Causality, interventions and counterfactuals in Structural Causal Models

Simon Fortier-Garceau

June 17, 2021

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

- It reflects our understanding that the natural world, and certain systems, operate under specific rules that follow a certain order.

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

- It reflects our understanding that the natural world, and certain systems, operate under specific rules that follow a certain order.
- There are various degrees at which we can exploit causal inference, depending on the strength of our model.

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

- It reflects our understanding that the natural world, and certain systems, operate under specific rules that follow a certain order.
- There are various degrees at which we can exploit causal inference, depending on the strength of our model.

For example, the statement "The event of missing the bus this morning is the cause of me arriving late at work." underlines several causal notions:

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

- It reflects our understanding that the natural world, and certain systems, operate under specific rules that follow a certain order.
- There are various degrees at which we can exploit causal inference, depending on the strength of our model.

For example, the statement "The event of missing the bus this morning is the cause of me arriving late at work." underlines several causal notions:

- **Causal sufficiency:** "missing the bus $\implies$ arriving late at work"

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

- It reflects our understanding that the natural world, and certain systems, operate under specific rules that follow a certain order.
- There are various degrees at which we can exploit causal inference, depending on the strength of our model.

For example, the statement "The event of missing the bus this morning is the cause of me arriving late at work." underlines several causal notions:

- **Causal sufficiency:** "missing the bus $\implies$ arriving late at work"
- **Causal necessity?:** "arriving late at work $\implies$ missing the bus"

Introduction to causality

Causality arises as a sense or attributed measure of responsibility of an event (a cause) in the occurrence of another (an effect), or in between the entities that drive those events (variables).

- It reflects our understanding that the natural world, and certain systems, operate under specific rules that follow a certain order.
- There are various degrees at which we can exploit causal inference, depending on the strength of our model.

For example, the statement "The event of missing the bus this morning is the cause of me arriving late at work." underlines several causal notions:

- **Causal sufficiency:** "missing the bus $\implies$ arriving late at work"
- **Causal necessity?:** "arriving late at work $\implies$ missing the bus"
- But arriving late at work cannot "cause" missing the bus, because arriving late at work can only happen **after** the event of missing the bus here.

Models of Evolving Systems and Causality

- In general, causal relations appear as an attribute of some underlying universe of events, ordered in the manner in which they can occur, an evolving system of sort.

Models of Evolving Systems and Causality

- In general, causal relations appear as an attribute of some underlying universe of events, ordered in the manner in which they can occur, an evolving system of sort.
- We think of an **evolving system** as a collection of interconnected events (or actions) that describe how that system changes along the order of time.

  Examples: Event structures, Bayesian networks, quantum systems, Markov chains, structural causal models, register machines, Petri nets, presheaves of labelled transition systems, presheaves on a directed space, etc.

Models of Evolving Systems and Causality

- In general, causal relations appear as an attribute of some underlying universe of events, ordered in the manner in which they can occur, an evolving system of sort.
- We think of an **evolving system** as a collection of interconnected events (or actions) that describe how that system changes along the order of time.

  Examples: Event structures, Bayesian networks, quantum systems, Markov chains, structural causal models, register machines, Petri nets, presheaves of labelled transition systems, presheaves on a directed space, etc.
- Causality manifests as various **cause-effect** relationships on events along the order of the system, where an effect is a result of or is influenced in an instrumental manner by the events that cause it. All of the above examples of evolving systems embody one or several forms of causal relations.

Models of Evolving Systems and Causality

- In general, causal relations appear as an attribute of some underlying universe of events, ordered in the manner in which they can occur, an evolving system of sort.
- We think of an **evolving system** as a collection of interconnected events (or actions) that describe how that system changes along the order of time.

  Examples: Event structures, Bayesian networks, quantum systems, Markov chains, structural causal models, register machines, Petri nets, presheaves of labelled transition systems, presheaves on a directed space, etc.
- Causality manifests as various **cause-effect** relationships on events along the order of the system, where an effect is a result of or is influenced in an instrumental manner by the events that cause it. All of the above examples of evolving systems embody one or several forms of causal relations.
- **Question:** What are these different types of cause-effect relationships, and which ones are essential to general causal models? How do we classify the strength of a causal model?

Causality in Probability and Statistics

The most significant parts of the theory of causal inference were developed within the context of probability and statistics; emphasizing that causation does not reduce to correlation.

Causality in Probability and Statistics

The most significant parts of the theory of causal inference were developed within the context of probability and statistics; emphasizing that causation does not reduce to correlation.

- Different schools of thought exist and debate as to what constitutes a proper foundation for causality (example: are deterministic mechanisms required at the base of a causal model or not?).

Causality in Probability and Statistics

The most significant parts of the theory of causal inference were developed within the context of probability and statistics; emphasizing that causation does not reduce to correlation.

- Different schools of thought exist and debate as to what constitutes a proper foundation for causality (example: are deterministic mechanisms required at the base of a causal model or not?).
- The most common assumptions for an order on variables to represent a causal order are **the Markov Condition**, **the Minimality Condition**, and **the Faithfulness Condition**. (P. Spirtes, C. Glymour, and R. Scheines, 1993).

Causality in Probability and Statistics

The most significant parts of the theory of causal inference were developed within the context of probability and statistics; emphasizing that causation does not reduce to correlation.

- Different schools of thought exist and debate as to what constitutes a proper foundation for causality (example: are deterministic mechanisms required at the base of a causal model or not?).
- The most common assumptions for an order on variables to represent a causal order are **the Markov Condition**, **the Minimality Condition**, and **the Faithfulness Condition**. (P. Spirtes, C. Glymour, and R. Scheines, 1993).
- But we will focus on more recent trends of causal inference research that follow Pearl's ladder of causation, in which **interventions** and **counterfactuals** are used as a basis of causal discourse.

Pearl's Ladder of Causation

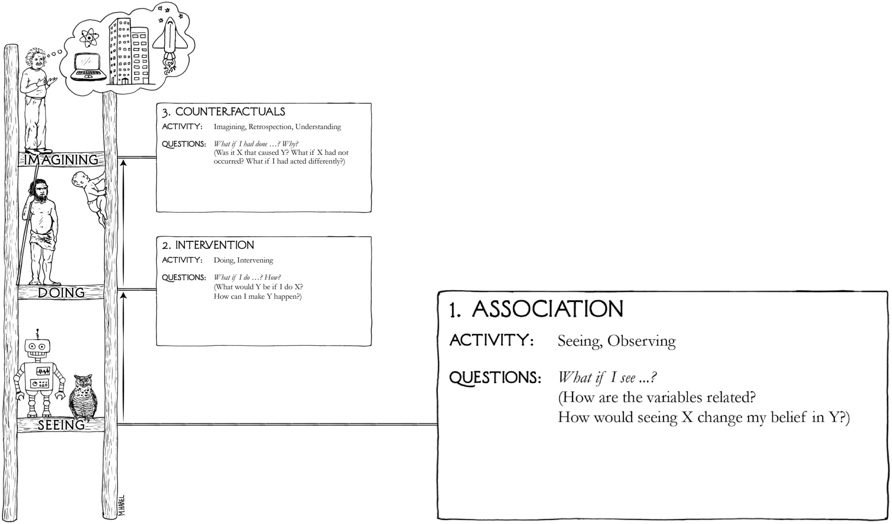

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

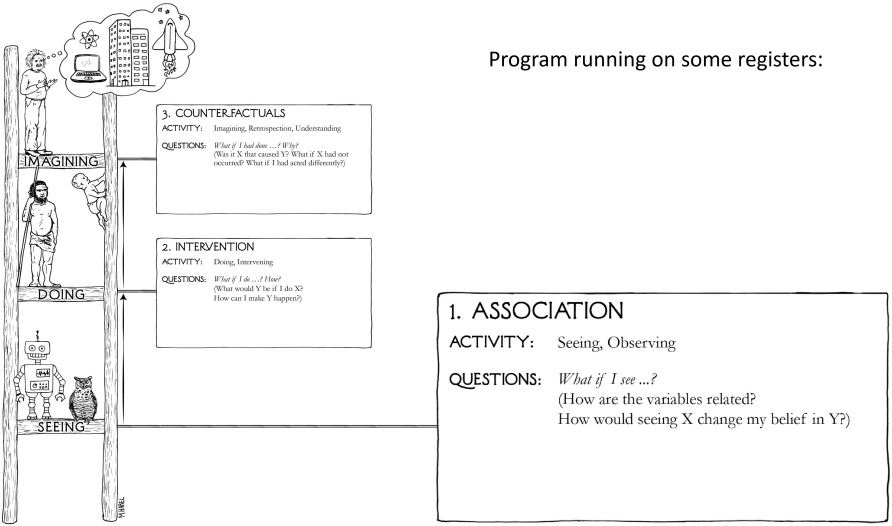

Program running on some registers:

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

Program running on some registers:

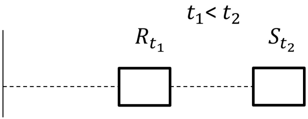

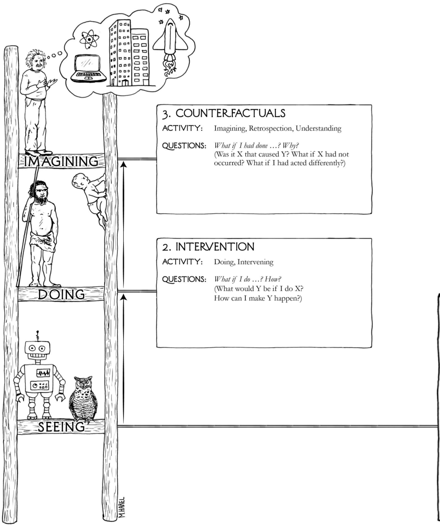
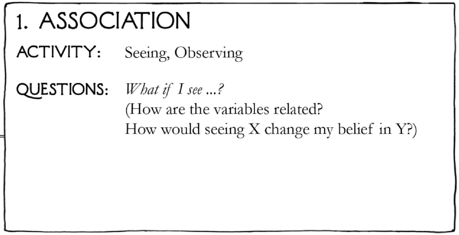

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

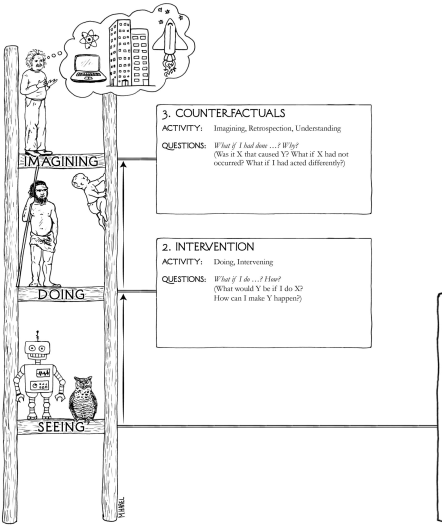

Program running on some registers:

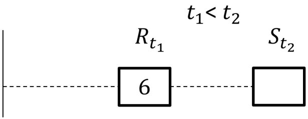

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

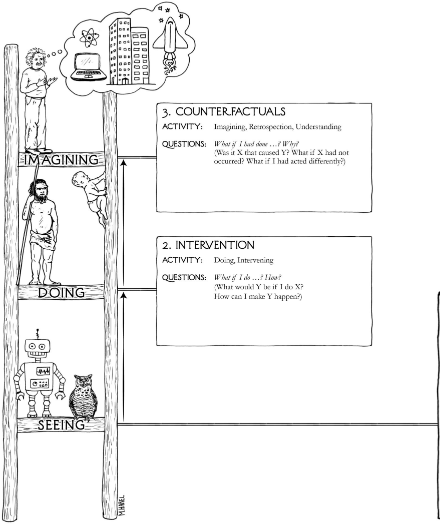

Program running on some registers:

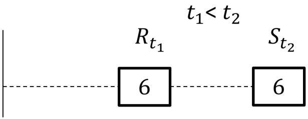

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

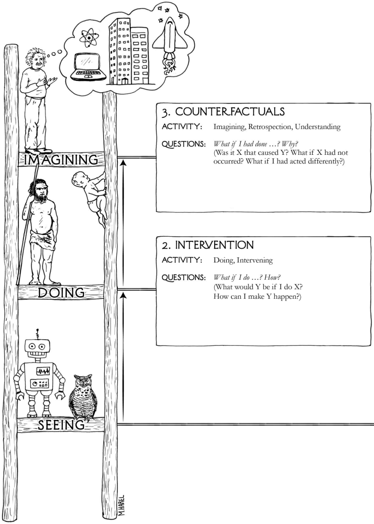

Program running on some registers:

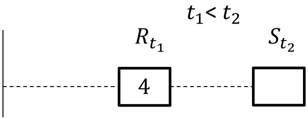

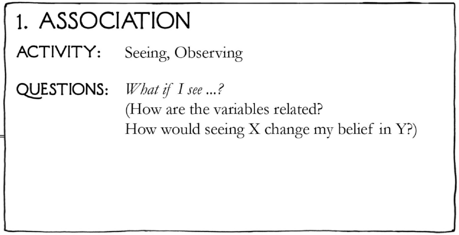

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

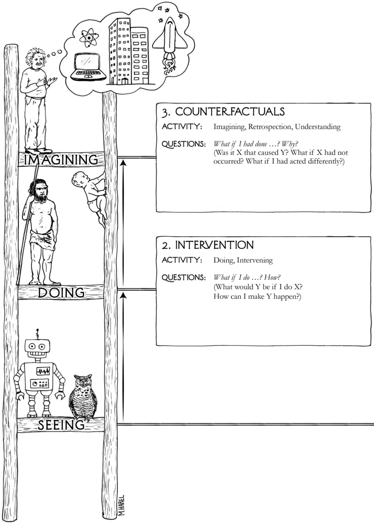

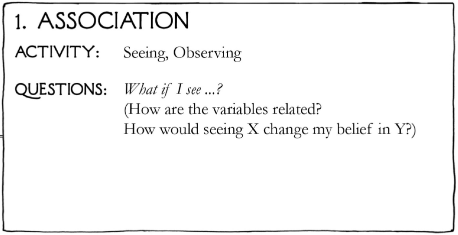

Program running on some registers:

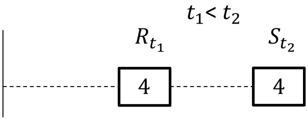

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

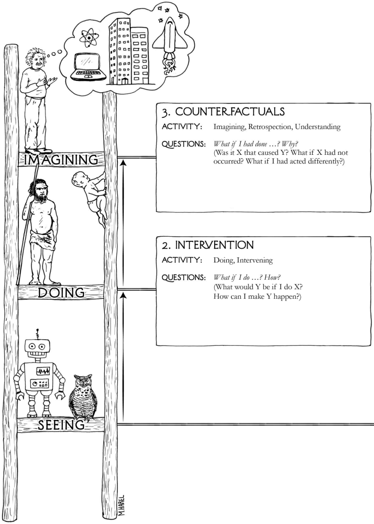

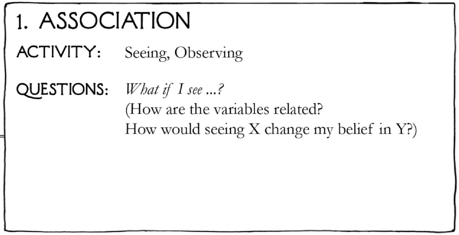

Program running on some registers:

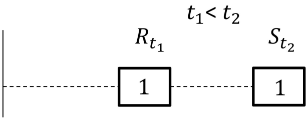

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

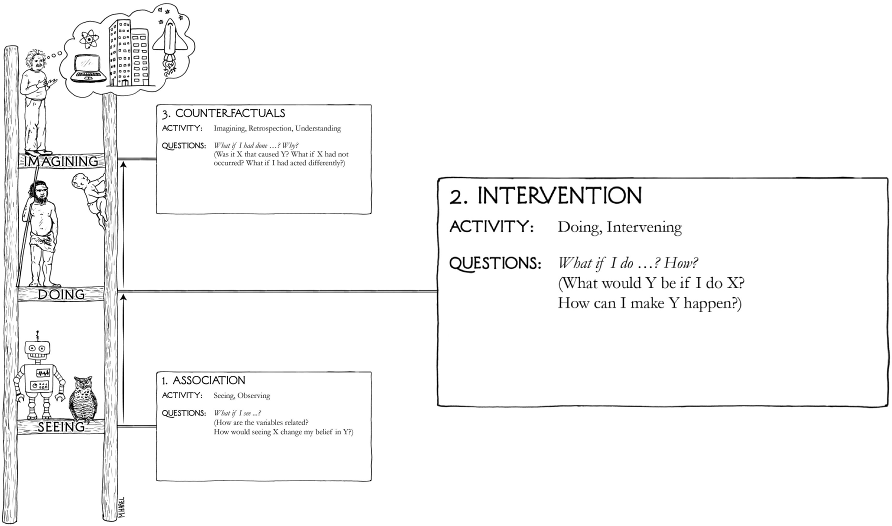

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

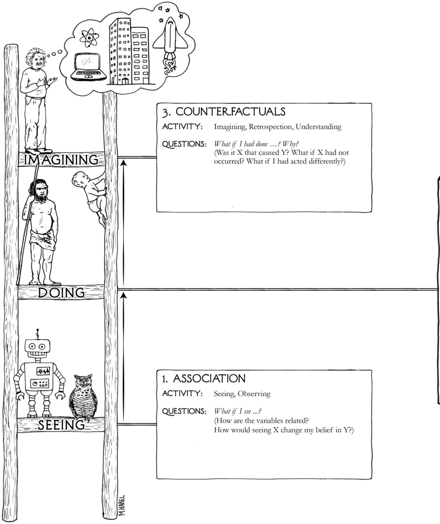

Program running on some registers:

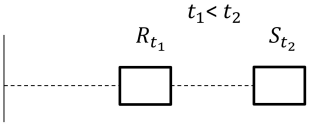

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

Program running on some registers:

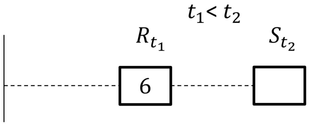

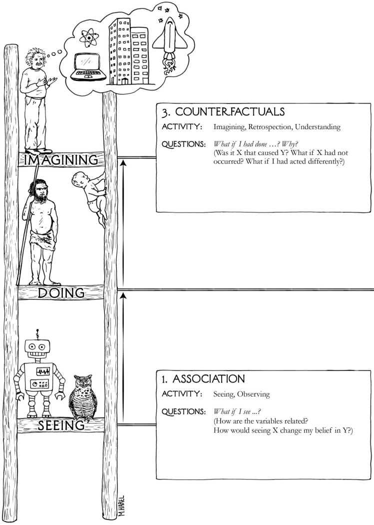

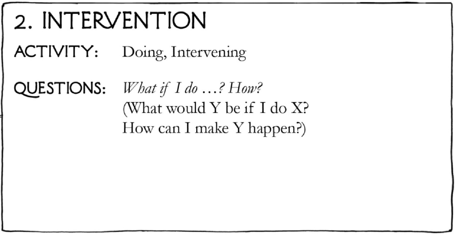

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl’s Ladder of Causation

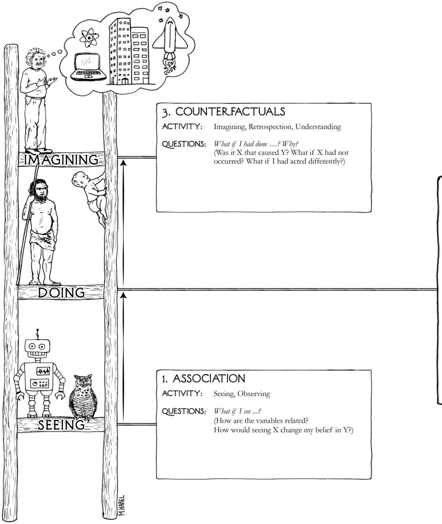

Program running on some registers:

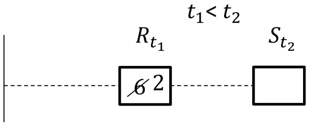

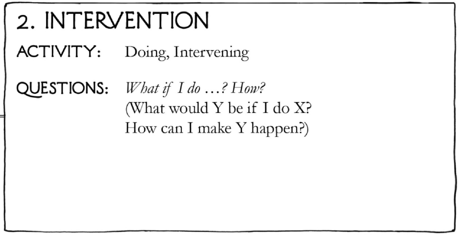

Figure: J. Pearl’s ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

Program running on some registers:

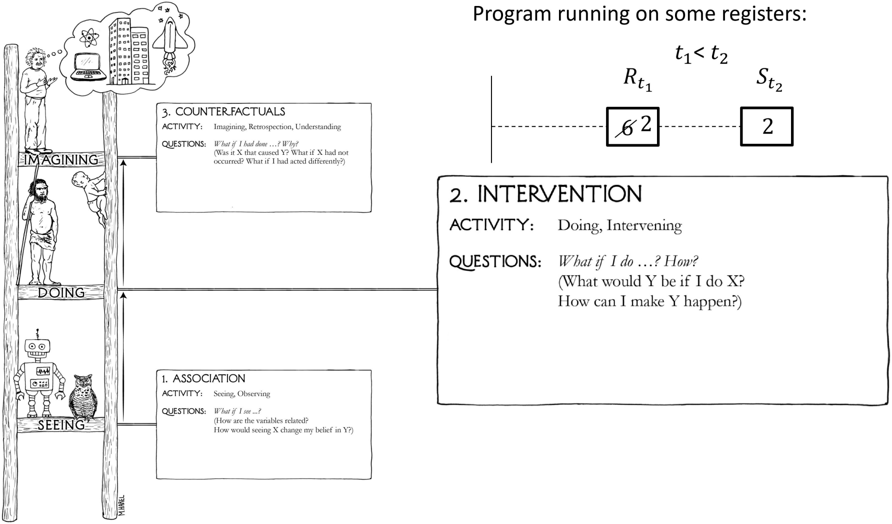

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

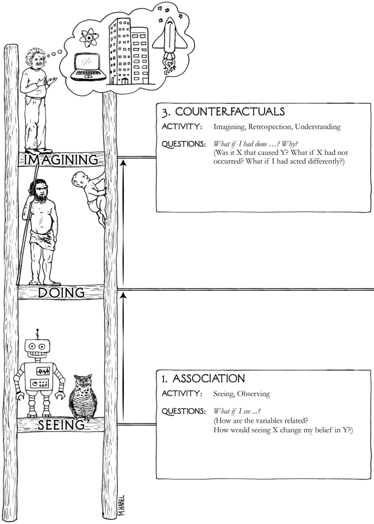

Program running on some registers:

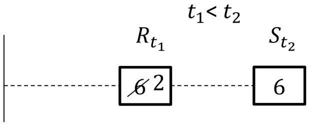

2. INTERVENTION

ACTIVITY: Doing, Intervening

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

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

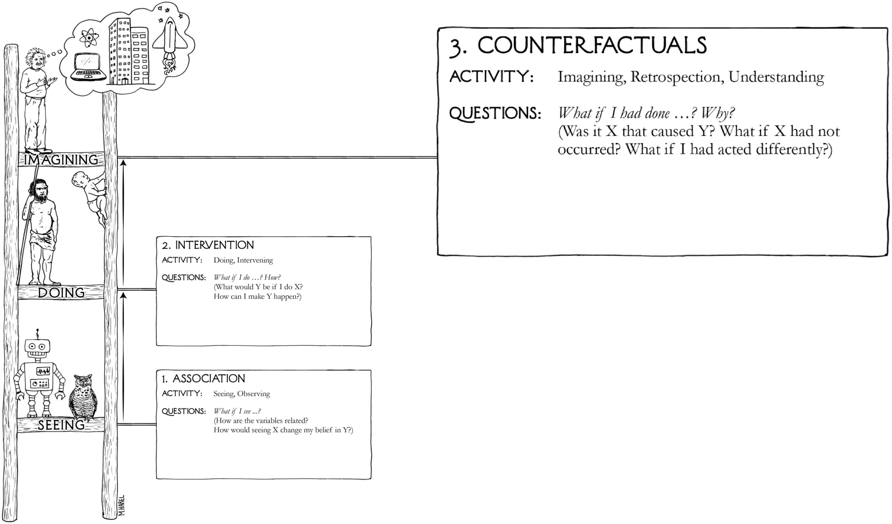

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

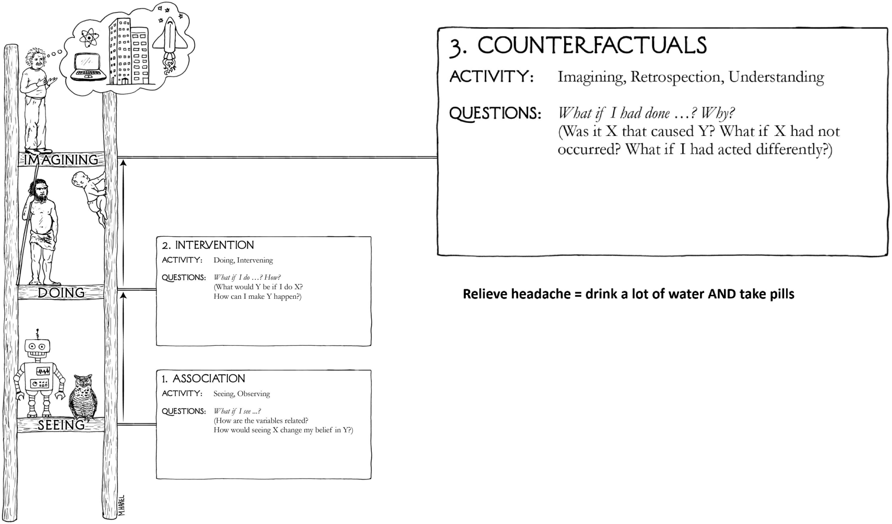

Relieve headache = drink a lot of water AND take pills

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

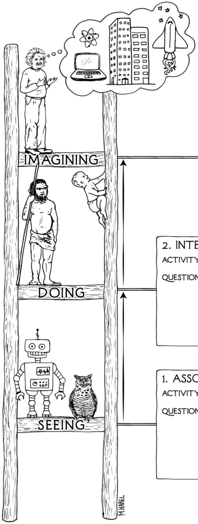

1. ASSOCIATION
ACTIVITY: Seeing, Observing
QUESTIONS: *What if I see ...?*
(How are the variables related?
How would seeing X change my belief in Y?)

2. INTERVENTION
ACTIVITY: Doing, Intervening
QUESTIONS: *What if I do ...? How?*
(What would Y be if I do X?
How can I make Y happen?)

3. COUNTERFACTUALS
ACTIVITY: Imagining, Retrospection, Understanding
QUESTIONS: *What if I had done ...? Why?*
(Was it X that caused Y? What if X had not occurred? What if I had acted differently?)

Relieve headache = drink a lot of water AND take pills

Observed Situation 1: No relief after drinking a lot water and no pill taken.

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

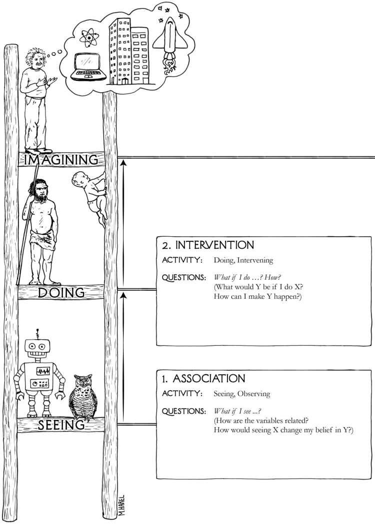

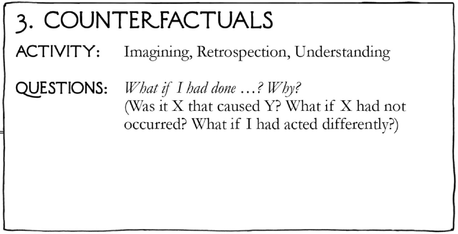

Relieve headache = drink a lot of water AND take pills

Observed Situation 1: No relief after drinking a lot water and no pill taken.

In this situation, if had taken a pill, I would've relieved the headache.
So taking a pill is causally relevant to providing relief in this situation.

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

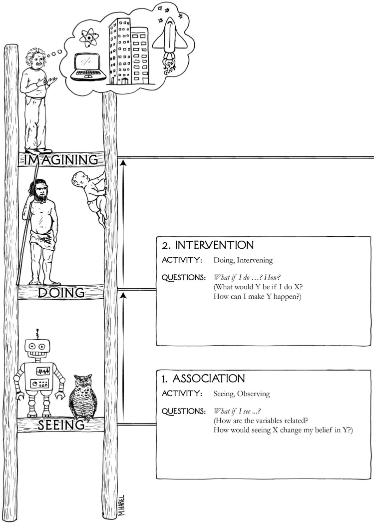

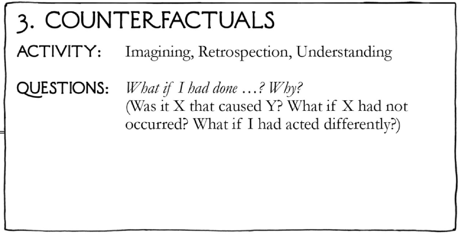

Relieve headache = drink a lot of water AND take pills

Observed Situation 1: No relief after drinking a lot water and no pill taken.

In this situation, if had taken a pill, I would've relieved the headache.
So taking a pill is causally relevant to providing relief in this situation.

Observed Situation 2: No relief after not drinking enough water and no pill taken.

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Pearl's Ladder of Causation

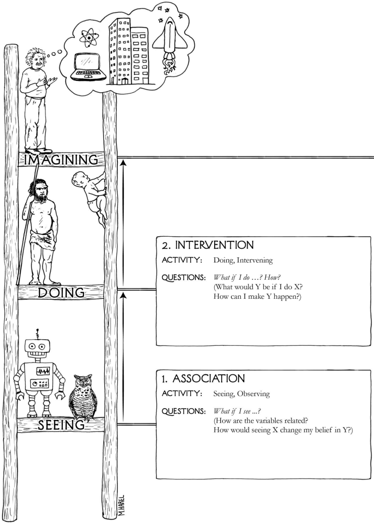

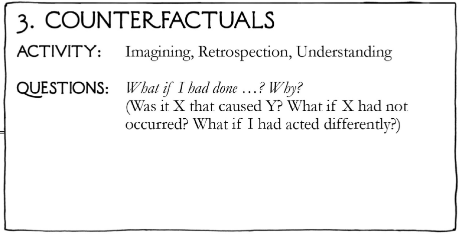

Relieve headache = drink a lot of water AND take pills

Observed Situation 1: No relief after drinking a lot water and no pill taken.

In this situation, if had taken a pill, I would've relieved the headache.
So taking a pill is causally relevant to providing relief in this situation.

Observed Situation 2: No relief after not drinking enough water and no pill taken.

In this situation, if had taken the pill, it would've changed nothing.
So taking a pill is causally irrelevant to providing relief in this situation.

Figure: J. Pearl's ladder of causation (*The Book of Why*, 2018), Illustrator: M. Hare

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

- $G = \langle \mathcal{X}, \mathcal{A} \rangle$ is a finite simple acyclic directed graph with set of vertices $\mathcal{X}$ representing **variables** of M and set of arrows $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. An arrow (X, Y) in $\mathcal{A}$ represents a potential for **direct influence** of X on Y.

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

- $G = \langle \mathcal{X}, \mathcal{A} \rangle$ is a finite simple acyclic directed graph with set of vertices $\mathcal{X}$ representing **variables** of M and set of arrows $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. An arrow (X, Y) in $\mathcal{A}$ represents a potential for **direct influence** of X on Y.
- Given $X \in \mathcal{X}$, set of parents of X: $\mathrm{PA}_X := \{Y \in \mathcal{X} \mid (Y, X) \in \mathcal{A}\}$.

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

- $G = \langle \mathcal{X}, \mathcal{A} \rangle$ is a finite simple acyclic directed graph with set of vertices $\mathcal{X}$ representing **variables** of M and set of arrows $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. An arrow (X, Y) in $\mathcal{A}$ represents a potential for **direct influence** of X on Y.
- Given $X \in \mathcal{X}$, set of parents of X: $\mathrm{PA}_X := \{Y \in \mathcal{X} \mid (Y, X) \in \mathcal{A}\}$.
- Background variables: $\mathcal{U} := \{U \in \mathcal{X} \mid \mathrm{PA}_U = \emptyset\}$.

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

- $G = \langle \mathcal{X}, \mathcal{A} \rangle$ is a finite simple acyclic directed graph with set of vertices $\mathcal{X}$ representing **variables** of M and set of arrows $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. An arrow (X, Y) in $\mathcal{A}$ represents a potential for **direct influence** of X on Y.
- Given $X \in \mathcal{X}$, set of parents of X: $\mathrm{PA}_X := \{Y \in \mathcal{X} \mid (Y, X) \in \mathcal{A}\}$.
- Background variables: $\mathcal{U} := \{U \in \mathcal{X} \mid \mathrm{PA}_U = \emptyset\}$.
- Endogeneous variables: $\mathcal{V} := \mathcal{X} \setminus \mathcal{U}$.

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

- $G = \langle \mathcal{X}, \mathcal{A} \rangle$ is a finite simple acyclic directed graph with set of vertices $\mathcal{X}$ representing **variables** of M and set of arrows $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. An arrow (X, Y) in $\mathcal{A}$ represents a potential for **direct influence** of X on Y.
- Given $X \in \mathcal{X}$, set of parents of X: $\mathrm{PA}_X := \{Y \in \mathcal{X} \mid (Y, X) \in \mathcal{A}\}$.
- Background variables: $\mathcal{U} := \{U \in \mathcal{X} \mid \mathrm{PA}_U = \emptyset\}$.
- Endogeneous variables: $\mathcal{V} := \mathcal{X} \setminus \mathcal{U}$.
- $\{D_X\}_{X \in \mathcal{X}}$ is a family of nonempty finite sets indexed by the variables $\mathcal{X}$, where D_X represents the **domain** of the variable X. For $\mathcal{Y} \subseteq \mathcal{X}$, write $D(\mathcal{Y}) := \prod_{Y \in \mathcal{Y}} D_Y$.

Definition of Structural Causal Models (SCM)

An acyclic **structural causal model** is a tuple

$$M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle \quad \text{where:}$$

- $G = \langle \mathcal{X}, \mathcal{A} \rangle$ is a finite simple acyclic directed graph with set of vertices $\mathcal{X}$ representing **variables** of M and set of arrows $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$. An arrow (X, Y) in $\mathcal{A}$ represents a potential for **direct influence** of X on Y.
- Given $X \in \mathcal{X}$, set of parents of X: $\mathrm{PA}_X := \{Y \in \mathcal{X} \mid (Y, X) \in \mathcal{A}\}$.
- Background variables: $\mathcal{U} := \{U \in \mathcal{X} \mid \mathrm{PA}_U = \emptyset\}$.
- Endogeneous variables: $\mathcal{V} := \mathcal{X} \setminus \mathcal{U}$.
- $\{D_X\}_{X \in \mathcal{X}}$ is a family of nonempty finite sets indexed by the variables $\mathcal{X}$, where D_X represents the **domain** of the variable X. For $\mathcal{Y} \subseteq \mathcal{X}$, write $D(\mathcal{Y}) := \prod_{Y \in \mathcal{Y}} D_Y$.
- $\{f_V\}_{V \in \mathcal{V}}$ is a family of set functions $f_V : \prod_{X \in \mathrm{PA}_V} D_X \to D_V$ indexed by the endogeneous variables $\mathcal{V}$, where f_V represents a local rule that determines the outcome at V solely in terms of input from the parent variables PA_V.

Properties of SCM

$\mathcal{U} = \{U_1, U_2, U_3\}$

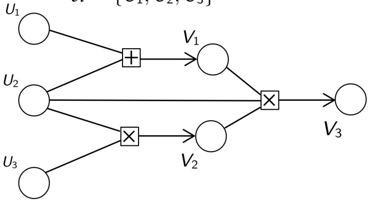

Properties of SCM

$$\mathcal{U} = \{U_1, U_2, U_3\}$$

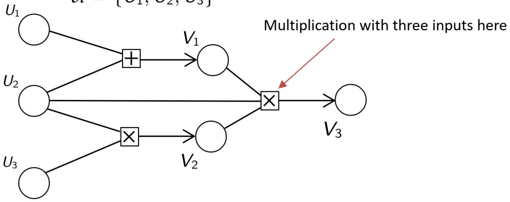

Properties of SCM

$\mathcal{U} = \{U_1, U_2, U_3\}$

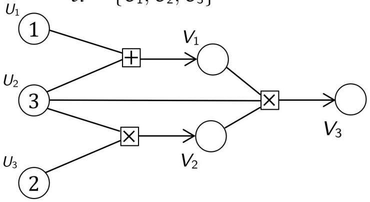

Properties of SCM

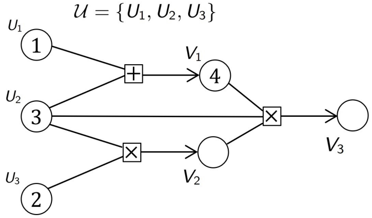

Properties of SCM

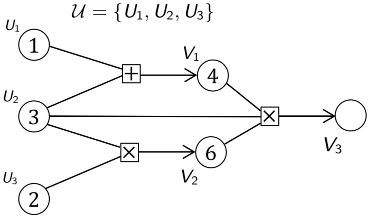

Properties of SCM

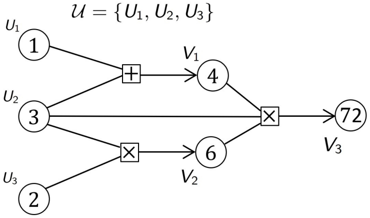

Properties of SCM

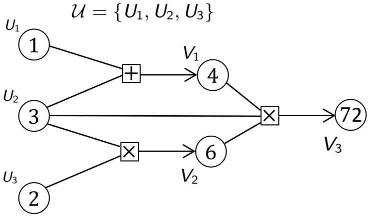

In any SCM, for every input $u \in D(\mathcal{U})$ on the background variables, there exists a unique global outcome $x \in D(\mathcal{X})$ s.t. $x|_{\mathcal{U}} = u$ and $x|_V = f_V(x|_{\mathrm{PA}_V})$ for all $V \in \mathcal{V}$.

Properties of SCM

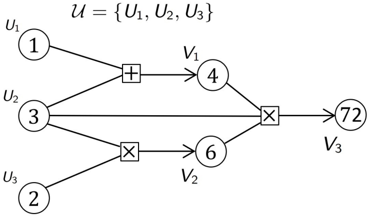

In any SCM, for every input $u \in D(\mathcal{U})$ on the background variables, there exists a unique global outcome $x \in D(\mathcal{X})$ s.t. $x|_{\mathcal{U}} = u$ and $x|_V = f_V(x|_{\mathsf{PA}_V})$ for all $V \in \mathcal{V}$.

Writing $\tilde{\mathcal{X}}(u)$ to denote this unique outcome, we get a global solution map $\tilde{\mathcal{X}} : D(\mathcal{U}) \to D(\mathcal{X})$. For any subset $\mathcal{Y} \subseteq \mathcal{X}$, restriction to $D(\mathcal{Y})$ also gives a $\mathcal{Y}$-solution map $\tilde{\mathcal{Y}} : D(\mathcal{U}) \to D(\mathcal{X})$.

Properties of SCM

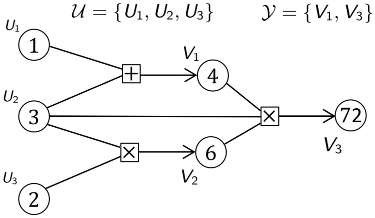

In any SCM, for every input $u \in D(\mathcal{U})$ on the background variables, there exists a unique global outcome $x \in D(\mathcal{X})$ s.t. $x|_{\mathcal{U}} = u$ and $x|_V = f_V(x|_{\mathsf{PA}_V})$ for all $V \in \mathcal{V}$.
Writing $\tilde{\mathcal{X}}(u)$ to denote this unique outcome, we get a global solution map $\tilde{\mathcal{X}} : D(\mathcal{U}) \to D(\mathcal{X})$. For any subset $\mathcal{Y} \subseteq \mathcal{X}$, restriction to $D(\mathcal{Y})$ also gives a $\mathcal{Y}$-solution map $\tilde{\mathcal{Y}} : D(\mathcal{U}) \to D(\mathcal{X})$.

Properties of SCM

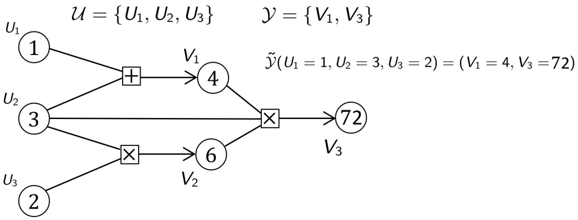

In any SCM, for every input $u \in D(\mathcal{U})$ on the background variables, there exists a unique global outcome $x \in D(\mathcal{X})$ s.t. $x|_{\mathcal{U}} = u$ and $x|_V = f_V(x|_{\mathrm{PA}_V})$ for all $V \in \mathcal{V}$. Writing $\tilde{\mathcal{X}}(u)$ to denote this unique outcome, we get a global solution map $\tilde{\mathcal{X}} : D(\mathcal{U}) \to D(\mathcal{X})$. For any subset $\mathcal{Y} \subseteq \mathcal{X}$, restriction to $D(\mathcal{Y})$ also gives a $\mathcal{Y}$-solution map $\tilde{\mathcal{Y}} : D(\mathcal{U}) \to D(\mathcal{X})$.

Properties of SCM

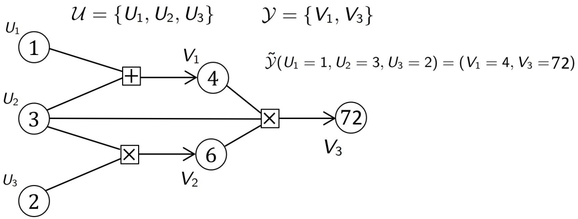

In any SCM, for every input $u \in D(\mathcal{U})$ on the background variables, there exists a unique global outcome $x \in D(\mathcal{X})$ s.t. $x|_{\mathcal{U}} = u$ and $x|_V = f_V(x|_{\mathrm{PA}_V})$ for all $V \in \mathcal{V}$. Writing $\tilde{\mathcal{X}}(u)$ to denote this unique outcome, we get a global solution map $\tilde{\mathcal{X}} : D(\mathcal{U}) \to D(\mathcal{X})$. For any subset $\mathcal{Y} \subseteq \mathcal{X}$, restriction to $D(\mathcal{Y})$ also gives a $\mathcal{Y}$-solution map $\tilde{\mathcal{Y}} : D(\mathcal{U}) \to D(\mathcal{X})$.

If M is additionally equipped with a probability distribution $P(u)$ on its background variables $\mathcal{U}$, the pair $\langle P, M\rangle$ is called **probabilistic (structural) causal model**. There is a unique extension of P to all of $\mathcal{X}$ by pushing out along the map $\tilde{\mathcal{X}} : D(\mathcal{U}) \to D(\mathcal{X})$, i.e.

$$P(x) = \sum_{u \in \tilde{\mathcal{X}}^{-1}(x)} P(u)$$

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0, 1\}$-valued variables X, Y, Z, W:

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0, 1\}$-valued variables X, Y, Z, W:

Order in time:

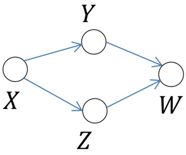

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0,1\}$-valued variables X, Y, Z, W:

Order in time:

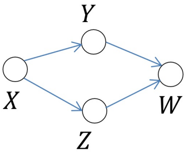

We observe each of the following global outcomes with relative frequency 1/2:

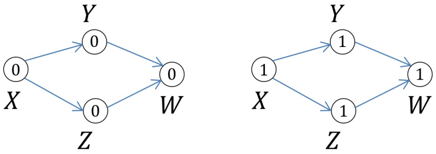

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0,1\}$-valued variables X, Y, Z, W:

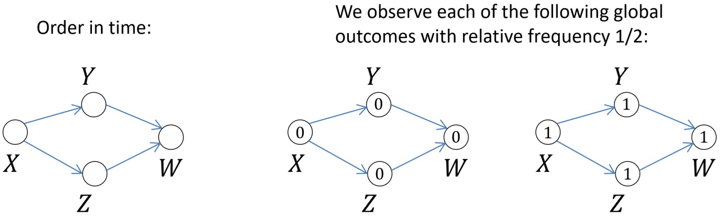

There are several distinct structural causal models that could represent this situation.

For example, we could construct three processes that each consist of three steps, and that share the same steps except the last one as follows:

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0,1\}$-valued variables X, Y, Z, W:

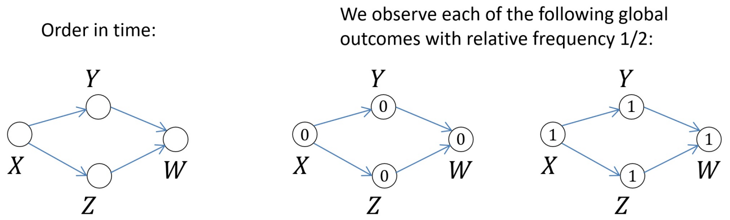

There are several distinct structural causal models that could represent this situation.

For example, we could construct three processes that each consist of three steps, and that share the same steps except the last one as follows:

Step 1: The background variable X receives a random input;

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0, 1\}$-valued variables X, Y, Z, W:

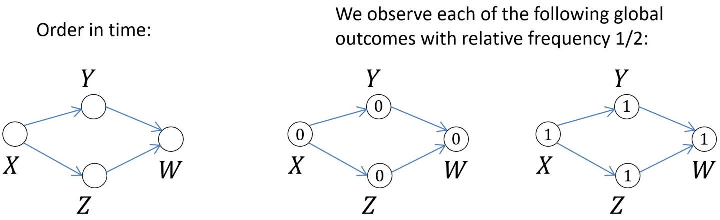

There are several distinct structural causal models that could represent this situation.

For example, we could construct three processes that each consist of three steps, and that share the same steps except the last one as follows:

Step 1: The background variable X receives a random input;
Step 2: after, Y and Z receive a copy of X's value; $f_Y(x) = f_Z(x) = x$;

Example: exposing local structure through interventions

Suppose we see the following process on four $\{0, 1\}$-valued variables X, Y, Z, W:

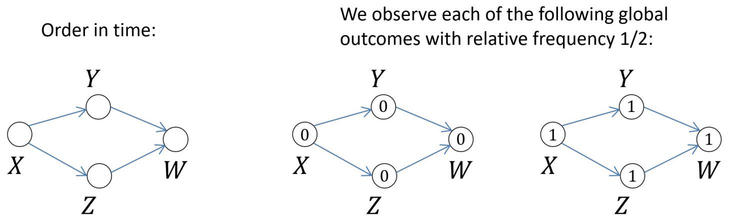

There are several distinct structural causal models that could represent this situation.

For example, we could construct three processes that each consist of three steps, and that share the same steps except the last one as follows:

Step 1: The background variable X receives a random input;
Step 2: after, Y and Z receive a copy of X's value; $f_Y(x) = f_Z(x) = x$;
Step 3: after,

Process 1	Process 2	Process 3
$f^1_W(y, z) = y$	$f^2_W(y, z) = z$	$f^3_W(y, z) = y \cdot z$

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y,z) = y$:

Interventions in Y have an impact on W:

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

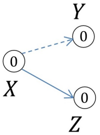

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

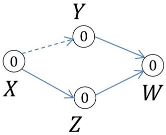

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f^1_W(y, z) = y$:

Interventions in Y have an impact on W:

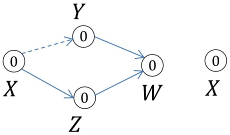

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

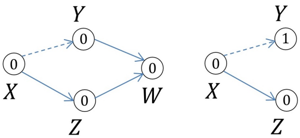

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

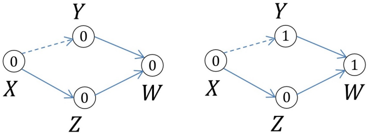

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y,z) = y$:

Interventions in Y have an impact on W:

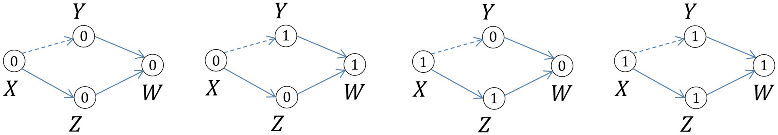

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f^1_W(y, z) = y$:

Interventions in Y have an impact on W:

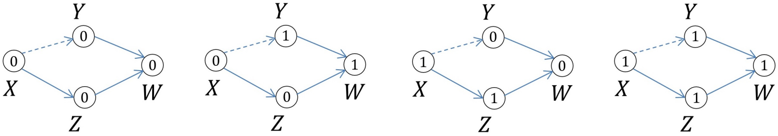

Interventions in Z have no impact on W:

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

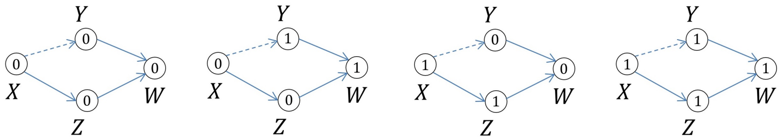

Interventions in Z have no impact on W:

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y,z) = y$:

Interventions in Y have an impact on W:

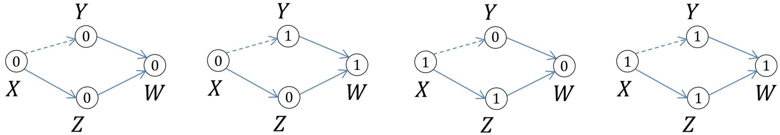

Interventions in Z have no impact on W:

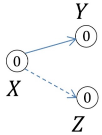

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

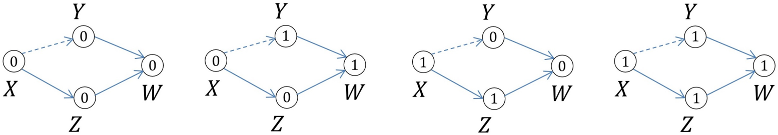

Interventions in Z have no impact on W:

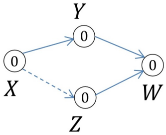

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

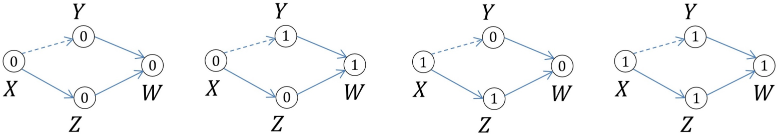

Interventions in Z have no impact on W:

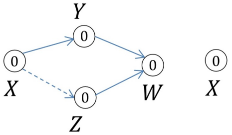

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

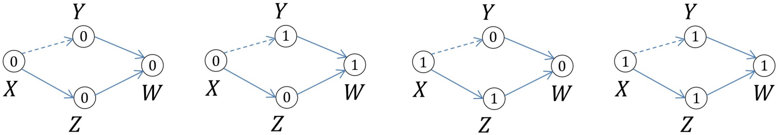

Interventions in Z have no impact on W:

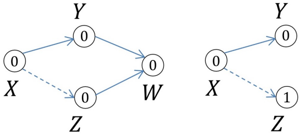

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y, z) = y$:

Interventions in Y have an impact on W:

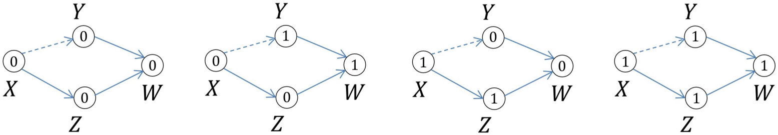

Interventions in Z have no impact on W:

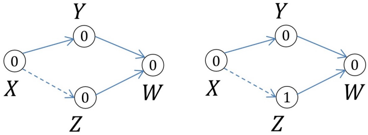

Example: exposing local structure through interventions

All three processes are globally indistiguishable in their solution set, and any distribution $P(u)$ will yield indistinguishable global joint distributions as well.

$\rightarrow$ Local structure is the key to telling these models apart, which can be exposed by interventions in the middle of the process.

For example, in Process 1 with $f_W^1(y,z) = y$:

Interventions in Y have an impact on W:

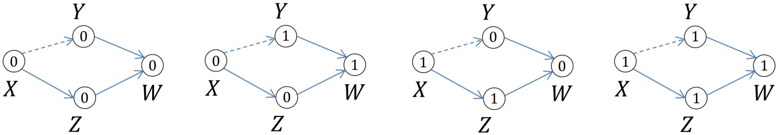

Interventions in Z have no impact on W:

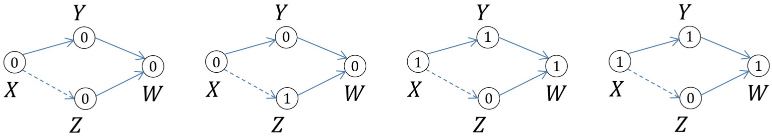

Causal relevance diagrams

Causal relevance diagrams for all three processes:

Process 1:

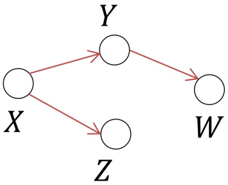

Process 2:

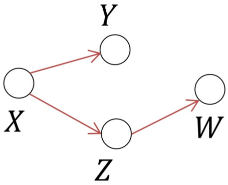

Process 3:

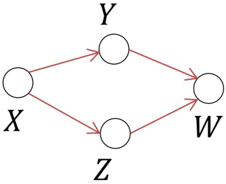

Interventions formally

Let $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$ be an SCM.

Interventions formally

Let $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$ be an SCM.

Given any set of endogeneous variables $\mathcal{Y} \subseteq \mathcal{V}$ and any realization $y \in D(\mathcal{Y})$.
The **y-intervention submodel** M_y of M is the structural causal model

$$M_y = \langle G = \langle \mathcal{X}, \mathcal{A}_{\mathcal{Y}} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V} \setminus \mathcal{Y}} \rangle$$

$$\text{where } \ \mathcal{A}_{\mathcal{Y}} = \mathcal{A} \setminus \{(X, Y) \in \mathcal{A} \mid Y \in \mathcal{Y}\}.$$

Interventions formally

Let $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$ be an SCM.

Given any set of endogeneous variables $\mathcal{Y} \subseteq \mathcal{V}$ and any realization $y \in D(\mathcal{Y})$.
The **y-intervention submodel** M_y of M is the structural causal model

$$M_y = \langle G = \langle \mathcal{X}, \mathcal{A}_{\mathcal{Y}} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V} \setminus \mathcal{Y}} \rangle$$
$$\text{where } \ \mathcal{A}_{\mathcal{Y}} = \mathcal{A} \setminus \{(X, Y) \in \mathcal{A} \mid Y \in \mathcal{Y}\}.$$

But we also set each $V \in \mathcal{Y}$ to the constant $* \mapsto y|_V$.

Interventions formally

Let $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$ be an SCM.

Given any set of endogeneous variables $\mathcal{Y} \subseteq \mathcal{V}$ and any realization $y \in D(\mathcal{Y})$. The **y-intervention submodel** M_y of M is the structural causal model

$$M_y = \langle G = \langle \mathcal{X}, \mathcal{A}_\mathcal{Y} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V} \setminus \mathcal{Y}} \rangle$$

$$\text{where} \ \ \mathcal{A}_\mathcal{Y} = \mathcal{A} \setminus \{(X, Y) \in \mathcal{A} \mid Y \in \mathcal{Y}\}.$$

But we also set each $V \in \mathcal{Y}$ to the constant $* \mapsto y|_V$.

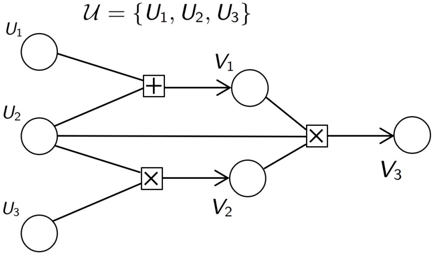

Interventions formally

Let $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$ be an SCM.

Given any set of endogeneous variables $\mathcal{Y} \subseteq \mathcal{V}$ and any realization $y \in D(\mathcal{Y})$.
The **y-intervention submodel** M_y of M is the structural causal model

$$M_y = \langle G = \langle \mathcal{X}, \mathcal{A}_\mathcal{Y} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V} \setminus \mathcal{Y}} \rangle$$

$$\text{where}\ \ \mathcal{A}_\mathcal{Y} = \mathcal{A} \setminus \{(X, Y) \in \mathcal{A} \mid Y \in \mathcal{Y}\}.$$

But we also set each $V \in \mathcal{Y}$ to the constant $* \mapsto y|_V$.

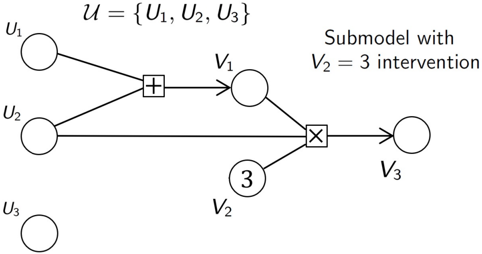

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{W}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{W}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

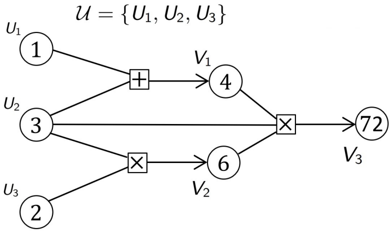

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{W}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

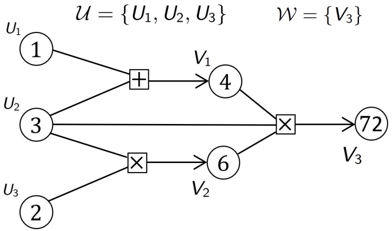

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

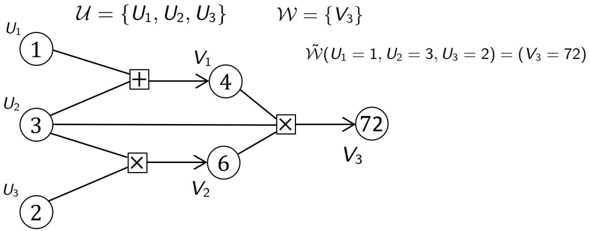

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

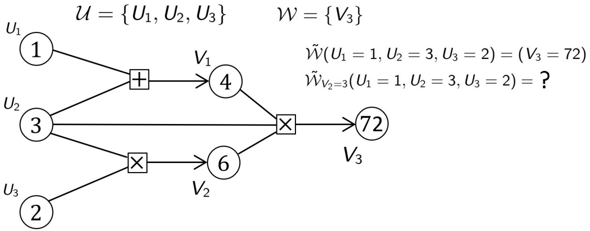

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

$\mathcal{U} = \{U_1, U_2, U_3\}$ $\quad$ $\mathcal{W} = \{V_3\}$

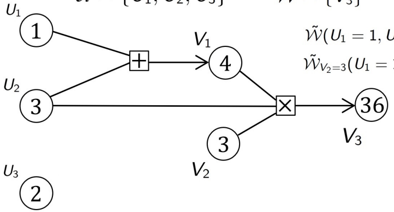

$\tilde{\mathcal{W}}(U_1 = 1, U_2 = 3, U_3 = 2) = (V_3 = 72)$

$\tilde{\mathcal{W}}_{V_2=3}(U_1 = 1, U_2 = 3, U_3 = 2) = (V_3 = 36)$

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

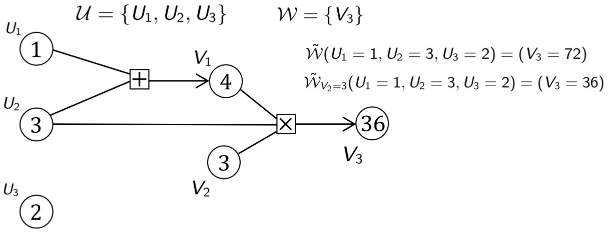

- The original value $V_2 = 6$ is **causally relevant** to $V_3 = 72$ in situation $U_1 = 1, 3, 2$;

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{W}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

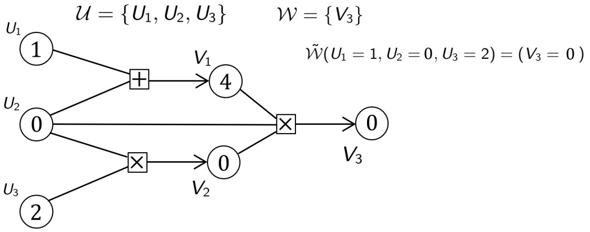

- The original value $V_2 = 6$ is **causally relevant** to $V_3 = 72$ in situation $U_1 = 1, 3, 2$;

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{W}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

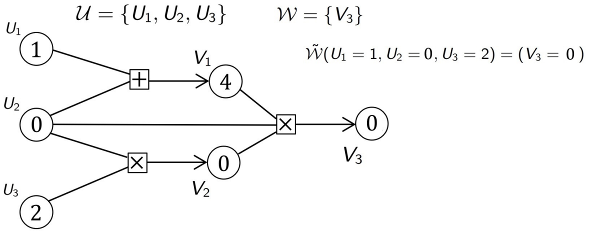

- The original value $V_2 = 6$ is **causally relevant** to $V_3 = 72$ in situation $U_1 = 1, 3, 2$;
- But in situation $U_1 = 1, 0, 2$, $V_2 = 0$ is **causally irrelevant** to $V_3 = 0$.

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

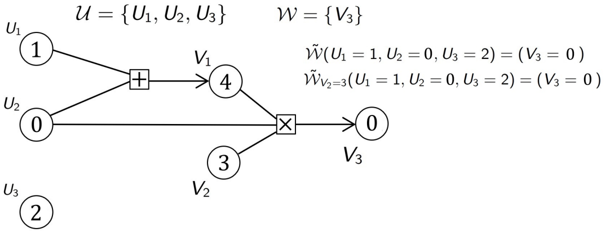

- The original value $V_2 = 6$ is **causally relevant** to $V_3 = 72$ in situation $U_1 = 1, 3, 2$;
- But in situation $U_1 = 1, 0, 2$, $V_2 = 0$ is **causally irrelevant** to $V_3 = 0$.

Counterfactuals and causal relevance

- Given a set of variables $\mathcal{W} \subseteq \mathcal{X}$, the **potential response** of $\mathcal{W}$ to the $\mathcal{Y} = y$-intervention is represented by the $\mathcal{W}$-solution map $\tilde{\mathcal{W}}_{\mathcal{Y}=y} : D(\mathcal{U}) \to D(\mathcal{Z})$ in the M_y model.

  Given $u \in D(\mathcal{U})$, $\tilde{W}_{\mathcal{Y}=y}(u)$ represents the **counterfactual value** that $\mathcal{W}$ would have in situation u, if $\mathcal{Y}$ had been y.

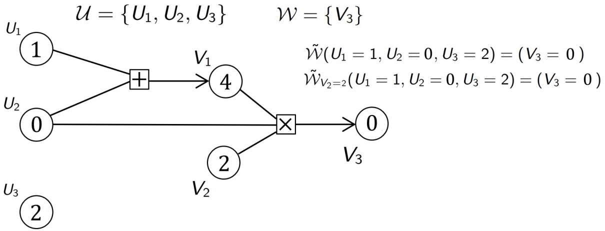

- The original value $V_2 = 6$ is **causally relevant** to $V_3 = 72$ in situation $U_1 = 1, 3, 2$;
- But in situation $U_1 = 1, 0, 2$, $V_2 = 0$ is **causally irrelevant** to $V_3 = 0$.

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

$$\{(X, V) \mid X \in \mathrm{PA}_V\}$$

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

$$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\})$$

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

$$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

  $$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

  The presheaf F has at least one amalgamation for each matching family in this topology. (In fact, $\mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$.)

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

$$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

  The presheaf F has at least one amalgamation for each matching family in this topology. (In fact, $\mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$.)
- Yet, each local behavior functions $f_V : \prod_{X \in \mathrm{PA}_V} D_X \to D_V$ of the SCM essentially provide a gluing operation on matching families over S_V and establish a map:

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \{D_X\}_{X \in \mathcal{X}}, \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

  $$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

  The presheaf F has at least one amalgamation for each matching family in this topology. (In fact, $\mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$.)
- Yet, each local behavior functions $f_V : \prod_{X \in \mathrm{PA}_V} D_X \to D_V$ of the SCM essentially provide a gluing operation on matching families over S_V and establish a map:

  $$\frac{S_V \longrightarrow F}{t_V \longrightarrow F}$$

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle,\ \{D_X\}_{X \in \mathcal{X}},\ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

  $$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

  The presheaf F has at least one amalgamation for each matching family in this topology. (In fact, $\mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$.)
- Yet, each local behavior functions $f_V : \prod_{X \in \mathrm{PA}_V} D_X \to D_V$ of the SCM essentially provide a gluing operation on matching families over S_V and establish a map:

$$\frac{S_V \longrightarrow F}{t_V \longrightarrow F} \qquad \bar{f}_V : \mathrm{Match}(S_V, F) \to \mathrm{Match}(t_V, F)$$

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \{D_X\}_{X \in \mathcal{X}}, \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

  $$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

  The presheaf F has at least one amalgamation for each matching family in this topology. (In fact, $\mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$.)
- Yet, each local behavior functions $f_V : \prod_{X \in \mathrm{PA}_V} D_X \to D_V$ of the SCM essentially provide a gluing operation on matching families over S_V and establish a map:

  $$\frac{S_V \longrightarrow F}{t_V \longrightarrow F} \qquad \bar{f}_V : \mathrm{Match}(S_V, F) \to \mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$$

Seeking Categorical Models of SCM

Given a structural causal model $M = \langle G = \langle \mathcal{X}, \mathcal{A} \rangle, \ \{D_X\}_{X \in \mathcal{X}}, \ \{f_V\}_{V \in \mathcal{V}} \rangle$, we started by looking for a representation within the category of presheaves $[C(G)^{op}, Set]$ on the freely generated category of the graph G. This goes as follows:

- Treating the vertices $\mathcal{X}$ as a discrete category, we get a functor $D : \mathcal{X}^{op} \to Set$ (i.e. $D(X) = D_X$), and the right Kan extension of D along the inclusion in $C(G)^{op}$ gives a functor $F = Ran(D) : C(G)^{op} \to Set$.
- There is an obvious topology on $C(G)$, where parents of a node form a cover:

$$S_V = (\{(X, V) \mid X \in \mathrm{PA}_V\}) = t_V \setminus \{id_V\}$$

  The presheaf F has at least one amalgamation for each matching family in this topology. (In fact, $\mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$.)
- Yet, each local behavior functions $f_V : \prod_{X \in \mathrm{PA}_V} D_X \to D_V$ of the SCM essentially provide a gluing operation on matching families over S_V and establish a map:

$$\frac{S_V \longrightarrow F}{t_V \longrightarrow F} \qquad \bar{f}_V : \mathrm{Match}(S_V, F) \to \mathrm{Match}(t_V, F) \cong \mathrm{Match}(S_V, F) \times D_V$$

  (This map is a section of the restriction map the other way around.)

Categorical SCM

- This categorical SCM representation $\langle F, \{\bar{f_V}\}_{V \in \mathcal{V}} \rangle$ of a SCM comes with a notion of morphism for SCM models (i.e. natural transformations on F that preserve the $\bar{f_V}$ maps), and a notion of isomorphism as well.

Categorical SCM

- This categorical SCM representation $\langle F, \{\bar{f_V}\}_{V \in \mathcal{V}} \rangle$ of a SCM comes with a notion of morphism for SCM models (i.e. natural transformations on F that preserve the $\bar{f_V}$ maps), and a notion of isomorphism as well.
- In fact, each SCM gives rise to a categorical SCM in such a way that distincts SCMs give non-isomorphic categorical SCMs.

Categorical SCM

- This categorical SCM representation $\langle F, \{\bar{f_V}\}_{V \in \mathcal{V}} \rangle$ of a SCM comes with a notion of morphism for SCM models (i.e. natural transformations on F that preserve the $\bar{f_V}$ maps), and a notion of isomorphism as well.
- In fact, each SCM gives rise to a categorical SCM in such a way that distincts SCMs give non-isomorphic categorical SCMs.
- This justifies moving our analysis of SCMs to the categorical setting, where it will be possible to study interventions, counterfactuals and other relations involved in causal modeling.

Categorical SCM

- This categorical SCM representation $\langle F, \{\bar{f_V}\}_{V \in \mathcal{V}} \rangle$ of a SCM comes with a notion of morphism for SCM models (i.e. natural transformations on F that preserve the $\bar{f_V}$ maps), and a notion of isomorphism as well.
- In fact, each SCM gives rise to a categorical SCM in such a way that distincts SCMs give non-isomorphic categorical SCMs.
- This justifies moving our analysis of SCMs to the categorical setting, where it will be possible to study interventions, counterfactuals and other relations involved in causal modeling.

Thank you!