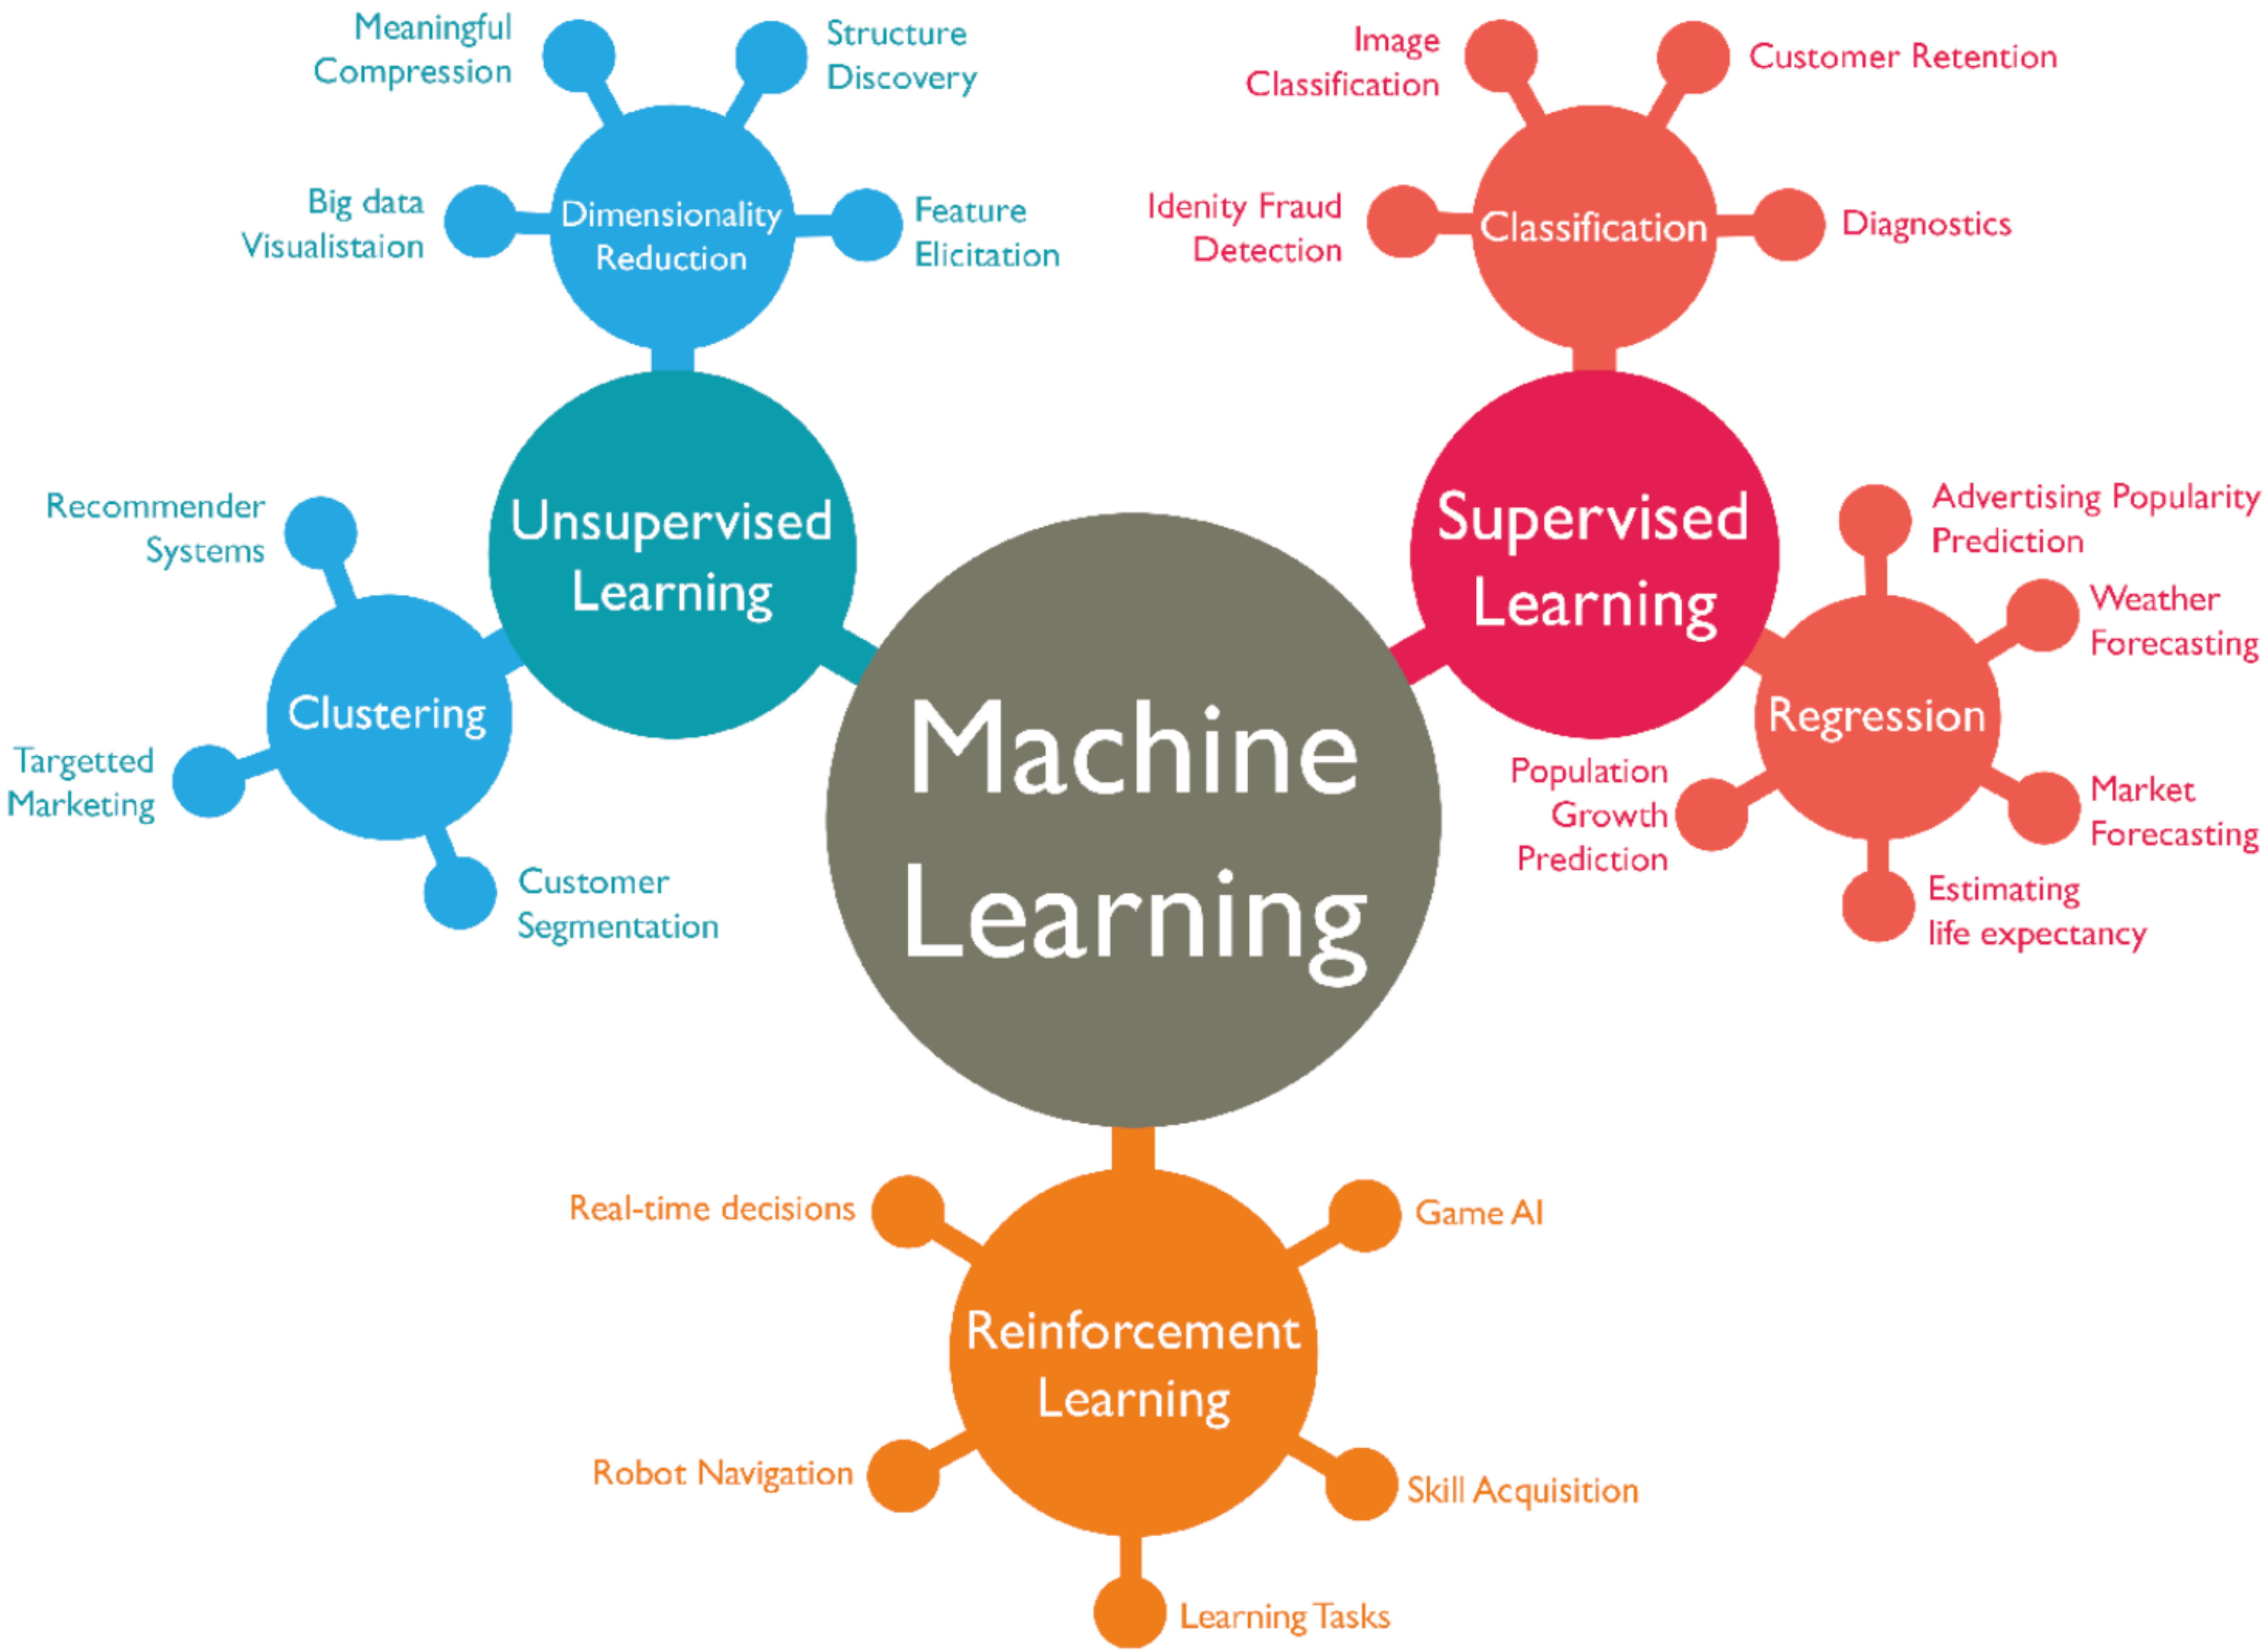

Image via Abdul Rahid

Supervised Learning

- Classification

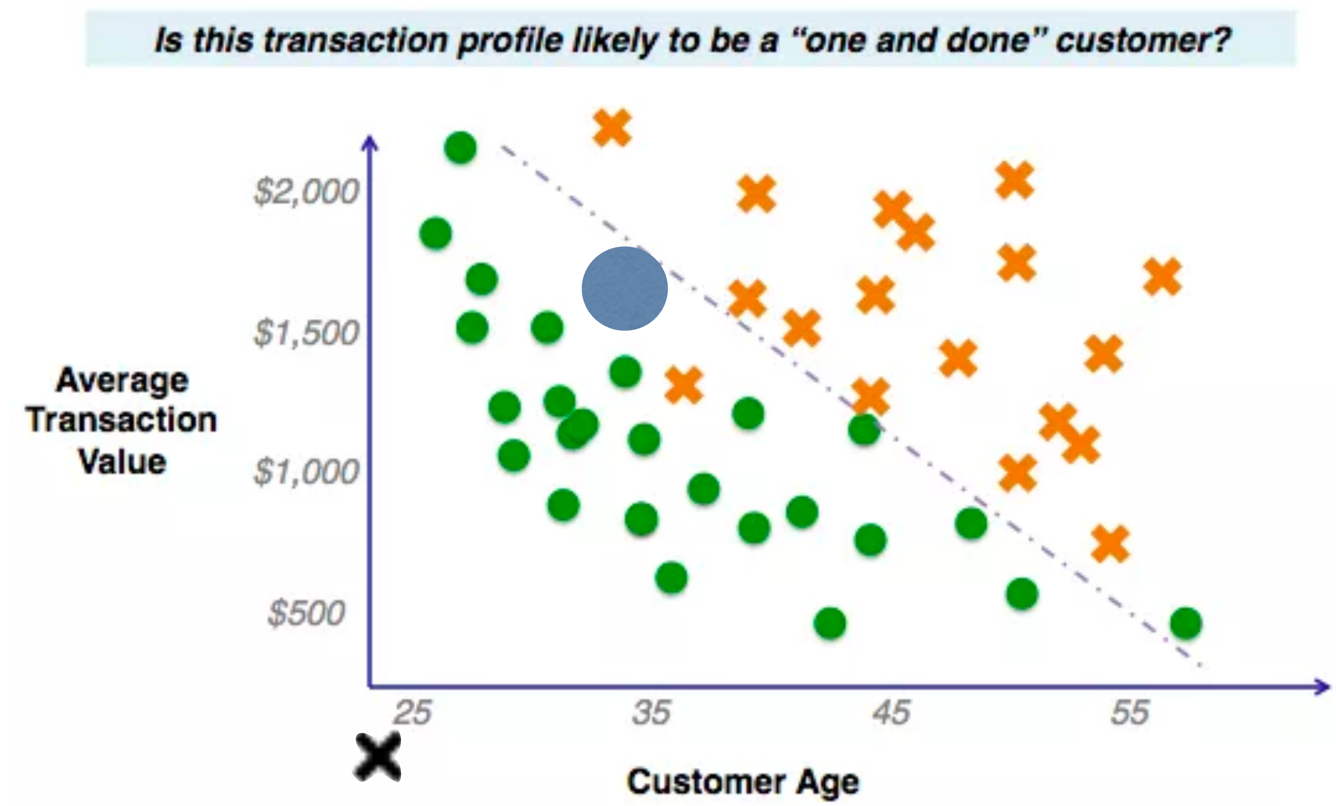

- Regression

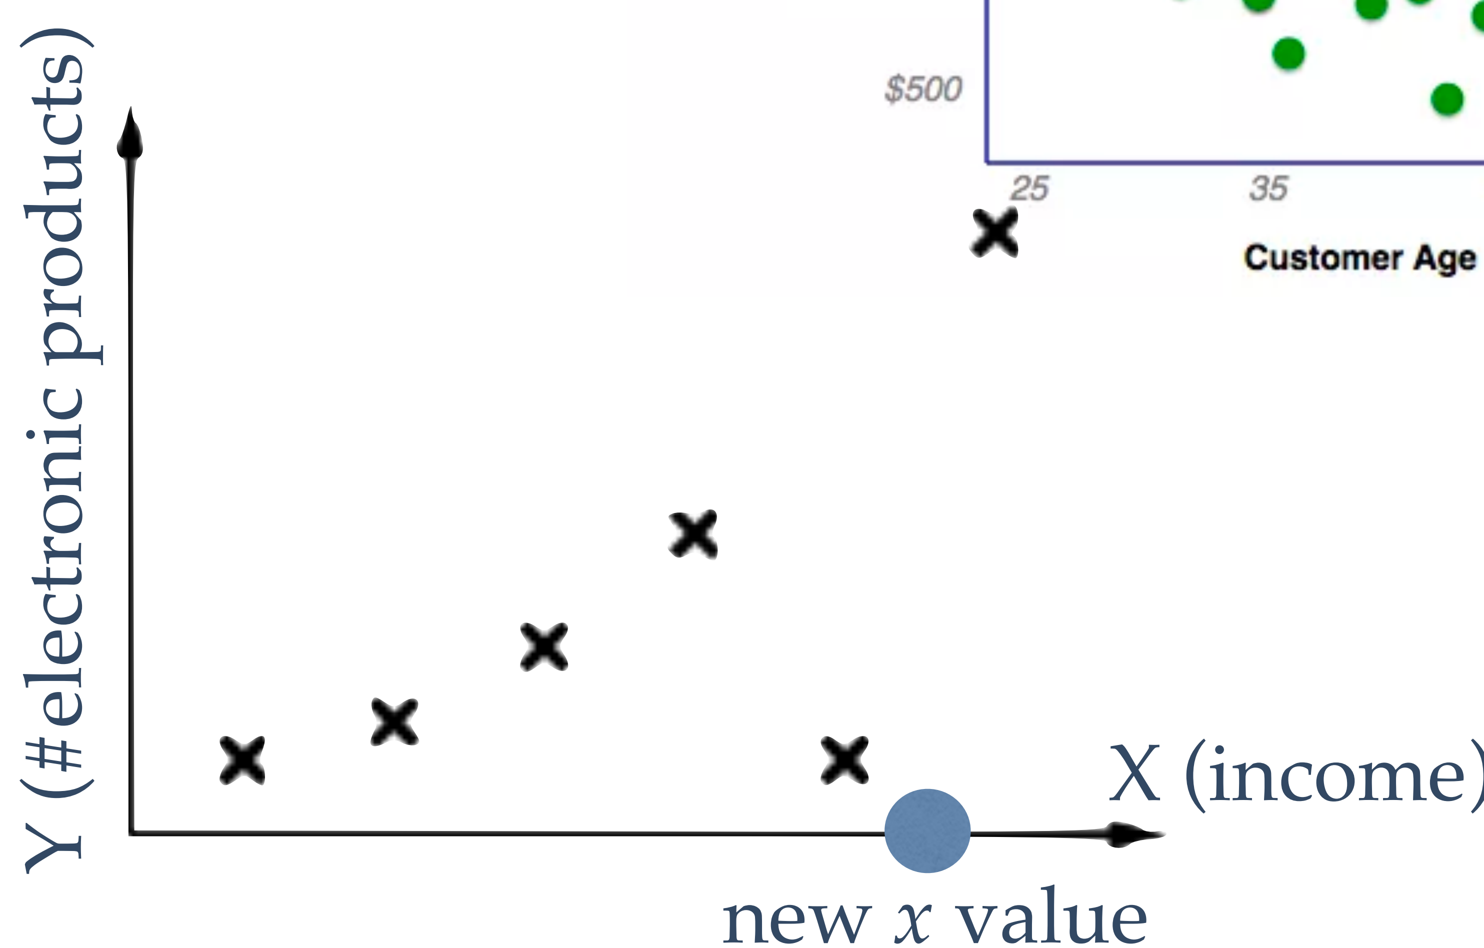

An Essential Problem in Supervised Learning

- Data: data pairs (x_i, y_i), i=1,…,n
- Learning: fitting curve $f(x)$ to "agree with" data
- Key: a good $f(x)$ generalizes well (i.e., predict unseen examples well)

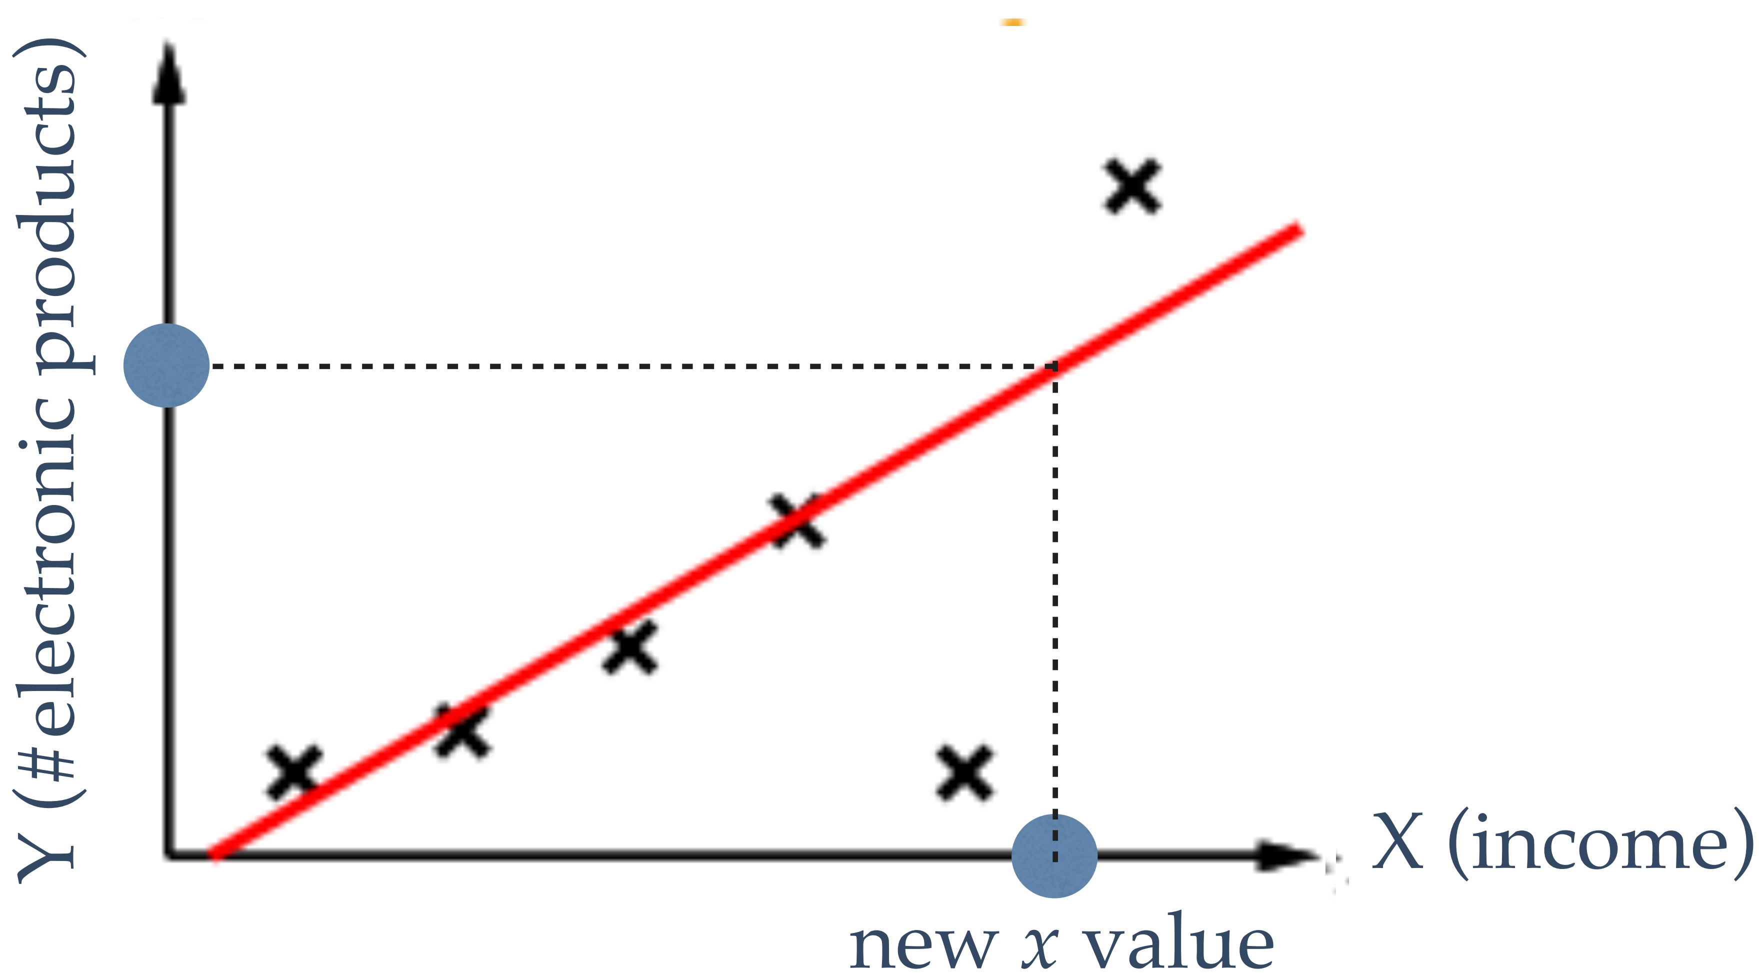

Supervised Learning Algorithms

- Nearest-neighbor
- Decision trees
- Linear/nonlinear regression
- Neural networks/deep learning
- …

Supervised Learning Algorithms

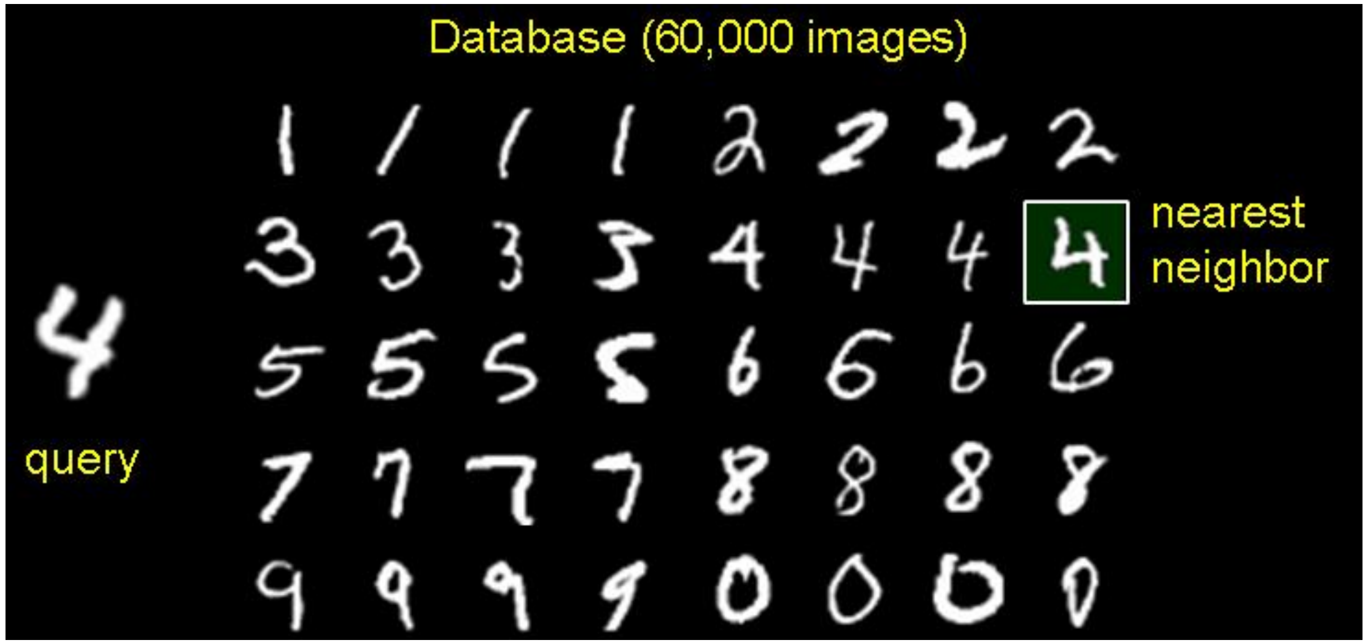

- **Nearest-neighbor**
- Decision trees
- Linear / nonlinear regression
- Neural networks / deep learning
- …

Unsupervised Learning

- Clustering
- Visualization
- Dimensional reduction

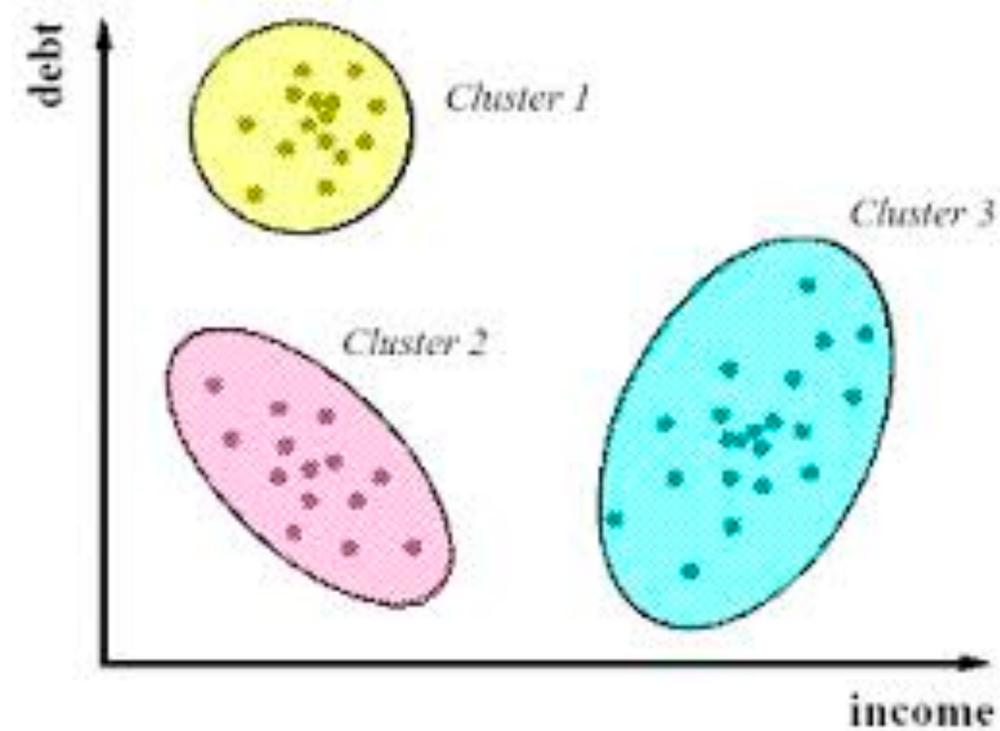

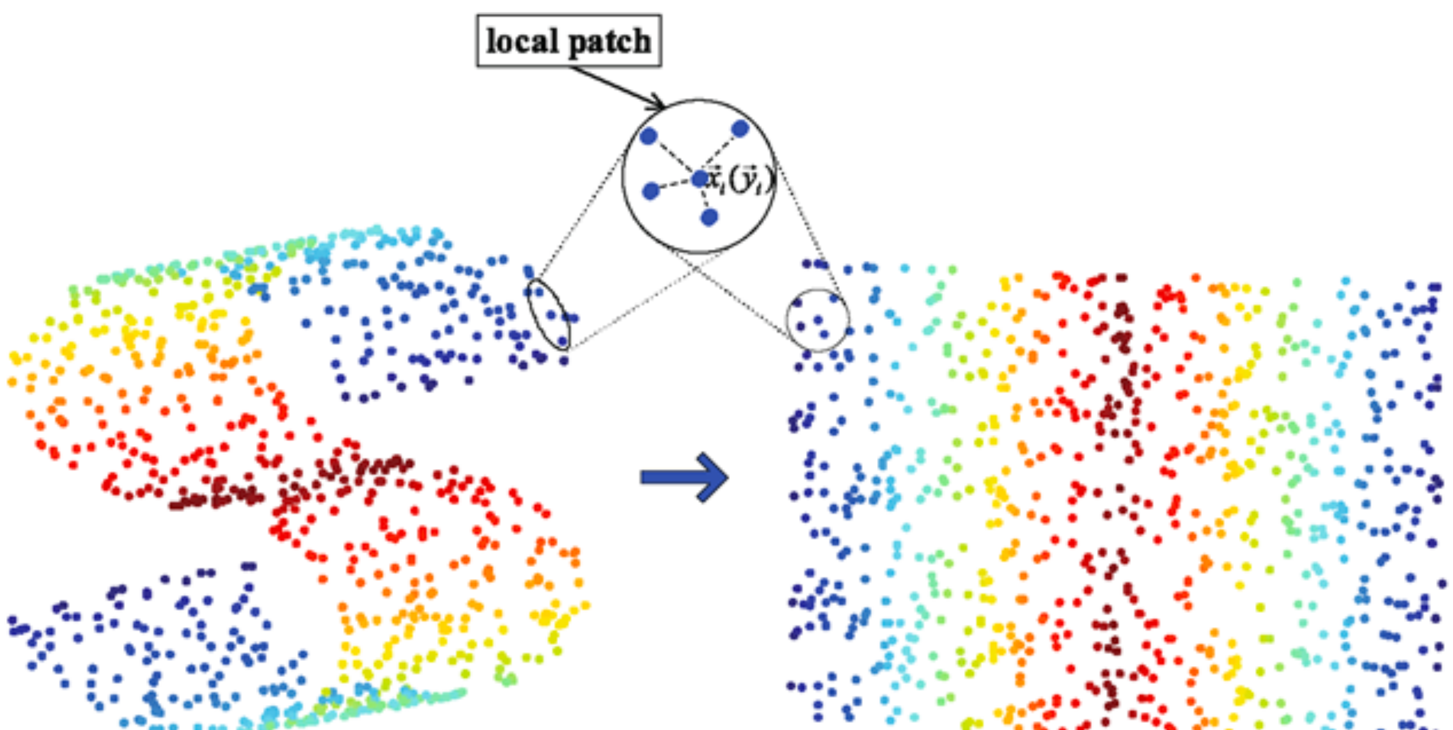

Two Ways of Finding Simpler Data Representations

- Fewer “data points” vs. ***fewer dimensions (#variables)***?

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1	Id	Population	Sex	Cranial size	Diet or subsistence					Paramastic	Dental wear		Geographic location per population			Climate per population					
2			(Male, fem	(Centroid S	Gathering	Hunting	Fishing	Pastoralism	Agriculture	Yes=1, no=	Average att	Attrition pe	Distance to	Longitude	Latitude	Tmean	Tmin	Tmax	Vpmean	Vpmin	Vpmax
3	AINU31_1	Ainu	Unknown	713.2942	2	3	4	0	1	0	1.5	2	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
4	AINU7_1	Ainu	Unknown	676.148	2	3	4	0	1	0	1.5	1	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
5	AINU7_2	Ainu	Unknown	675.4924	2	3	4	0	1	0	1.5	1	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
6	AINU_1016	Ainu	Male	684.3304	2	3	4	0	1	0	1.5	2.5	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
7	AINU_1016	Ainu	Female	686.285	2	3	4	0	1	0	1.5	4	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
8	AUSM245	Australia	Male	673.8749	6	4	0	0	0	1	2.5	1	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
9	AUSM246	Australia	Male	647.4586	6	4	0	0	0	1	2.5	4	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
10	AUSM8217	Australia	Male	658.6616	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
11	AUSM8177	Australia	Male	667.5444	6	4	0	0	0	1	2.5	4	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
12	AUSM8173	Australia	Male	629.7138	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
13	AUSM8173	Australia	Male	648.7064	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
14	AUSM8171	Australia	Male	643.0378	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
15	AUSM8165	Australia	Male	616.55	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
16	AUSM8154	Australia	Male	635.0605	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
17	AUSM8153	Australia	Male	650.6959	6	4	0	0	0	1	2.5	3	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
18	AUSF1412	Australia	Female	618.4781	6	4	0	0	0	1	2.5	1	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
19	AUSF8179	Australia	Female	634.3122	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
20	AUSF8175	Australia	Female	605.1759	6	4	0	0	0	1	2.5	1.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
21	AUSF8172	Australia	Female	613.8324	6	4	0	0	0	1	2.5	3	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
22	AUSF8169	Australia	Female	619.1206	6	4	0	0	0	1	2.5	2.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
23	AUSF8157	Australia	Female	628.2819	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
24	AUSF8155	Australia	Female	628.4609	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
25	AUSF1578	Australia	Female	640.6311	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
26	AUSF243	Australia	Female	606.164	6	4	0	0	0	1	2.5	2.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
27	AUSF8158	Australia	Female	631.6258	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
28	DENM1432	Denmark	Male	663.6198	0	0	1	3	6	0	2.1	2	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
29	DENM1011	Denmark	Male	651.4847	0	0	1	3	6	0	2.1	3	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
30	DENM1205	Denmark	Male	636.9831	0	0	1	3	6	0	2.1	1.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
31	DENM116_	Denmark	Male	642.9192	0	0	1	3	6	0	2.1	3	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
32	DENM116_	Denmark	Male	646.6609	0	0	1	3	6	0	2.1	2.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
33	DENM116_	Denmark	Male	674.9799	0	0	1	3	6	0	2.1	2	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
34	DENM7_77	Denmark	Male	666.53	0	0	1	3	6	0	2.1	2.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
35	DENM1_58	Denmark	Male	627.4583	0	0	1	3	6	0	2.1	1.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
36	DENM903	Denmark	Male	662.5953	0	0	1	3	6	0	2.1	2	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
37	DENM901	Denmark	Male	672.8408	0	0	1	3	6	0	2.1	NaN	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
38	DENF1559	Denmark	Female	604.4864	0	0	1	3	6	0	2.1	0.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27

INFO | Shape_rawcoordinates | Shape_symmProcCoord. | Factors | Background | Geodistance_extra | +

Unsupervised Learning Algorithms

- K-means clustering
- Principal component analysis
- Autoencoders
- …

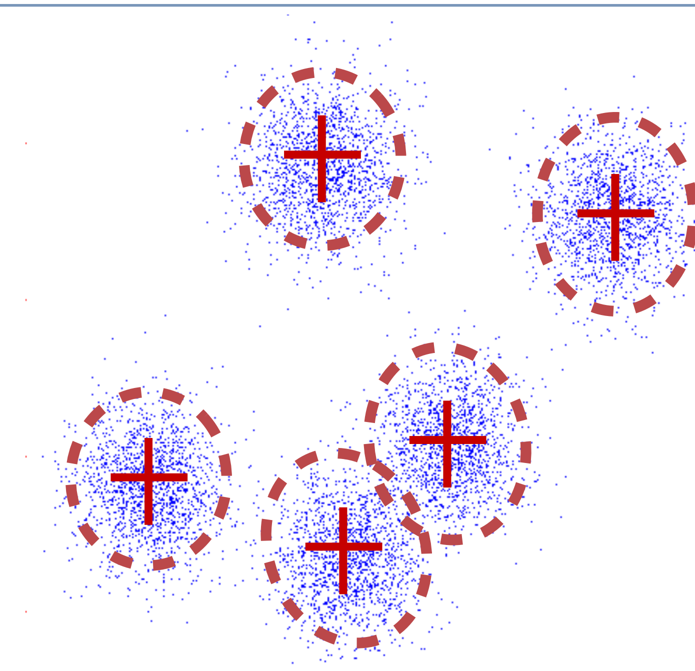

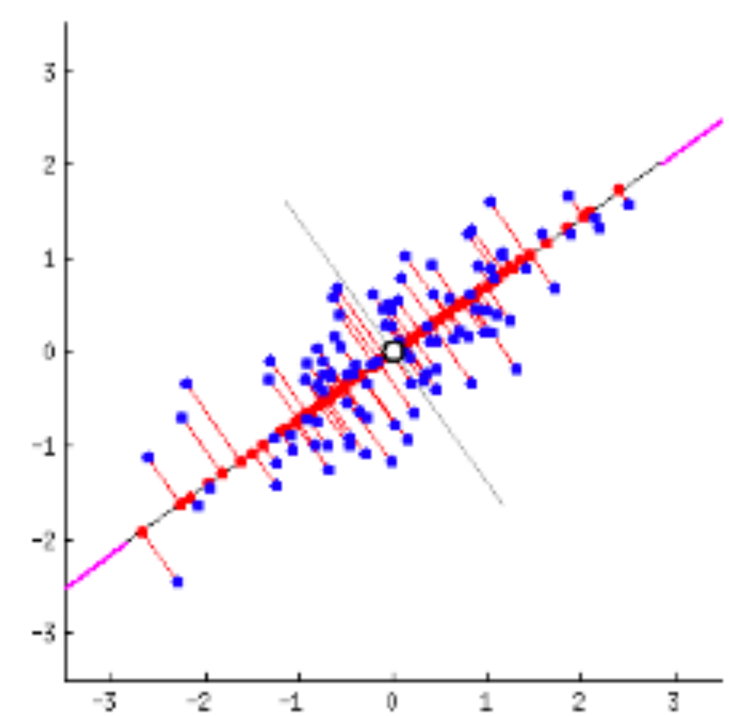

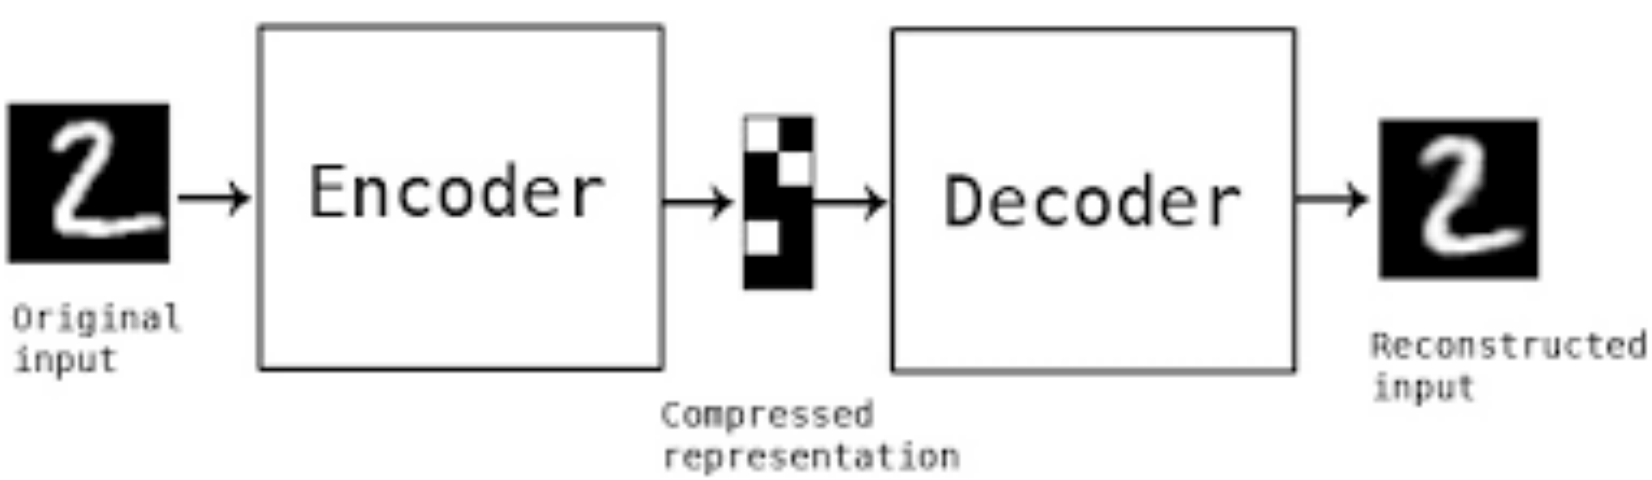

Unsupervised Learning: Clustering

Group similar things e.g. images

[Goldberger et al.]

Meaningful Clustering?

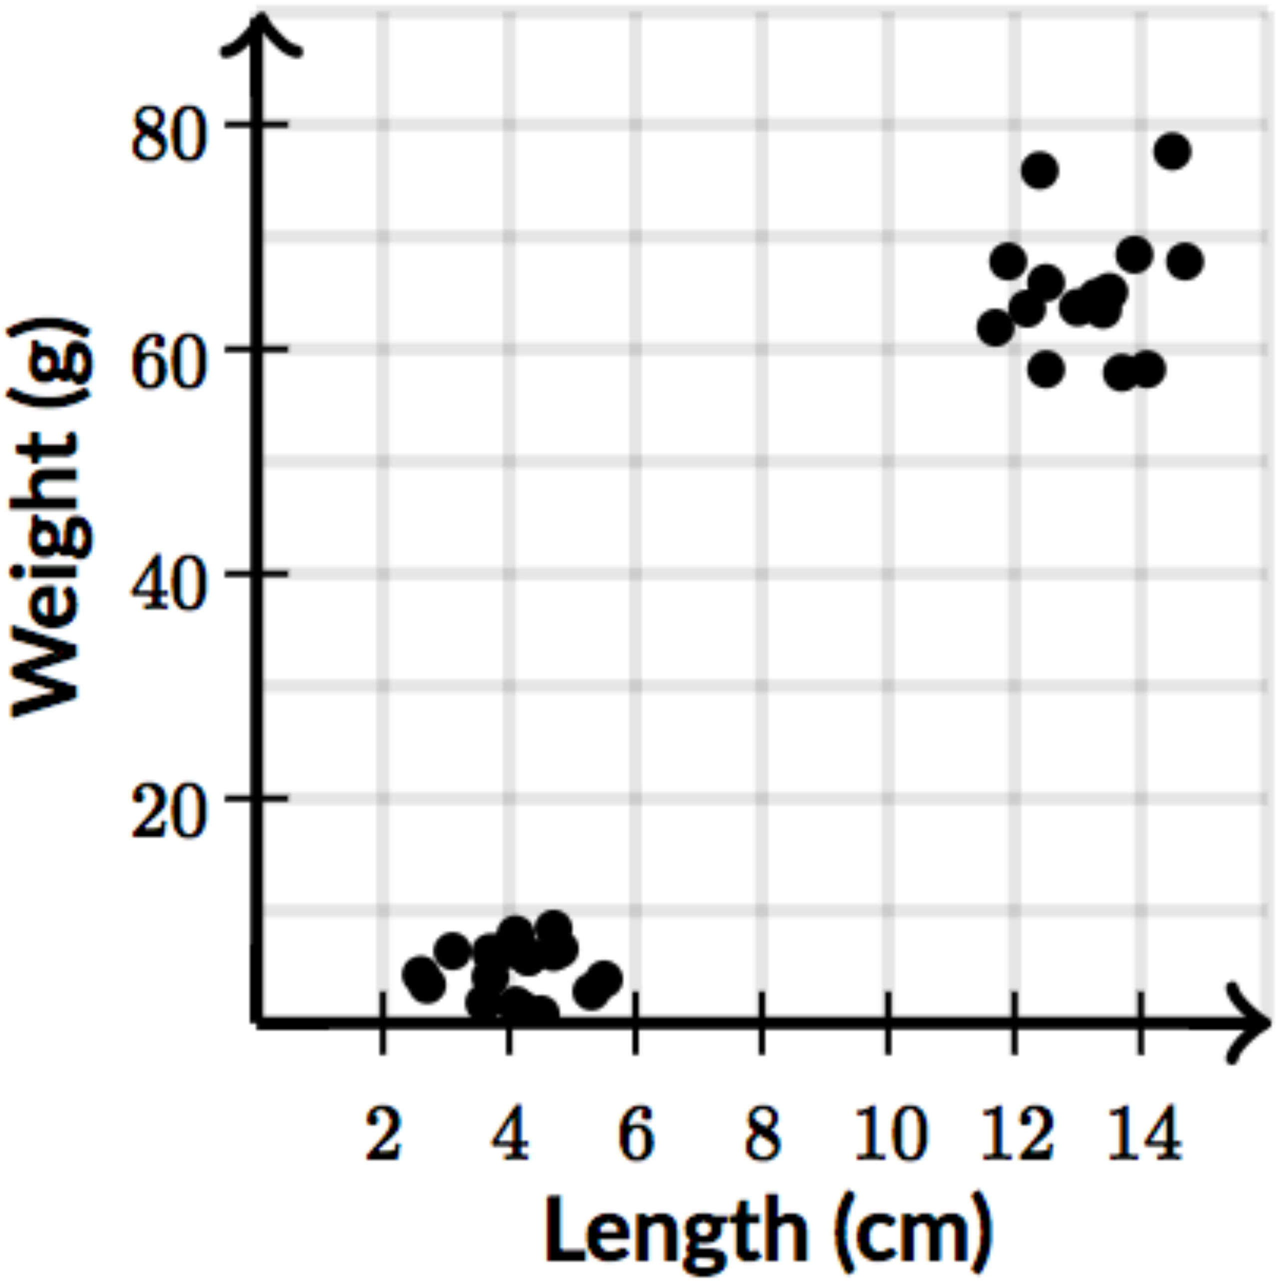

- Can you cluster the percentage grade to obtain letter grade?
- Can you recover 'red wine quality' from measured features (fixed acidity, volatile acidity, citric acid, residual sugar, chlorides, free sulfur dioxide, total sulfur dioxide, density, pH, etc.)?

Causality and Invariance, Robustness, etc.

- Consider prediction (with regression) in different time periods

data

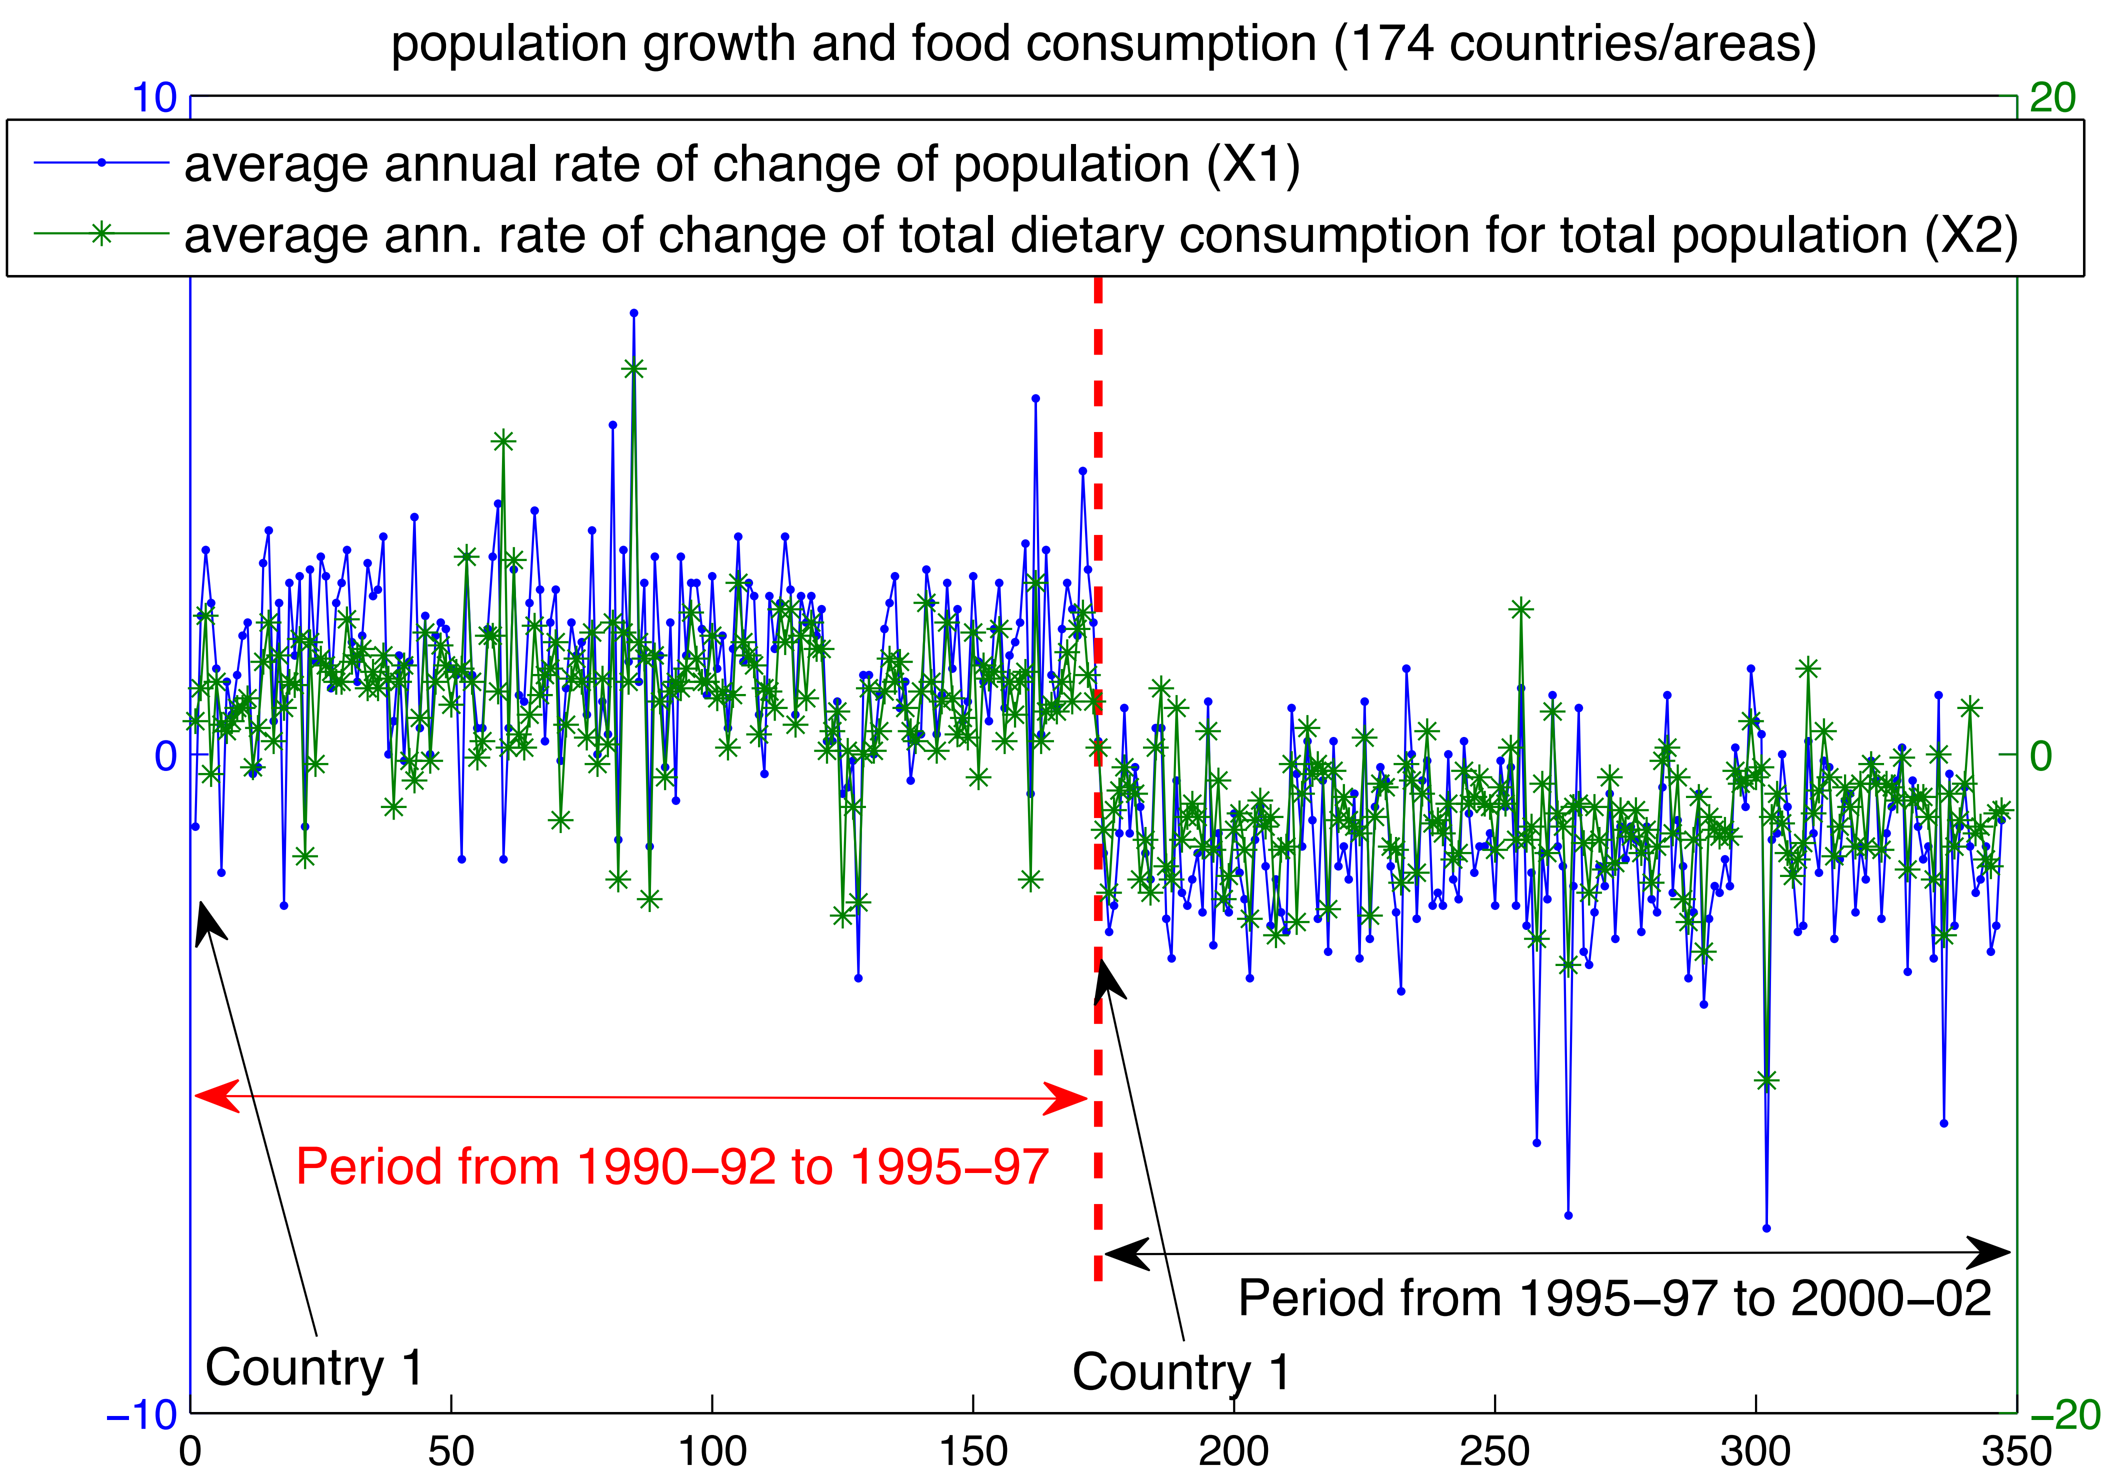

Data

- Population growth and food consumption:
  - data for 174 countries or areas, during the period from 1990-92 to 1995-97 (former 174 data points) and that from 1995-97 to 2000-02 (latter 174 points).
- X1: the average annual rate of change of population; X2: the average annual rate of change of total dietary consumption for total population (kcal/day)

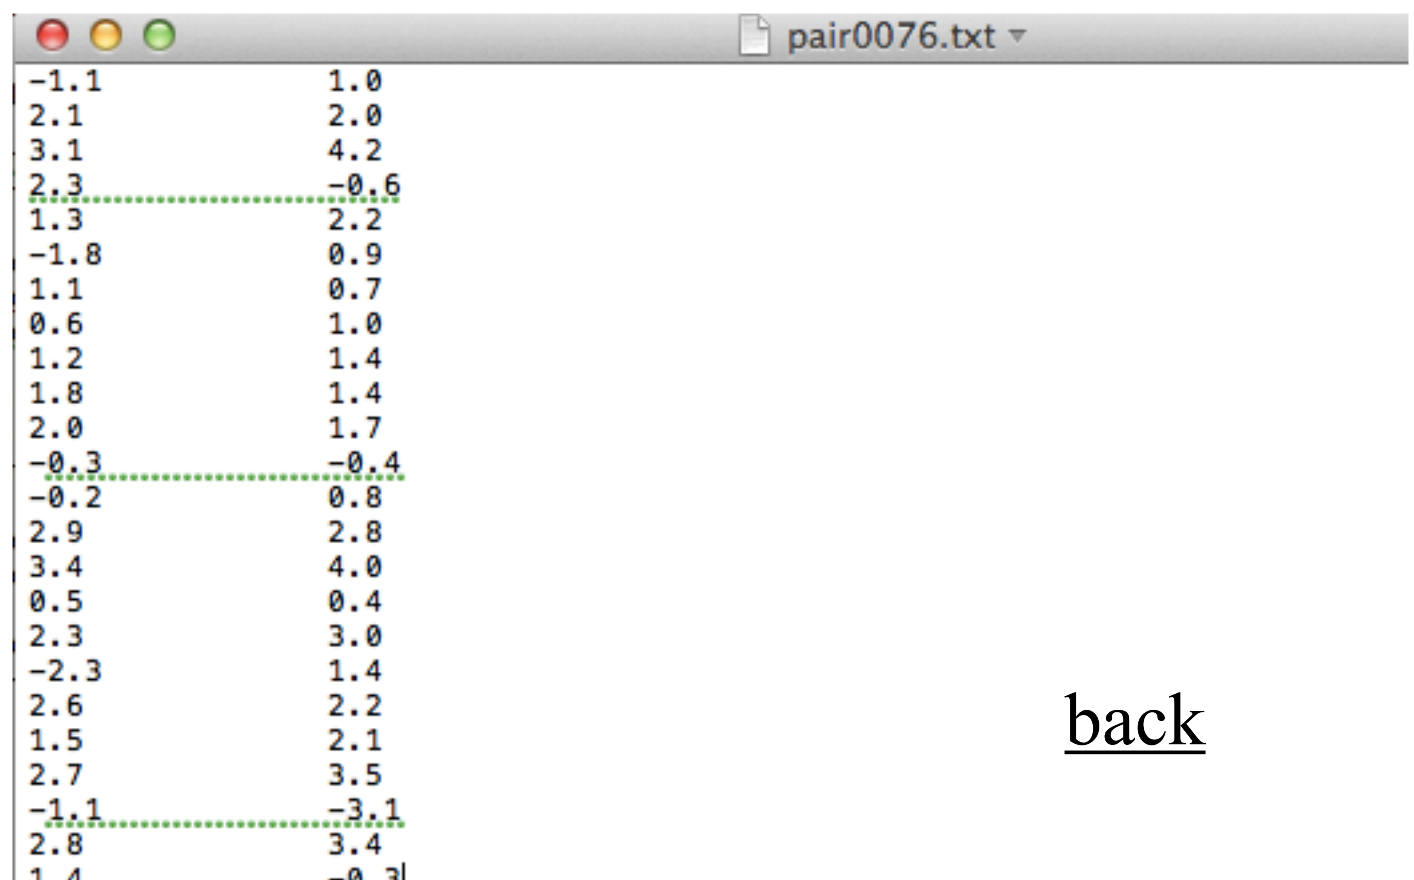
pair0076.txt

-1.1	1.0
2.1	2.0
3.1	4.2
2.3	-0.6
1.3	2.2
-1.8	0.9
1.1	0.7
0.6	1.0
1.2	1.4
1.8	1.4
2.0	1.7
-0.3	-0.4
-0.2	0.8
2.9	2.8
3.4	4.0
0.5	0.4
2.3	3.0
-2.3	1.4
2.6	2.2
1.5	2.1
2.7	3.5
-1.1	-3.1
2.8	3.4
1.4	-0.3

back

Causality and Invariance, Robustness, etc.

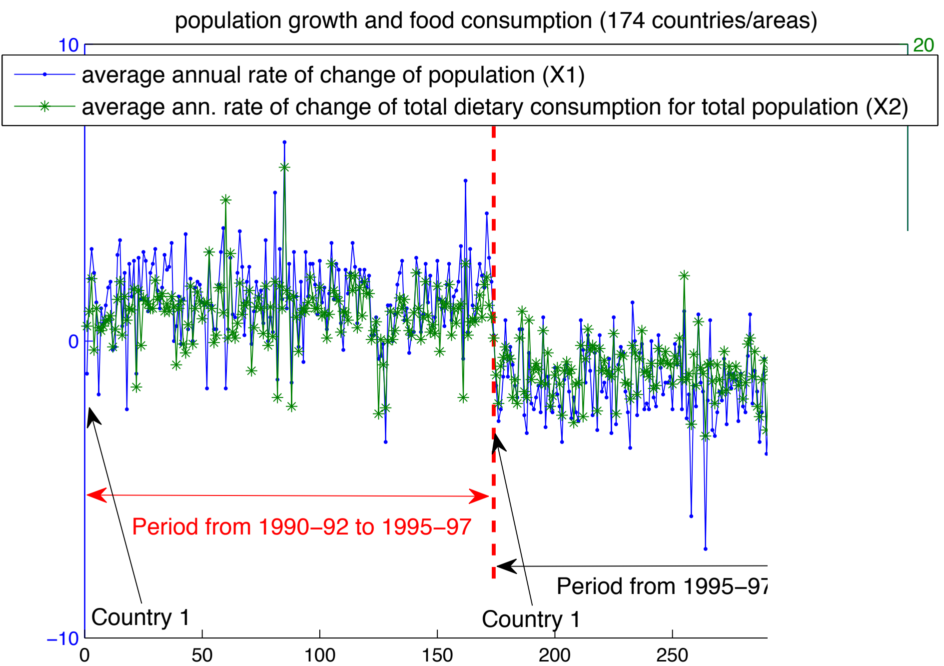

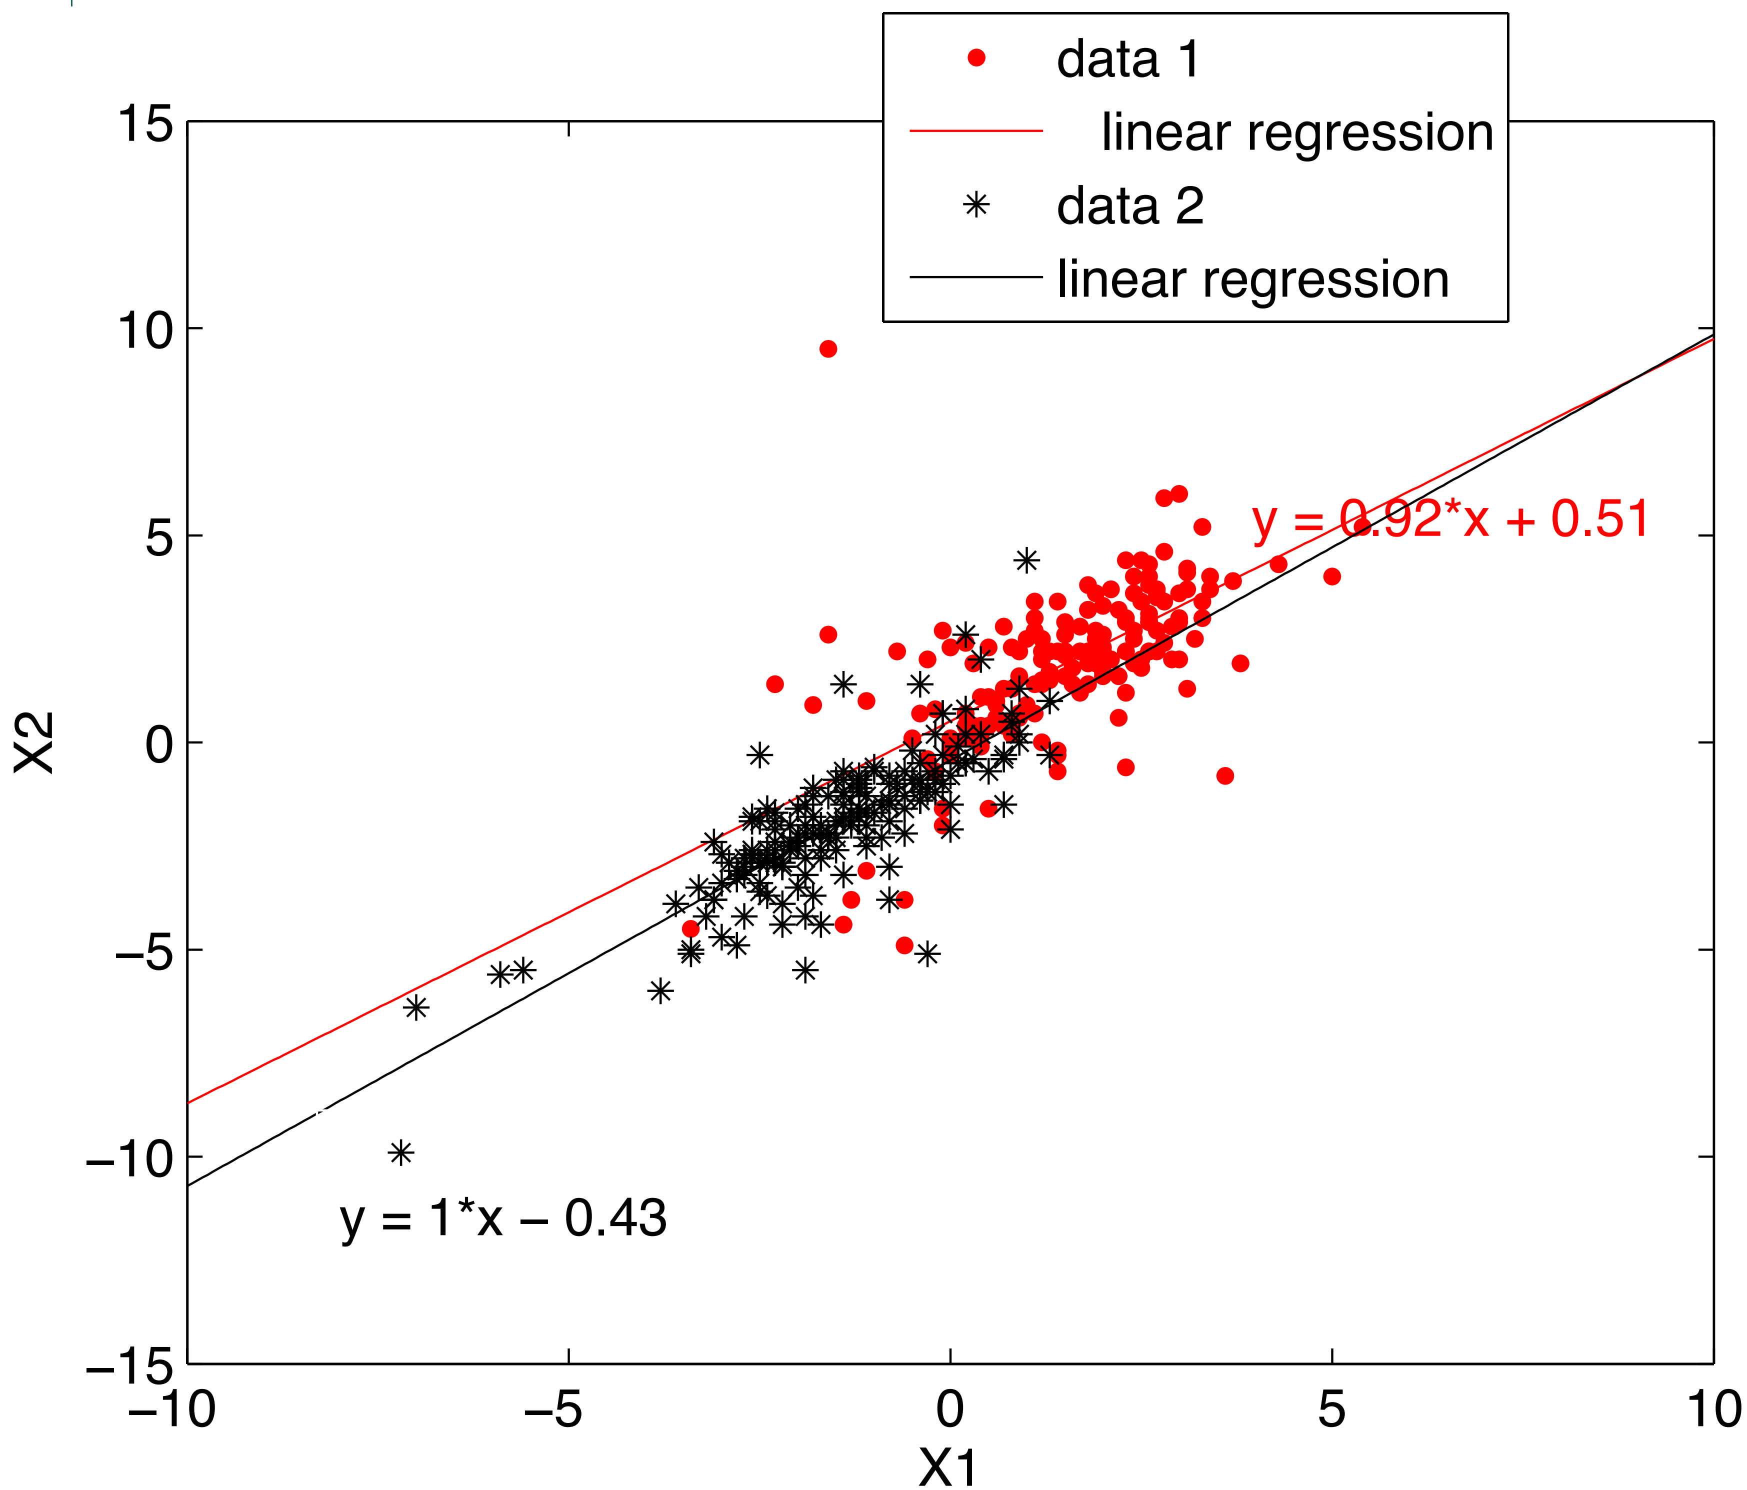

Causality and Invariance, Robustness, etc.

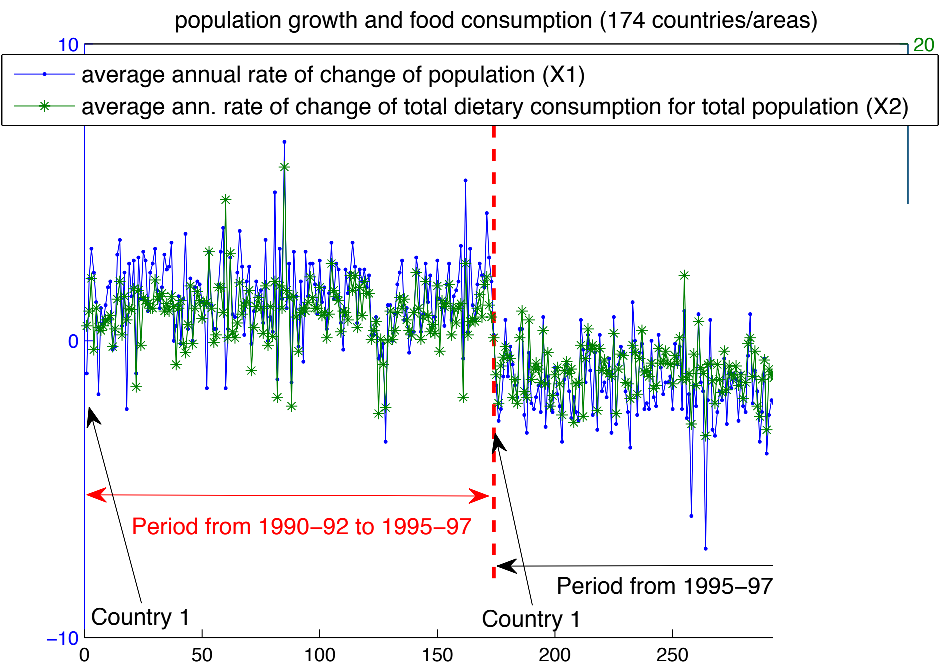

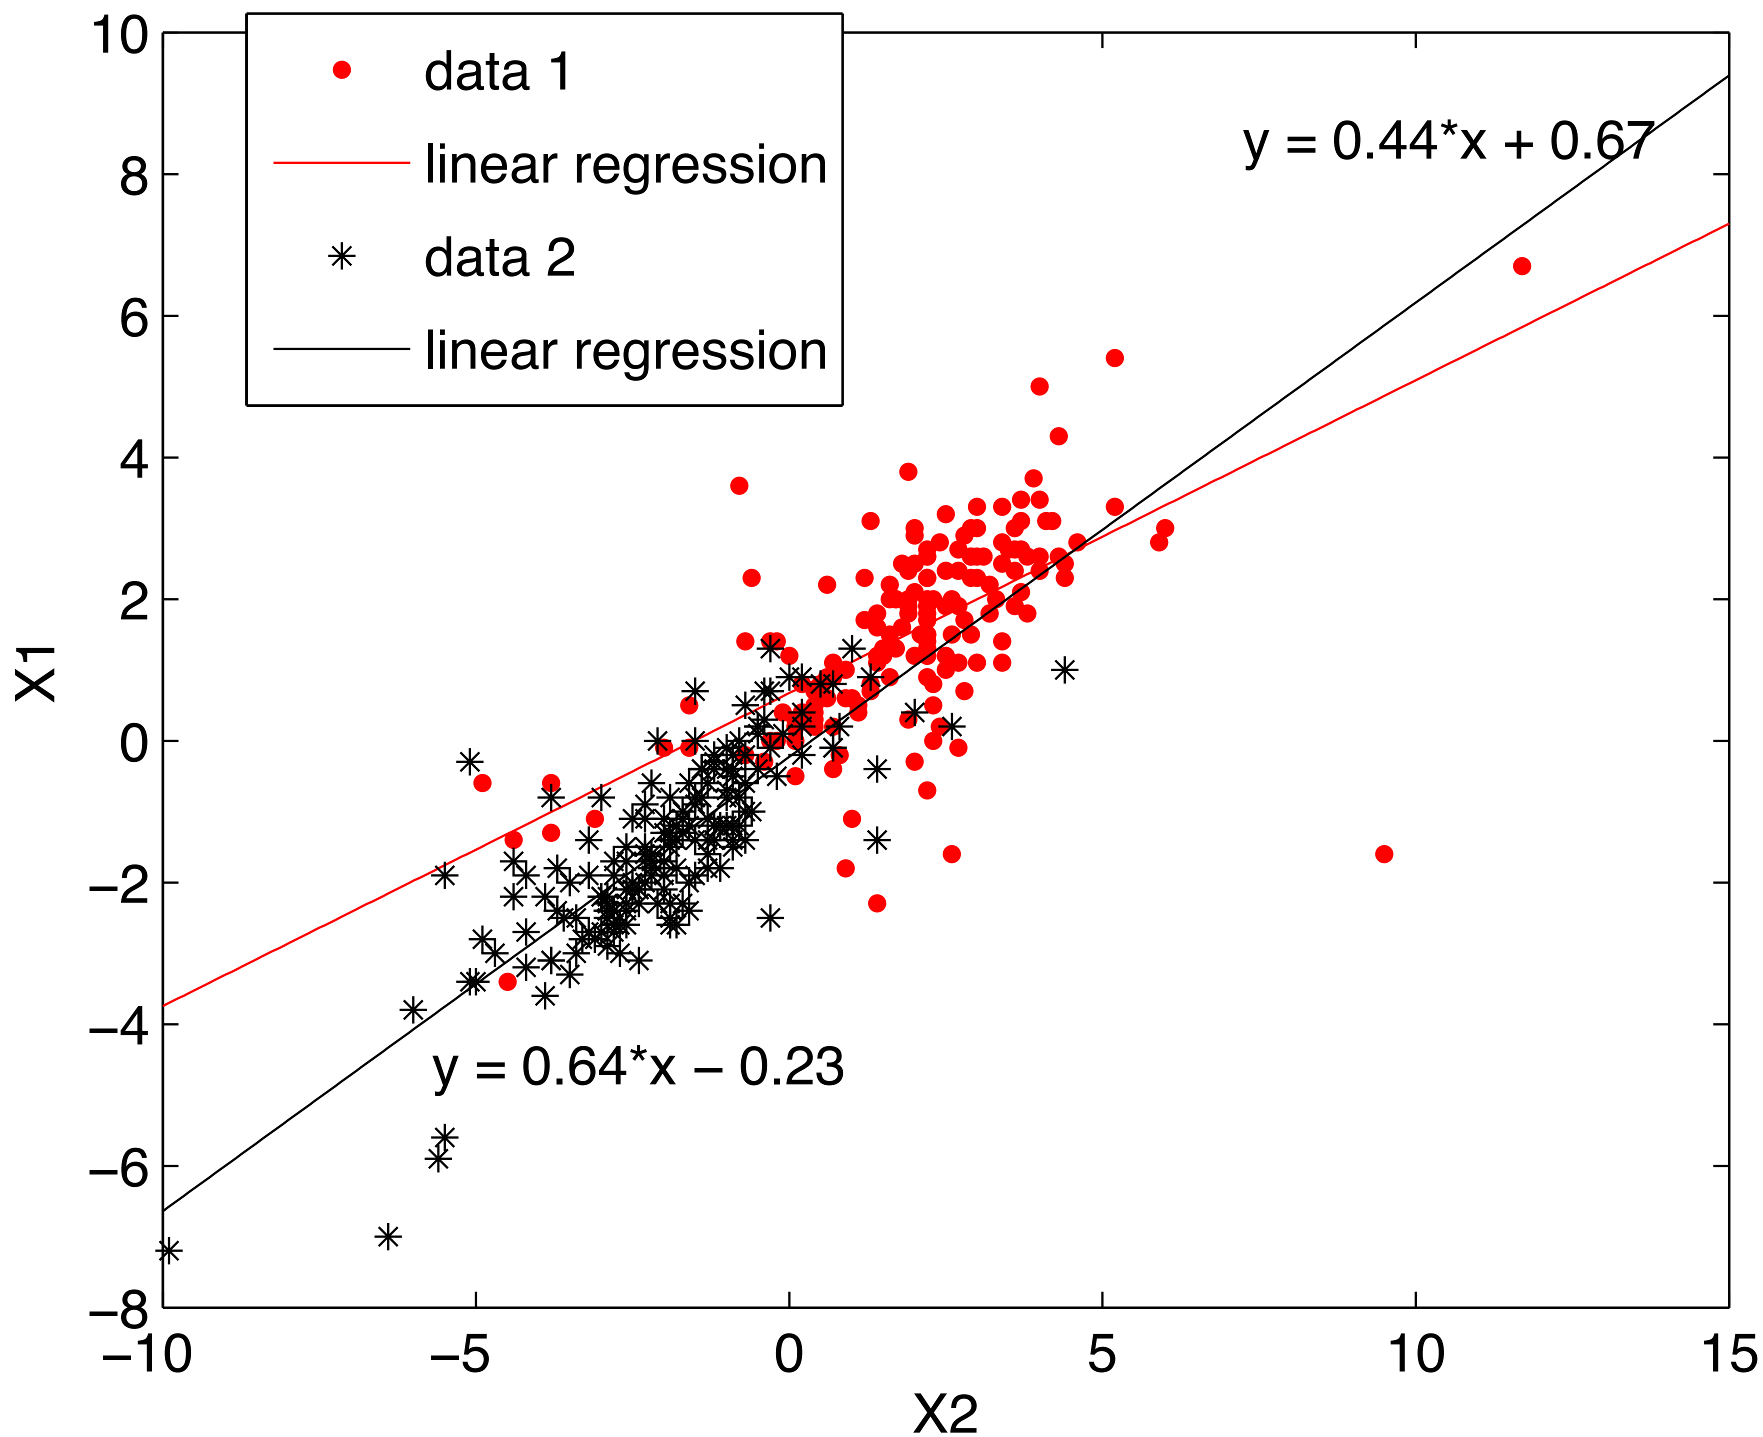

Causality and Invariance, Robustness, etc.

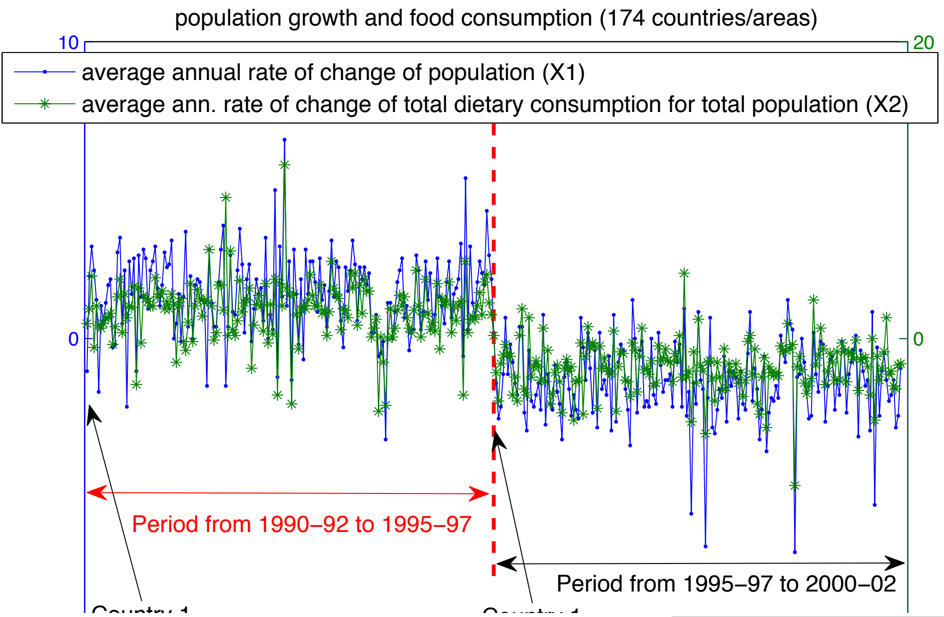

- *Invariance!*
- *More generally, independent changes*

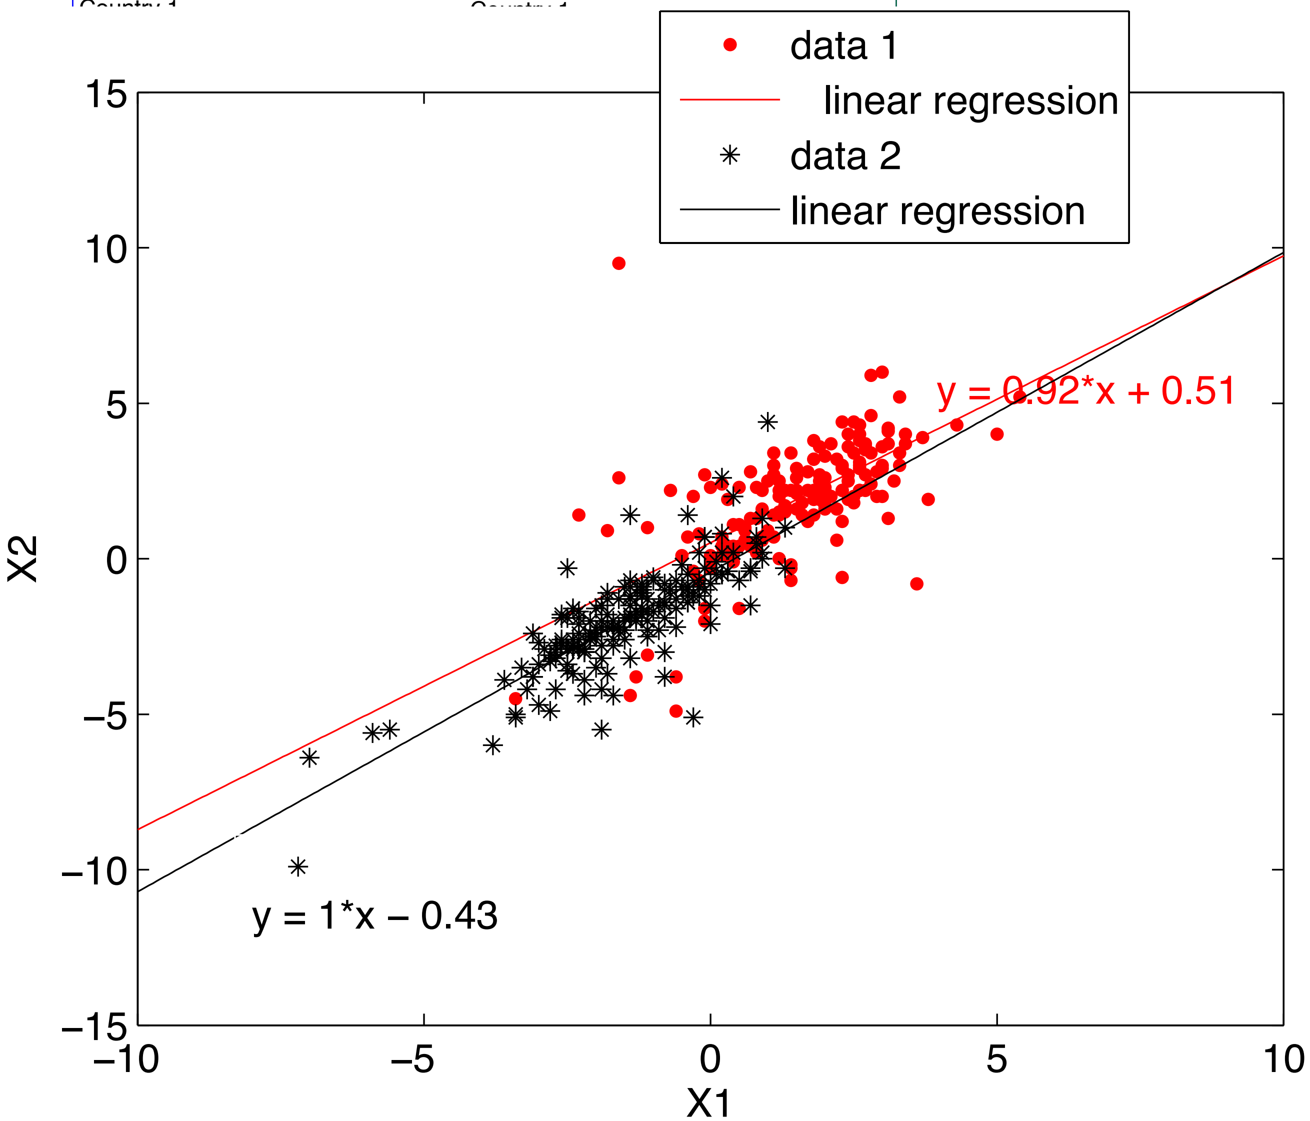

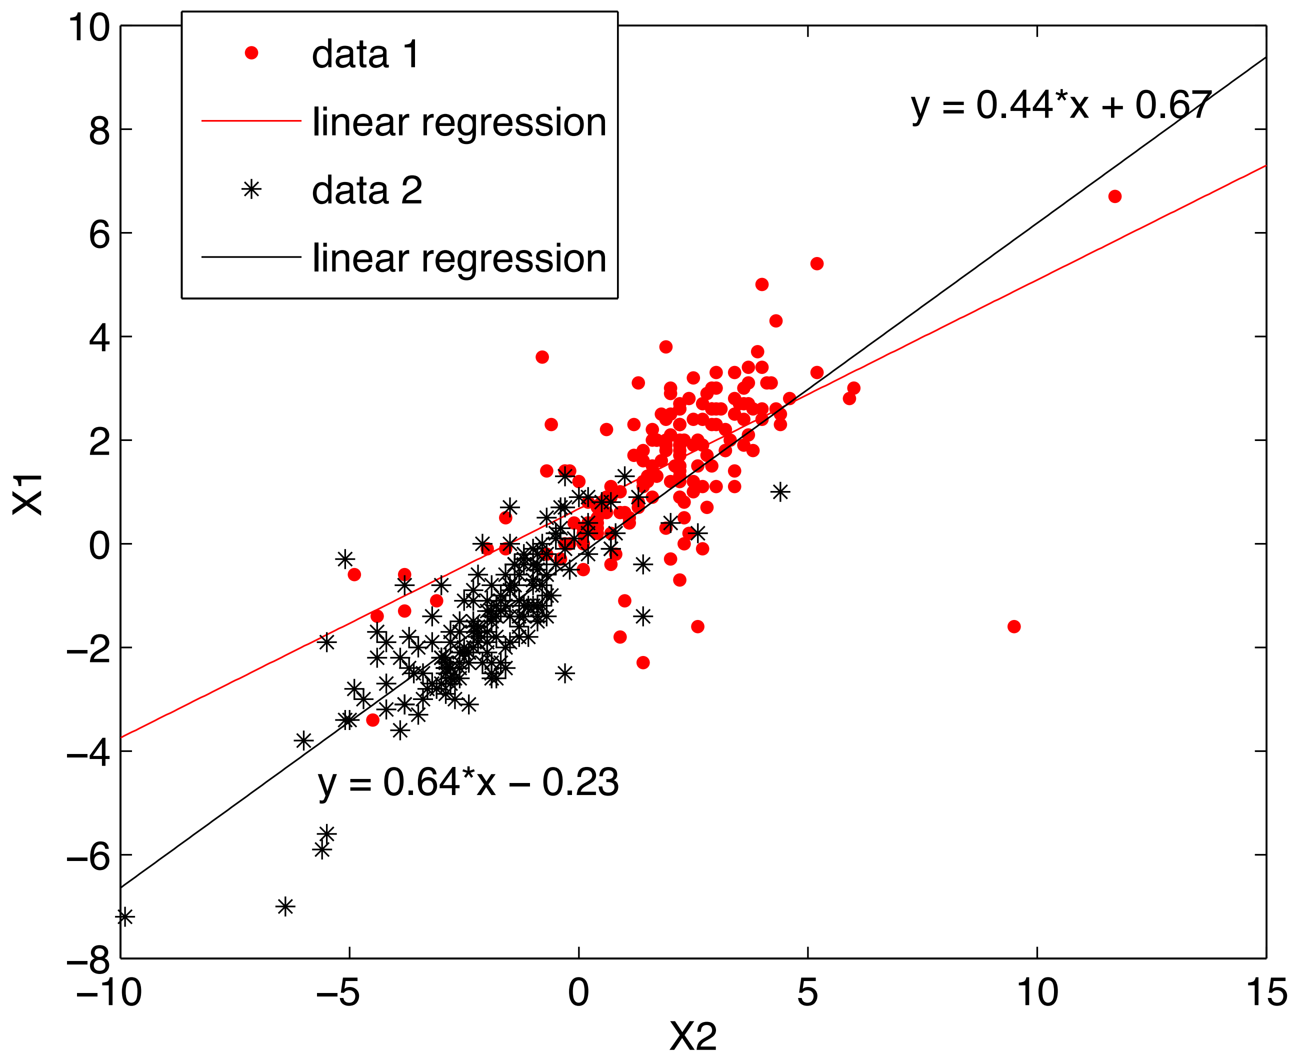

Recommender Systems: Why Make Recommendations?

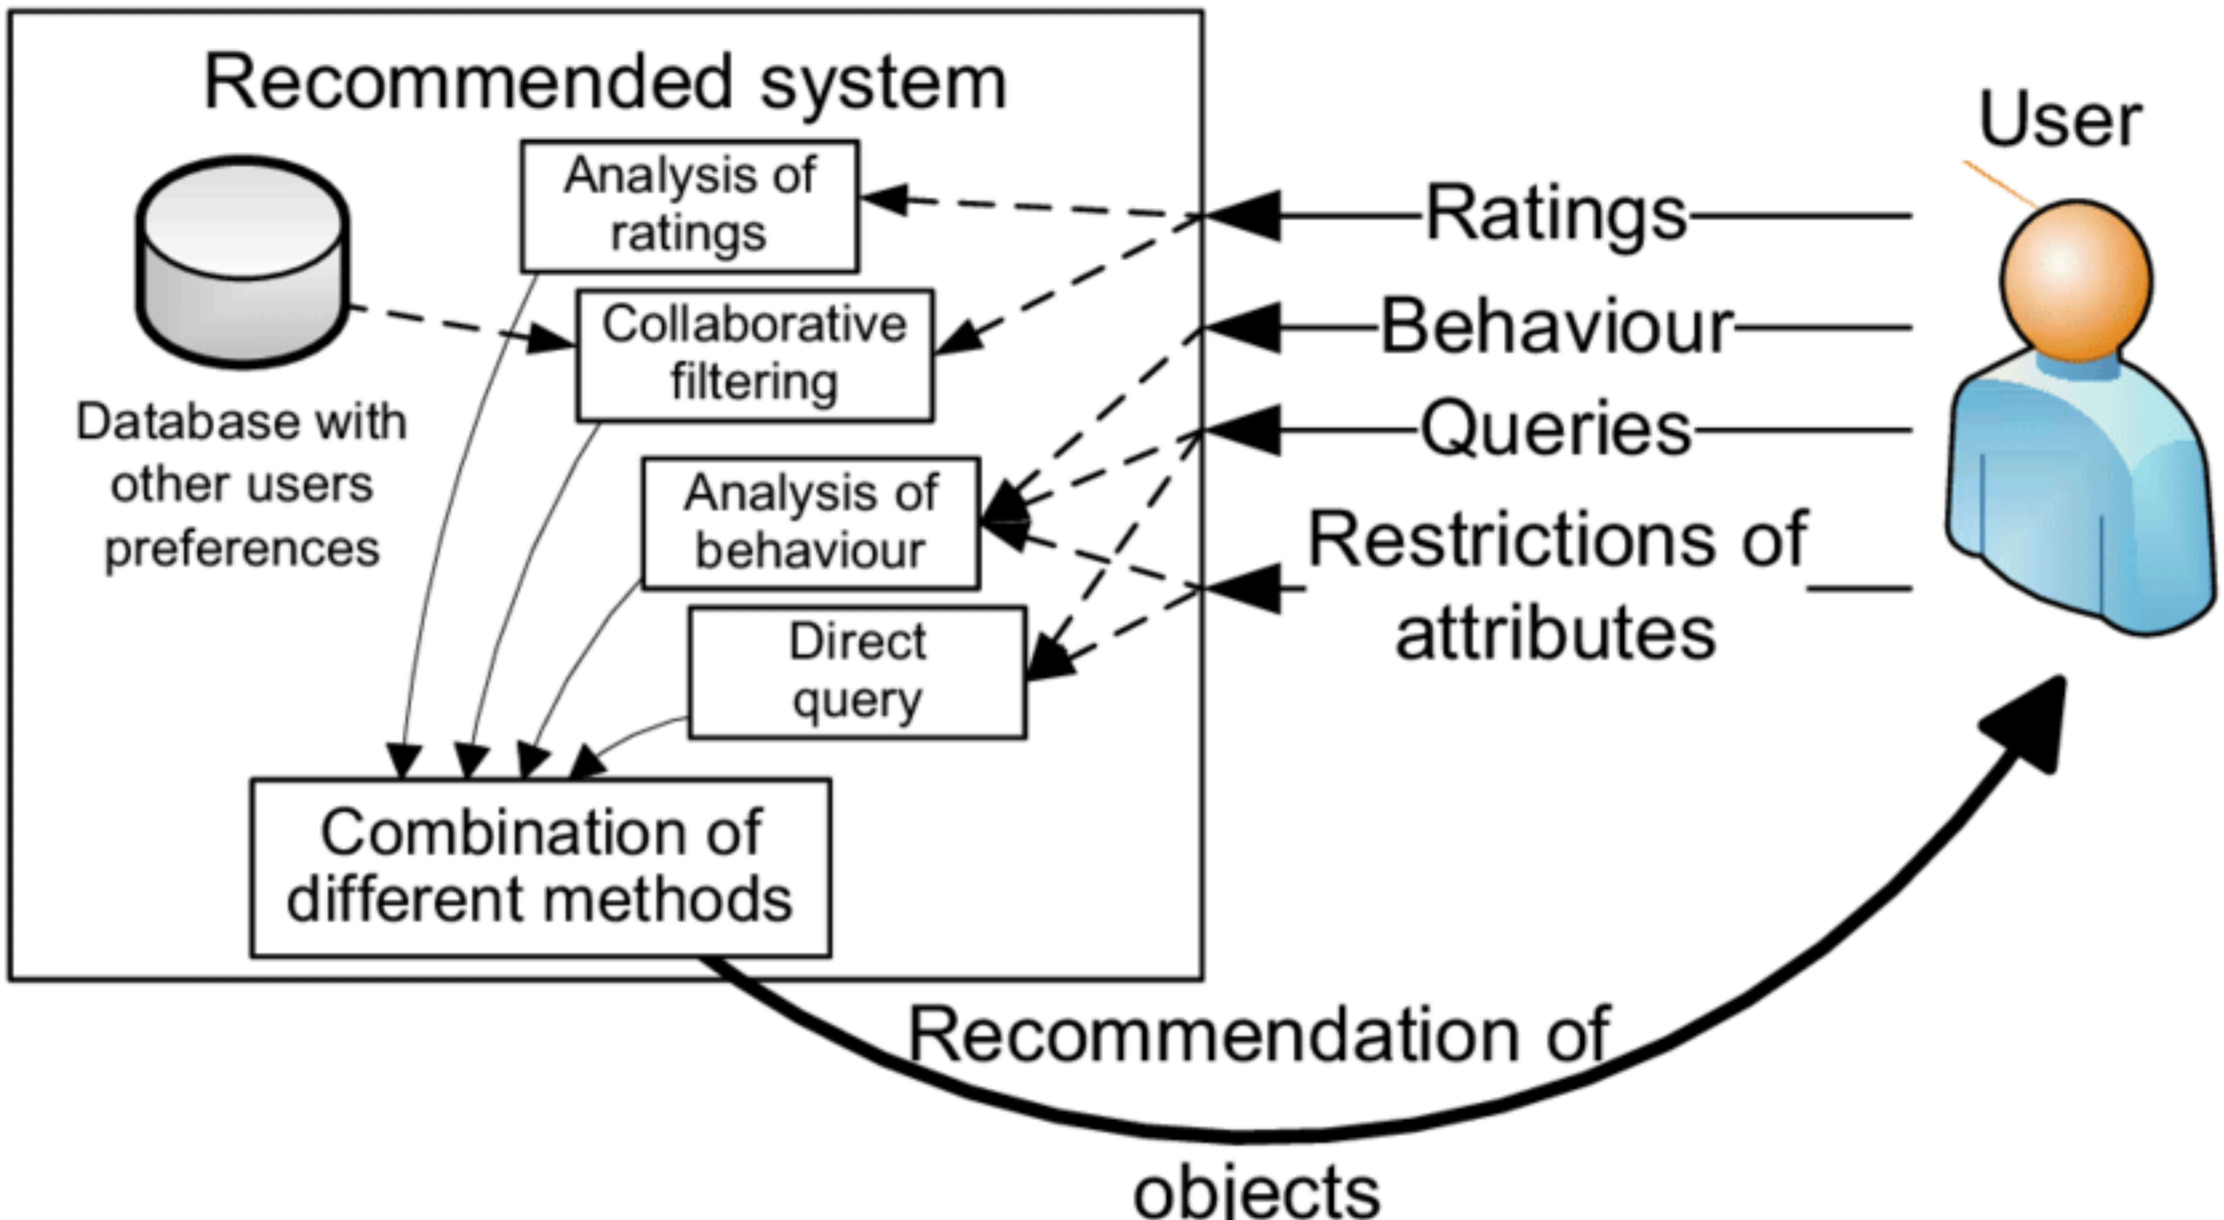

Introduction

- What is causality?
  - Classic ways to find causal information
- Introduction to ML and AI, and some connections with causality
- Causal thinking

1. Making "changes"
2. Understanding & information fusion
3. Prediction in complex environments
4. Artificial "intelligence"…

- Typical problems in causality research
  - Identification of causal effects
  - Counterfactual reasoning
  - Causal discovery & causal representation learning

Carl Jung's Forward to the English Translation of "Book of Changes"

I Ching

From Wikipedia, the free encyclopedia

For other uses, see I Ching (disambiguation).

The ***I Ching*** ([î tɕíŋ]; Chinese: 易經; pinyin: *Yìjīng*), also known as the ***Classic of Changes*** or ***Book of Changes*** in English, is an ancient divination text and the oldest of the Chinese classics. The *I Ching* was originally a divination manual in the Western Zhou period (1000–750 BC), but over the course of the Warring States period and early imperial period (500–200 BC) was transformed into a cosmological text with a series of

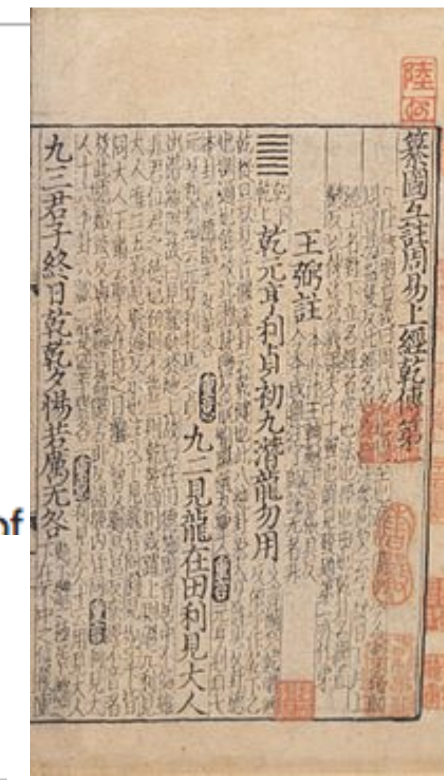

Carl Jung

From Wikipedia, the free encyclopedia

"Jung" redirects here. For other uses, see Jung (disambiguation).

Carl Gustav Jung (/jʊŋ/; German: [ˈkarl ˈɡʊstaf jʊŋ]; 26 July 1875 – 6 June 1961), often referred to as **C. G. Jung**, was a Swiss psychiatrist and psychotherapist who founded analytical psychology. His work has been influential not only in psychiatry but also in philosophy, anthropology, archaeology, literature, and religious studies. He was a prolific writer, though many of his works were not published until after his death.

Foreword to the I Ching

by Carl Gustav Jung

HTML Edition by Dan Baruth

developed what we call science. Our science, however, is based upon the principle of causality, and causality is considered to be an axiomatic truth. But a great change in ou ... *e of Pure Reason* failed to do, is being accomplished by modern physics. The axioms of causality are being shaken to their foundations: we know now that what we term natural laws are merely statistical truths and thus must necessarily allow for exceptions. We have not sufficiently taken into account as yet that we need the laboratory with its incisive restrictions in order to demonstrate the invariable validity of natural law. If we leave things to nature, we see a very different picture: every process is partially or totally interfered with by chance, so much so that under natural circumstances a course of events absolutely conforming to specific laws is almost an exception.

The Chinese mind, as I see it at work in the *I Ching*, seems to be exclusively preoccupied with the chance aspect of events. What we call coincidence seems to be the chief concern of this peculiar mind, and what we worship as causality passes almost unnoticed. We must admit that there is something to be said for the immense importance of chance. An incalculable amount of human effort is directed to combating and restricting the nuisance or danger represented by chance. Theoretical considerations of cause and effect often look pale and dusty in comparison to the practical results of chance. It is all very well to say that the crystal of quartz is a hexagonal prism. The statement is quite true in so far as an ideal crystal is envisaged. But in nature one finds no two crystals exactly alike, although all are unmistakably hexagonal. The actual form, however, seems to appeal more to the Chinese sage than the ideal one. The jumble of natural laws constituting empirical reality holds more significance for him than a causal explanation of events that, moreover, must usually be separated from one another in order to be properly dealt with.

The manner in which the *I Ching* tends to look upon reality seems to disfavor our causalistic procedures. The moment under actual observation appears to the ancient Chinese view more of a chance hit than a clearly defined result of concurring causal chain processes. The matter of interest seems to be the configuration formed by chance events in the moment of observation, and not at all the hypothetical reasons that seemingly account for the coincidence. While the Western mind carefully sifts, weighs, selects, classifies, isolates, the Chinese picture of the moment encompasses everything down to the minutest nonsensical detail, because all of the ingredients make up the observed moment.

Thus it happens that when one throws the three coins, or counts through the forty-nine yarrow stalks, these chance details enter into the picture of the moment of observation and form a part of it -- a part that is insignificant to us, yet most meaningful to the Chinese mind. With us it would be a banal and almost meaningless statement (at least on the face of it) to say that whatever happens in a given moment possesses inevitably the quality peculiar to that moment. This is not an abstract argument but a very practical one. There are certain connoisseurs who can tell you merely from the appearance, taste, and behavior of a wine the site of its vineyard and the year of its origin. There are antiquarians who with almost uncanny accuracy will name the time and place of origin and the maker of an *objet d'art* or piece of furniture on merely looking at it. And there are even astrologers who can tell you, without any previous knowledge of your nativity, what the position of sun and moon was and what zodiacal sign rose above the horizon in the moment of your birth. In the face of such facts, it must be admitted that moments can leave long-lasting traces.

In other words, whoever invented the *I Ching* was convinced that the hexagram worked out in a certain moment coincided with the latter in quality no less than in time. To him the hexagram was the exponent of the moment in which it was cast -- even more so than the hours of the clock or the divisions of the calendar could be -- inasmuch as the hexagram was understood to be an indicator of the essential situation prevailing in the moment of its origin.

This assumption involves a certain curious principle that I have termed synchronicity,[2] a concept that formulates a point of view diametrically opposed to that of causality. Since the latter is a merely statistical truth and not absolute, it is a sort of working hypothesis of how events evolve one out of another, whereas synchronicity takes the coincidence of events in space and time as meaning something more than mere chance, namely, a peculiar interdependence of objective events among themselves as well as with the subjective (psychic) states of the observer or observers.

The ancient Chinese mind contemplates the cosmos in a way comparable to that of the modern physicist, who cannot deny that his model of the world is a decidedly psychophysical structure. The microphysical event includes the observer just as much as the reality underlying the *I Ching* comprises subjective, i.e., psychic conditions in the totality of the momentary situation. Just as causality describes the sequence of events, so synchronicity to the Chinese mind deals with the coincidence of events. The causal point of view tells us a dramatic story about how *D* came into existence: it took its origin from *C*, which existed before *D*, and *C* in its turn had a father, *B*, etc. The synchronistic view on the other hand tries to produce an equally meaningful picture of coincidence. How does it happen that *A'*, *B'*, *C'*, *D'*, etc., appear all in the same moment and in the same place? It happens in the first place because the physical events *A'* and *B'* are of the same quality as the psychic events *C'* and *D'*, and further because all are the exponents of one and the same momentary situation. The situation is assumed to represent a legible or understandable picture.

Now the sixty-four hexagrams of the *I Ching* are the instrument by which the meaning of sixty-four different yet typical situations can be determined. These interpretations are equivalent to causal explanations. Causal connection is statistically necessary and can therefore be subjected to experiment. Inasmuch as situations are unique and cannot be repeated, experimenting with synchronicity seems to be impossible under ordinary

Why Causal Models? Changes!

- Infer effect of interventions:

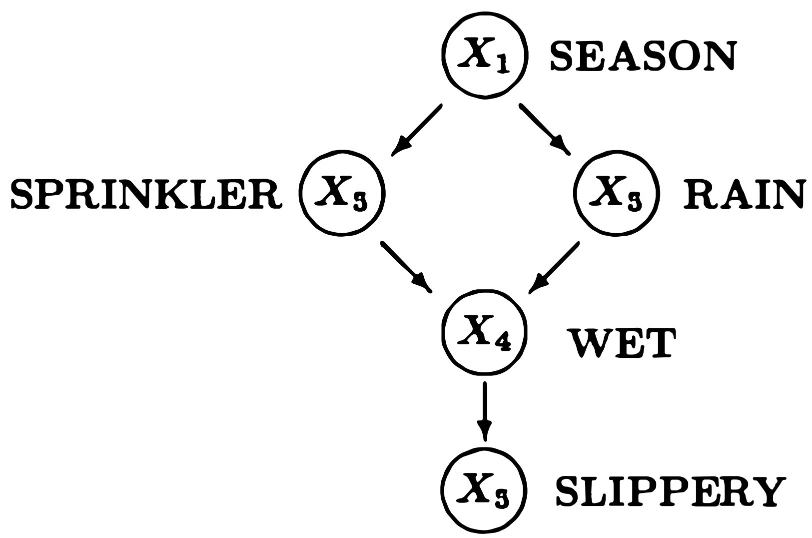

What is
$P_{X3=ON}(X_1,X_2,X_4,X_5)$?

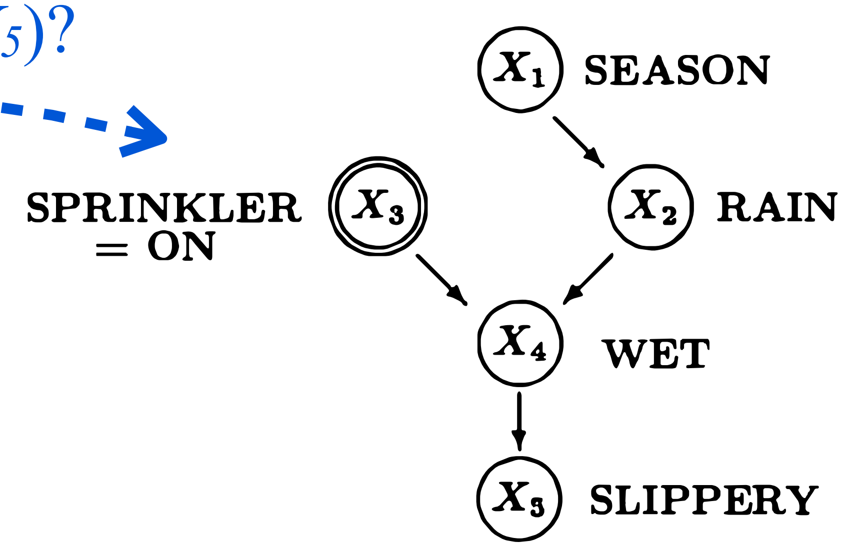

Causal Thinking: Making Changes?

- Dependence vs. causality

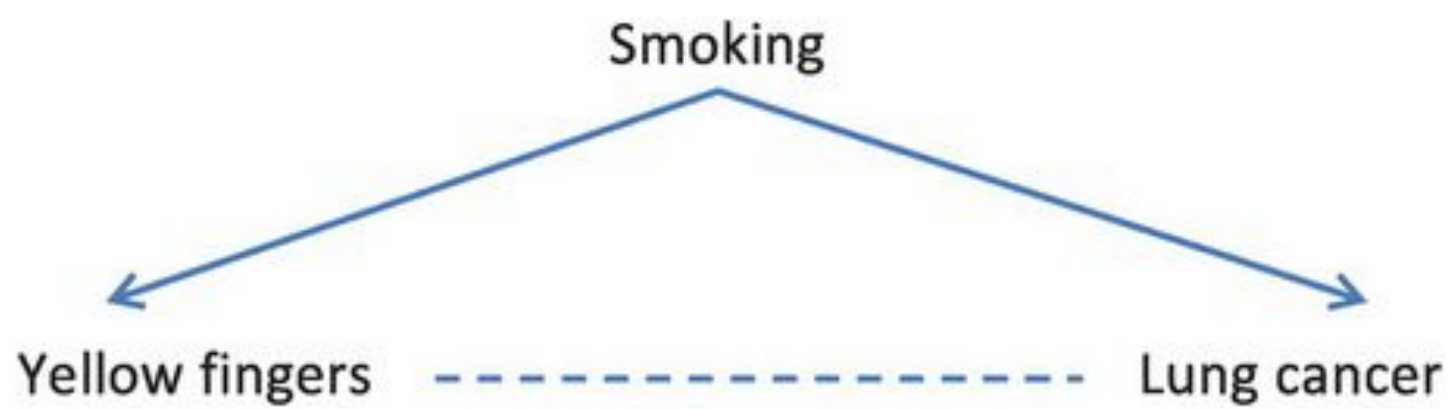

Causal Thinking: Why "Paradox"?

- Dependence vs. causality

- Simpson's paradox

	Treatment A	**Treatment B**
Small Stones	*Group 1* **93% (81/87)**	*Group 2* 87% (234/270)
Large Stones	*Group 3* **73% (192/263)**	*Group 4* 69% (55/80)
Both	78% (273/350)	**83% (289/350)**

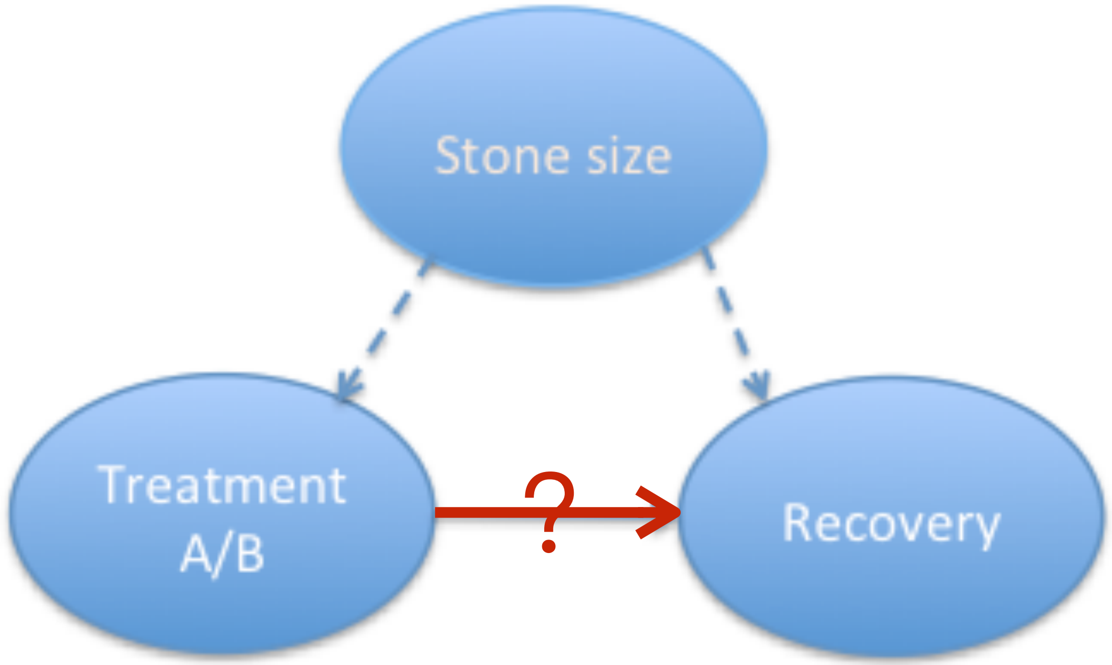

Causal Thinking: Why “Paradox”?

- Dependence vs. causality

- Simpson’s paradox

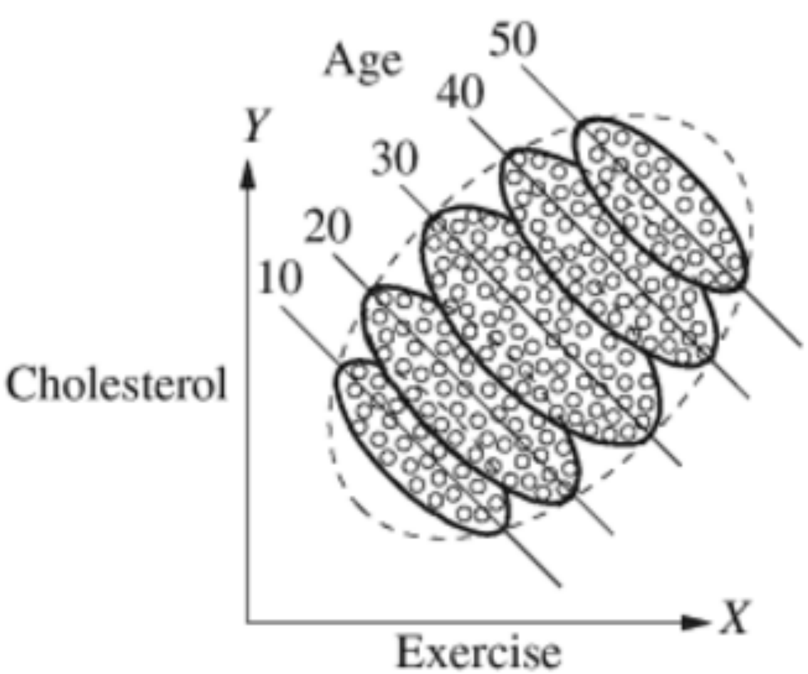

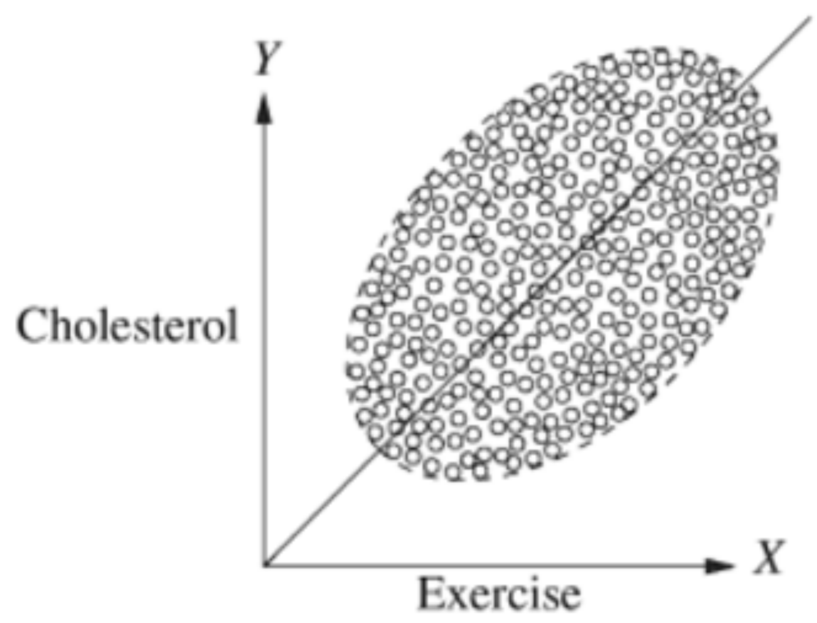

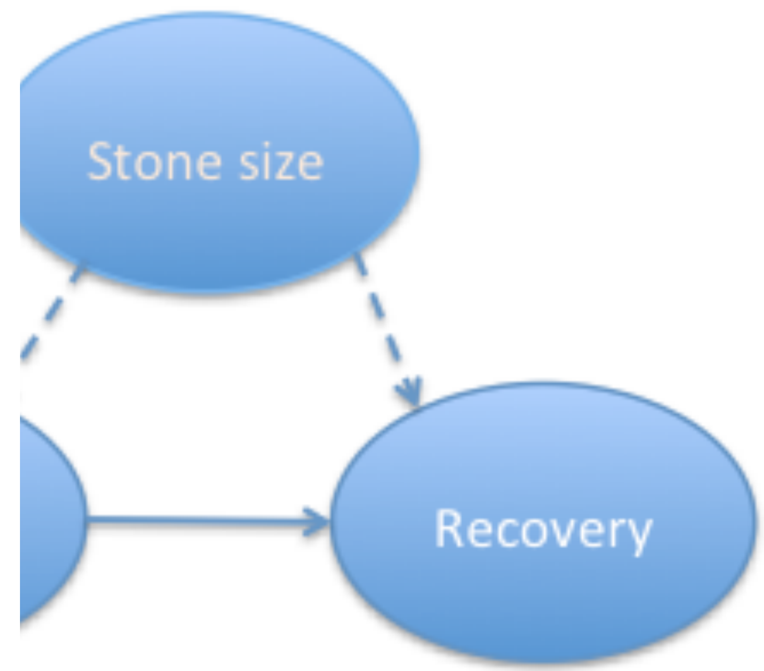

Causal Thinking: Sample vs. Population

- Dependence vs. causality
- Simpson's paradox
- "Strange" dependence
  - Go back 50 years; female college students were smarter than male ones on average. Why?

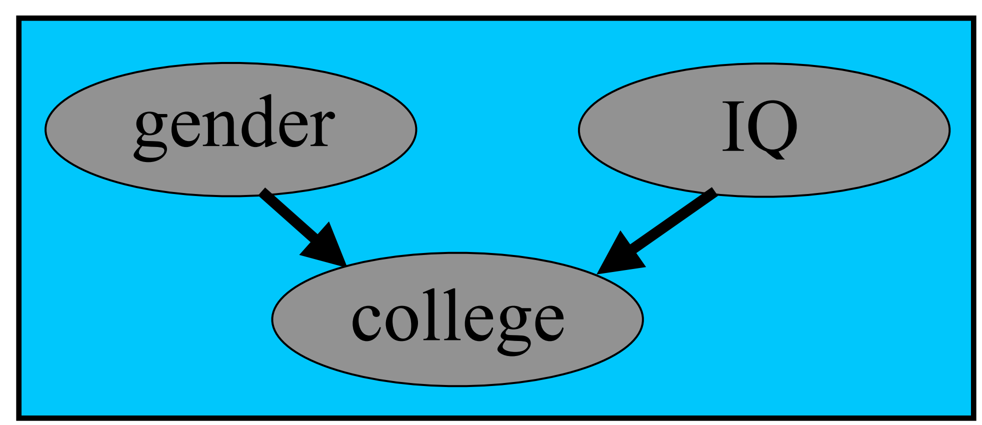

Causal Thinking: Sample vs. Population

- Dependence vs. causality
- Simpson's paradox
- "Strange" dependence

Causal Thinking: Sample vs. Population

- Dependence vs. causality

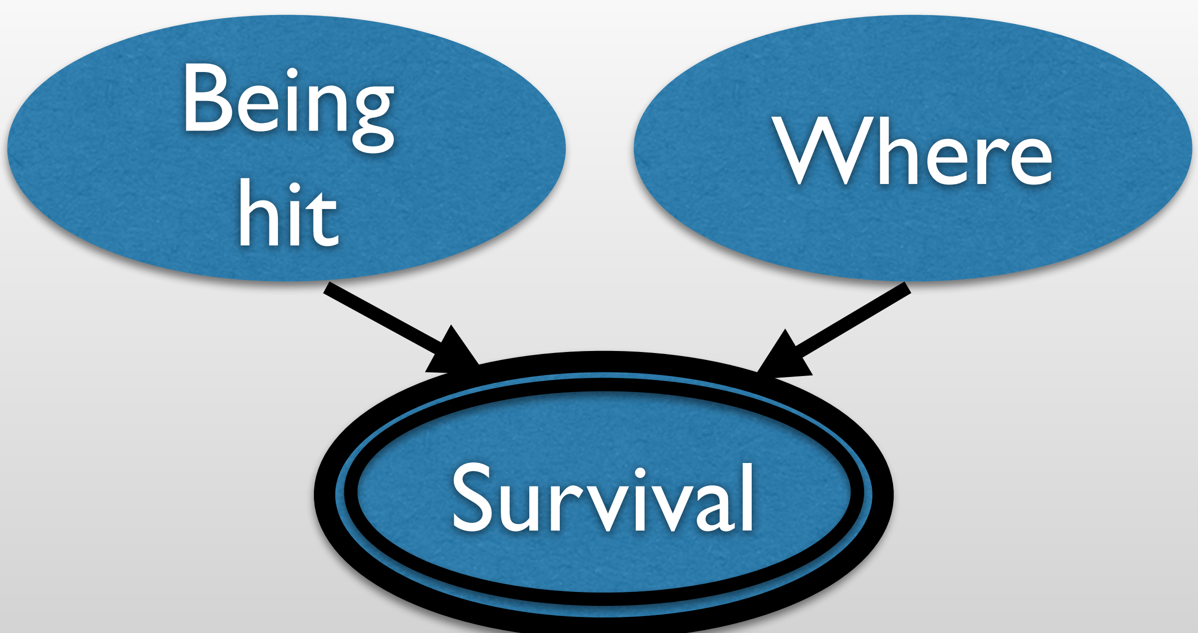

- Observe P(B,W | S=1)
- Infer P(B,W | S=0)

With the Causal Story…

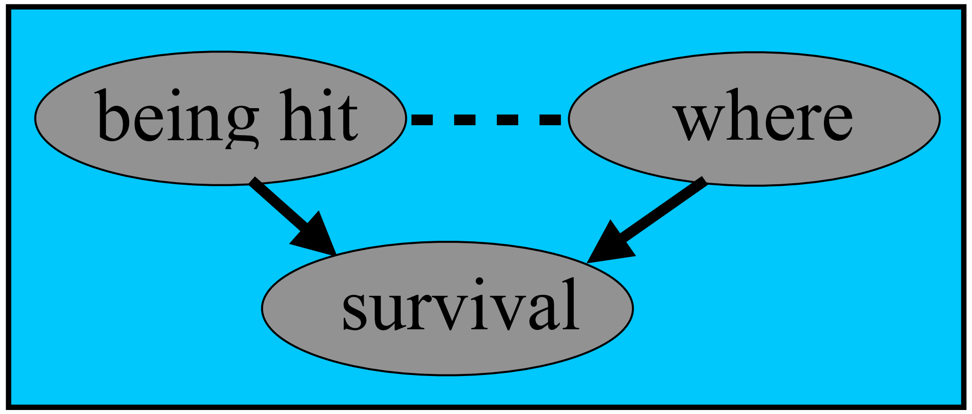

- We know P(where | being hit = 1, survival=1), and aim to find P(where | being hit = 1, survival=0)
- $P(W \mid H = 1) = P(W \mid H = 1, S = 1)P(S=1 \mid H=1) + P(W \mid H = 1, S = 0)P(S=0 \mid H=1)$

Addressing These AI Problems

- Generalization/adaptation/robustness, decision making, recommendations, fairness, generative AI…

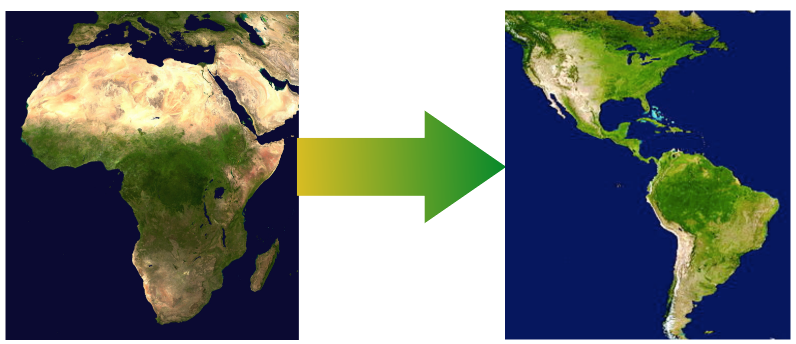

- Dealing with adversarial attacks?

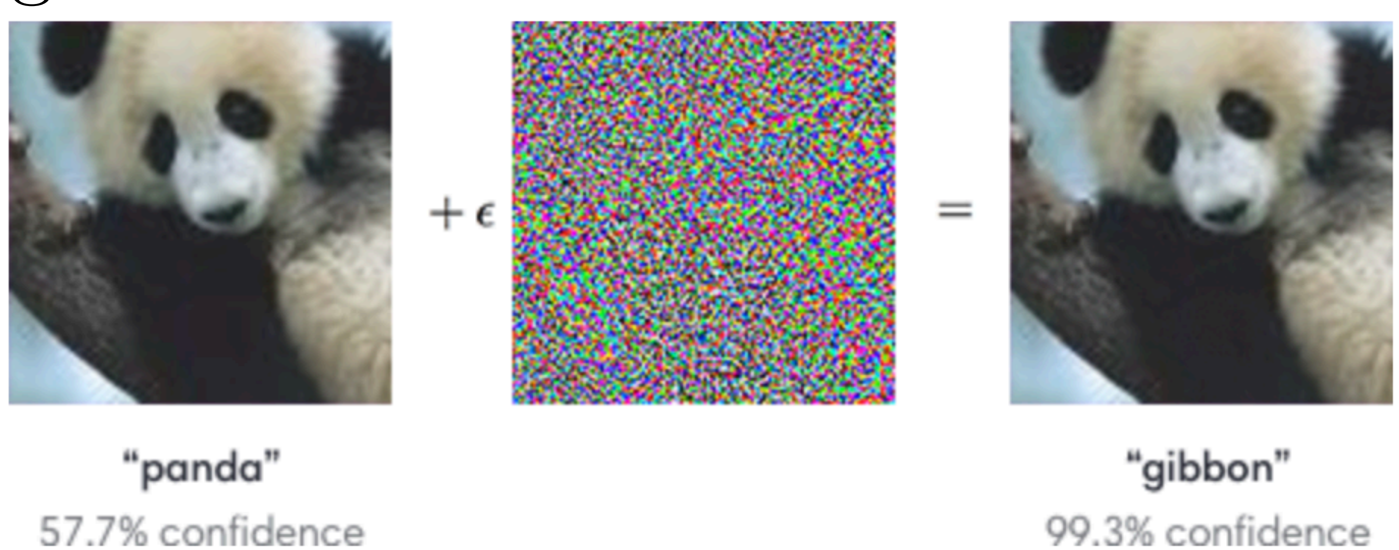

(Goodfellow et al., 2014)

An adversarial input, overlaid on a typical image, can cause a classifier to miscategorize a panda as a gibbon.

Let's Look at Stable Diffusion for Image Generation

- Prompt 1: a human with the eyes of a dragonfly

Let's Look at Stable Diffusion for Image Generation

- Prompt 1: a human with the eyes of a dragonfly

Let's Look at Stable Diffusion for Image Generation

- Prompt 2: a peacock eating ice cream

Let's Look at Stable Diffusion for Image Generation

- Prompt 2: a peacock eating ice cream

Unsupervised Image-to-Image Translation

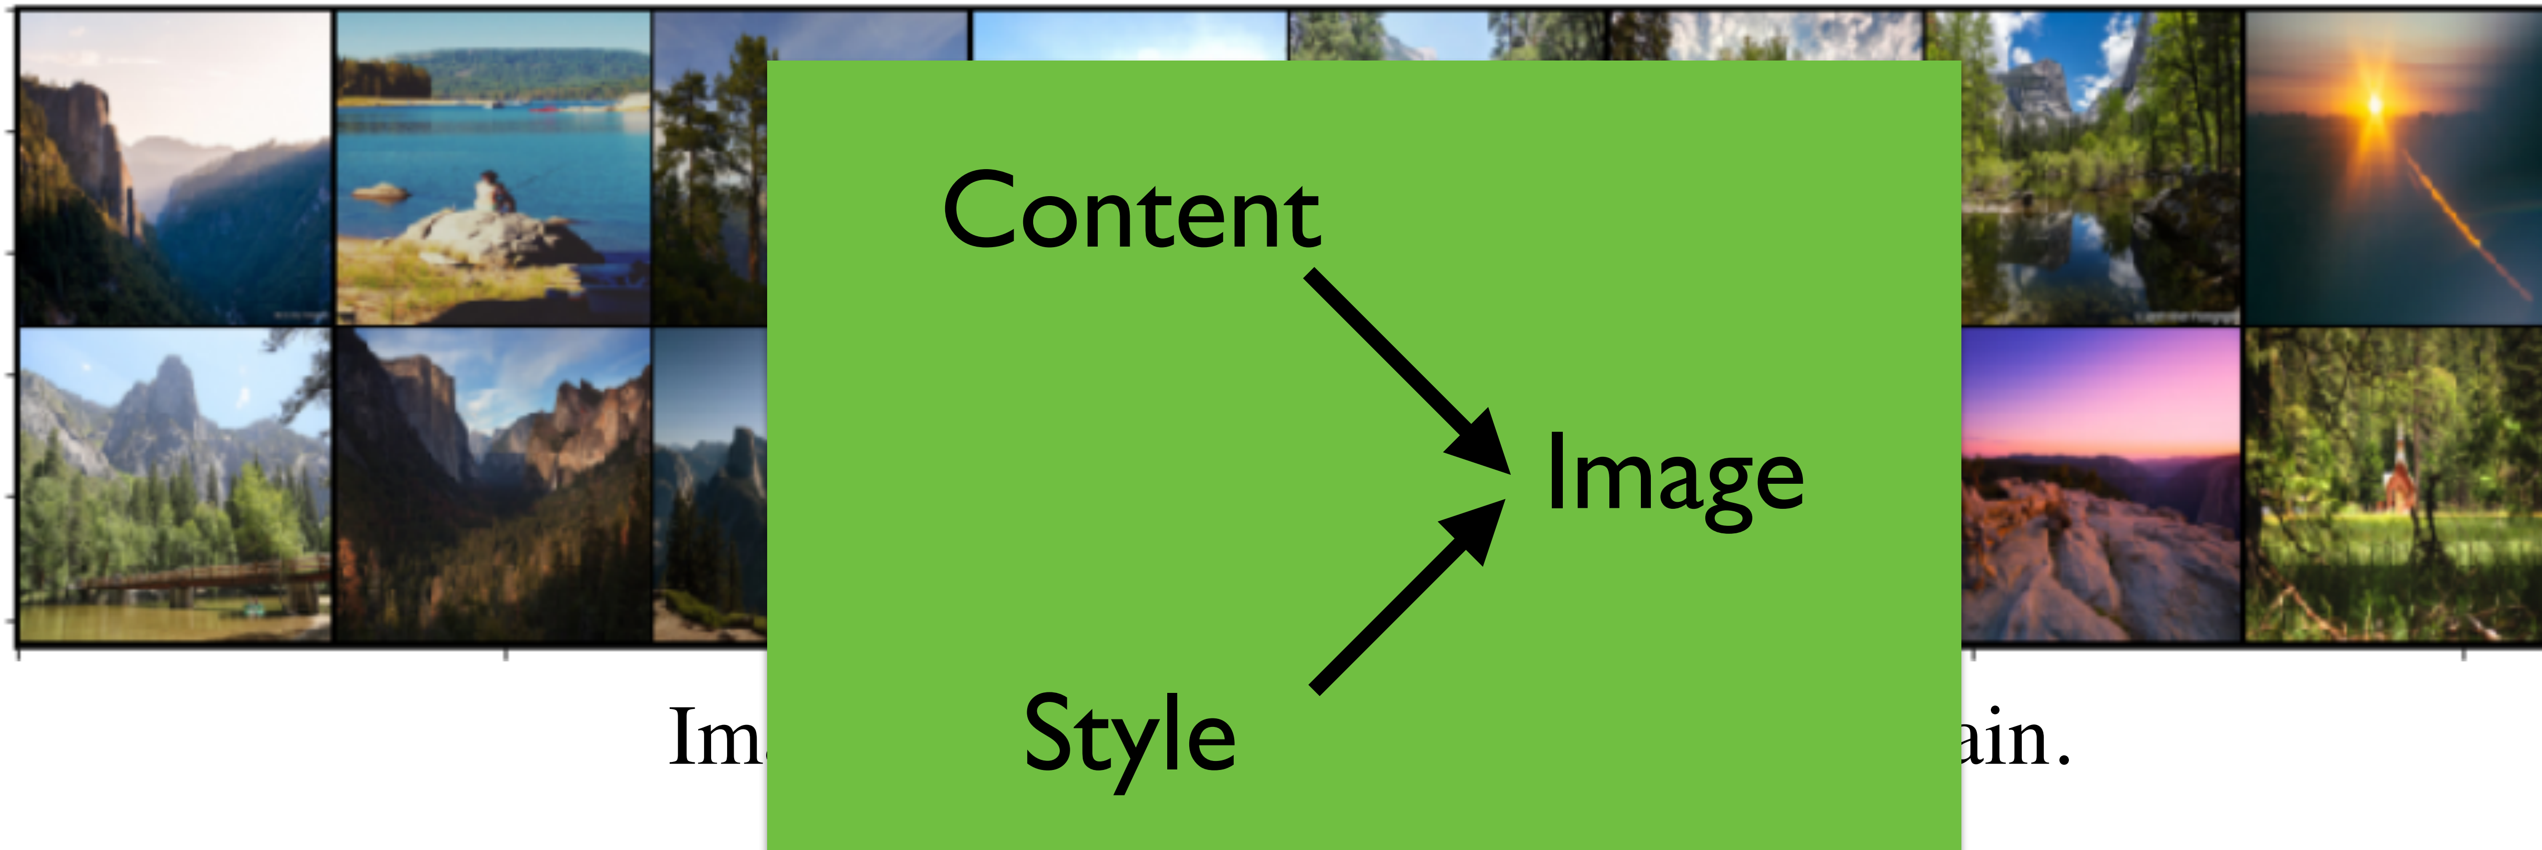

*Minimize the **influence** of 'Style' on 'Image' during translation.*

*How? A **minimal number** of changing components?*

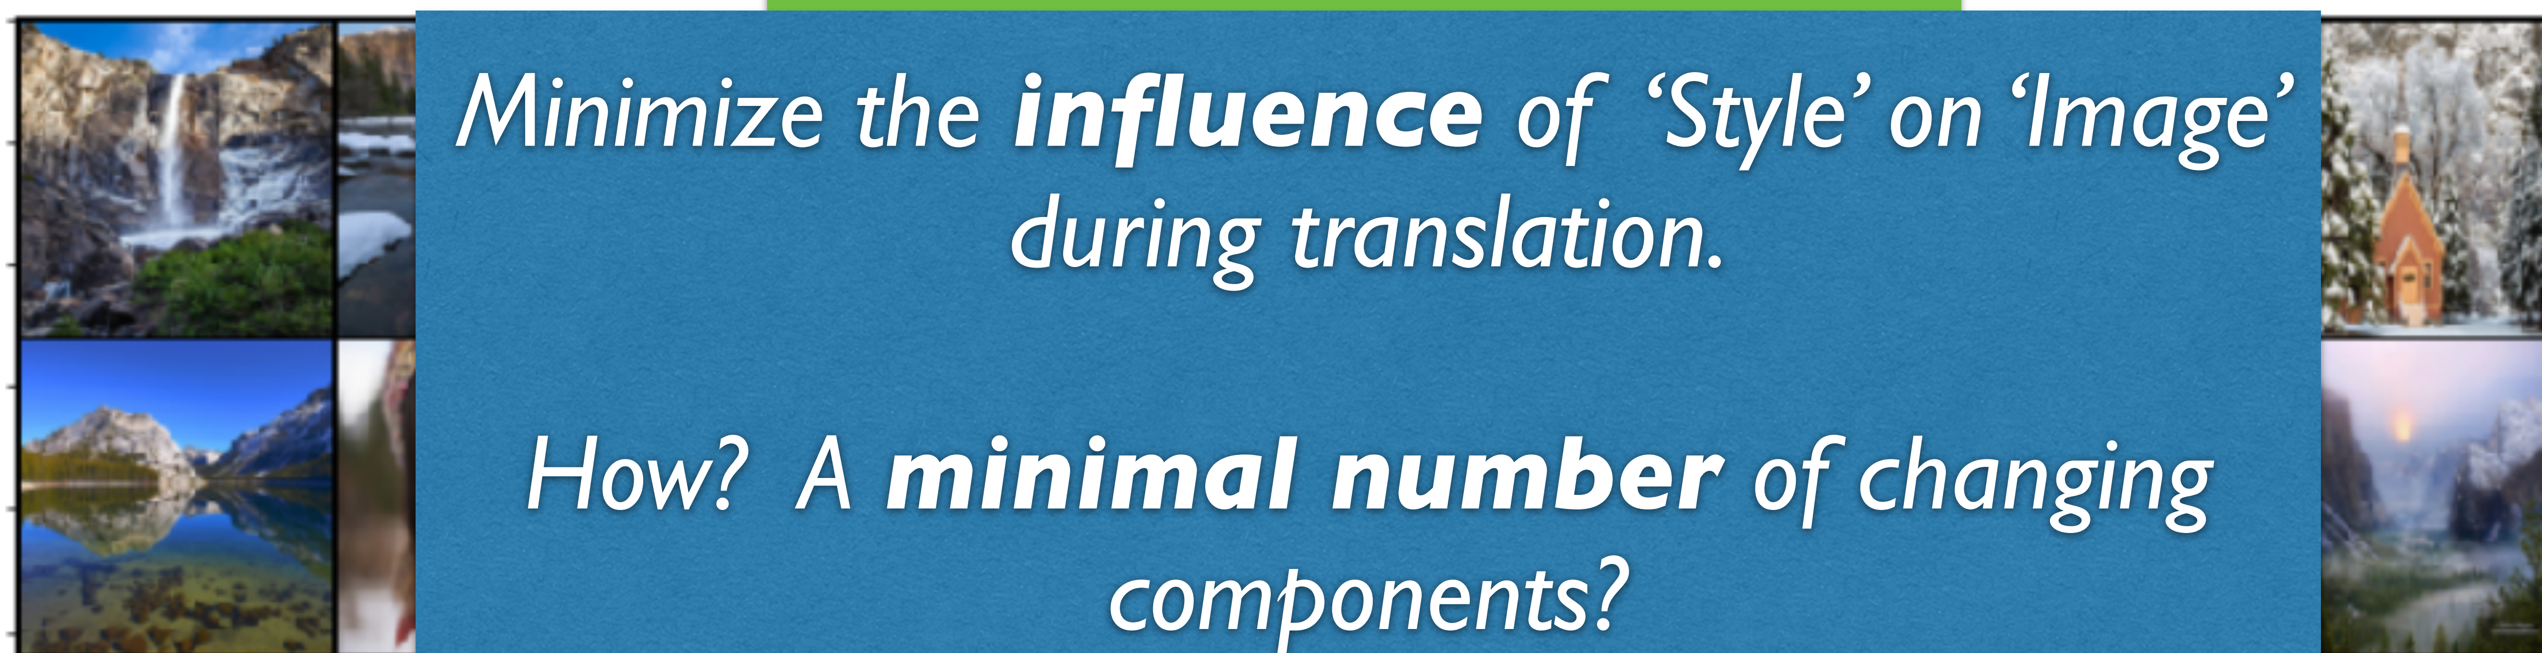

Images from the winter season domain.

Causality may Matter in Prediction: An Illustration

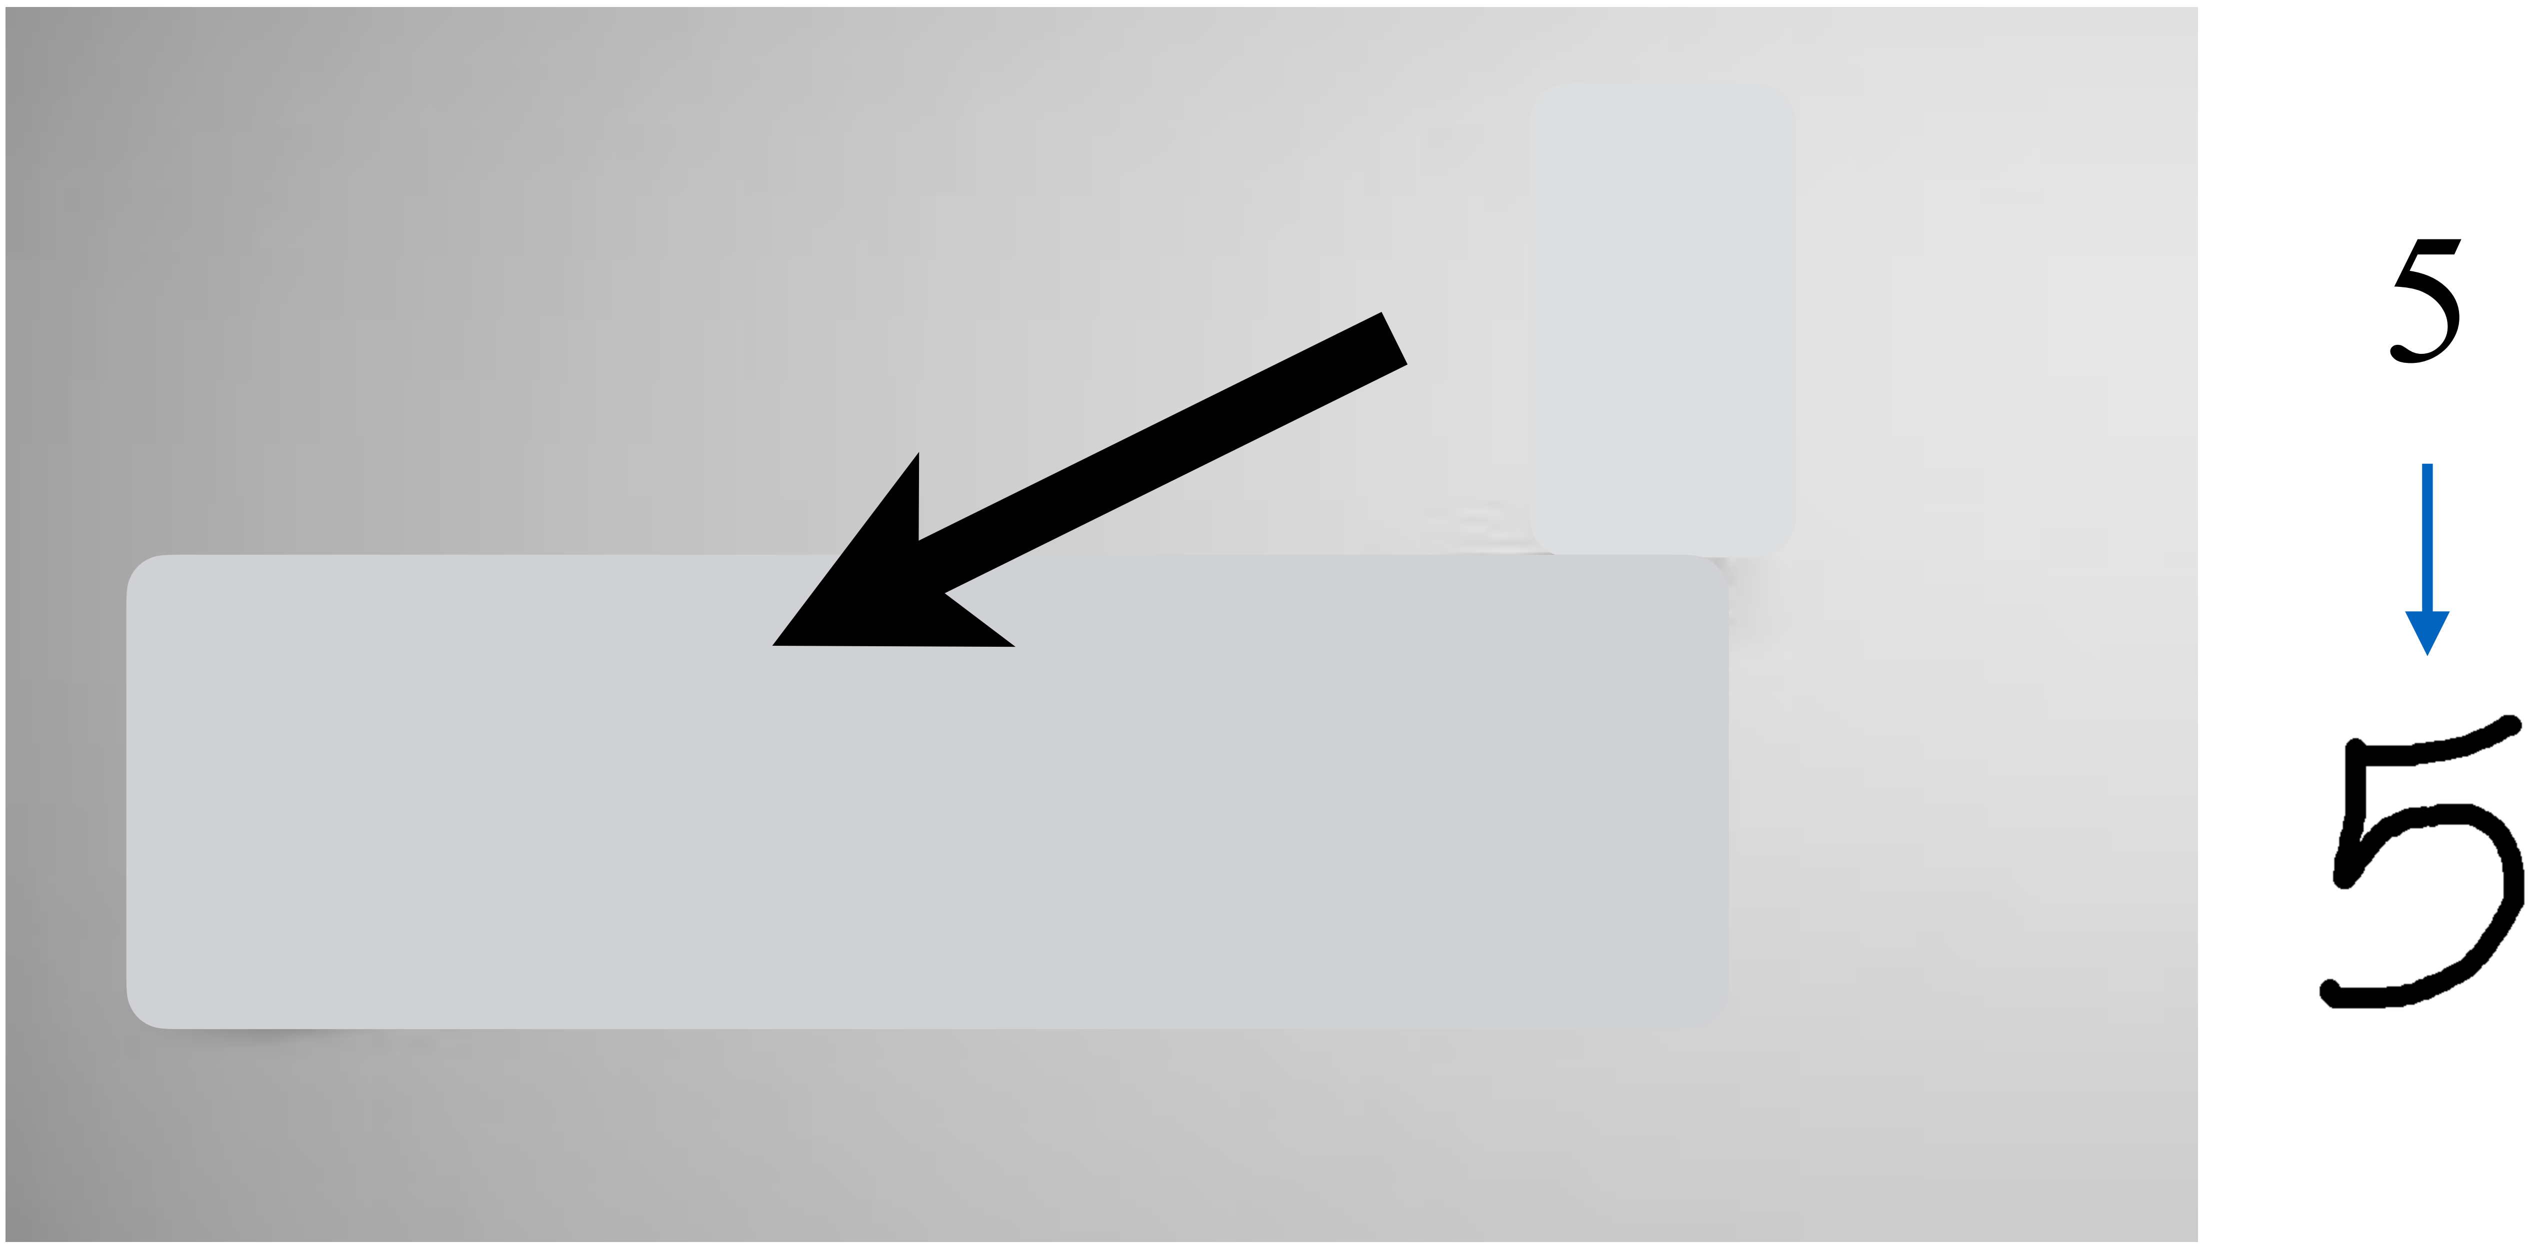

Understanding connections between different scenarios & *modeling* differences

Causal Thinking Makes a Difference

- Active manipulation /control vs. passive prediction
- Generalization / adaptation ability in new environments?
- Integration of causal information: what is the causal model for X, Y, and Z if
  - $X \rightarrow Y$, $Y \rightarrow Z$ (expansion) or $X \rightarrow Z$, $Y \rightarrow Z$ (refinement)...
- Creativity
  - Thoughts consist of the "What if?" and the "If I had only..." + knowledge integration + ...

Artificial “Intelligence”

- Traditional machine learning usually assumes a fixed data distribution; avoids overfitting

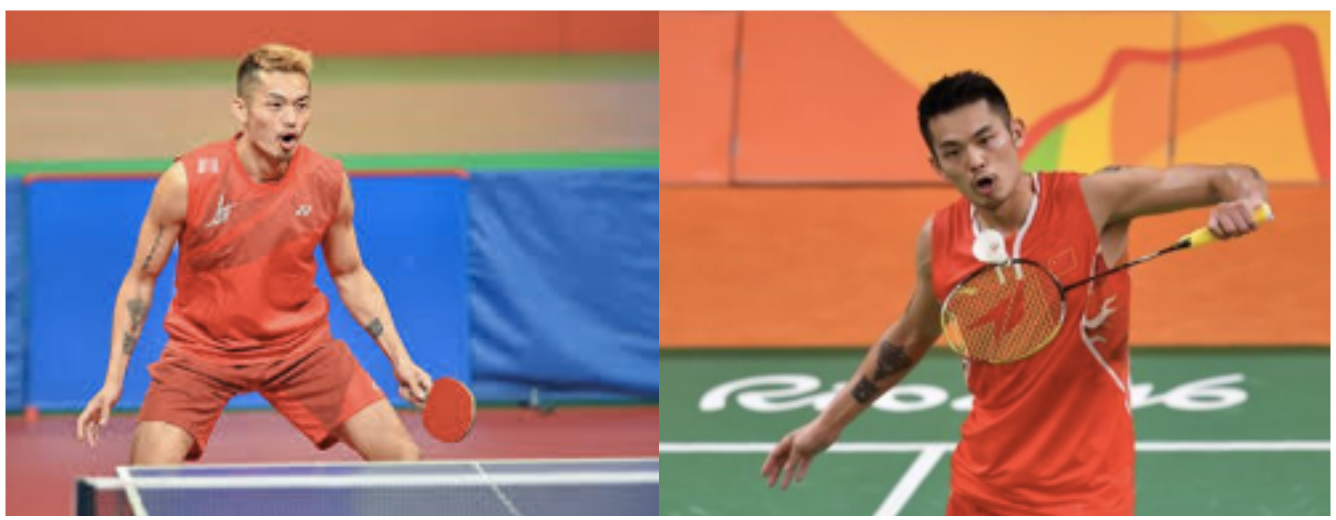

- Intelligence: understanding; control/intervention; decomposability; information fusion, learning with few examples, extrapolation

Introduction

- What is causality?
  - Classic ways to find causal information
- Introduction to ML and AI, and some connections with causality
- Causal thinking

1. Making "changes"
2. Understanding & information fusion
3. Prediction in complex environments
4. Artificial "intelligence"…

- Typical problems in causality research
  - Identification of causal effects
  - Counterfactual reasoning
  - Causal discovery & causal representation learning

Three Types of Problems in current AI

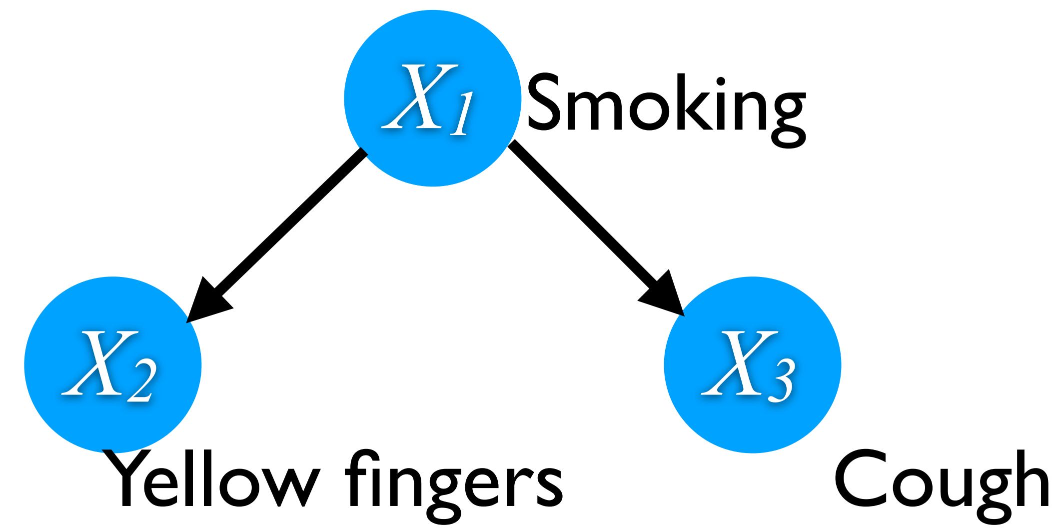

- Three questions:

X_1	X_2	X_3
1	0	0
0	0	1
0	1	1
1	1	1
0	0	0
0	1	0
1	1	1
1	1	1
0	0	0
1	0	0
…	…	…

- **Prediction**: Would the person cough if we *find* he/she has yellow fingers?

$$P(X3 \mid X2=1)$$

- **Intervention**: Would the person cough if we *make sure* that he/she has yellow fingers?

$$P(X3 \mid \mathrm{do}\,(X2=1))$$

- **Counterfactual**: Would George cough *had* he had yellow fingers, *given that he does not have yellow fingers and coughs*?

$$P(X3_{X2=1} \mid X2 = 0, X3 = 1)$$

Identification of Causal Effects

Effect of *do(Treatment=A or B)* on *Recovery*

- "Golden standard": randomized controlled experiments
  - **All the other factors** that influence the outcome variable are either fixed or vary at random, so any changes in the outcome variable must be due to the controlled variable

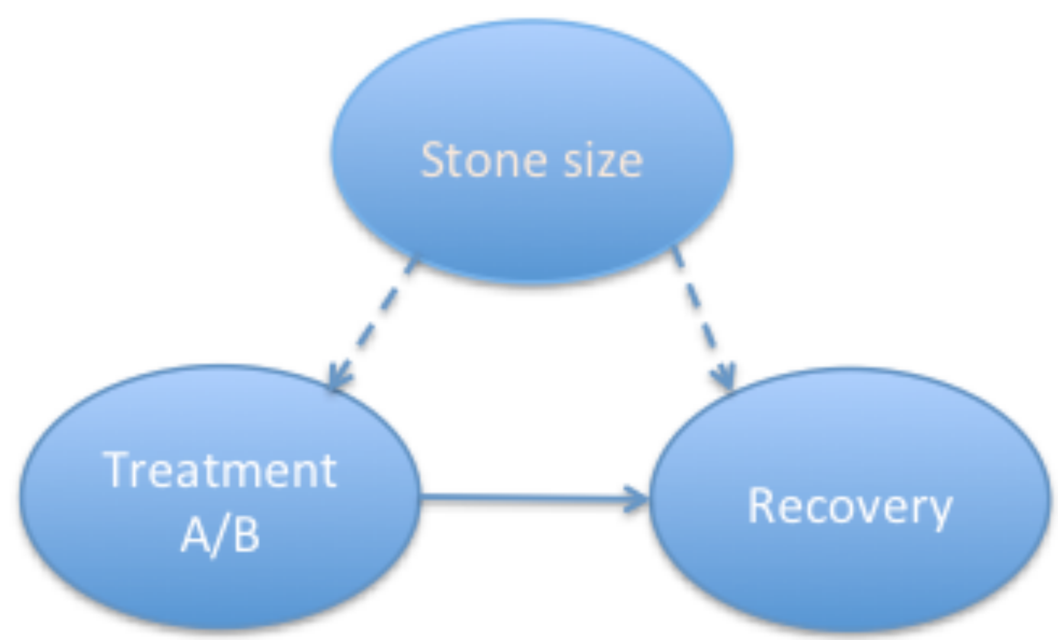

  - Usually expensive or impossible to do!

Identification of Causal Effects: Example

	Treatment A	Treatment B
Small Stones	*Group 1* **93% (81/87)**	*Group 2* 87% (234/270)
Large Stones	*Group 3* **73% (192/263)**	*Group 4* 69% (55/80)
Both	78% (273/350)	**83% (289/350)**

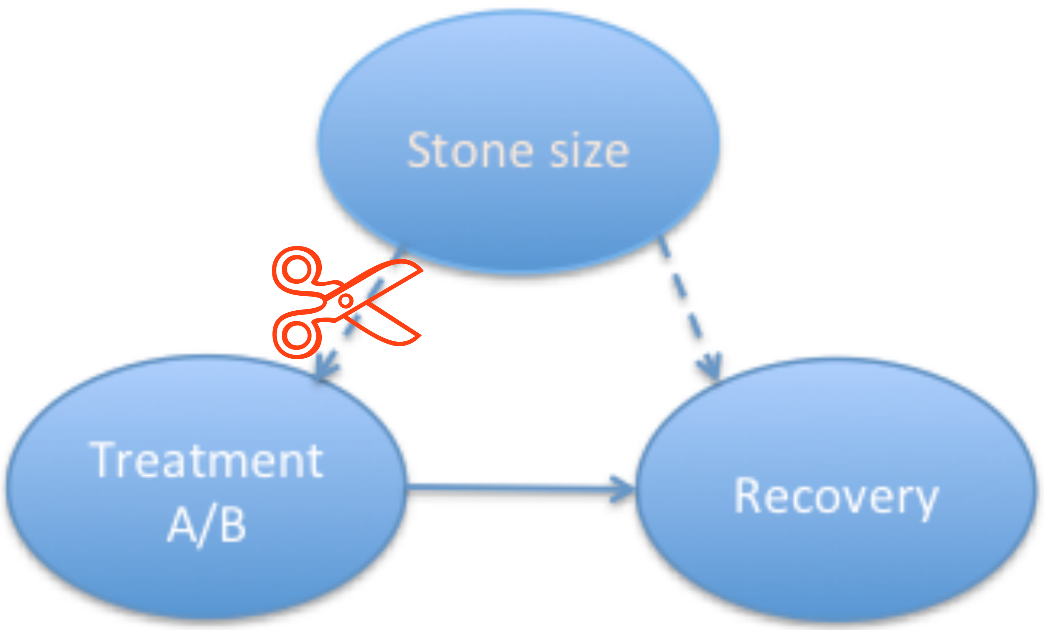

In the future you will learn:

$$P(R|T) = \sum_S P(R|T,S)P(S|T)$$

$$P(R|do(T)\,) = \sum_S P(R|T,S)P(S)$$

conditioning vs. manipulating

Identification of Causal Effects: Example

	Treatment A	Treatment B
Small Stones	*Group 1* **93% (81/87)**	*Group 2* 87% (234/270)
Large Stones	*Group 3* **73% (192/263)**	*Group 4* 69% (55/80)
Both	78% (273/350)	**83% (289/350)**

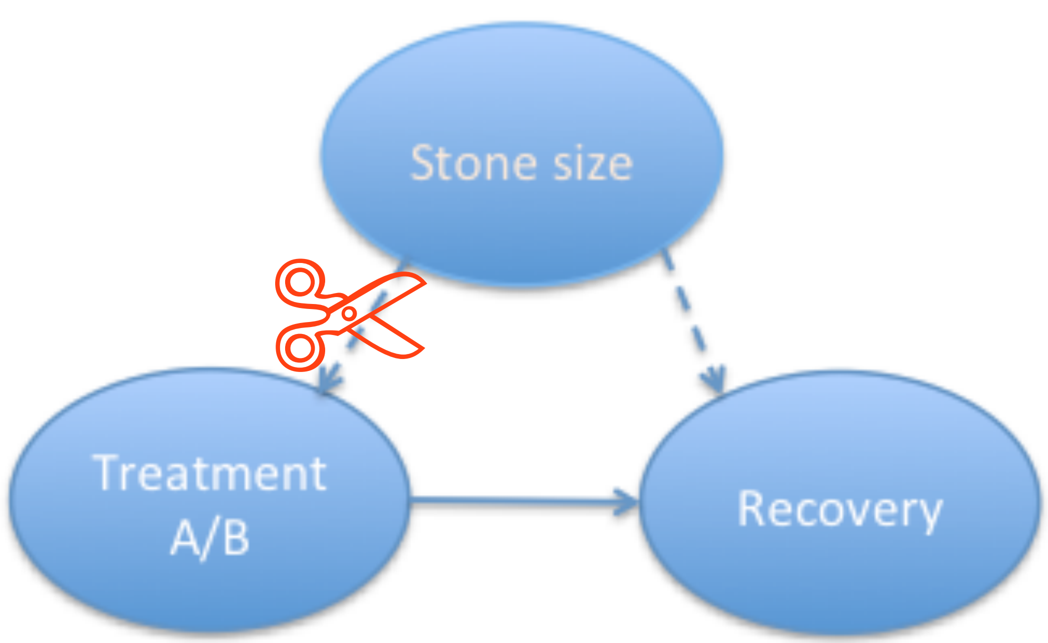

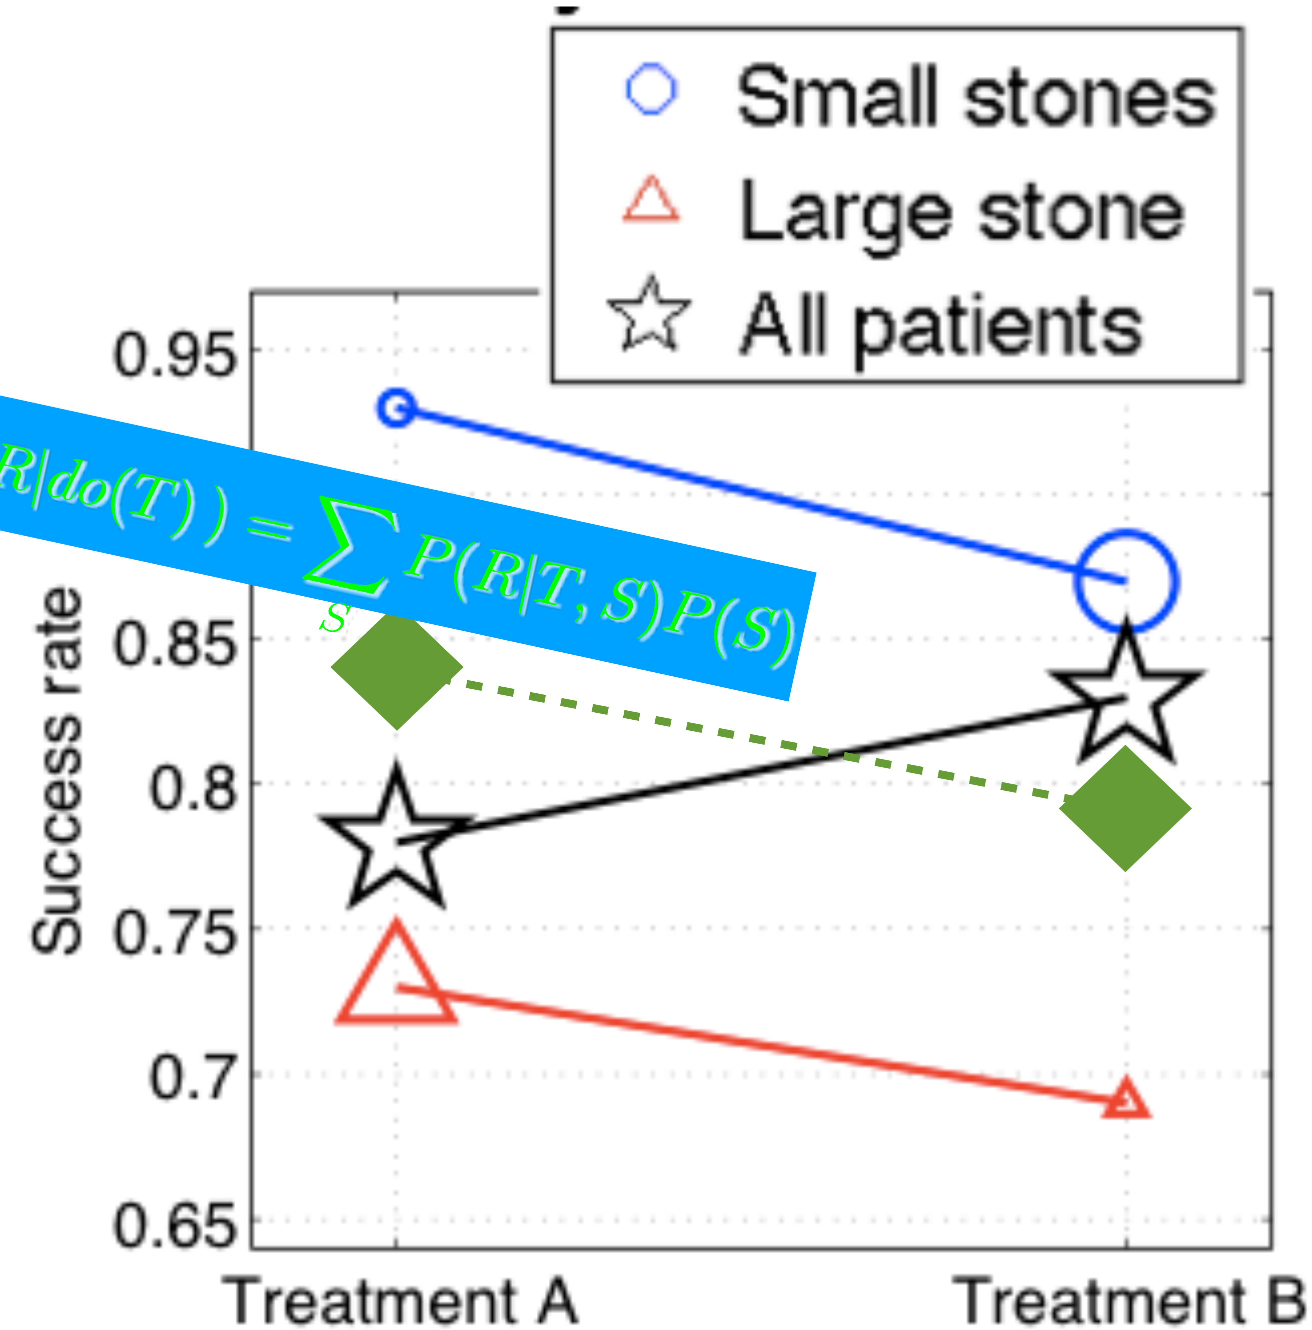

$$P(R|do(T)) = \sum_S P(R|T,S)P(S)$$

Prediction vs. causal effect

Counterfactual Inference vs. Prediction

- Suppose X→Y (*attendance* → *grade*) with Y = log(X + U + 3). For an individual with (x,y), what would Y be if X had been x' ?

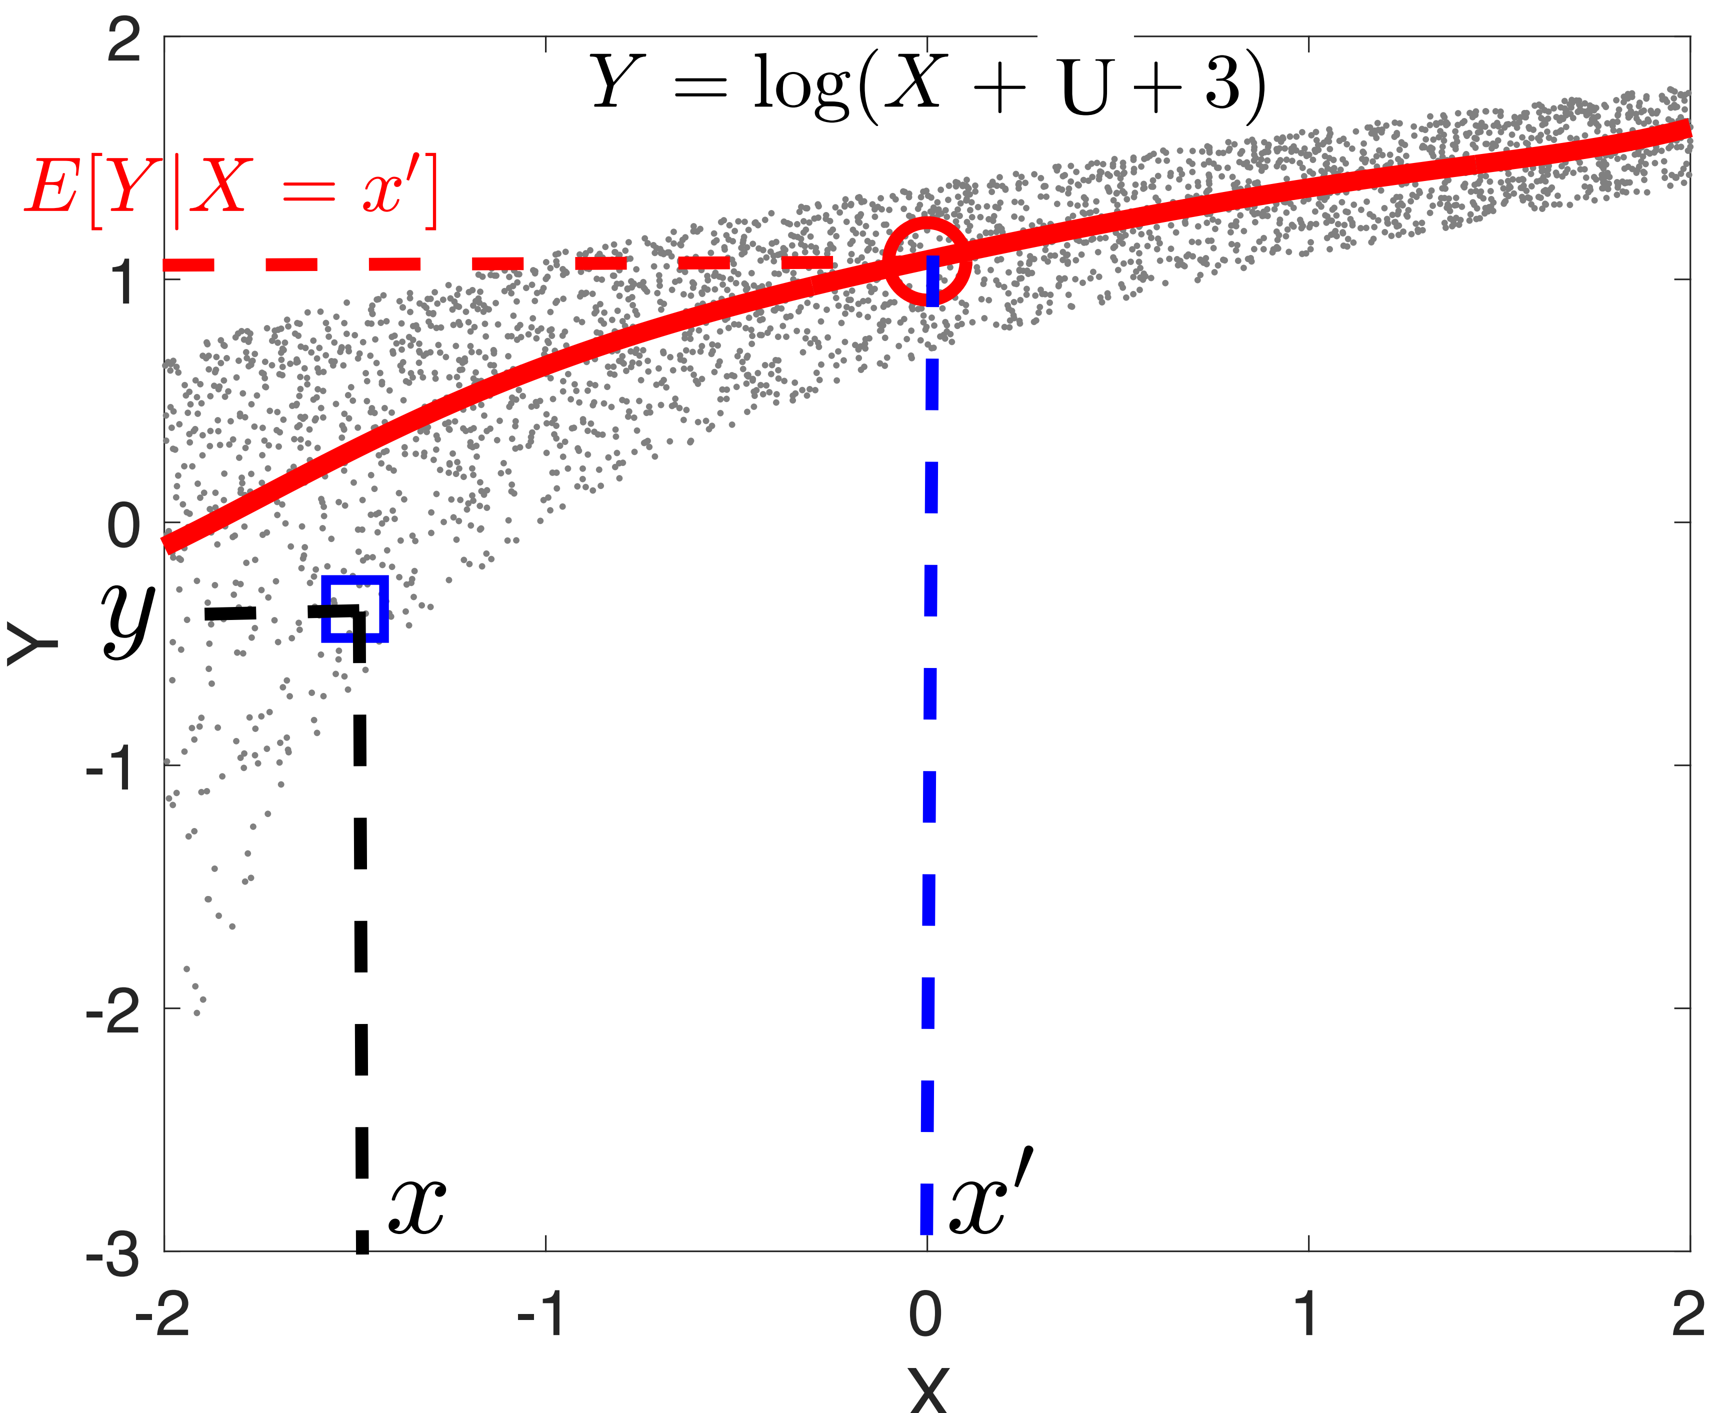

Counterfactual Inference vs. Prediction

- Suppose X→Y with Y = log(X + U + 3). For an individual with (*x*,*y*), what would Y be if X had been *x'* ?

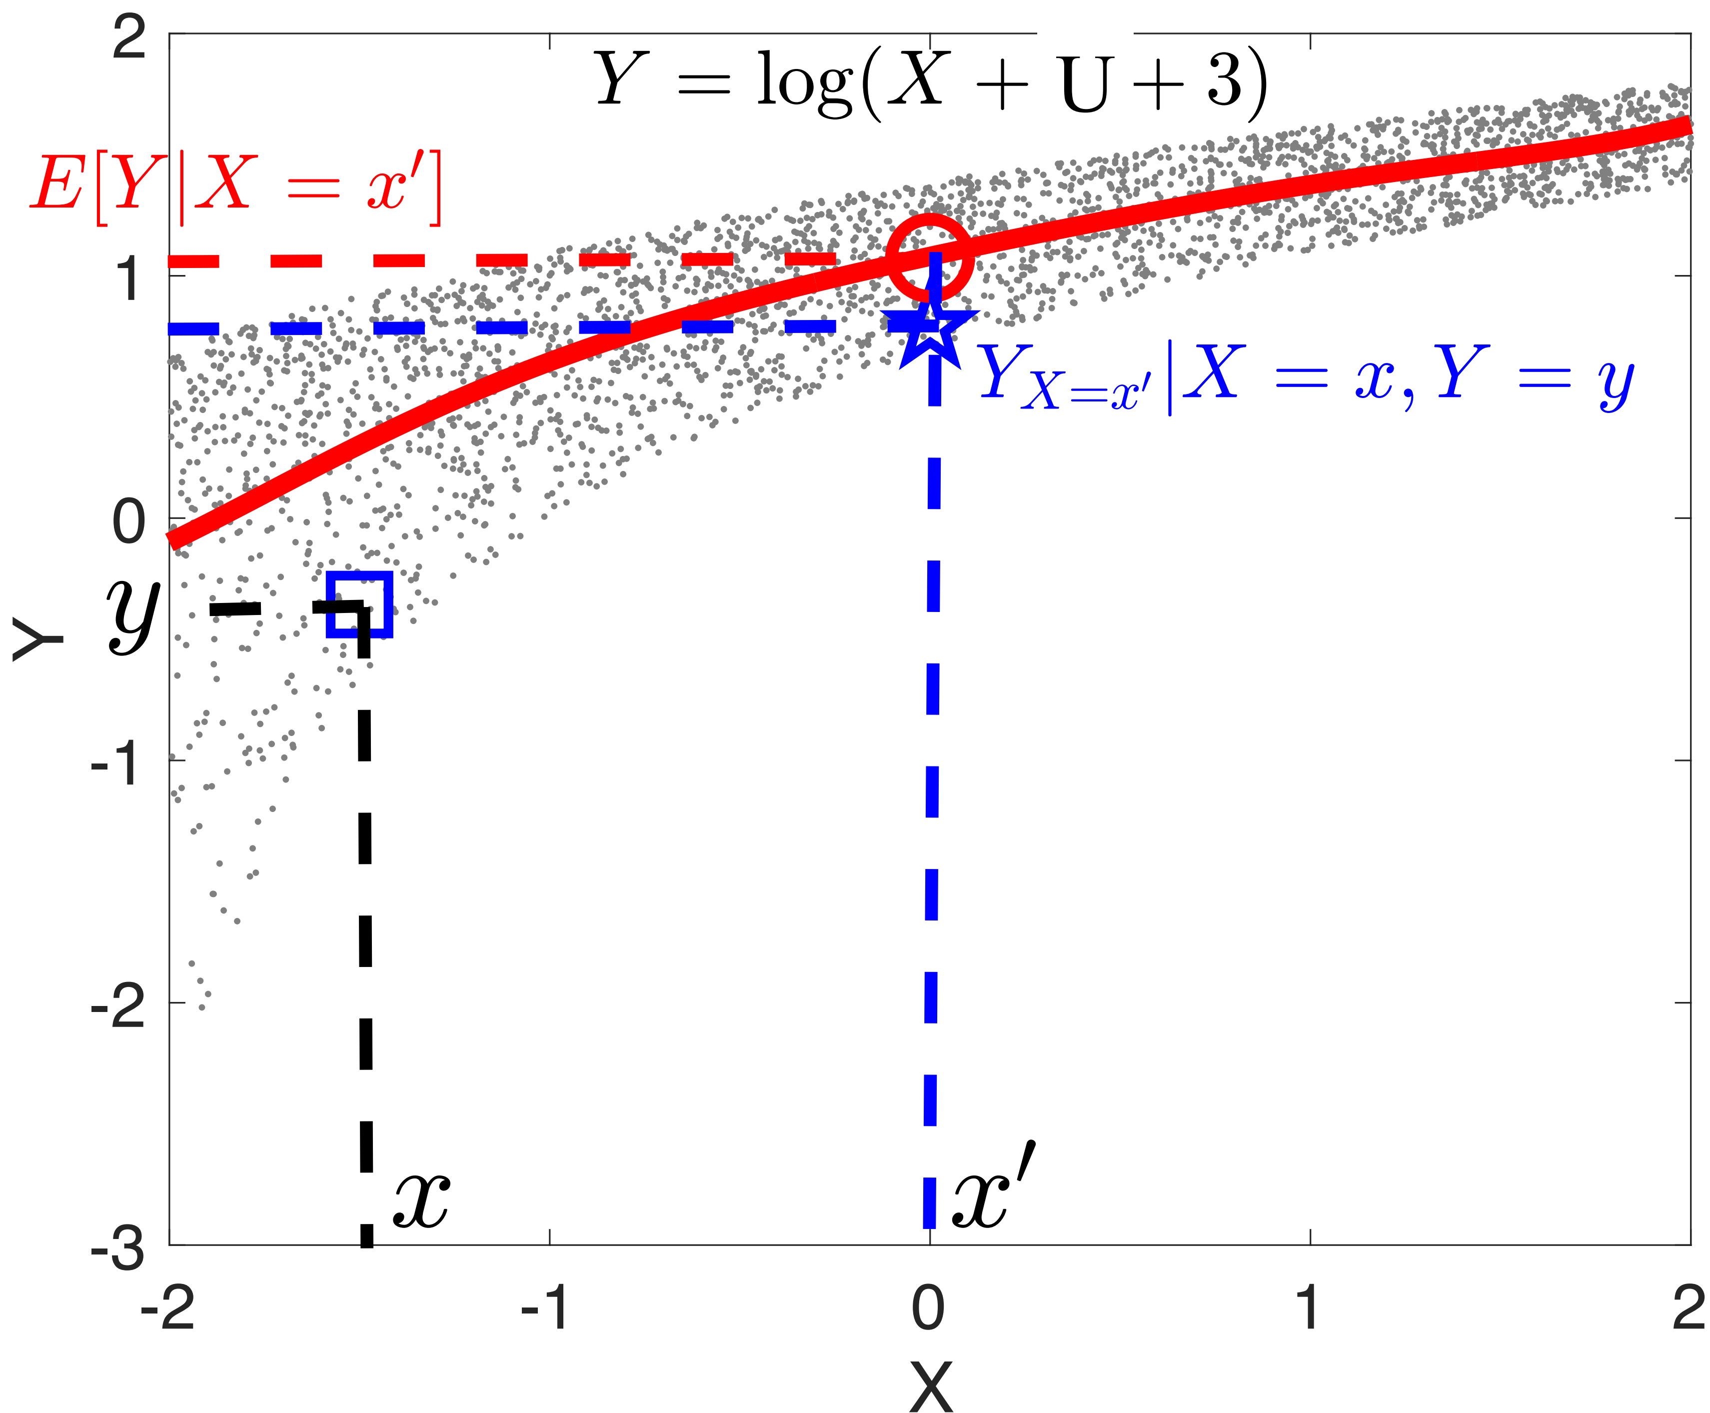

Causal Discovery

Possible to

discover causal information (*specific properties of the true process*)

from purely observational data ?

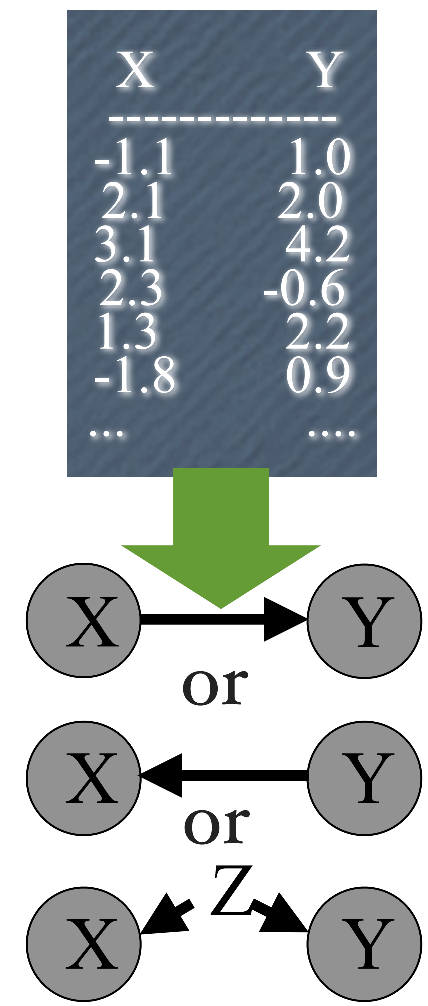

Can we go beyond the data?

Causality Examples

The Telegraph

Home | News | World | Sport | Finance | Comment | Culture | Travel | Life | Women | Fashion

USA | Asia | China | Europe | Middle East | Australasia | Africa | South America | Central Asia

France | Francois Hollande | Germany | Angela Merkel | Russia | Vladimir Putin | Greece

HOME » NEWS » WORLD NEWS » EUROPE

Couples who share the housework are more likely to divorce, study finds

Divorce rates are far higher among "modern" couples who share the housework than in those where the woman does the lion's share of the chores, a Norwegian study has found.

Causality Examples

Causality Examples

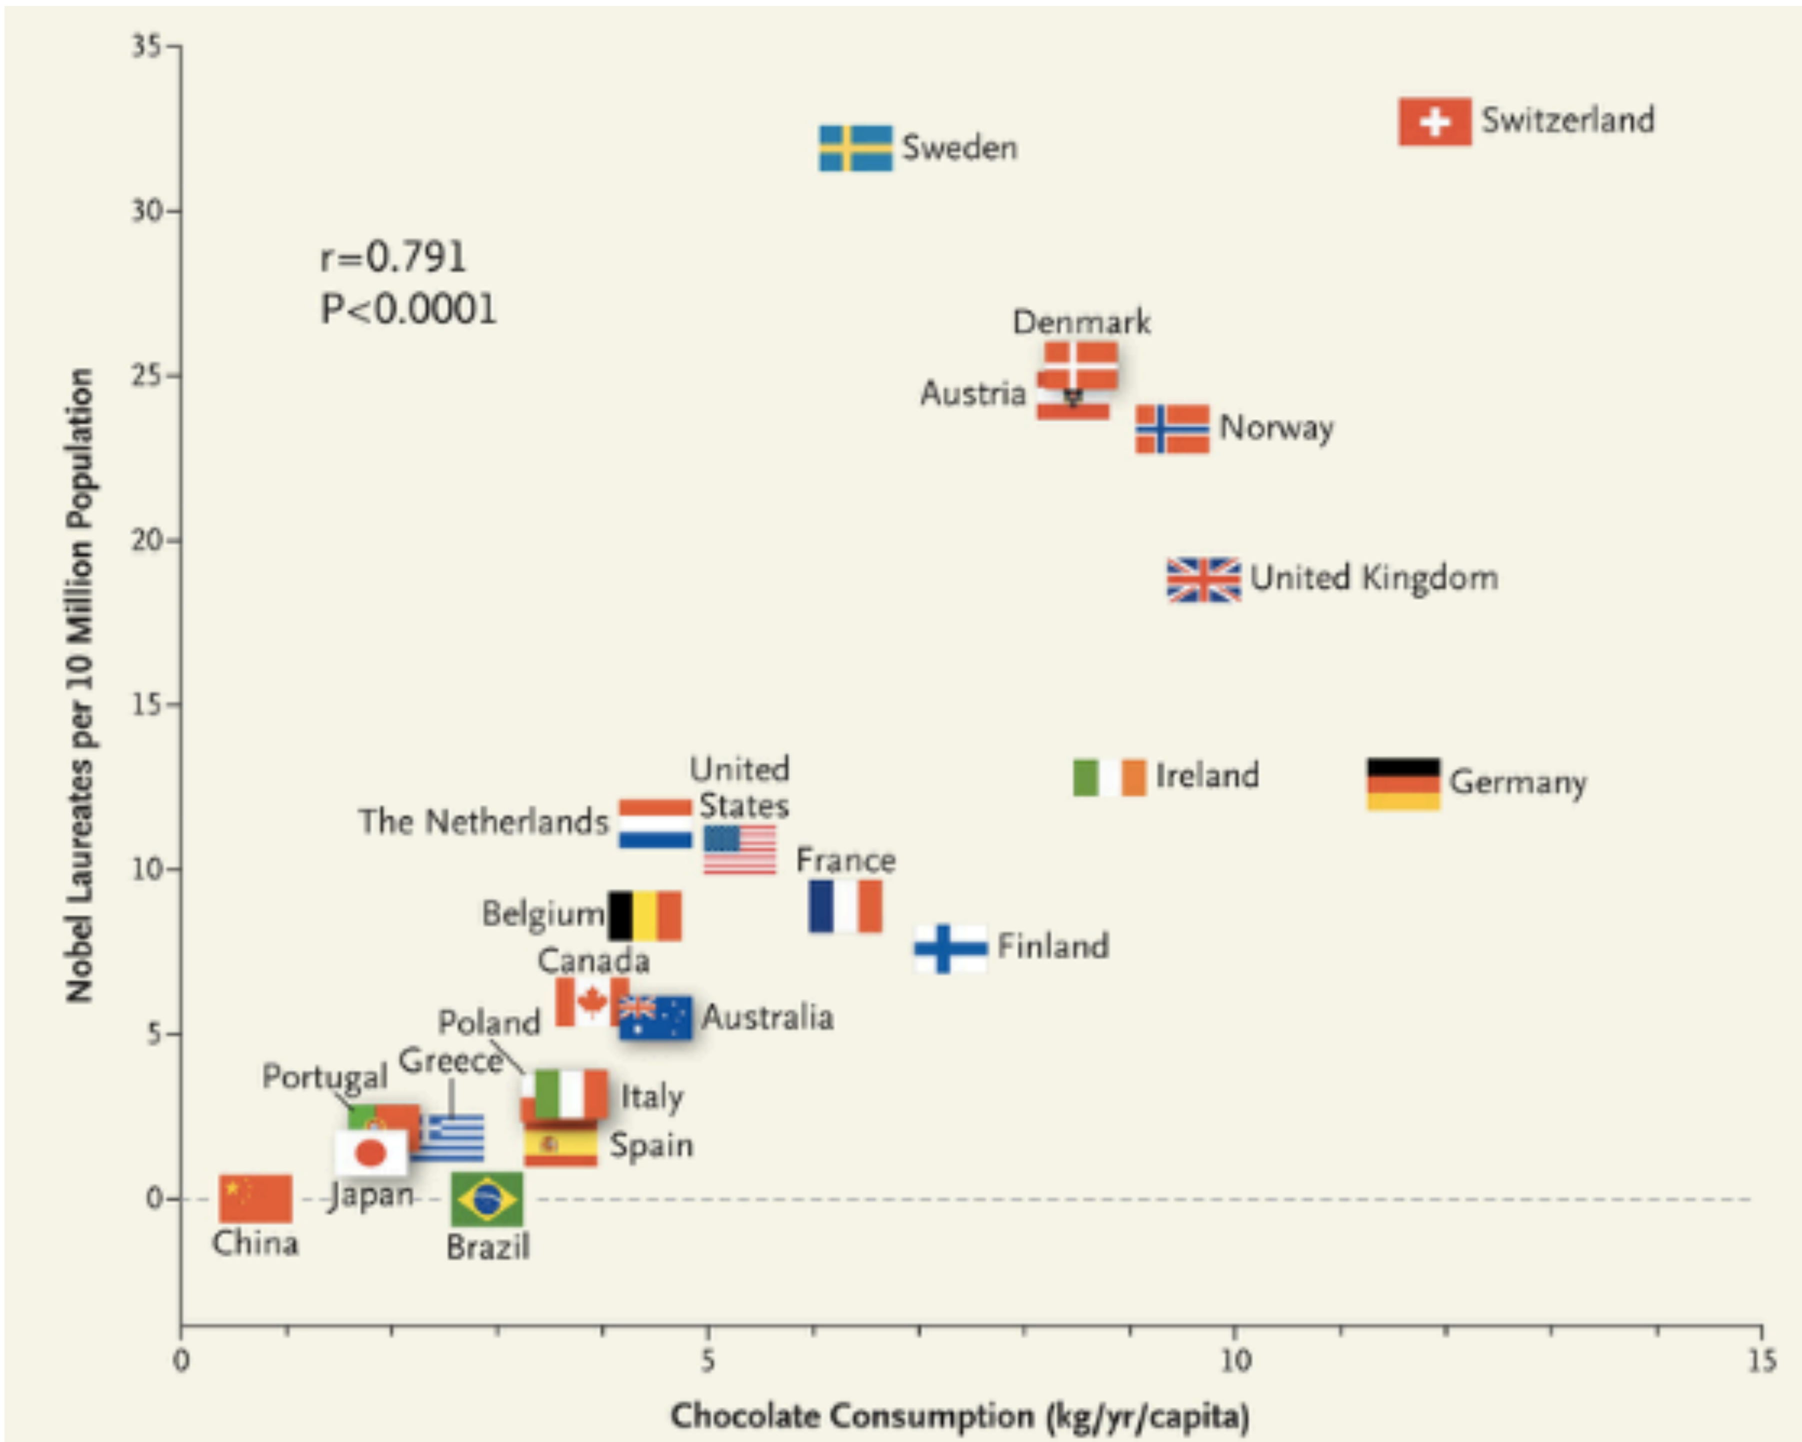

(Simple) Causal Discovery as an Estimation Problem

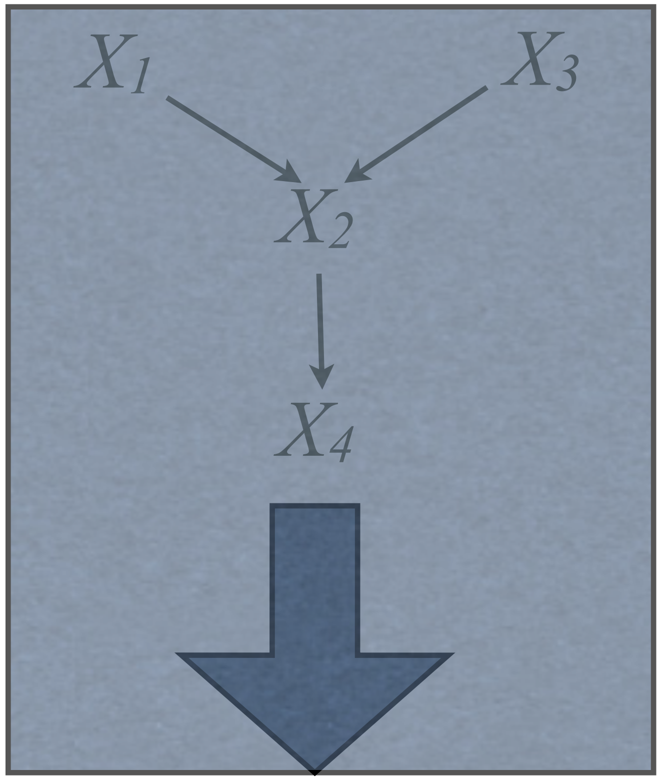

Data

X_1	X_2	X_3	X_4
-1.1	1.0	1.3	0.2
2.1	2.0	3.1	-1.3
3.1	4.2	2.6	0.6
2.3	-0.6	-3.5	0.8
1.3	2.2	0.9	2.4
-1.8	0.9	-1.3	0.9
...	...	...	...

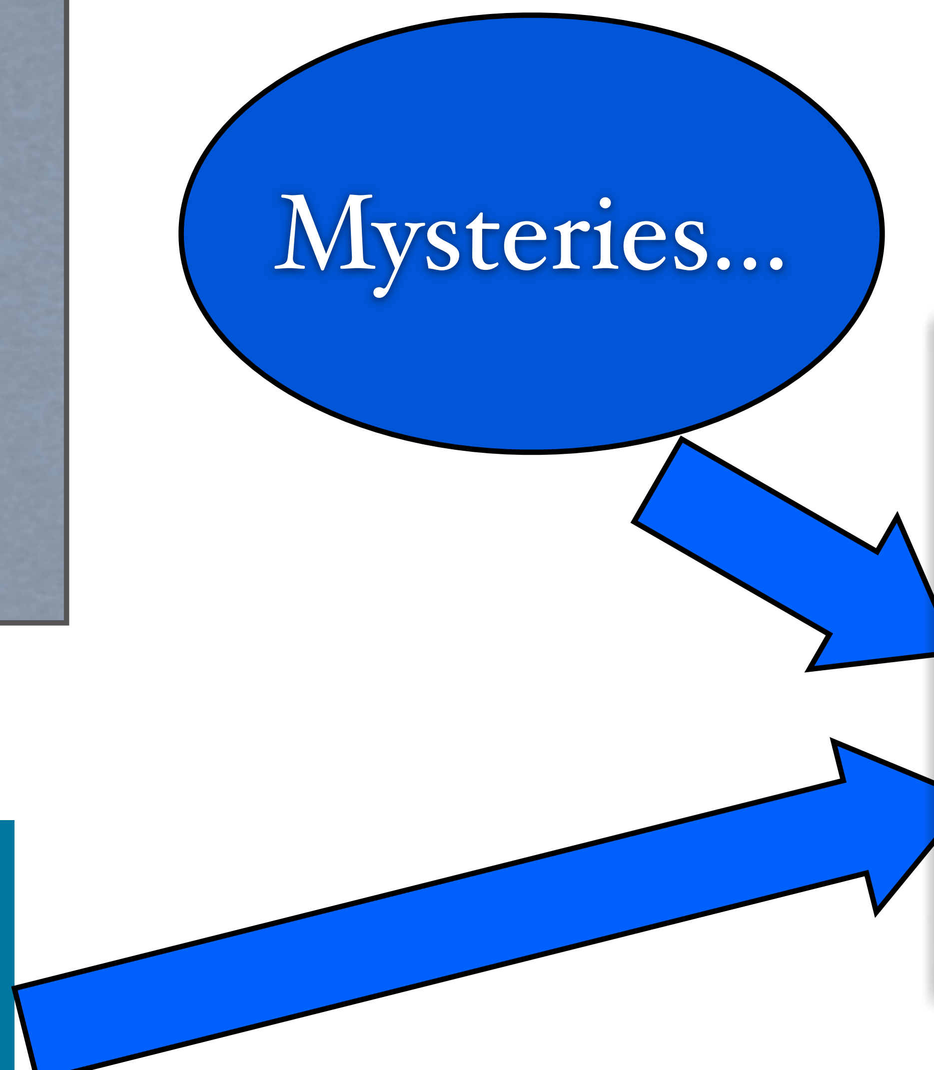

	X_1	X_2	X_3	X_4
X_1	0	0	0	0
X_4	0	1	0	0

Linear identifiable cases, find: $\mathbf{X} = \mathrm{B}\cdot\mathbf{X} + \mathbf{E}$

Nonlinear identifiable cases, find $X_i = f_i(\mathrm{PA}_i, E_i)$

What if there are latent confounders?

Temporal Order Often Helpful.
I.I.D. Case More Difficult.

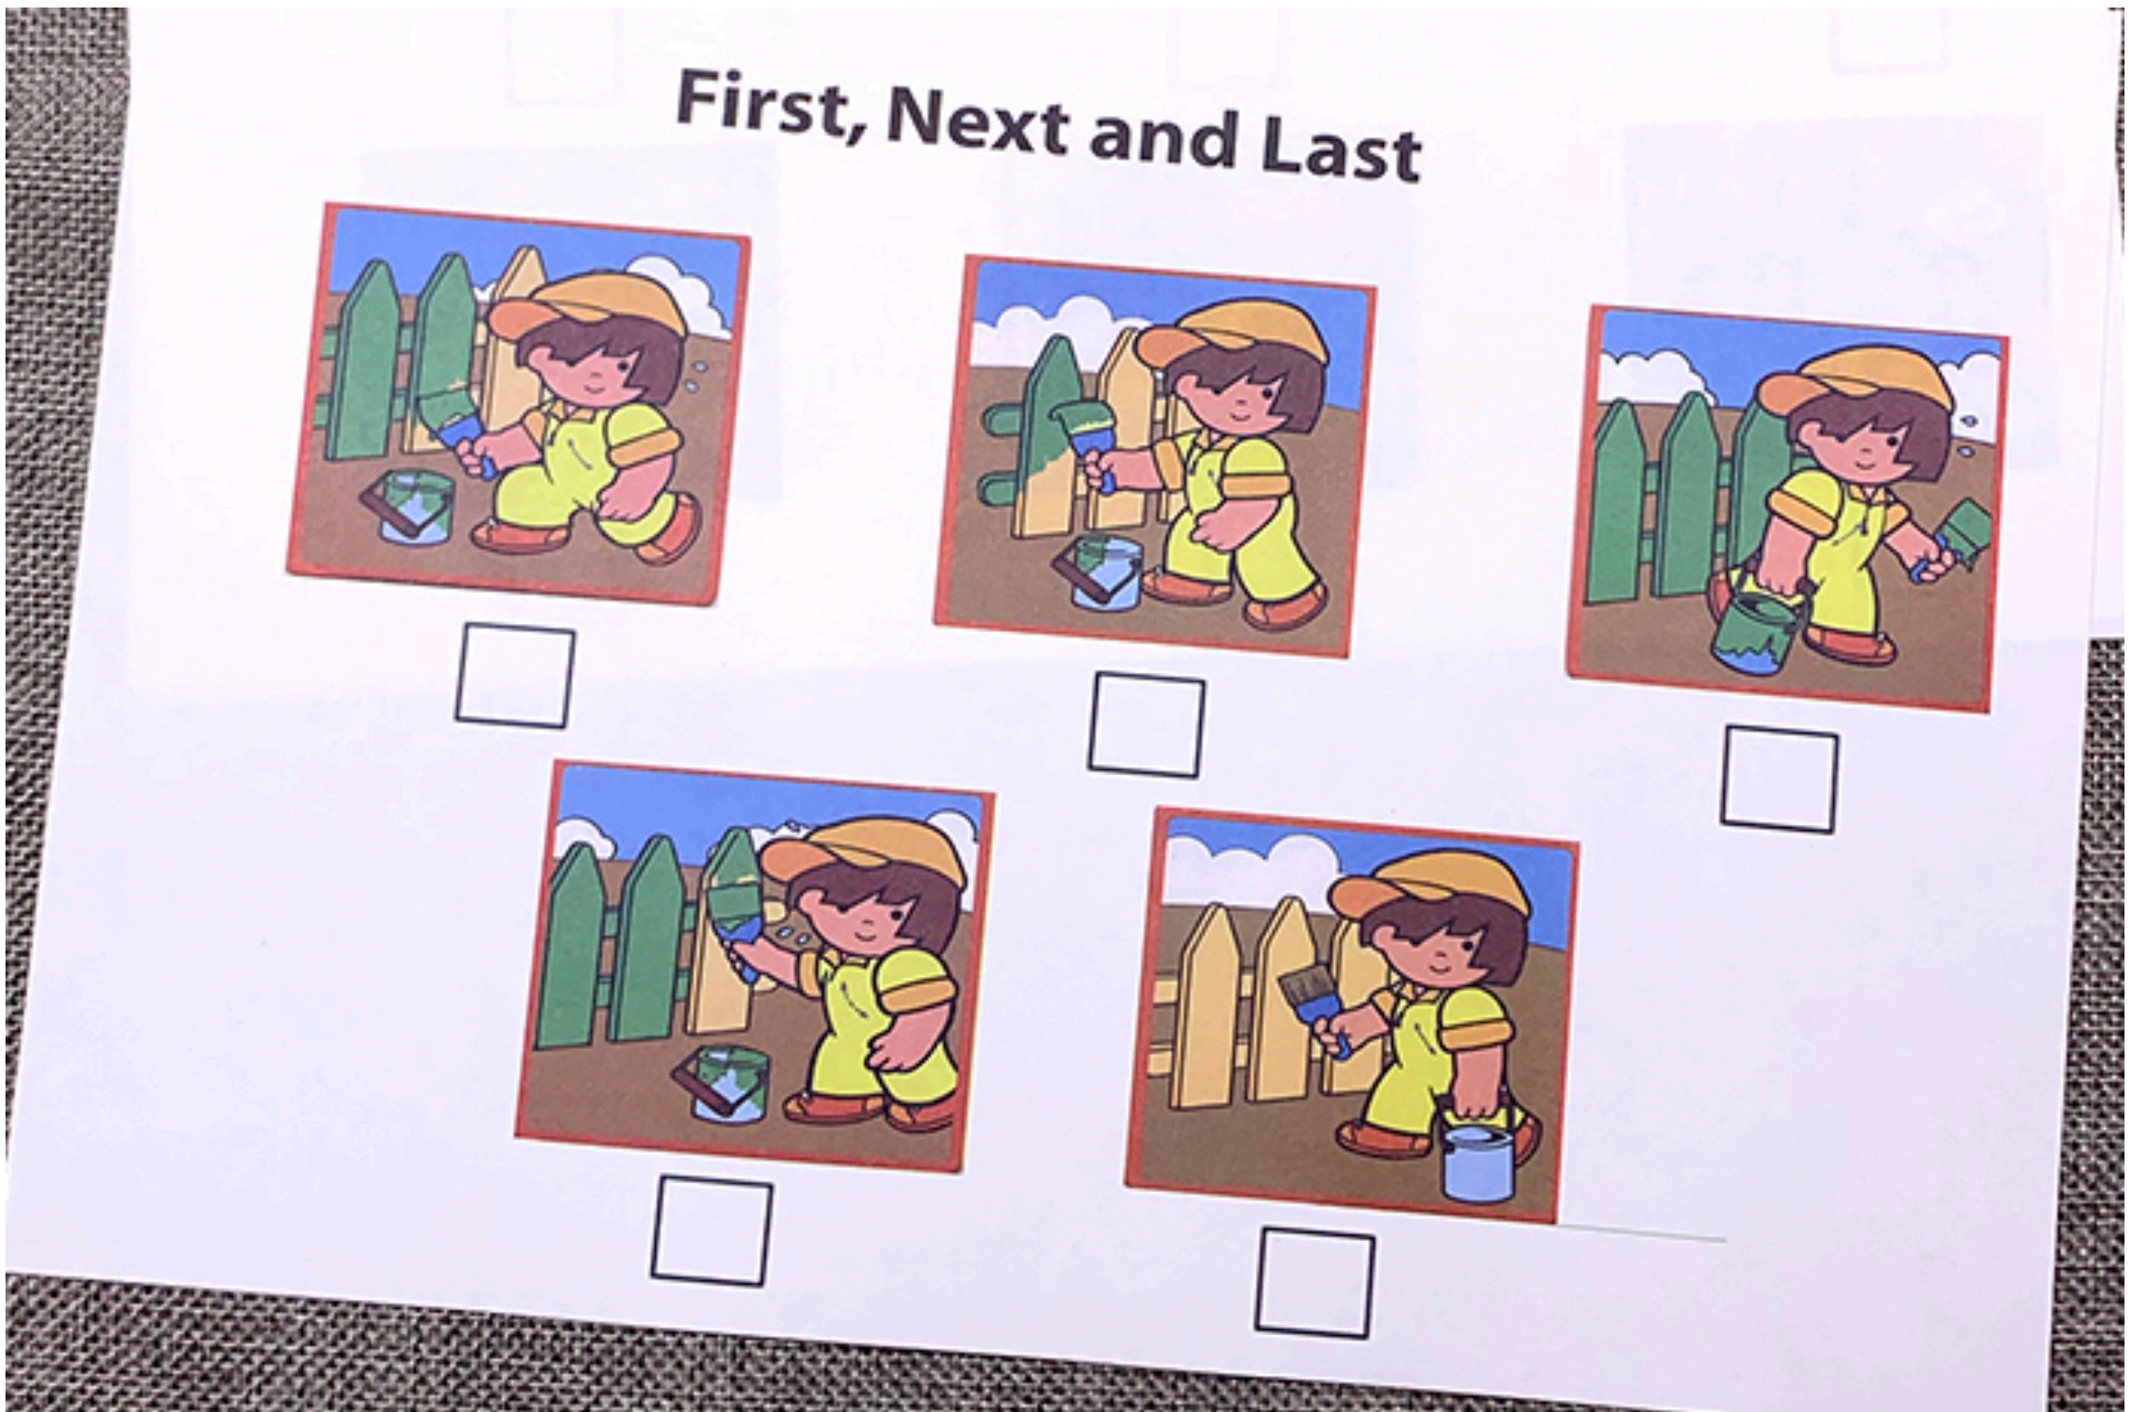

Causal Discovery: A Bit of History

- Reichenbach's common cause principle ("The Direction of Time", 1956)
- Markov condition (Kiiveri et al., 1984)
- "Causation, Prediction, and Search" (Spirtes, Glymour, & Scheines, 1993)
    - Faithfulness condition, PC algorithm, SGS, FCI, Tetrad program..
- "Causality: Models, Reasoning and Inference" (Pearl, 2000)
- Greedy equivalence search (GES) (Chickering, 2003)
- Functional causal model-based methods (LiNGAM, PNL… since 2005)
- Latent variable recovery: Factor analysis (Spearman, 1904), Tetrad condition (Spearman & CMU), Latent tree structure (Pearl et al., 1989), measurement model (CMU 2006), GIN (GDUT & CMU), rank deficiency (CMU)…

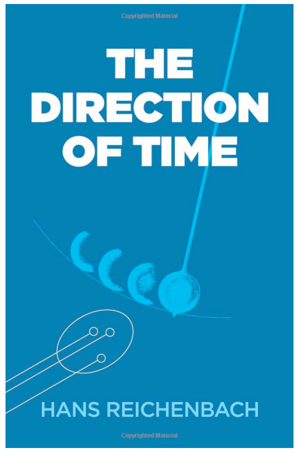

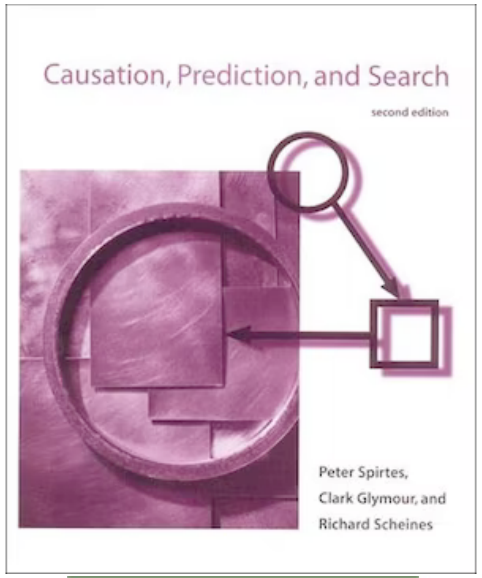

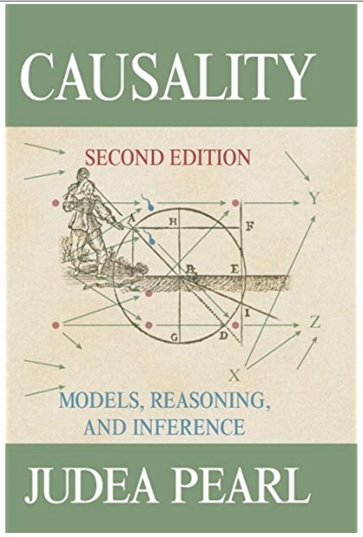

Uncover Causality from Observational Data?

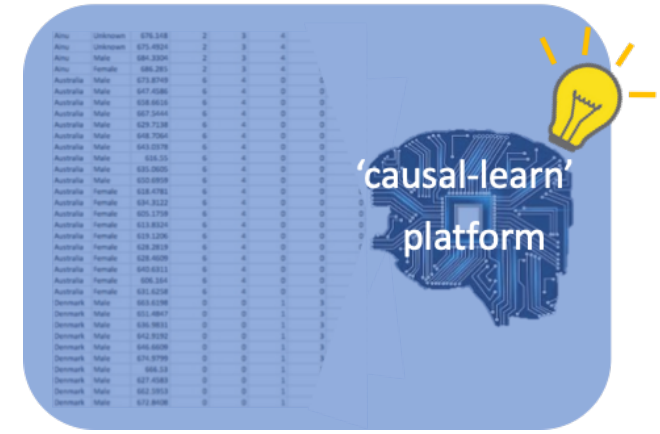

- Causal system has "irrelevant" modules (Pearl, 2000; Spirtes et al., 1993)

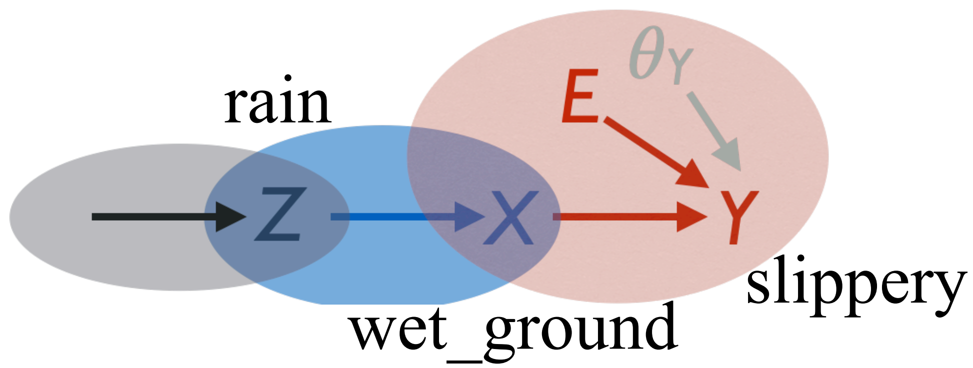

- conditional independence among variables;
- independent noise condition;
- minimal (and independent) changes...

Footprint of causality in data

- Causal discovery (Spirtes et al., 1993)/ causal representation learning (Schölkopf et al., 2021): find such representations with identifiability guarantees

- Three dimensions of the problem:

i.i.d. data?	**Parametric constraints?**	**Latent confounders?**
Yes	No	No
No	Yes	Yes

Causal Discovery in Archeology: An Example

i.i.d. data?	Parametric constraints?	Latent confounders?
Yes	No	No
No	Yes	Yes

Thanks to Marlijn Noback

- 8 variables of 250 skeletons collected from different locations

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
1	Id	Population	Sex	Cranial size	Diet or subsistence					Paramastic	Dental wear		Geographic location per population			Climate per population					
2			(Male, fem	(Centroid S	Gathering	Hunting	Fishing	Pastoralism	Agriculture	Yes=1, no=	Average at	Attrition pe	Distance to	Longitude	Latitude	Tmean	Tmin	Tmax	Vpmean	Vpmin	Vpmax
3	AINU31_1	Ainu	Unknown	713.2942	2	3	4	0	1	0	1.5	2	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
4	AINU7_1	Ainu	Unknown	676.148	2	3	4	0	1	0	1.5	1	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
5	AINU7_2	Ainu	Unknown	675.4924	2	3	4	0	1	0	1.5	1	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
6	AINU_1016	Ainu	Male	684.3304	2	3	4	0	1	0	1.5	2.5	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
7	AINU_1016	Ainu	Female	686.285	2	3	4	0	1	0	1.5	4	16464	43.548548	142.639159	2.86	-11.19	17.01	7.43	2.27	16.83
8	AUSM245	Australia	Male	673.8749	6	4	0	0	0	1	2.5	1	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
9	AUSM246	Australia	Male	647.4586	6	4	0	0	0	1	2.5	4	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
10	AUSM8217	Australia	Male	658.6616	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
11	AUSM8177	Australia	Male	667.5444	6	4	0	0	0	1	2.5	4	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
12	AUSM8173	Australia	Male	629.7138	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
13	AUSM8173	Australia	Male	648.7064	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
14	AUSM8171	Australia	Male	643.0378	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
15	AUSM8165	Australia	Male	616.55	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
16	AUSM8154	Australia	Male	635.0605	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
17	AUSM8153	Australia	Male	650.6959	6	4	0	0	0	1	2.5	3	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
18	AUSF1412	Australia	Female	618.4781	6	4	0	0	0	1	2.5	1	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
19	AUSF8179	Australia	Female	634.3122	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
20	AUSF8175	Australia	Female	605.1759	6	4	0	0	0	1	2.5	1.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
21	AUSF8172	Australia	Female	613.8324	6	4	0	0	0	1	2.5	3	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
22	AUSF8169	Australia	Female	619.1206	6	4	0	0	0	1	2.5	2.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
23	AUSF8157	Australia	Female	628.2819	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
24	AUSF8155	Australia	Female	628.4609	6	4	0	0	0	1	2.5	3.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
25	AUSF1578	Australia	Female	640.6311	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
26	AUSF243	Australia	Female	606.164	6	4	0	0	0	1	2.5	2.5	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
27	AUSF8158	Australia	Female	631.6258	6	4	0	0	0	1	2.5	2	20164	-24.287027	135.615234	22.46	13.33	30.27	11.10	7.55	15.96
28	DENM1432	Denmark	Male	663.6198	0	0	1	3	6	0	2.1	2	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
29	DENM1011	Denmark	Male	651.4847	0	0	1	3	6	0	2.1	3	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
30	DENM1205	Denmark	Male	636.9831	0	0	1	3	6	0	2.1	1.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
31	DENM116_	Denmark	Male	642.9192	0	0	1	3	6	0	2.1	3	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
32	DENM116_	Denmark	Male	646.6609	0	0	1	3	6	0	2.1	2.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
33	DENM116_	Denmark	Male	674.9799	0	0	1	3	6	0	2.1	2	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27
34	DENM7_77	Denmark	Male	666.53	0	0	1	3	6	0	2.1	2.5	10440	55.717055	11.711426	8.01	-0.02	16.66	9.67	5.59	15.27

Result of PC on the Archeology Data

Thanks to collaborator Marlijn Noback

- By PC algorithm (Spirtes et al., 1993) + kernel-based conditional independence test (Zhang et al., 2011)

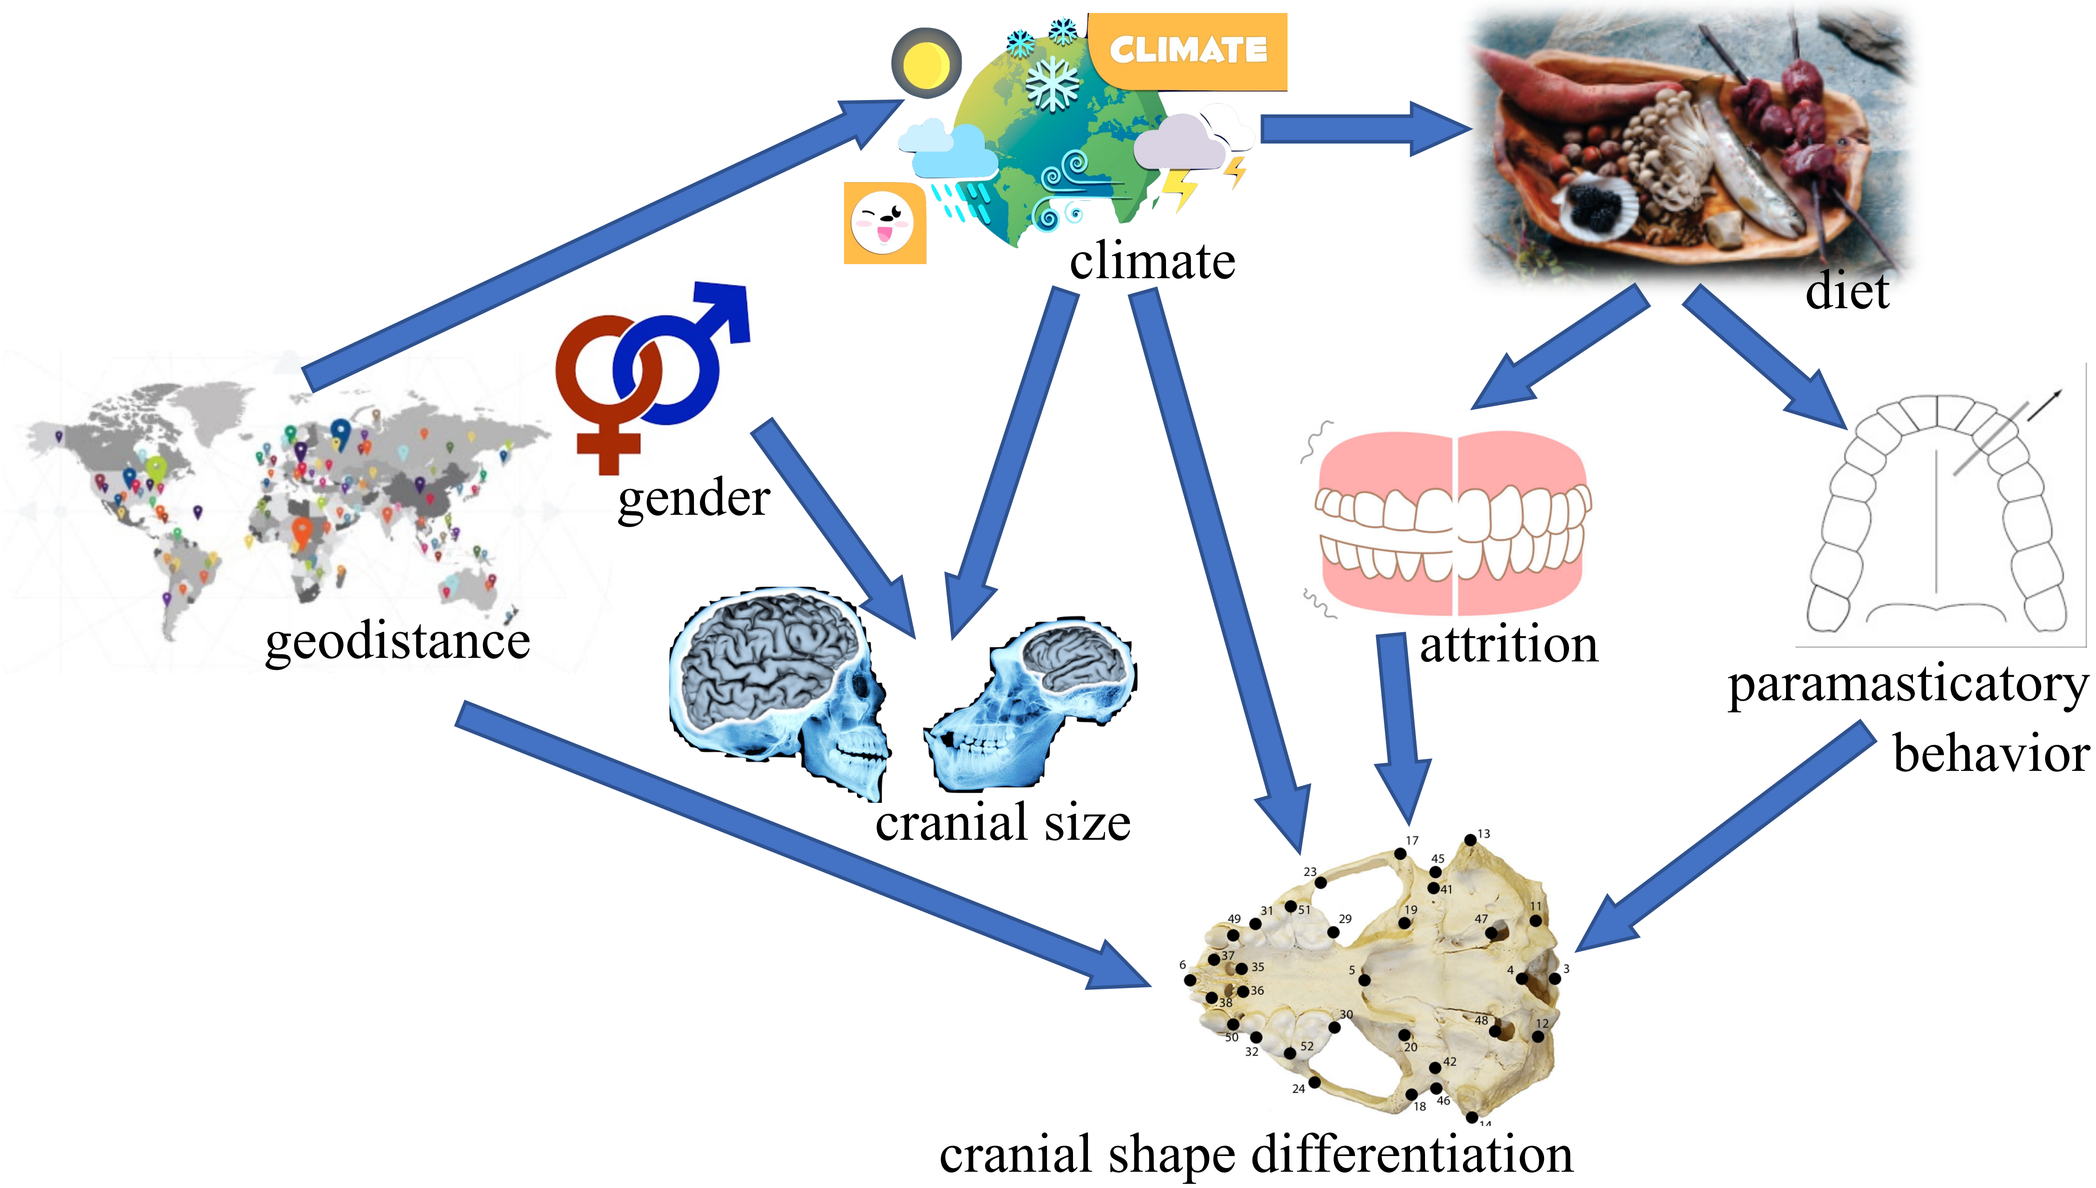

Functional Causal Model-Based Causal Discovery

i.i.d. data?	Parametric constraints?	Latent confounders?
Yes	No	No
No	Yes	Yes

"Independent changes" renders causal direction identifiable

- Linear non-Gaussian model (Shimizu et al., 2006):
  $Y = aX + E$

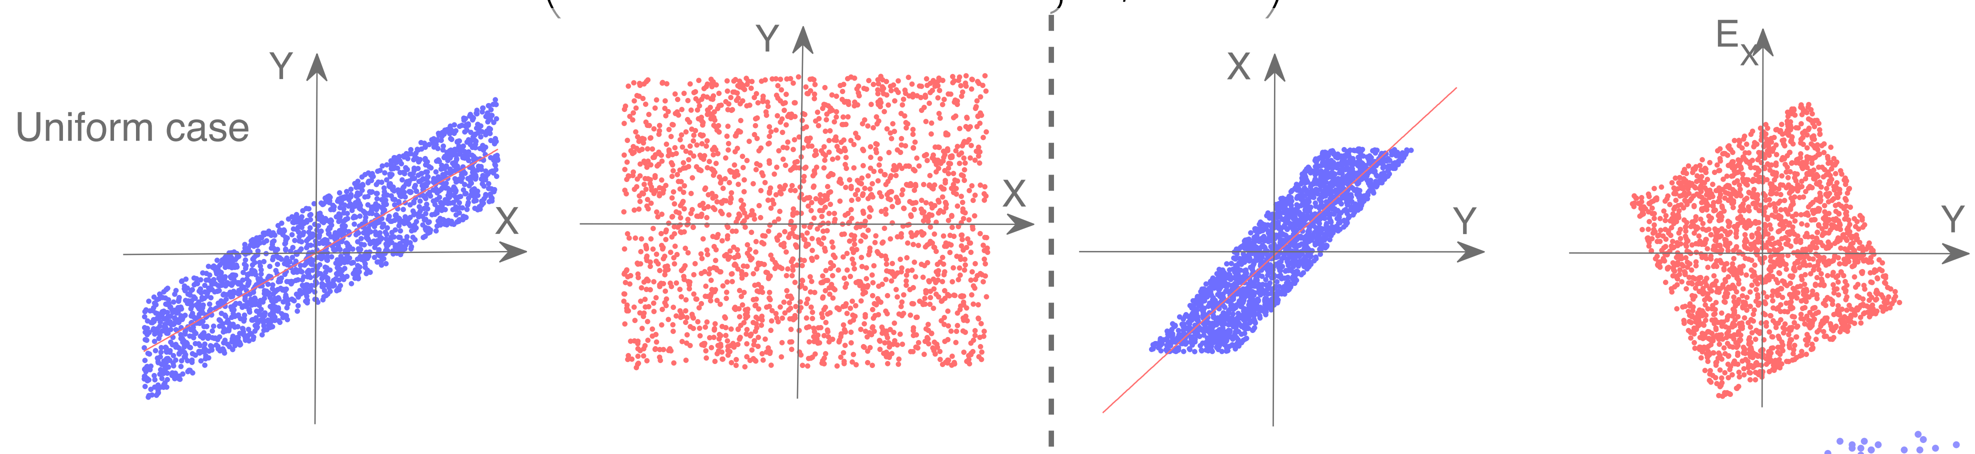

- Post-nonlinear causal model (Zhang & Chan, 2006):

$$Y = f_2\,(\,f_1(X) + E\,)$$

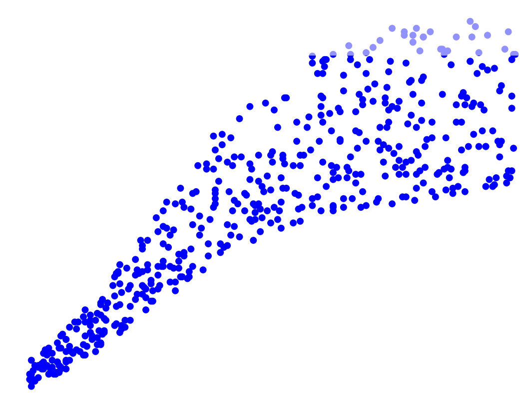

- Additive noise model (Hoyer et al, 2009)

$$Y = f(X) + E$$

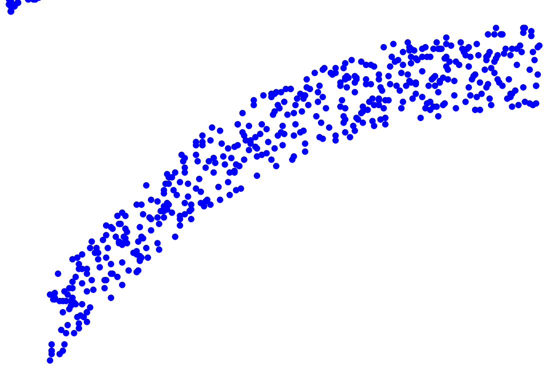

A Problem in Psychology: Finding Underlying Mental Conditions?

- 50 questions for big 5 personality test

race	age	engnat	gender	hand	source	country	E1	E2	E3	E4	E5	E6	E7	E8	E9	E10	N1	N2	N3	N4	N5	N6	N7	N8	N9	N10	A1	A2	A3	A4	A5
3	53	1	1	1	1	US	4	2	5	2	5	1	4	3	5	1	1	5	2	5	1	1	1	1	1	1	1	5	1	5	2
13	46	1	2	1	1	US	2	2	3	3	3	3	1	5	1	5	2	3	4	2	3	4	3	2	2	4	1	3	3	4	4
1	14	2	2	1	1	PK	5	1	1	4	5	1	1	5	5	1	5	1	5	5	5	5	5	5	5	5	5	1	5	5	1
3	19	2	2	1	1	RO	2	5	2	4	3	4	3	4	4	5	5	4	4	2	4	5	5	5	4	5	2	5	4	4	3
11	25	2	2	1	2	US	3	1	3	3	3	1	3	1	3	5	3	3	3	4	3	3	3	3	3	4	5	5	3	5	1
13	31	1	2	1	2	US	1	5	2	4	1	3	2	4	1	5	1	5	4	5	1	4	4	1	5	2	2	2	3	4	3
5	20	1	2	1	5	US	5	1	5	1	5	1	5	4	4	1	2	4	2	4	2	2	3	2	2	2	5	5	1	5	1
4	23	2	1	1	2	IN	4	3	5	3	5	1	4	3	4	3	1	4	4	4	1	1	1	1	1	1	2	5	1	4	3
5	39	1	2	3	4	US	3	1	5	1	5	1	5	2	5	3	2	4	5	3	3	5	5	4	3	3	1	5	1	5	1
3	18	1	2	1	5	US	1	4	2	5	2	4	1	4	1	5	5	2	5	2	3	4	3	2	3	4	2	3	1	4	2
3	17	2	2	1	1	IT	1	5	2	5	1	4	1	4	1	5	5	3	5	3	2	5	3	3	4	3	2	4	2	4	1
13	15	2	1	1	1	IN	3	3	5	3	3	3	2	4	3	3	1	5	3	3	2	3	2	3	2	4	4	4	2	2	5
13	22	1	2	1	2	US	3	3	4	2	4	2	2	3	4	3	3	3	3	3	2	2	4	4	2	3	1	4	1	5	1
3	21	1	2	1	5	US	1	3	2	5	1	1	1	5	1	5	5	3	5	2	5	5	3	2	5	3	1	1	1	4	2
3	28	2	2	1	2	US	3	3	3	4	3	2	2	4	3	5	2	4	4	4	4	4	2	2	3	2	1	4	2	4	2
3	21	1	1	1	5	US	2	3	2	3	3	1	1	3	4	4	2	4	2	4	1	2	2	2	2	2	4	2	4	2	5
13	19	1	2	1	2	FR	1	3	2	4	2	4	1	4	3	4	4	2	3	2	1	3	1	2	2	3	4	2	3	1	4
3	21	1	2	1	5	US	4	1	5	2	5	1	5	3	5	1	5	2	5	2	3	3	3	3	4	2	1	5	2	5	2
3	26	1	2	3	5	GB	2	3	4	3	1	4	1	4	1	5	4	2	5	2	1	4	2	2	2	2	2	2	2	2	2
3	26	1	2	1	1	US	2	2	3	3	3	3	1	3	3	3	4	4	3	1	3	2	2	2	4	4	1	3	2	4	3

Learning Hidden Variables & Their Relations

i.i.d. data?	Parametric constraints?	Latent confounders?
Yes	No	No
No	Yes	Yes

- Measured variables (e.g., answer scores in psychometric questionnaires) were generated by causally related latent variables

X1	X2	X3	X4	X5	X6	X7	X8
4.2	3.6	6.5	6.8	9.6	7.6	2.7	4.8
3.8	1.9	6.5	7.3	8.9	6.9	1.1	4.6
4.2	3.4	6.5	6.9	9.5	7.4	2.5	4.6
4.2	2.2	6.2	6.9	9.6	7.2	1.9	4.8
3.9	1.9	6.5	6.8	9.0	6.8	1.7	4.4
4.0	2.0	6.4	7.2	9.1	7.0	1.0	4.6
3.8	1.7	6.4	7.3	9.0	6.7	0.8	4.3
4.1	2.8	6.5	6.9	9.3	6.7	2.7	4.6
...	...	...	...	...	...	...	...

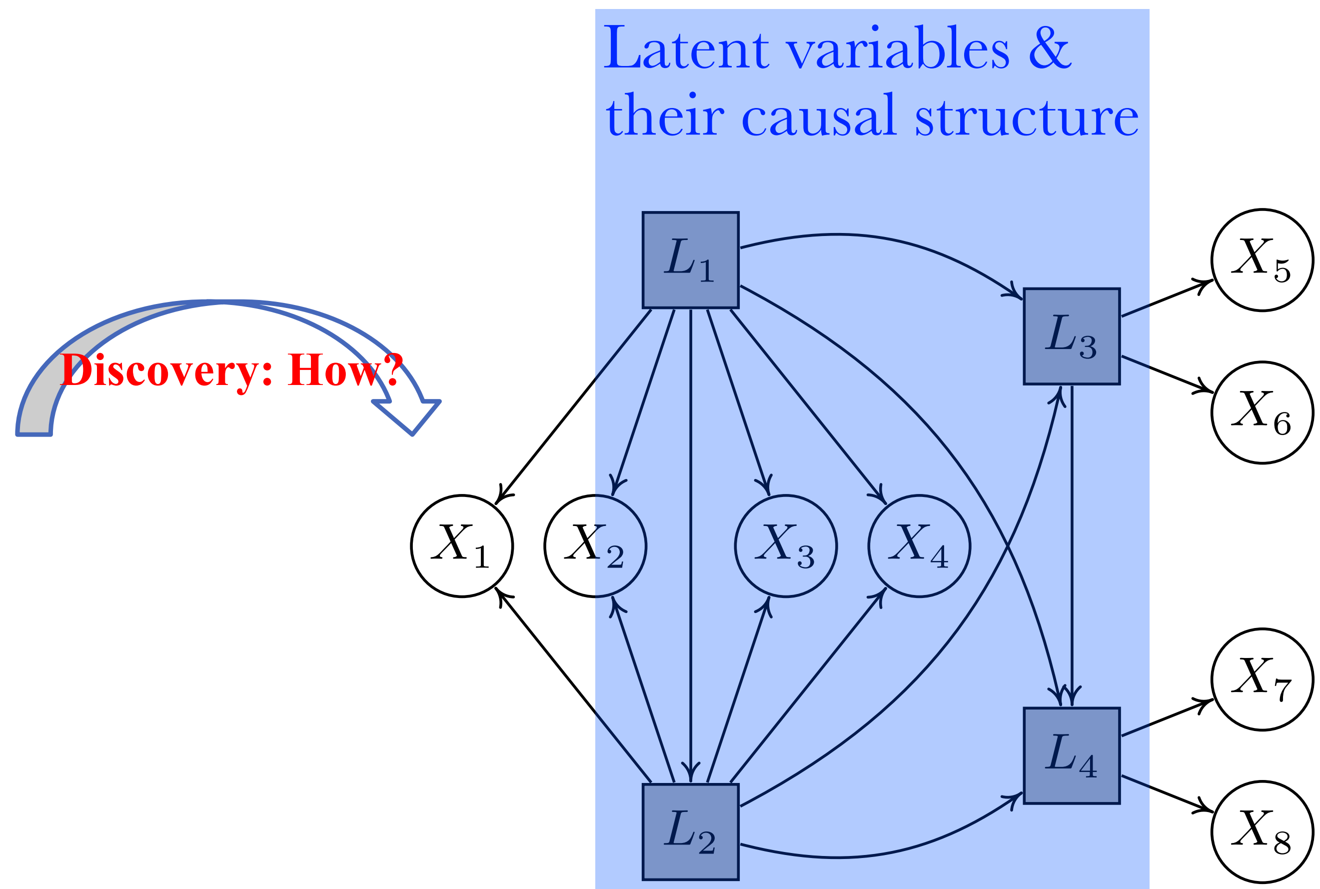

- Find latent variables L_i and their causal relations ?
- Rank deficiency or GIN helps solve the problem

Example: Big 5 Questions Are Well Designed but…

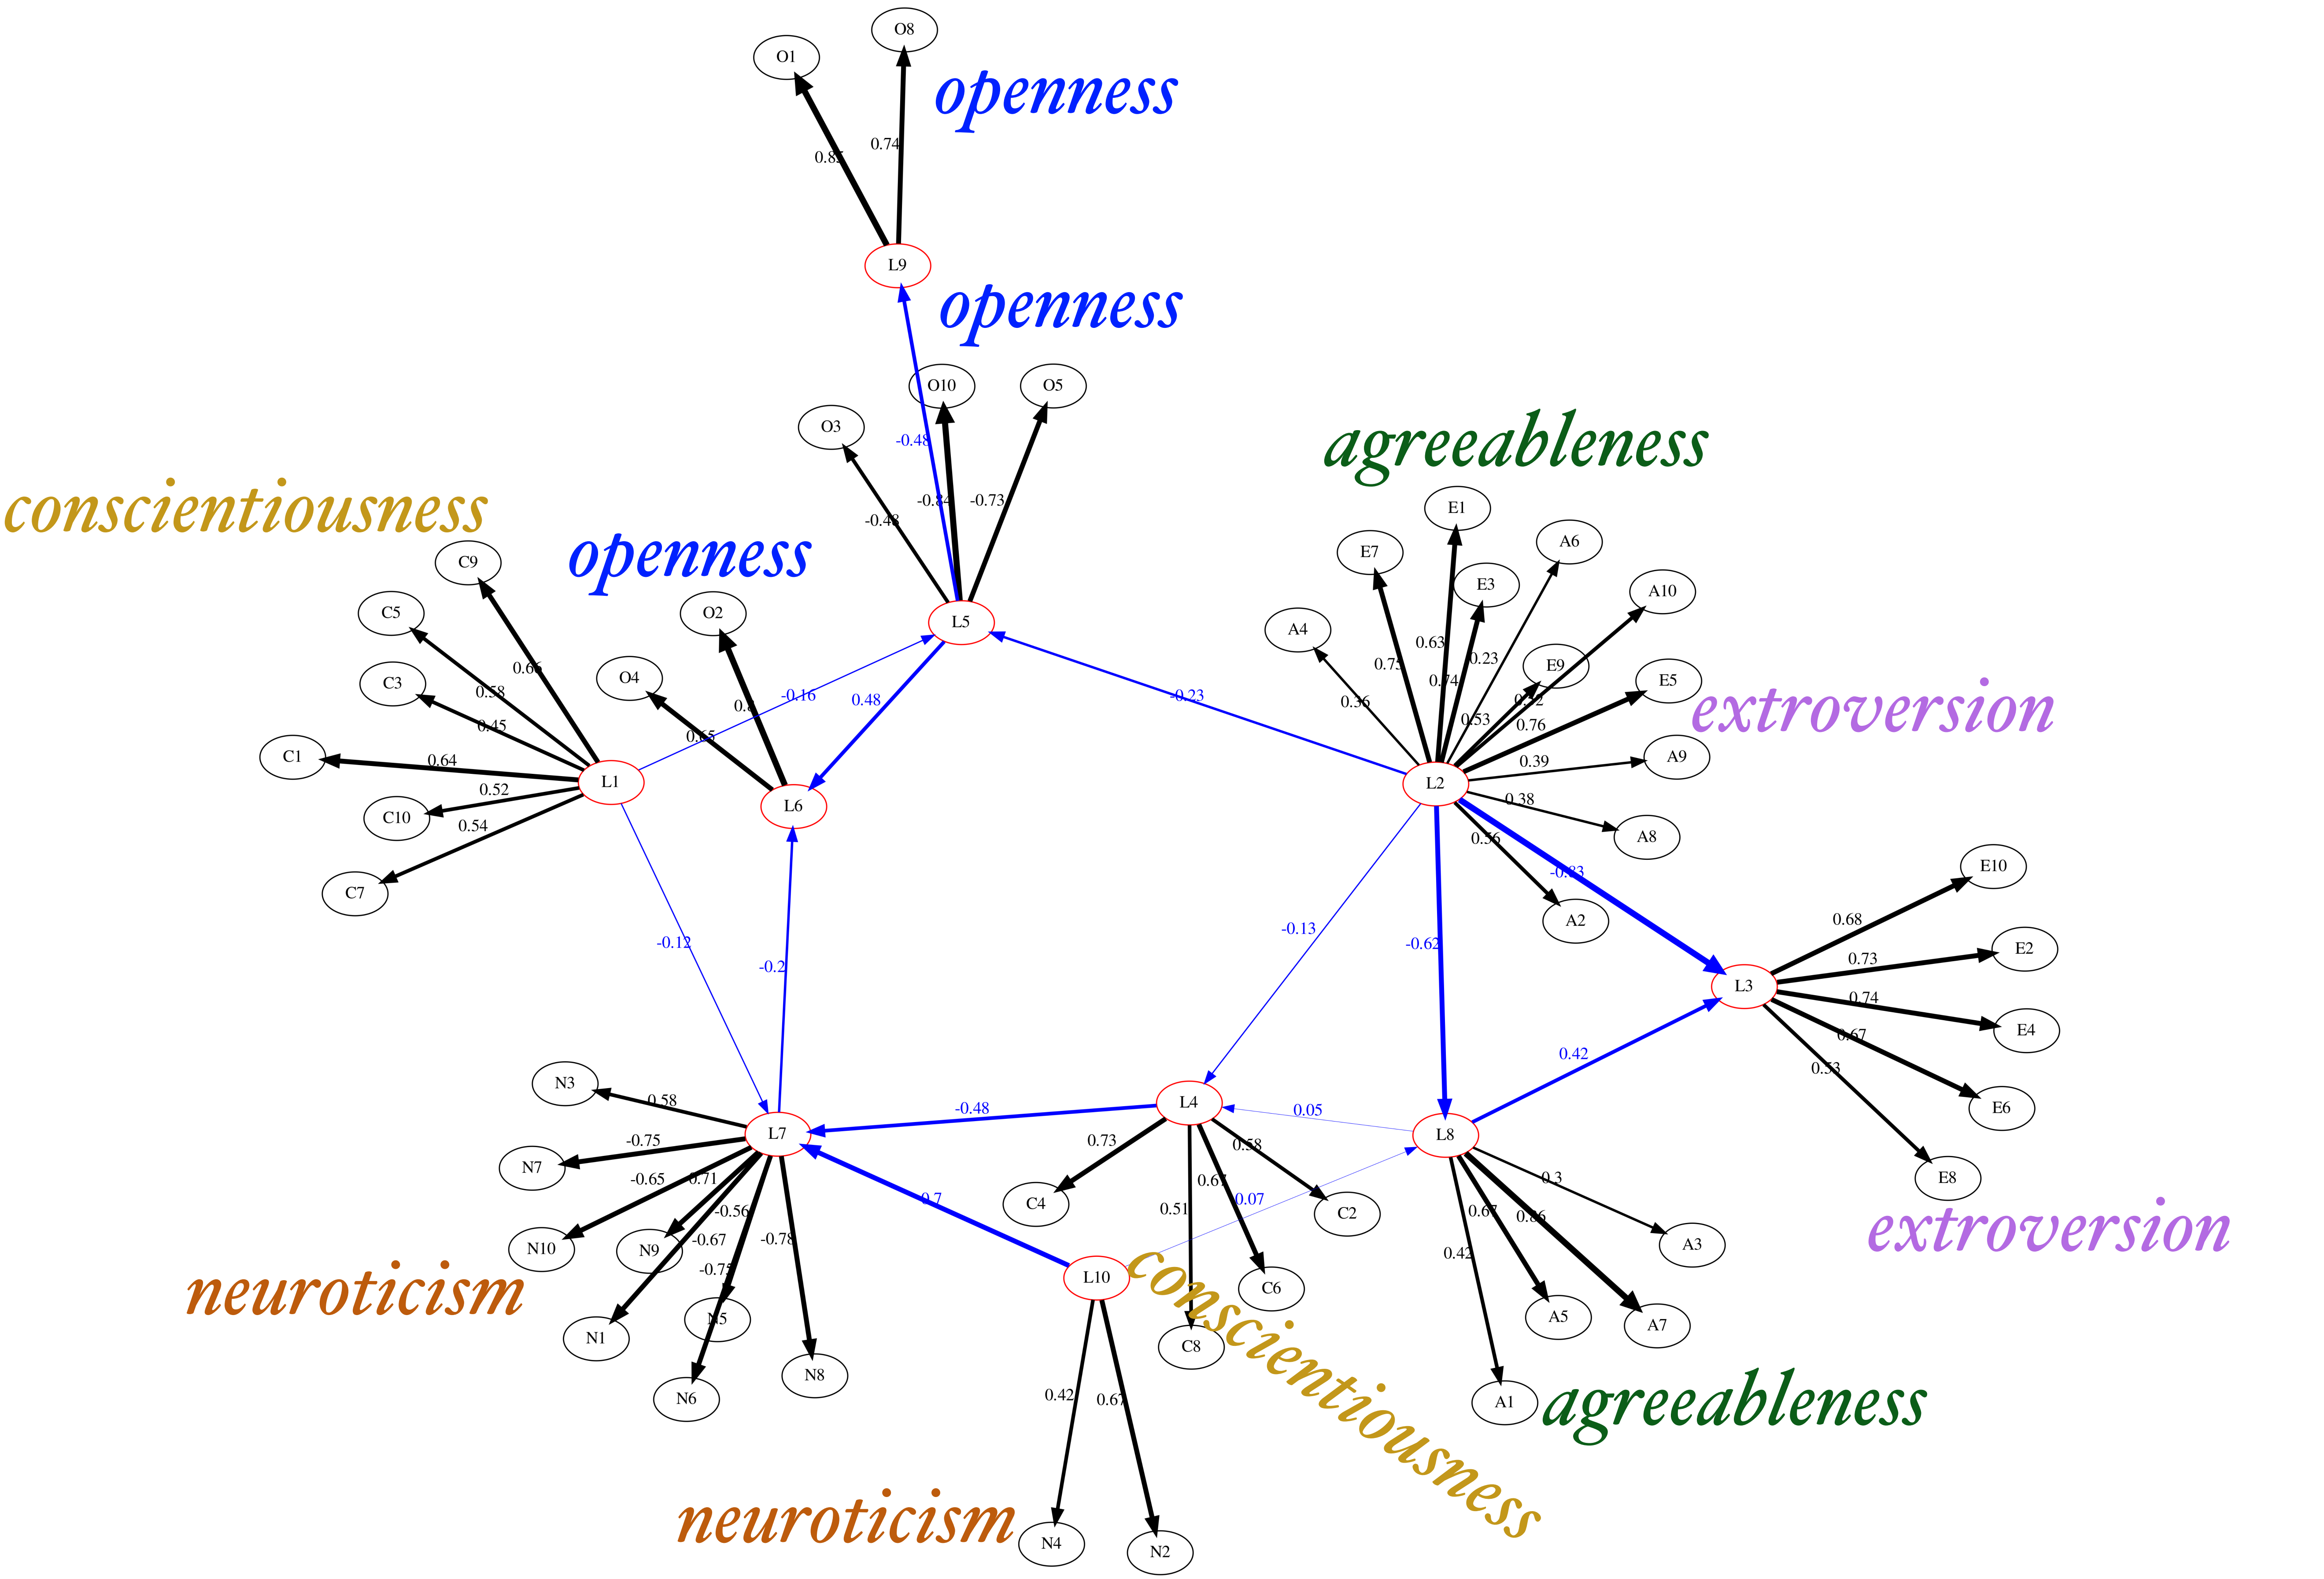

Example: Big 5 Questions Are Well Designed but…

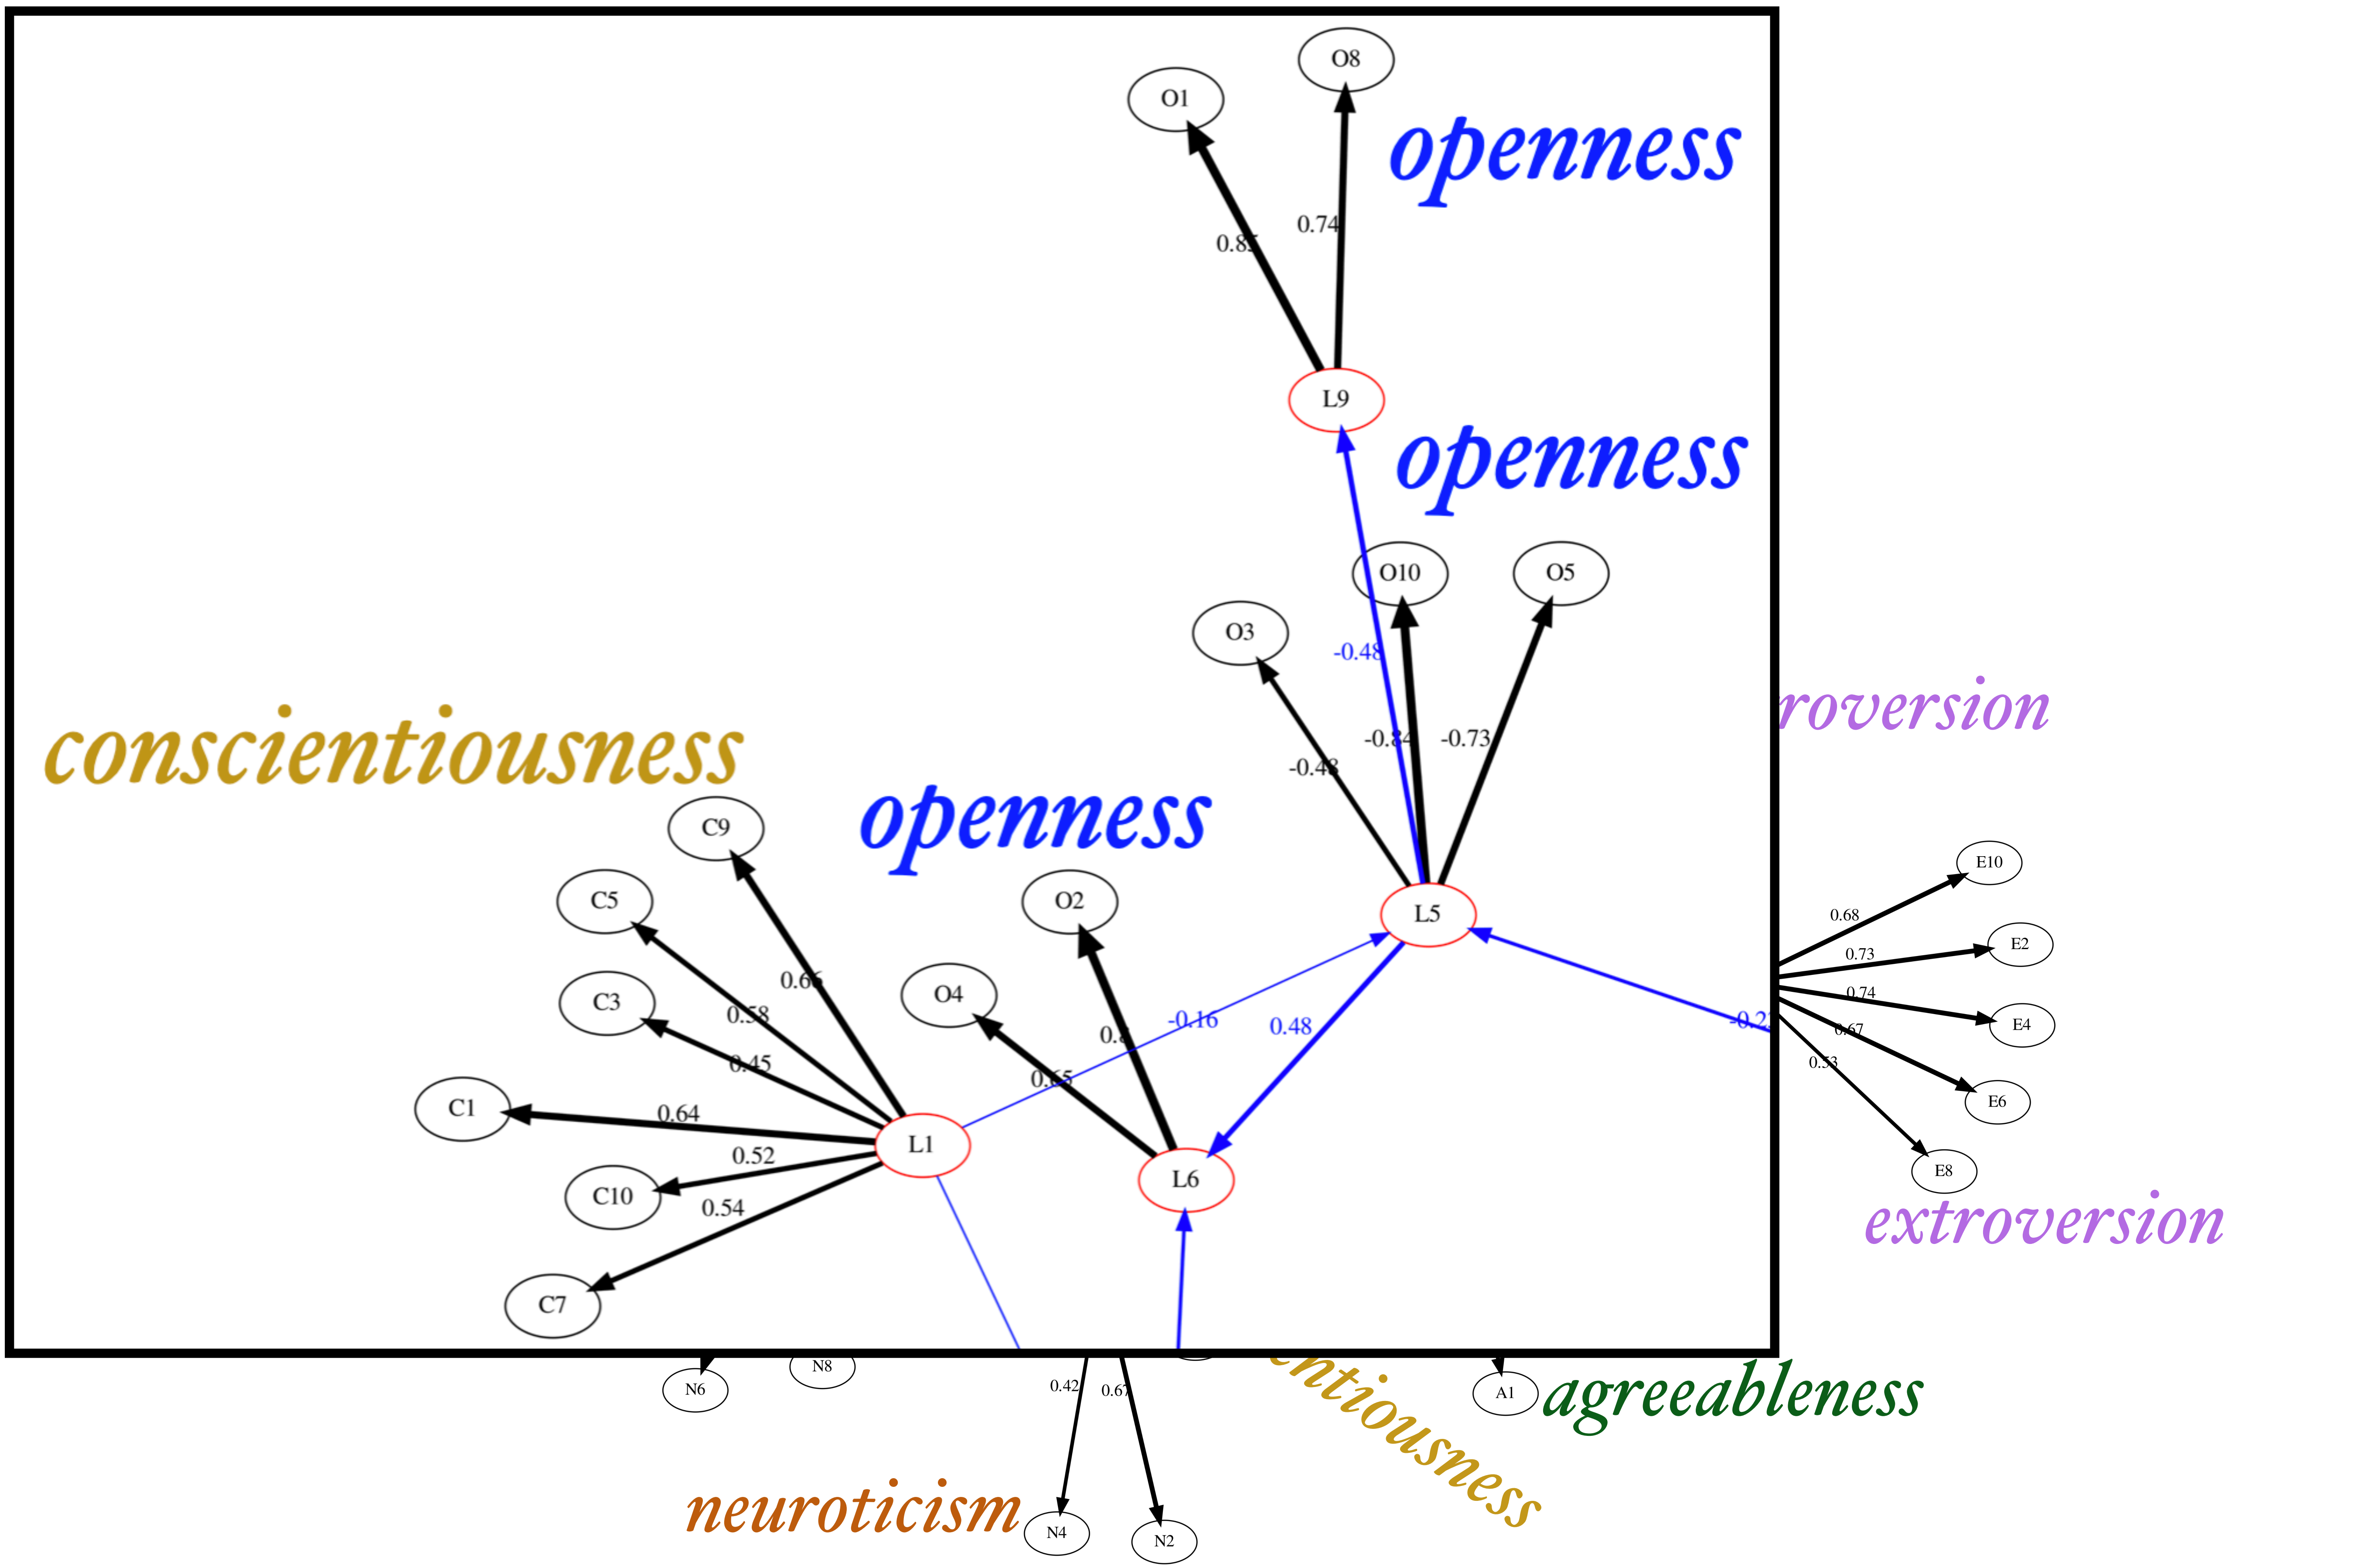

Example: Big 5 Questions Are Well Designed but…

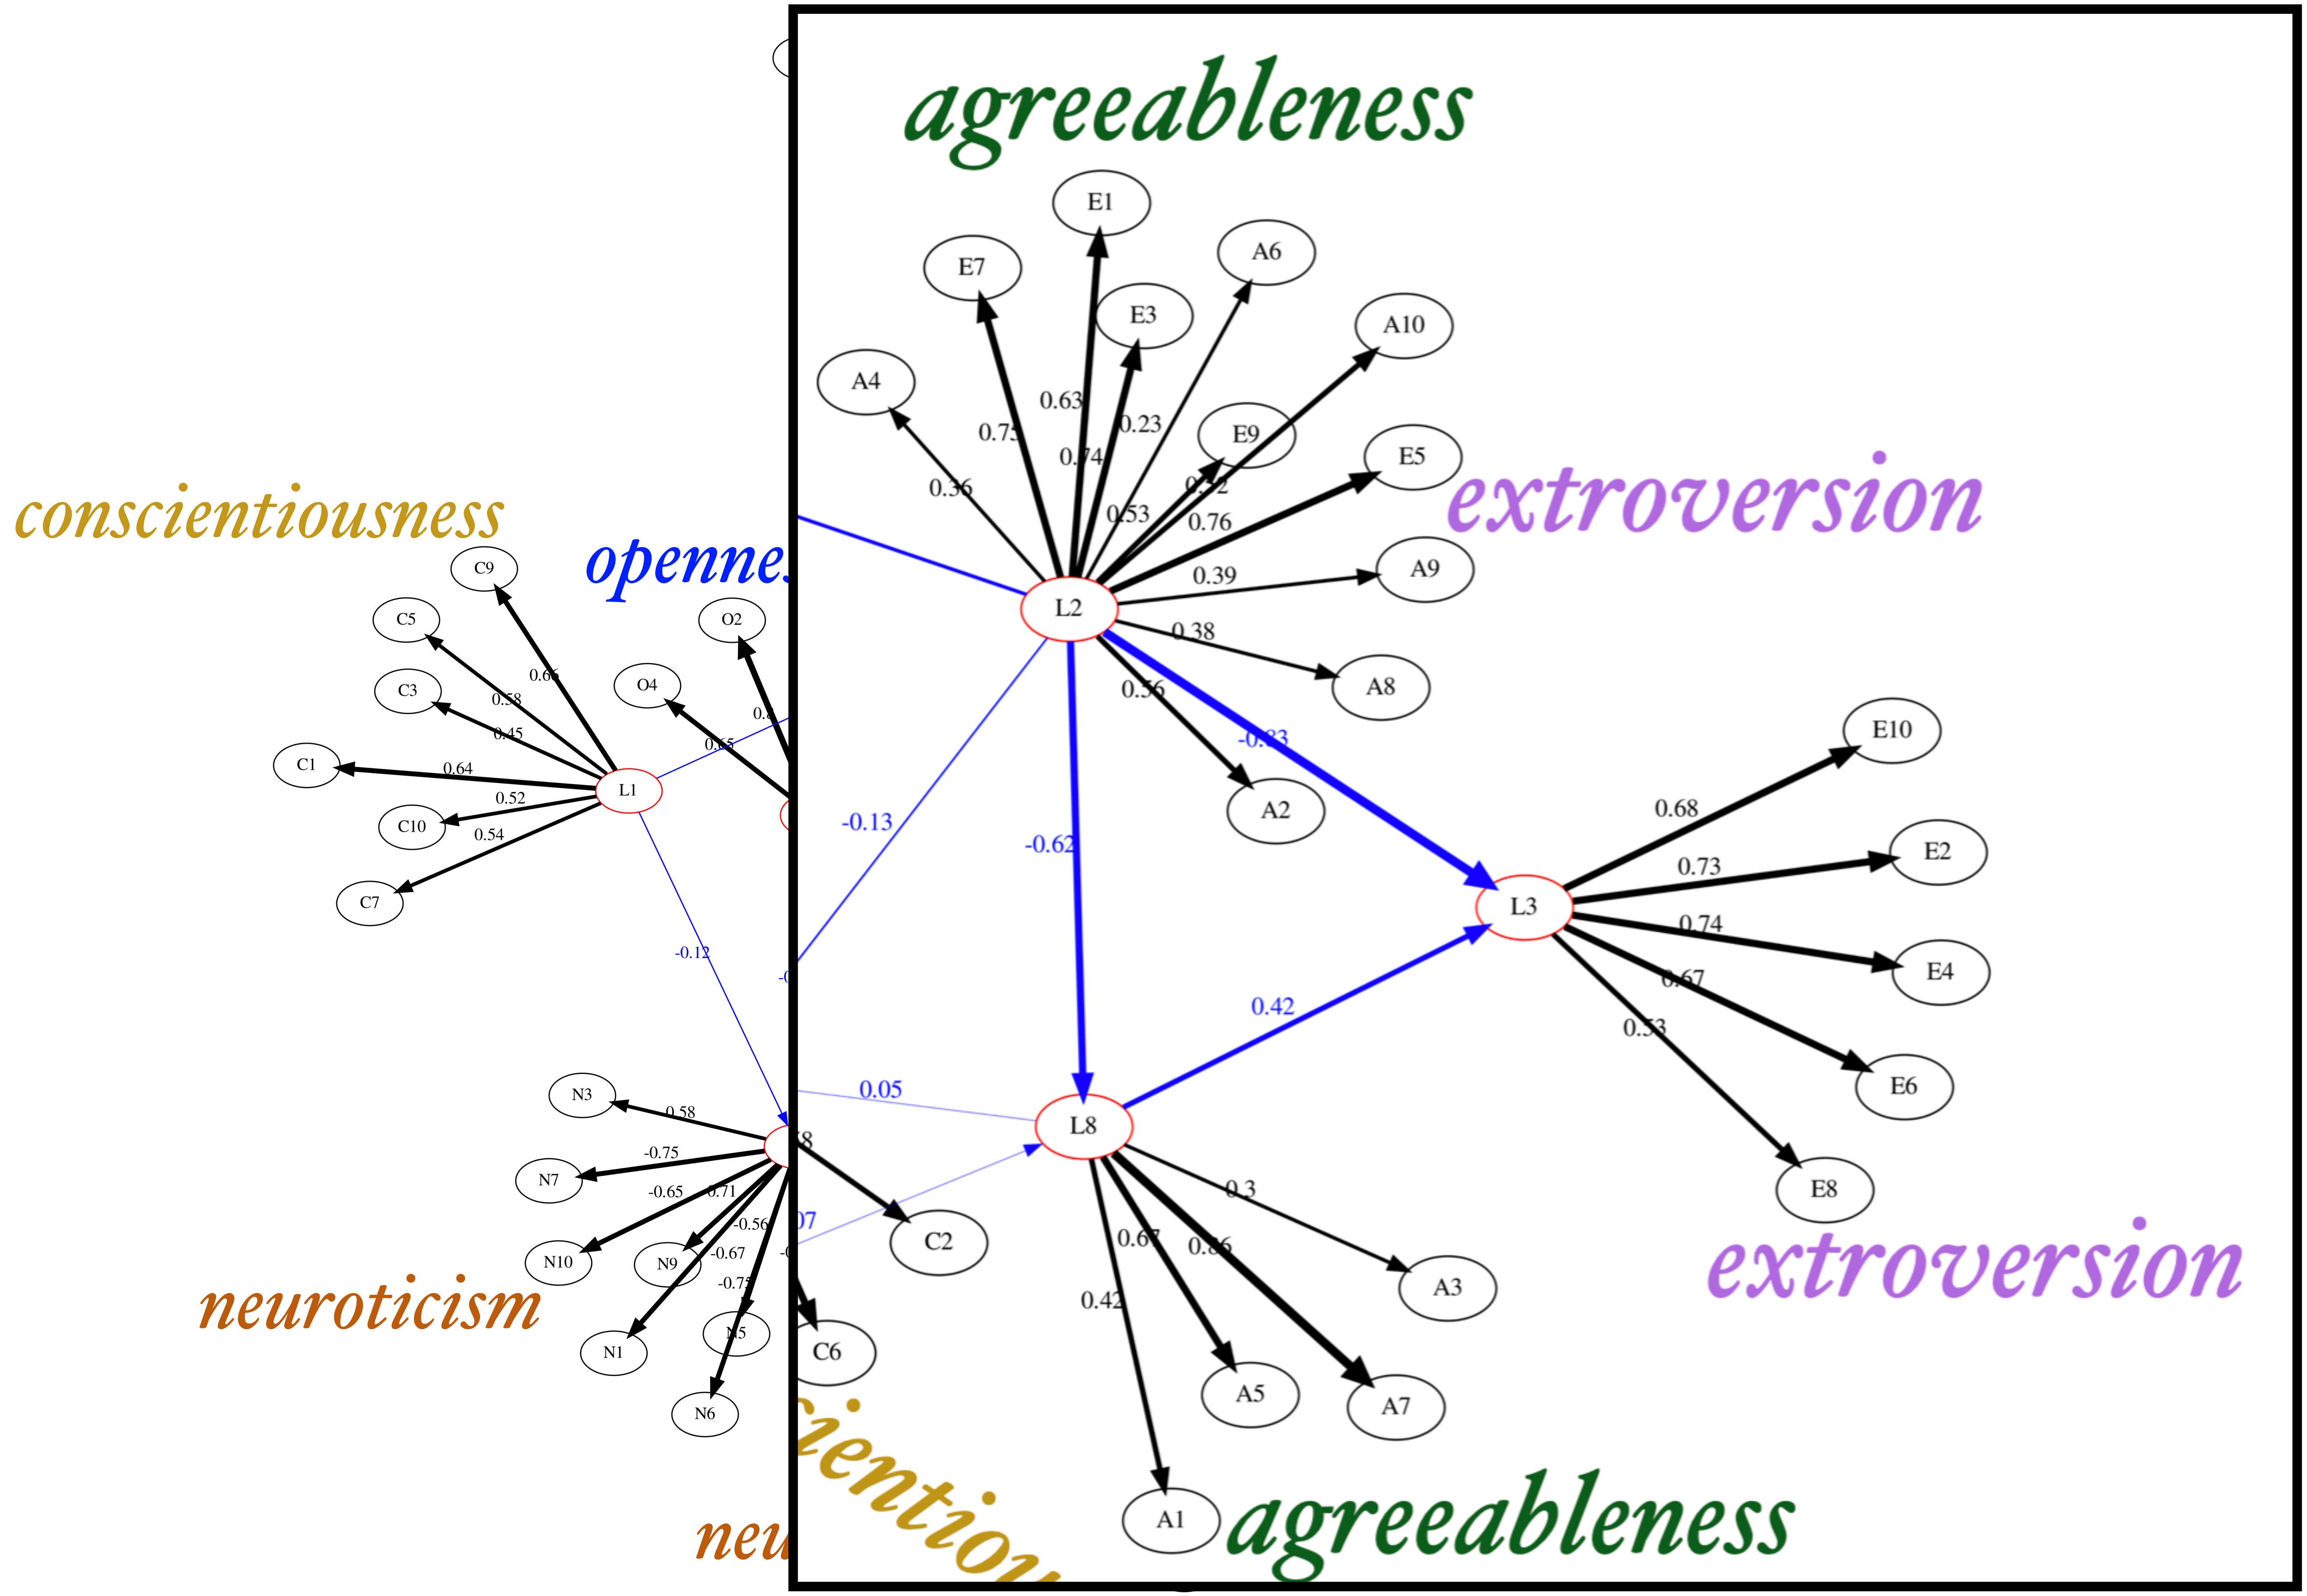

Learning Latent Causal Dynamics

i.i.d. data?	Parametric constraints?	Latent confounders?
Yes	No	No
No	Yes	Yes

Learn the underlying causal dynamics from their mixtures?
"Time-delayed" influence renders latent processes & their relations identifiable

Time-series Inputs $\{x_t\}_{t=0}^{T}$

Unsupervised Representation Learning

$$\mathbf{x}_t = \mathbf{g}(\mathbf{z}_t)$$

Latent processes

Latent temporal causal processes z_{it} can be recovered if they follow

- completely **nonparametric** model; or furthermore,
- **non-stationary** noise; or
- **non-stationary** causal influence, or
- **Parametric** constraints

Causal Skeleton Recovery

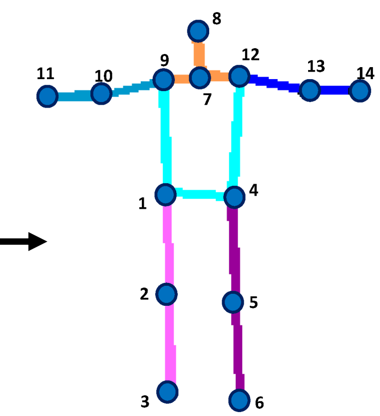

Recovered latent processes

- *Yao, Chen, Zhang, "Causal Disentanglement for Time Series," NeurIPS 2022*
- *Yao, Sun, Ho, Sun, Zhang, "Learning Temporally causal latent processes from general temporal data," ICLR 2022*

Results on Video Data

- For easy interpretation, consider two simple video data sets
  - KiTTiMask: a video dataset of binary pedestrian masks
  - Mass-spring system: a video dataset with ball movement and invisible springs

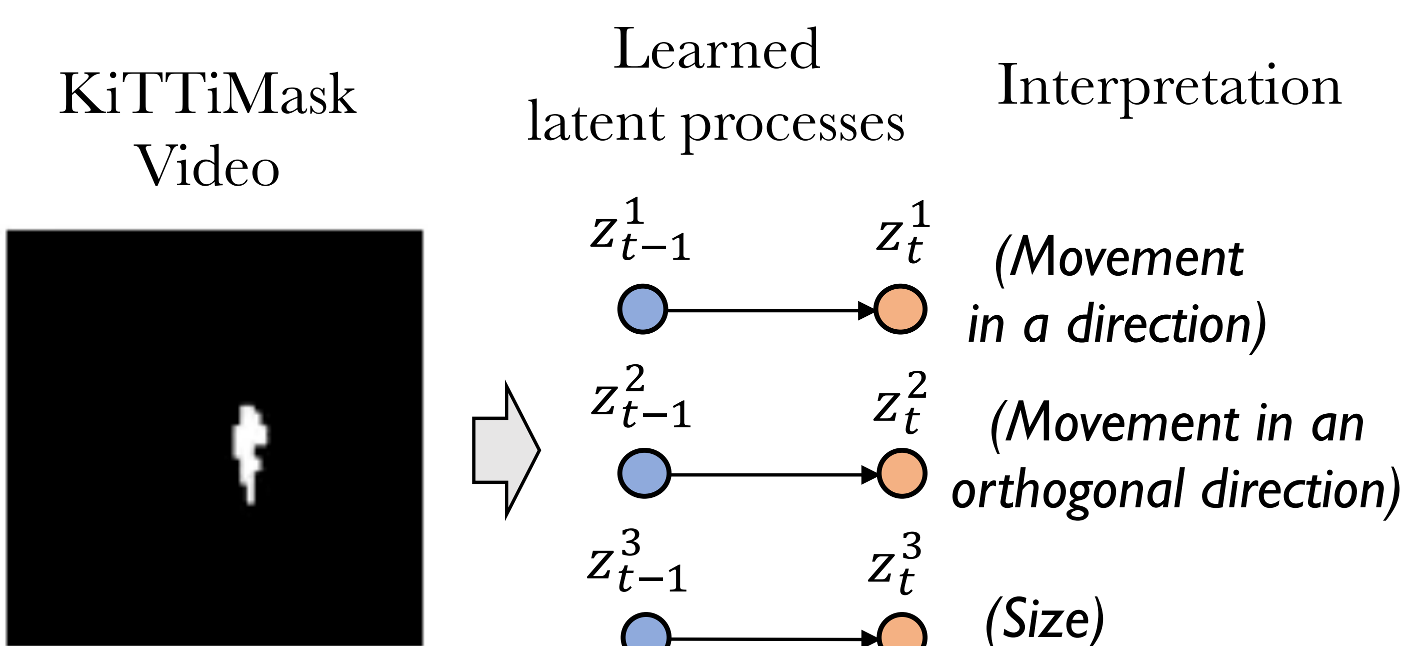

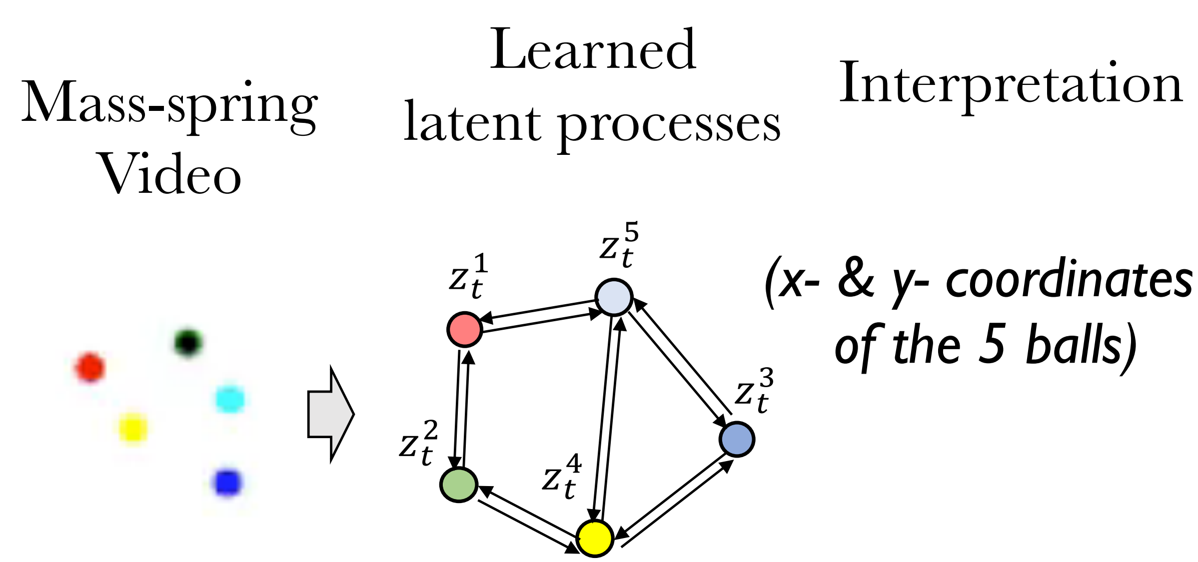

- *Yao, Chen, Zhang, "Learning Latent Causal Dynamics," NeurIPS 2022*
- *Yao, Sun, Ho, Sun, Zhang, "Learning Temporally causal latent processes from general temporal data," ICLR 2022*

Summary

- Definition of causality based on interventions
- Review of ML
- A number of ML or AI problems are related to causality
- What we can benefit from causal thinking
- Typical problems in modern causality research