COCOA tutorial handout

Bridget Copley

3 June 2021

Welcome to COCOA! Handouts, schedules, and a survey for the monthly zoominar are available at `https://oasis.cnrs.fr/cocoa`.

If you don't know anything/don't feel confident about causal models: During this tutorial and the Q&A session, feel free to ask questions in the chat. When I've gotten through this handout, you are welcome to ask your questions by voice or chat.

If you do know something/feel confident about causal models: During this tutorial and the Q&A session, feel free to use the chat to answer/clarify/add references/give nuance to anything. I will ask for this at certain points but feel free to chime in via chat at any point.

1 Introduction

1.1 What even is a causal model?

For our purposes, a causal model is a **formal representation** of the **structure** that **causal relations** give to our **conceptual model** of the world.

Basics (bedtime reading, sort of statisticky in parts): Pearl and Mackenzie (2018)

Basics (mathy): Pearl (2000); Halpern (2000); Halpern and Pearl (2005) ...

In linguistic semantics: Work since about 2010 by Rebekah Baglini, Elitzur Bar-Asher Siegal, Cleo Condoravdi, Bridget Copley, Sven Lauer, Fabienne Martin, Prerna Nadathur, ...

- Causal structures are formally represented by means of a *directed acyclic graph* (DAG).
- There is a set V of variables that are the vertices (or nodes) of the graph.
- These are connected by a set of edges (or arrows) E.

- The edges are *directed* and represent the dependency of one value on another. For instance, (A) → (B) represents that the value of (B) is dependent in some way on the value of (A).
- Absence of an edge between two variables means that the values are independent of each other.

(1) U: season ({spring, summer, winter, fall})
X: sprinkler ({on, off})
Z: rain ({yes, no})
Y: wet ({yes, no})
W: slippery ({yes, no})

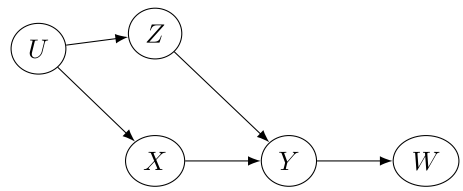

A plausible valuation for this model: U = summer, X = on, Z = no, Y = yes, W = yes.

U	X	Z	Y	W
summer	on	no	yes	yes
summer	off	no	no	no
summer	on	no	yes	yes
spring	off	yes	yes	yes
...	...	...	...	...

- Variables without arrows pointing at them are *exogenous* variables; their value depends only on circumstances that are not represented in the model (*background* variables, also *exogenous*). Variables with arrows pointing at them are *endogenous* variables.
- There is an asymmetry between exogenous variables and endogenous variables.
- Values of endogenous variables can be expressed by means of an equation over the variables that they depend on. When we do this, we can call the model a "structural equation model".

(2) An example from Halpern and Pearl (2005) with truth values

$F = 1$ if *There is fire* is true, $F = 0$ if it is false.

$M = 1$ if *The match is lit* is true, $M = 0$ if it is false.

$L = 1$ if *There is lightning* is true, $L = 0$ if it is false.

Function that takes L, M and yields value of F: $F = 1$ if either $L = 1$ or $M = 1$.

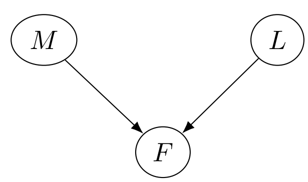

(3) Taxonomy of basic 3-node structures:

Collider Chain Fork

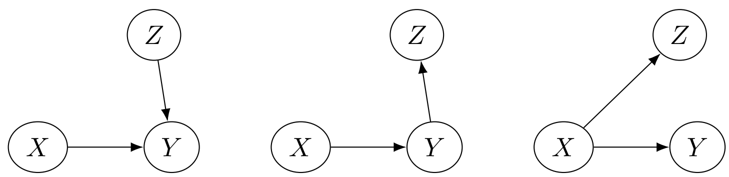

1.2 FAQs

How are the arrows to be read?

- The arrows are NOT to be read as CAUSE, *leads to*, material implication, or Talmian forces.
- The arrows ARE to be read as
  "the value of A affects the value of B",
  "the value of A influences the value of B", or "the value of B listens to the value of A".

Wait, don't causal models have to use probabilities and Bayes' Theorem?

- Short answer: No, they can use any value.
- Long answer: I'll talk more about this next week.

Which variable is THE cause?

- Short answer: Attend the next talk for the long answer!

How does one decide which variables to include?

- Short answer: Include everything that's relevant.
- Long answer: See next long answer.

How does one decide which arrows to include?

- Short answer: We have intuitions about where the causal relations are supposed to be.

- Long answer: We need to make sure that …

  … spurious dependencies do not get arrows (e.g. shoe size and reading ability)

  … if (A) $\rightarrow$ (B) $\rightarrow$ (C) is in our model, the value of (C) only listens to the value of (A) through the value of (B)

  … for any variable (X) in the structure, given (X)'s immediate causal ancestors, (X) is independent of its (other) non-descendants (Causal Markov condition)

  …

1.3 Where do causal models come from?

The short answer is that they come from the field of statistics in the early 20th century, with a huge ramp up of formalization and popularity in the last 30 years or so. See `https://plato.stanford.edu/entries/causal-models/`. Pearl and Mackenzie (2018) also has some history.

1.4 Why use causal models?

Pearl and Mackenzie (2018) Chapter 1 (`http://bayes.cs.ucla.edu/WHY/why-ch1.pdf`) in particular gives some context:

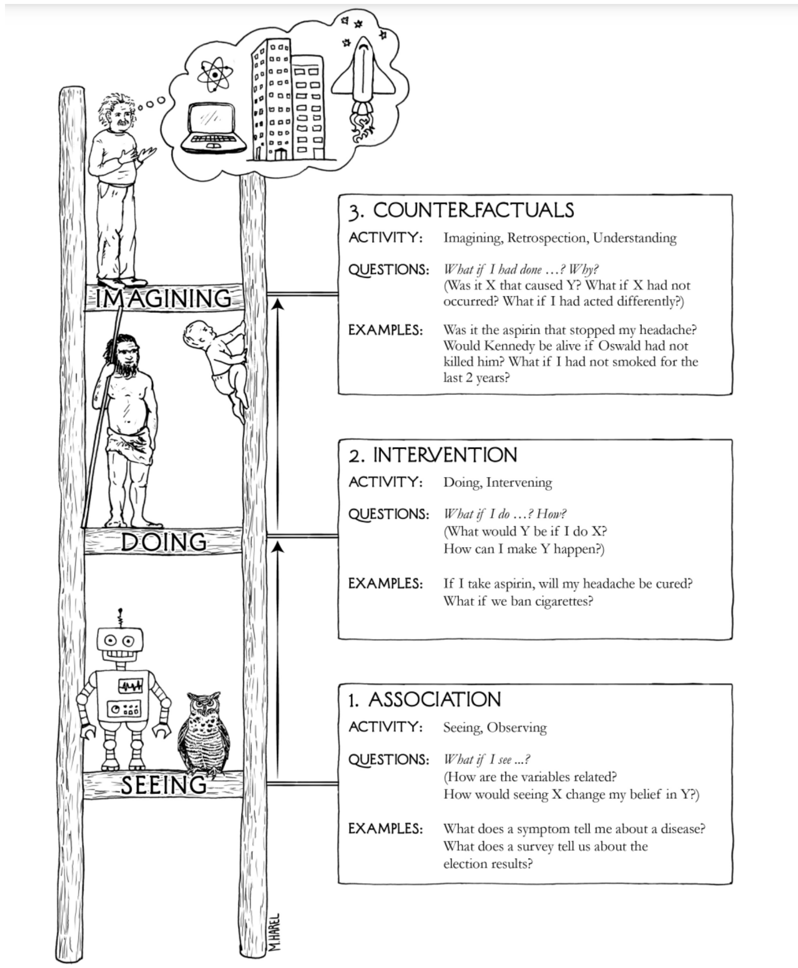

(4) Level 2: Intervention (Pearl (2000), see also Woodward (2006))
$do(X = x)$: Set the value of X to x and erase all arrows that point to X

(5) Level 3: Counterfactuals

a. Pearl (2000) uses probabilities for these, e.g. in a case where $X \rightarrow Y$ is in the model, and we know that $X = 1$, we want to know the likelihood of $Y = 1$ if we (or something) had intervened to set X to 1. $P(Y_{X=0} \mid X = 1, Y = 1)$

b. Kaufmann (2013): No need for probabilities if we use both Kratzer's causal premise semantics along with causal models. The premise background corresponding to the ordering source determines whether the value of Y is necessary or sufficient, given its parents, avoiding determinism. I want to take a closer look at this paper for a meeting in the fall semester.

2 How are causal models related to...

2.1 ...truth tables?

Traditional truth tables have Level 1 information only: **correlations** of the values, and both the **direction** (what values depend on what other values?) and **nature** (which relation/function is it?) of the dependency. The tables associated with causal models are similar, except that all of the above are assumed to be due to causality (kind of a Level 1.5?).

(6) Familiar truth tables:

And:

P	Q	$P \wedge Q$
1	1	1
1	0	0
0	1	0
0	0	0

Or:

P	Q	$P \vee Q$
1	1	1
1	0	1
0	1	1
0	0	0

Material conditional:

P	Q	$P \Rightarrow Q$
1	1	1
1	0	0
0	1	1
0	0	1

(7) Relations we will be dealing with:

Causal necessity

A	B
1	{(0),1}
0	0

Causal sufficiency

A	B
1	1
0	{0,(1)}

Stimulatory influence (+):

A	B
1	1
0	0

A	B
1	.1
2	.01
3	.001
4	.0001

Inhibitory influence (-):

A	B
1	0
0	1

A	B
.0006	55,679
.0008	16,450
.0010	9611
.0012	5001

2.2 ...event arguments?

Statisticians who use causal models talk about $X = x$ as an "event". We may or may not choose to represent a Davidsonian event as a node.

(8) a. $\text{I} \rightarrow \text{R}$
b. $\text{I} \rightarrow \text{E} \rightarrow \text{R}$

2.3 ...times/situations?

(9) Two ways of looking at a timeline

a. The variables are relativized to times (Halpern and Pearl, 2005, 18): $X_{i_1}, X_{i_2}, X_{i_3}$ ...

b. Variables have values, and these values are relativized to times. The valuation function (call it $\mathcal{R}$) that could be redefined to take a time argument in addition to the variable argument. So, $\mathcal{R}(X)(i) = x$, or in a more semantics-friendly notation, $\mathcal{R}(X) = \lambda i.\ [\![\text{p}]\!](i)$.

2.4 ...possible worlds?

Kaufmann (2013):

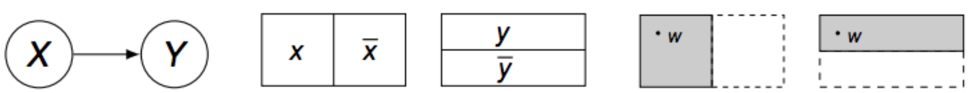

Causal Networks	Possible worlds
Outcome	Possible world
Event	Proposition
Variable	Partition

From there we have a choice: What do we need from the familiar Kratzer-style tools (atomic possible worlds, sets of atomic possible worlds, orderings on sets of atomic possible worlds)? And if we don't use them, do we need to use probabilities?

2.5 What can causal models do for semantics?

- We get the benefits of explicitly talking about causal relations, even when there are multiple influences and/or effects

Quantification over atomic possible worlds IS TO explicit causal relations
AS
Optimality Theory IS TO transformational phonology

- We get the benefits of talking about causation without talking about spatiotemporal locality
- We get the benefits of assuming a closed world
- We can easily represent goal-directed action
- We get the benefits of talking easily about causal powers (e.g. intentions, abilities, dispositions)

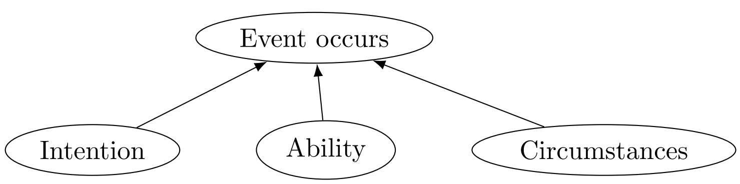

3 Appendix: A formal setup

From Halpern and Pearl (2005):

(10) A *signature* $\mathcal{S}$ is a tuple $(\mathcal{U}, \mathcal{V}, \mathcal{R})$ where $\mathcal{U}$ is a set of endogenous variables, $\mathcal{V}$ is a set of exogenous variables, and $\mathcal{R}$ associates with every variable $Y \in \mathcal{U} \cup \mathcal{V}$ a nonempty set $\mathcal{R}(\mathcal{Y})$ of possible values for Y.

(11) A *causal model* (or *structural model*) is a tuple $M = (\mathcal{S}, \mathcal{F})$, where $\mathcal{F}$ associates with each variable $X \in \mathcal{V}$ a function denoted F_X such that: $F_X = (\times_{U \in \mathcal{U}} \mathcal{R}(U)) \times (\times_{Y \in \mathcal{V}-X} \mathcal{R}(Y)) \rightarrow \mathcal{R}(X)$

a. In words: F_X is a function that determines the value of X given the values of all the other variables.

b. Notation: The set of all ordered pairs (a, b), where a is an element of A and b is an element of B, is called the Cartesian product of A and B and is denoted by $A \times B$.

c. More notation: $(\times_{U \in \mathcal{U}} \mathcal{R}(U))$ is the ordered n-tuple of all the values of all the variables in $\mathcal{U}$.

d. Exercise: How might we describe F_X in words, using "the arrows that..."?

(12) Example: if $F_X(Y, Z, U) = Y + U$, then if $Y = 3$ and $U = 2$, then $X = 5$ no matter how Z is set. (As shorthand, we can also write this equation as: $X = Y + U$) The function $\mathcal{F}$ defines a set of these, called *structural equations.*

References

Halpern, J. Y. (2000). Axiomatizing causal reasoning. *Journal of Artificial Intelligence Research 12*, 317–337.

Halpern, J. Y. and J. Pearl (2005). Causes and explanations: A structural-model approach. part i: Causes. *The British journal for the philosophy of science 56*(4), 843–887.

Kaufmann, S. (2013). Causal premise semantics. *Cognitive science 37*(6), 1136–1170.

Pearl, J. (2000). *Causality: Models, reasoning and inference.* Cambridge University Press.

Pearl, J. and D. Mackenzie (2018). *The book of why: the new science of cause and effect.* Basic Books.

Woodward, J. (2006). Sensitive and insensitive causation. *The Philosophical Review 115*(1), 1–50.