

Causality = time + modality + effective difference-making

Dean McHugh

Lira Seminar
Institute of Logic, Language and Computation
University of Amsterdam

5 March 2020

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

Two questions

- ① What kind of information do we use when we judge that a causal relation holds?
- ② What are the truth conditions of causal claims?

Causality = time + modality + effective difference-making

Today



Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

Plan

- 1 Graphical models of causation
 - **Primer on structural causal models**
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

Structural causal models

Definition (Structural causal model)

A structural causal model is a triple $M = (V, E, F)$ where

- V is a set of variables
- (V, E) is a directed acyclic graph
- F is a set of functions of the form

$$F_X : \mathcal{R}(pa_X) \rightarrow \mathcal{R}(X),$$

one for each endogenous (i.e. has a parent) variable $X \in V$,
where $pa_X := \{Y \in V : (Y, X) \in E\}$

Structural causal models

The value of an endogenous variable X is determined by the values of its parents, according to F_X

- Since F_X are **functions**, the dependence is **deterministic**
- Where $U = u$ is an assignment of values to the exogenous variables in V , we call u a *setting* or *context* for M
 - i.e. the values of the exogenous variables determine the values of all the variables

Interventions in structural causal models

Let $M = (V, E, F)$ be a structural causal model

Definition (Interventions as model surgery)

$M_{do(X=x)}$ is the model $(V_{do(X=x)}, E_{do(X=x)}, F_{X=x})$ where

- $V_{do(X=x)} = V$
- $E_{do(X=x)} = E \setminus \{(Y, X) : Y \in V\}$
- $F_{do(X=x), X}(u) = x$ for every setting u of M ,
and $F_{do(X=x), Y} = F_Y$ for all $Y \in V, Y \neq X$

Definition (Truth conditions for interventions)

Let M be a structural causal model and u a setting of the exogenous variables.

$$M, u \models [X \leftarrow x]Y = y \quad \text{iff} \quad M_{do(X=x)}, u \models Y = y$$

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - **Directness and density**
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

SCMs represent dependence as direct

In causal diagrams, an arrow represents a “direct effect” of the parent on the child, although this effect is direct only relative to a certain level of abstraction, in that the graph omits any variables that might mediate the effect.

— Greenland and Pearl (2011, pp. 208–09)

In the beginning, dependence was direct

CORRELATION AND CAUSATION

By SEWALL WRIGHT

Senior Animal Husbandman in Animal Genetics, Bureau of Animal Industry, United States Department of Agriculture

PART I. METHOD OF PATH COEFFICIENTS

INTRODUCTION

The ideal method of science is the study of the direct influence of one condition on another in experiments in which all other possible causes of variation are eliminated. Unfortunately, causes of variation often seem to be beyond control. In the biological sciences, especially, one

Figure: Wright (1921)

Dense causal chains

In daily life and physicists' models alike,
space and time are represented as *dense*

Density Between any two points there is a third

Intuitive belief: Causal influence travels through dense spacetime

There are chains of events $(C_i)_{i \in I}$ where,
for every C_i, C_k on the chain, there is a C_j on the chain such that

C_i causally influenced C_j and C_j causally influenced C_k .



Is density compatible with direct dependence?

Question Can dense dependence be represented in terms of direct dependence?

- Intuitively, dense dependence is indirect
- One can add more arrows
- But still, more arrows just means more instances of direct dependence



[Causation, Coherence, and Concepts](#) pp 99-111 | [Cite as](#)

Bayesian Nets Are All There Is to Causal Dependence

Chapter

510

Downloads

Part of the [Boston Studies in the Philosophy of Science](#) book series (BSPS, volume 256)

Wolfgang Spohn (2009)

A simple dense causal chain




Figure: Turning on the light at time t_0

How do structural causal models represent dependence here?

Answer: A version of counterfactual dependence

(Paul, 1998; Yablo, 2002; Halpern and Pearl, 2005; Halpern, 2016; Beckers, 2016)

- Let x be a point illuminated at t , and $t' = t + \frac{|y-x|}{c}$
- If x had not been illuminated at t , y would not be illuminated at t'

Formalizing dense dependence

Intuitively, dense causal chains exhibit *dense dependence*

Definition (Dense dependence)

$M = (V, E, F)$ features *dense dependence* at $Z \in V$ iff for any parent X of Z there is a parent Y of Z distinct from X such that

$$M, u \models [x, y]z \quad \text{iff} \quad M, u \models [x', y]z$$

for any setting u of the exogenous variables in M , and any values x, x' of X , value y of Y and z of Z .

Equivalently,

$$F_Z(x, y, o) = F_Z(x', y, o)$$

for any $x, x' \in \mathcal{R}(X)$, $y \in \mathcal{R}(Y)$ and $o \in \mathcal{R}(pa_Z \setminus \{X, Y\})$

Dummy parents

Definition (Dummy parent)

X is a dummy parent of Z in M iff $F_Z(x, o) = F_Z(x', o)$ for any $x \in \mathcal{R}(X)$ and $o \in \mathcal{R}(pa_Z \setminus \{X\})$.

Lemma

If every parent of Z is a dummy parent of Z , then F_Z is constant.

Proposition

If dependence in M is dense at Z then F_Z is a constant function.

Proof.

If dependence in M is dense at Z then for any parent X of Z , there is a parent Y of Z with

$$F_Z(x, y, o) = F_Z(x', y, o)$$

for any $(y, o) \in \mathcal{R}(pa_Z \setminus \{X\})$. Then X is a dummy parent of Z , and as X was arbitrary, every parent of Z is a dummy parent of Z . Hence F_Z is constant. \square

Intuitively, if F_Z is constant then Z does not depend on any variable for its value

Corollary

There is no structural causal model M with a variable Z such that:

- *Dependence in M is dense at Z*
- *Z depends on some variable(s) for its value*

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - **Determinacy**
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways


Counterfactual dependence and determinancy

A popular principle:

(Halpern and Pearl, 2005; Halpern, 2016; Beckers, 2016; Beckers and Vennekens, 2018)

Counterfactual dependence implies actual causation

If C had not occurred, E would not have occurred
 C is an actual cause of E



The robot has to get anywhere on Main Street, choosing between the four ways at random. The robot took 1st Street and then Road B.

(1) The fact that the robot took 1st Street caused it to take Road B

To many, (1) is unacceptable

Sufficiency in our intuitive talk about causation

It is often said that causes...

- ‘make’ their effects happen
- ‘produce’ them
- ‘generate’ them
- ‘bring’ them about

But taking 1st Street did not ‘make’ the robot take Road B

- Many analyses of causal claims are given in terms of SCMs
- Surprisingly, sufficiency does not play a role in these theories (Halpern, 2016; Beckers, 2016)
- Though sufficiency is part of some approaches (e.g. Mackie, 1965; Baldwin and Neufeld, 2004; Braham and van Hees, 2012)

Testing the explanation

(1) The fact that the robot took 1st Street caused it to take Road B

Hypothesis 1

(1) is unacceptable because the robot taking 1st Street was not **sufficient** for it to take Road B

Causation without sufficiency

- (2) The fact that the switch was flicked caused the light to turn on.



- Flicking the switch was not enough, by itself, to make the light turn on
 - There needed to be power in the building, a working wire from the switch to the light, etc.
- (2) is acceptable nonetheless
- **Intuitive explanation:** (2) is acceptable because we ‘assumed’ (‘took for granted’) facts which guarantee that, after flicking the switch, the light will turn on

Hypothesis 2


(1) is unacceptable because there was no set of facts such that, within that set, the robot taking 1st Street was **sufficient** for it to take Road B

Let's test this:

- Change the scenario to provide an enabling condition
- Check whether our judgements also change

Testing the explanation

Program the robot with the rule: Always change direction!
e.g. if the robot turns left, it must then turn right



Sufficiency: analysis

- Structural causal models are deterministic
- So they are forced to add enabling conditions to represent scenarios that we can intuitively represent as non-deterministic
- In our actual causal reasoning:
 - Sometimes enabling conditions are present
 - Sometimes they are not

The need for a uniform model of causality

It would be remarkably surprising if we employed

- One kind of model of causation for discrete systems, and another for dense systems
- One kind of model of causation of deterministic systems, and another kind for non-deterministic systems

To do: Create a model of causal reasoning that:

- 1 Allows for dense causal chains
- 2 Can provide non-deterministic representations of scenarios

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

The essential ingredients of causality

What kind of information do we use
when we judge that a causal relation holds?

In other words, what should causal models represent?
What kind of information should they contain?

Causal models, like all models, are representations
As such, we need some semantic objects to do the representing
A semantic object is a point at which sentences are evaluated

Ingredient 1

Semantic objects (states, possible worlds, situations, ...)

There is no causality without time.

Ingredient 2
Time

How should we represent time?

We will adopt the view from classical mechanics and everyday life: time is represented by a **linear order over states**

(see Maudlin 2014, *New Foundations for Physical Geometry*)

We index states to allow repetitions:

Definition (Token state)

A *token state* is a state with an index.

Definition (Path)

A *path* is a linearly ordered set of token states.


i.e. where S is a set of states, a path is a linearly ordered subset of $\{s_i : s \in S, i < \#(s)\}$, where $\# : S \rightarrow \text{card}$ assigns to each state a cardinal

Causal reasoning is a form of hypothetical reasoning, one involving counterfactual thinking (see e.g. Byrne 2007) (e.g. “if the cause had not occurred, ...”)

Evidence: The problem of causal inference


If causal reasoning was about actuality alone, inferring causation would not be as hard as it is

Is the actual world enough?



The New York Times

***Pushing That Crosswalk Button
May Make You Feel Better, but ...***



The actual world is part of two sets that disagree on causation

Ingredient 3

Multiple possibilities

states	+	time	\Rightarrow	path
path	+	multiple possibilities	\Rightarrow	set of paths

Definition (Dynamical system)

A dynamical system is a set of paths.

Hypothesis Dynamical systems provide an adequate model of causal reasoning

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems**
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

Modelling dense dependence




Figure: Turning on the light at time t_0

- The definition of a dynamical system allows for states to be densely ordered
- Consider the set of paths where, for any points x, y in space and any time t , if x is light at t then y is light at $t + \frac{|y-x|}{c}$

In dynamical systems, one can model the fact that


- For any points in space x, z , if x has a causal influence on z , then there is a point y such that, intervening on y , x no longer has a causal influence on z
- (Note: this requires a formalization of the truth conditions of causal claims, including a treatment of intervention)

Plan


- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems**
 - Dense dependence in dynamical systems
 - **Non-determinism in dynamical systems**
- 4 From graphical models to dynamical systems
- 5 Take aways

Modelling the robot scenario

- (1) The fact that the robot took 1st Street caused it to take Road B
- The robot's actions can be represented as **random** or as **deterministic**
 - **Random** After taking 1st Street, the robot could have taken Road A and could have taken Road B
 - **Deterministic** The robot has some mechanism to decide
 - We can represent both interpretations:



(a) Nondeterministic



(b) Deterministic

Figure: Models of the robot's action

Free versus bound variables

Should the enabling conditions be represented by free variable (i.e. contextually determined) or by a bound variable?

- 'Assuming', 'Taking something for granted' are **epistemic states**
- If one uses interpreter-dependent notions in the analysis of causal claims, it is hard to give objective answers to concrete questions of causation
 - Subjective parameters leave causal reasoning open to manipulation
- Perhaps there is an interpreter-independent characterisation
 - The existence of a proposition

Formalizing contextual sufficiency

Causes are contextually sufficient for their effects

If C caused E at w , then there is a set of worlds A such that

- $w \in A$, and
- For every world v in A and state s on v , if C is true at s then
 - E occurs on v at some state later than s , and
 - For all worlds u that agree with v up to and including state s , $u \in A$ (A is “closed under open futures”)

Scenario 1

- The robot chooses which street to take at random
- A must include the actual world (taking Road B)
- Then by closure under open futures, A must also include the world where the robot takes Road A

⇒ Taking 1st St did **not** cause the robot to take Road B

Scenario 2

- The robot must change direction at every turn
- A includes the fact that the robot changes direction

⇒ Taking 1st St **did** cause the robot to take Road B

Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

What **in the world** is a Bayesian network?

What must the world be like for X to be a parent of Y ?

Bayes nets

Definition (Representing a graph)

Let (V, E) be an acyclic graph and $P : \mathcal{R}(V) \rightarrow [0, 1]$ a probability distribution. We say P represents (V, E) just in case, for any $X \in V, x \in \mathcal{R}(X)$,

$$P(x \mid pa_X) = P(x \mid pa_X, nd_X)$$

Definition

A Bayes net is a triple $B = (V, E, P)$, where (V, E) is an acyclic graph and $P : \mathcal{R}(V) \rightarrow [0, 1]$ is a probability distribution representing (V, E) .

The chain rule in Bayes nets

- Bayes nets generalise the chain rule from linear to acyclic graphs

Chain rule

$$\begin{aligned}P(x_1, \dots, x_n) &= P(x_n \mid x_1, \dots, x_{n-1})P(x_1, \dots, x_{n-1}) \\&= \prod_{i=1}^n P(x_i \mid x_1, \dots, x_{i-1})\end{aligned}$$

An observation

The chain rule w.r.t. linear orders is a special case of the rule

$$P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i \mid pa_{X_i});$$

namely, when pa_{X_i} contains at most one element.

Let $E = \{(X_i, X_{i+1})\}_{0 \leq i < n}$ and suppose P represents (V, E) .

$$P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i \mid pa_{X_i}) \quad (\text{Rule above})$$

$$= \prod_{i=1}^n P(x_n \mid x_{n-1}) \quad (\text{Definition of } E)$$

$$= \prod_{i=1}^n P(x_n \mid x_1, \dots, x_{n-1}) \quad (P \text{ represents } (V, E))$$

Two chain rules

Let $S = \mathcal{R}(V)$ and $s, t \in S$

- Order the variables X_1, \dots, X_n
- Let $s(X_i)$ be the value of X_i at s
- $P_s(t)$ is the probability that t is the next state, given that s is the current state


Static chain rule

$$P(s) = \prod_{i=1}^n P(\textcolor{red}{s}(X_i) \mid s(X_1), \dots, s(X_{i-1}))$$

Dynamic chain rule

$$P_s(t) := \prod_{i=1}^n P(\textcolor{red}{t}(X_i) \mid s(X_1), \dots, s(X_{i-1}))$$

Example



Dynamic probability

Lifting from states to sets of states. Where $A, B \subseteq \mathcal{R}(V)$,

$$P_A(B) := \sum_{s \in A} \left(P_s(s) \sum_{t \in B} P_s(t) \right)$$

- $P_A(B)$ is the probability that the **next** state is a B -state, given that the current state is an A -state
- $P(B \mid A)$ is the probability that the **current** state is a B -state, given that the current state is an A -state

Theorem

For any Bayes net (V, E, P) and $s \in \mathcal{R}(V)$, P_s is a probability distribution.

Proof.

By induction on $|V|$.



Plan

- 1 Graphical models of causation
 - Primer on structural causal models
 - Directness and density
 - Determinacy
- 2 The essential ingredients of causality
- 3 Causal dependence in dynamical systems
 - Dense dependence in dynamical systems
 - Non-determinism in dynamical systems
- 4 From graphical models to dynamical systems
- 5 Take aways

Take aways

- Structural causal models and Bayesian networks offer a marvellously user-friendly way to represent causal dependence
- But they have their expressive limits
 - Dense dependence
 - Scenarios we represent as non-deterministic
- Thankfully, there is an alternative: dynamical systems
 - Dynamical systems can represent dense dependence and non-determinism
 - Every structural causal model and Bayesian net can be represented as a dynamical system

References I

- Richard A. Baldwin and Eric Neufeld. The structural model interpretation of the ness test. In Ahmed Y. Tawfik and Scott D. Goodwin, editors, *Advances in Artificial Intelligence*, pages 292–307. Springer, 2004. doi:10.1007/978-3-540-24840-8_21.
- Sander Beckers. *Actual Causation: Definitions and Principles*. PhD thesis, PhD thesis, KU Leuven, 2016. URL https://limo.libis.be/primo-explore/fulldisplay?docid=LIRIAS1656621&context=L&vid=Lirias&search_scope=Lirias&tab=default_tab&lang=en_US.
- Sander Beckers and Joost Vennekens. A principled approach to defining actual causation. *Synthese*, 195(2):835–862, Feb 2018. doi:10.1007/s11229-016-1247-1.
- Matthew Braham and Martin van Hees. An anatomy of moral responsibility. *Mind*, 121 (483):601–634, 7 2012. doi:10.1093/mind/fzs081.
- Ruth MJ Byrne. *The Rational Imagination: How people create alternatives to reality*. MIT press, 2007. URL <https://mitpress.mit.edu/books/rational-imagination>.
- Sander Greenland and Judea Pearl. *Causal Diagrams*, pages 208–216. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-04898-2_162.
- Joseph Y Halpern. *Actual Causality*. MIT Press, 2016.
- Joseph Y Halpern and Judea Pearl. Causes and explanations: A structural-model approach. part i: Causes. *The British journal for the philosophy of science*, 56(4): 843–887, 2005.

References II

- John L Mackie. Causes and conditions. *American philosophical quarterly*, 2(4): 245–264, 1965.
- Tim Maudlin. *New foundations for physical geometry: the theory of linear structures*. Oxford University Press, 2014.
- L. A. Paul. Keeping track of the time: Emending the counterfactual analysis of causation. *Analysis*, 58(3):191–198, 1998. doi:10.1111/1467-8284.00121.
- Wolfgang Spohn, editor. *Bayesian Nets Are All There Is to Causal Dependence*, pages 99–111. Springer, 2009. doi:10.1007/978-1-4020-5474-7_4.
- Sewall Wright. Correlation and causation. *Journal of Agricultural Research*, 20(7): 557–580, 1921.
- Stephen Yablo. De facto dependence. *The Journal of philosophy*, 99(3):130–148, 2002. doi:10.2307/3655640.