The word *cause* is not in the vocabulary of standard probability theory. It is an embarrassing yet inescapable fact that probability theory, the official mathematical language of many empirical sciences, does not permit us to express sentences such as "Mud does not cause rain"; all we can say are that the two events are mutually correlated, or dependent – meaning that if we find one, we can expect to encounter the other. Scientists seeking causal explanations for complex phenomenon or rationales for policy decisions must therefore supplement the language of probability with a vocabulary for causality, one in which the symbolic representation for "Mud does not cause rain" is distinct from the symbolic representation for "Mud is independent of rain". Oddly, such distinctions have yet to be incorporated into standard scientific analysis.

– Judea Pearl, *Causality*, p 134.

Causality

- An intervention on a variable changes its value by some mechanism outside of the model.

Causality

- An intervention on a variable changes its value by some mechanism outside of the model.
- A causal model is a model which predicts the effects of interventions.

Causality

- An intervention on a variable changes its value by some mechanism outside of the model.
- A causal model is a model which predicts the effects of interventions.
- A direct cause of variable Y is a variable X such that intervening on X, holding all other variables constant, can affect Y.

Causality

- An intervention on a variable changes its value by some mechanism outside of the model.
- A causal model is a model which predicts the effects of interventions.
- A direct cause of variable Y is a variable X such that intervening on X, holding all other variables constant, can affect Y. e.g., mud and rain.

Causality

- An intervention on a variable changes its value by some mechanism outside of the model.
- A causal model is a model which predicts the effects of interventions.
- A direct cause of variable Y is a variable X such that intervening on X, holding all other variables constant, can affect Y. e.g., mud and rain.
- Assume no causal cycles; apparent cycles, e.g., *poverty* $\rightarrow$ *sickness* and *sickness* $\rightarrow$ *poverty*, are modeled using time.

Causality

- In a causal network, the parents of a variable are its direct causes.

Causality

- In a causal network, the parents of a variable are its direct causes.
- We would expect that a causal model to obey the independence assumption of a belief network – each variable is independent of its descendants given its parents.

Causality

- In a causal network, the parents of a variable are its direct causes.
- We would expect that a causal model to obey the independence assumption of a belief network – each variable is independent of its descendants given its parents.
  - All causal networks are belief networks.

Causality

- In a causal network, the parents of a variable are its direct causes.
- We would expect that a causal model to obey the independence assumption of a belief network – each variable is independent of its descendants given its parents.
  - All causal networks are belief networks.
  - Not all belief networks are causal networks.
    E.g. $rain \rightarrow mud$ versus $mud \rightarrow rain$.

Causality

- In a causal network, the parents of a variable are its direct causes.
- We would expect that a causal model to obey the independence assumption of a belief network – each variable is independent of its descendants given its parents.
  - All causal networks are belief networks.
  - Not all belief networks are causal networks.
    E.g. $rain \rightarrow mud$ versus $mud \rightarrow rain$.
- In a causal network $P(X \mid parents(X))$ is the same whether the parents are observed or intervened on.

Causality

- In a causal network, the parents of a variable are its direct causes.
- We would expect that a causal model to obey the independence assumption of a belief network – each variable is independent of its descendants given its parents.
  - All causal networks are belief networks.
  - Not all belief networks are causal networks.
    E.g. $rain \rightarrow mud$ versus $mud \rightarrow rain$.
- In a causal network $P(X \mid parents(X))$ is the same whether the parents are observed or intervened on.
- A causal mechanism specifies the $P(X \mid parents(X))$ when the parents are intervened on.

Causality

- In a causal network, the parents of a variable are its direct causes.
- We would expect that a causal model to obey the independence assumption of a belief network – each variable is independent of its descendants given its parents.
  - All causal networks are belief networks.
  - Not all belief networks are causal networks.
    E.g. $rain \rightarrow mud$ versus $mud \rightarrow rain$.
- In a causal network $P(X \mid parents(X))$ is the same whether the parents are observed or intervened on.
- A causal mechanism specifies the $P(X \mid parents(X))$ when the parents are intervened on.
- A structural causal model defines a causal mechanism for each modeled variable.

Sprinkler Example

Variables:

- Season dry or wet
- Rained last night
- Sprinkler was on last night
- Grass wet
- Grass shiny and appears to be wet
- Shoes wet after walking on grass

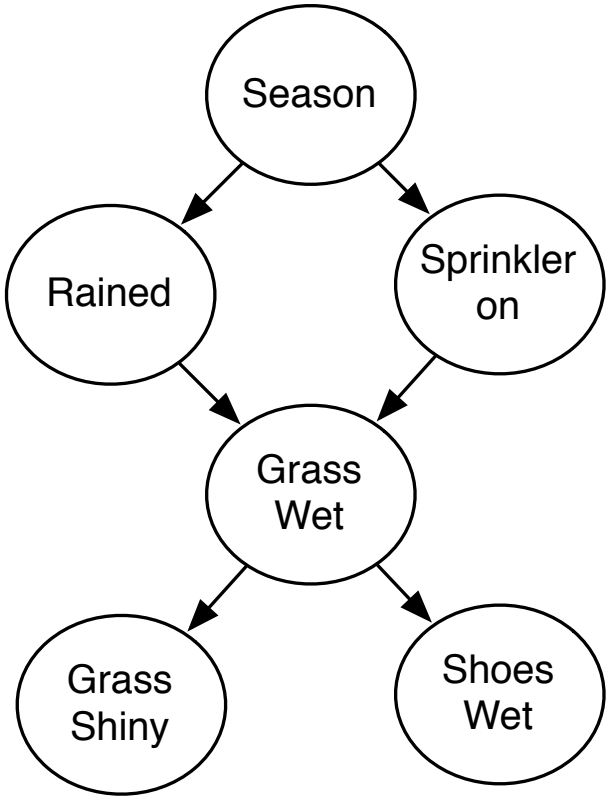

Sprinkler Example

Variables:

- Season dry or wet
- Rained last night
- Sprinkler was on last night
- Grass wet
- Grass shiny and appears to be wet
- Shoes wet after walking on grass

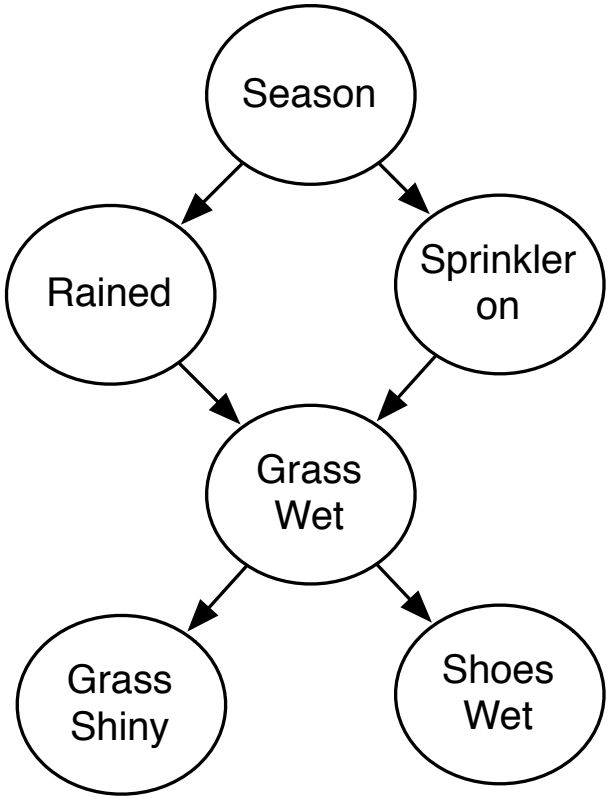

- Which probabilities change if you observe sprinkler on?

Sprinkler Example

Variables:

- Season dry or wet
- Rained last night
- Sprinkler was on last night
- Grass wet
- Grass shiny and appears to be wet
- Shoes wet after walking on grass

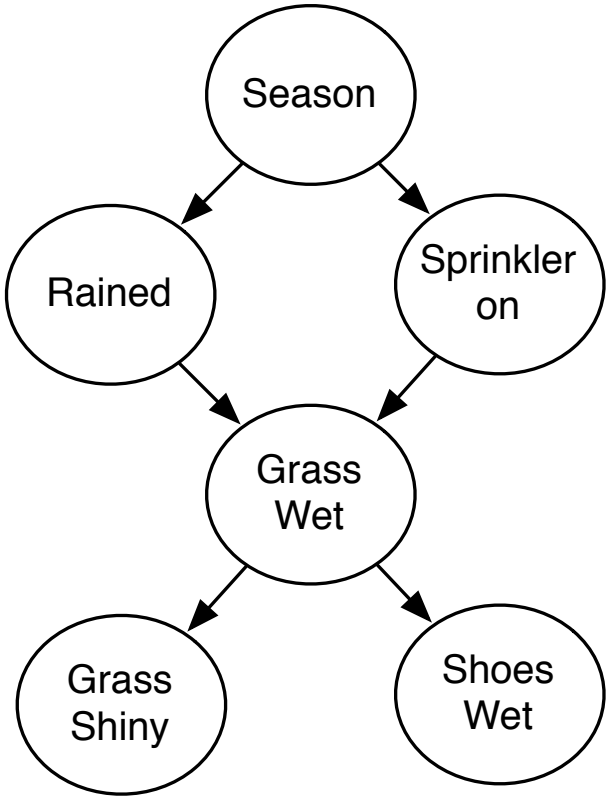

- Which probabilities change if you observe sprinkler on?
- Which probabilities change if you turn the sprinkler on?

Example: drowning and eating ice cream.

- Ice cream consumption and drowning are correlated.

Example: drowning and eating ice cream.

- Ice cream consumption and drowning are correlated.
- The top two can be made to fit the data
- Which is a better causal model?
- What experiments could be used to test the models?

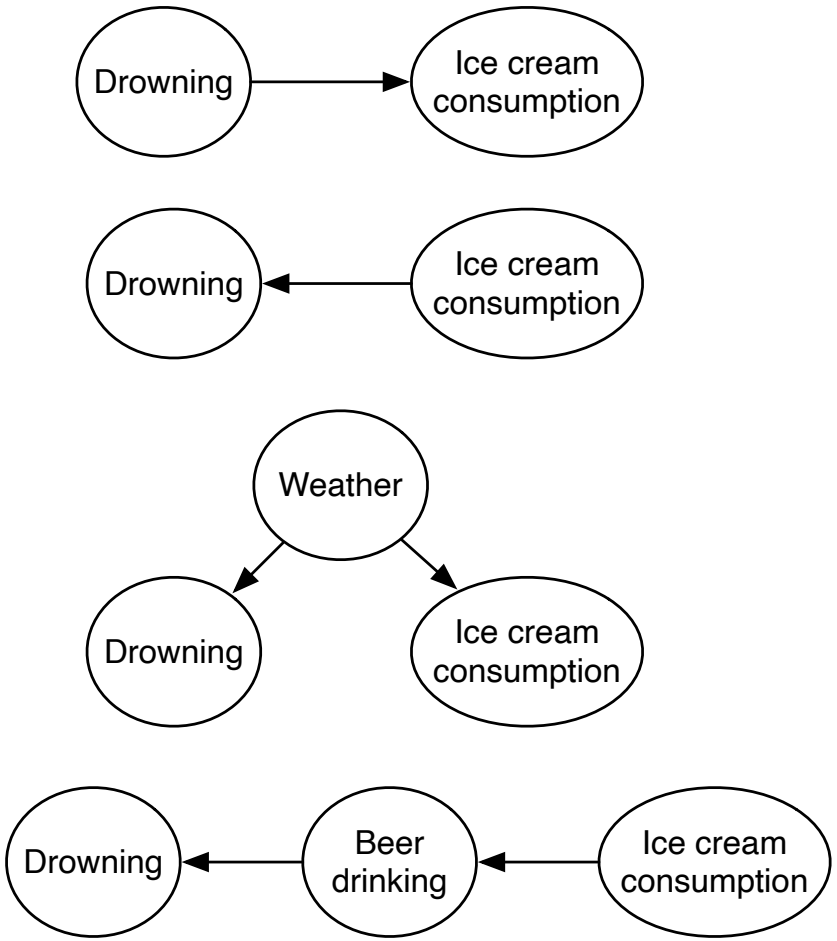

Clicker Question

Which of the following is not true:

A All belief networks are causal networks

B All causal networks are belief networks

C A causal network predicts the effect of an intervention

D An intervention changes the value of a variable by some mechanism external to the model

E Intervening on a variable only affects the descendents of the variable

Do notation

In a causal model:

- $P(x \mid do(z), y)$ is the probability that x is true after doing z and then observing y.

Do notation

In a causal model:

- $P(x \mid do(z), y)$ is the probability that x is true after doing z and then observing y.
- Are these different?
  $P(Season{=}wet \mid do(Sprinkler\_on{=}true))$
  $P(Season{=}wet \mid Sprinkler\_on{=}true)$?

Do notation

In a causal model:

- $P(x \mid do(z), y)$ is the probability that x is true after doing z and then observing y.
- Are these different?
  $P(Season{=}wet \mid do(Sprinkler\_on{=}true))$
  $P(Season{=}wet \mid Sprinkler\_on{=}true)$?
- Are these different?
  $P(grassWet \mid do(Sprinkler\_on{=}true), Rained{=}false)$
  $P(grassWet \mid Sprinkler\_on{=}true, Rained{=}false)$?

Do notation

In a causal model:

- $P(x \mid do(z), y)$ is the probability that x is true after doing z and then observing y.
- Are these different?
  $P(Season{=}wet \mid do(Sprinkler\_on{=}true))$
  $P(Season{=}wet \mid Sprinkler\_on{=}true)$?
- Are these different?
  $P(grassWet \mid do(Sprinkler\_on{=}true), Rained{=}false)$
  $P(grassWet \mid Sprinkler\_on{=}true, Rained{=}false)$?

Which is a reasonable causal model?

A fan is connected to a switch...

A)

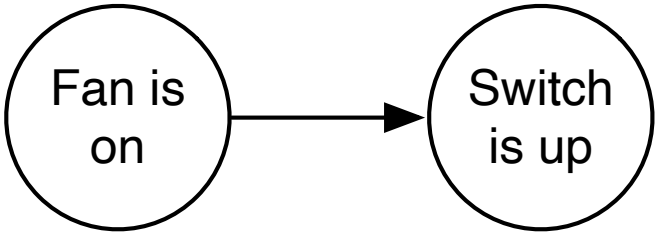

B)

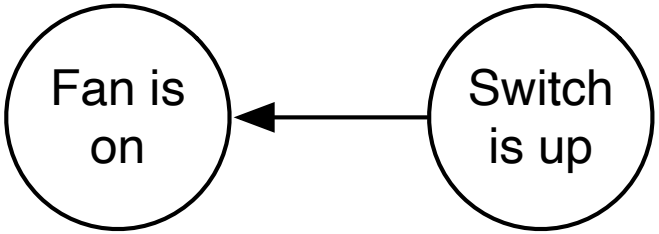

C) both

D) neither

Which is a reasonable causal model?

A)

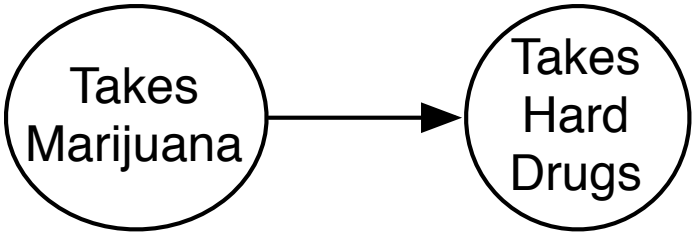

B)

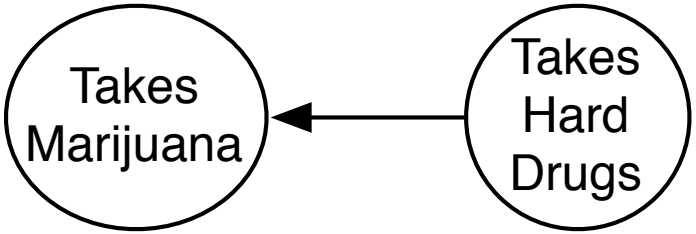

C) both

D) neither

Which is a reasonable causal model?

A)

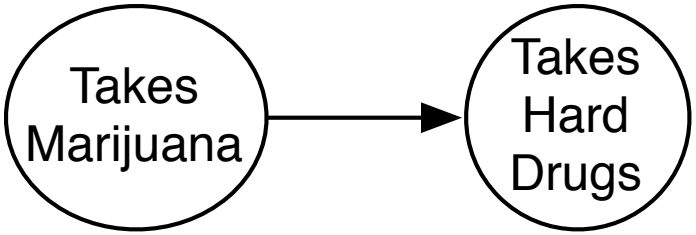

B)

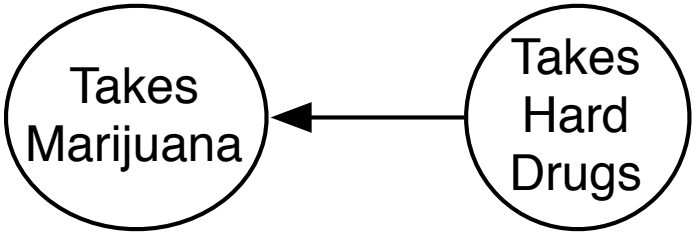

C) both

D) neither

These models cannot be distinguished by observations — but can be distinguished by interventions in controlled studies.

Which is a reasonable causal model?

A)

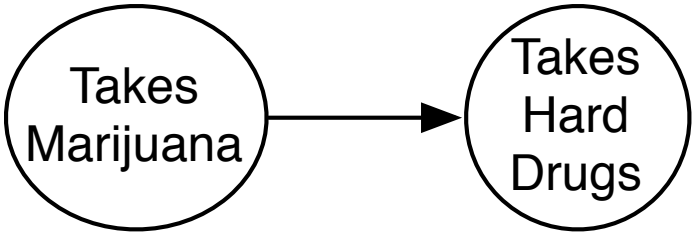

B)

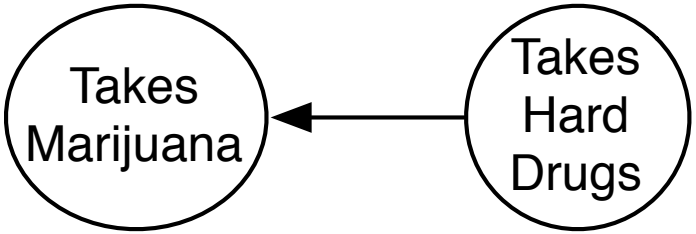

C) both

D) neither

These models cannot be distinguished by observations — but can be distinguished by interventions in controlled studies.
Try: `drugs` example in `probDo.py` of `AIPython`.

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.
- Intervening on a variable only affects its descendants.

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.
- Intervening on a variable only affects its descendants.
- Can be modeled by variable X with parents Zs having a new parent, "$ForceX$"

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.
- Intervening on a variable only affects its descendants.
- Can be modeled by variable X with parents Zs having a new parent, "$ForceX$"
  - $domain(ForceX) = domain(X) \cup \{\perp\}$

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.
- Intervening on a variable only affects its descendants.
- Can be modeled by variable X with parents Zs having a new parent, "$ForceX$"
  - $domain(ForceX) = domain(X) \cup \{\bot\}$
  - $ForceX{=}\bot$ when X is not intervened on; otherwise it is the value X is set to.

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.
- Intervening on a variable only affects its descendants.
- Can be modeled by variable X with parents Zs having a new parent, "$ForceX$"
  - $domain(ForceX) = domain(X) \cup \{\perp\}$
  - $ForceX{=}\perp$ when X is not intervened on; otherwise it is the value X is set to.
  - $$P(X \mid ForceX, Zs) = \begin{cases} P(X \mid Zs) & \text{if } ForceX{=}\perp \\ ForceX & \text{otherwise} \end{cases}$$

Causality

In a causal model:

- To intervene on a variable:
  - remove the arcs into the variable from its parents
  - set the value of the variable
- An intervention has a different effect than an observation.
- Intervening on a variable only affects its descendants.
- Can be modeled by variable X with parents Zs having a new parent, "$ForceX$"
  - $domain(ForceX) = domain(X) \cup \{\bot\}$
  - $ForceX{=}\bot$ when X is not intervened on; otherwise it is the value X is set to.
  - $$P(X \mid ForceX, Zs) = \begin{cases} P(X \mid Zs) & \text{if } ForceX{=}\bot \\ ForceX & \text{otherwise} \end{cases}$$
  - $do(X{=}v)$ becomes $ForceX{=}v$

Causality

- One of the following is a better causal model of the world:

...same as belief networks, but different as causal networks

Causality

- One of the following is a better causal model of the world:

  ...same as belief networks, but different as causal networks

- We can't learn causal models from observational data unless we are prepared to make modeling assumptions.

Causality

- One of the following is a better causal model of the world:

  ...same as belief networks, but different as causal networks
- We can't learn causal models from observational data unless we are prepared to make modeling assumptions.
- Causal models can be learned from randomized experiments — assuming the randomization isn't correlated with other variables.

Causality

- One of the following is a better causal model of the world:

  ...same as belief networks, but different as causal networks

- We can't learn causal models from observational data unless we are prepared to make modeling assumptions.
- Causal models can be learned from randomized experiments — assuming the randomization isn't correlated with other variables.
- Conjecture: causal belief networks are more natural and more concise than non-causal networks.
- Conjecture: causal model are more stable to changing circumstances (transportability)