Causal Bayesian networks

Peter Antal
antal@mit.bme.hu

Outline

- Can we represent exactly (in)dependencies by a BN?
  - From a causal model? Suff.&nec.?
- Can we interpret
  - edges as causal relations
    - with no hidden variables?
    - in the presence of hidden variables?
  - local models as autonomous mechanisms?
- Can we infer the effect of interventions?
- Optimal study design to infer the effect of interventions?

Motivation: from observational inference…

- In a Bayesian network, any query can be answered corresponding to passive observations: $p(Q=q|E=e)$.
  - What is the (conditional) probability of $Q=q$ given that $E=e$.
  - Note that Q can preceed temporally E.

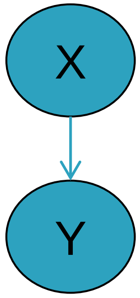

- Specification: $p(X)$, $p(Y|X)$
- Joint distribution: $p(X,Y)$
- Inferences: $p(X)$, $p(Y)$, $p(Y|X)$, $p(X|Y)$

Motivation: to interventional inference...

- Perfect intervention: do(X=x) as set X to x.
- What is the relation of $p(Q=q|E=e)$ and $p(Q=q|do(E=e))$?

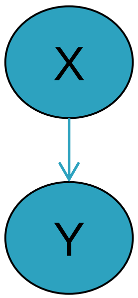

- Specification: $p(X)$, $p(Y|X)$
- Joint distribution: $p(X,Y)$
- Inferences:
  - $p(Y|X=x)=p(Y|do(X=x))$
  - $p(X|Y=y)\neq p(X|do(Y=y))$

- What is a formal knowledge representation of a causal model?
- What is the formal inference method?

Motivation: and to counterfactual inference

- Imagery observations and interventions:
  - We observed X=x, but imagine that x' would have been observed: denoted as X'=x'.
  - We set X=x, but imagine that x' would have been set: denoted as do(X'=x').
- What is the relation of
  - Observational p(Q=q|E=e, X=x')
  - Interventional p(Q=q|E=e, do(X=x'))
  - Counterfactual p(Q'=q'|Q=q, E=e, do(X=x), do(X'=x'))

- O: What is the probability that the patient recovers if he takes the drug *x'*.
- I:What is the probability that the patient recovers if we prescribe* the drug *x'*.
- C: Given that the patient did not recovered for the drug x, what would have been the probability that patient recovers if we had prescribed* the drug *x'*, instead of *x*.
- ~C (time-shifted): Given that patient did not recovered for the drug *x* and he has not respond well**, what is the probability that patient will recover if we change the prescribed* drug *x to x'*.

- *: Assume that the patient is fully compliant.
- **" expected to neither he will.

Challenges in a complex domain

The domain is defined by the joint distribution

$$P(X_1,\ldots, X_n|\text{Structure},\text{parameters})$$

1. Representation of parameteres

   „small number of parameters”

2. Representation of independencies

   „what is relevant for diagnosis”

3. Representation of causal relations

   „what is the effect of a treatment”

4. Representation of possible worlds

?

quantitave

qualitative

passive (observational)

Active (interventional)

Imagery (counterfactual)

Principles of causality

- strong association,
- X precedes temporally Y,
- plausible explanation without alternative explanations based on confounding,
- necessity (generally: if cause is removed, effect is decreased or actually: y would not have been occurred with that much probability if x had not been present),
- sufficiency (generally: if exposure to cause is increased, effect is increased or actually: y would have been occurred with larger probability if x had been present).

- Autonomous, transportable mechanism.

- The probabilistic definition of causation formalizes many, but for example not the counterfactual aspects.

Conditional independence

$I_P(X;Y|Z)$ or $(X \perp\!\!\!\perp Y|Z)_P$ denotes that X is independent of Y given Z: $P(X;Y|z)=P(Y|z)\,P(X|z)$ for all z with $P(z)>0$.

(Almost) alternatively, $I_P(X;Y|Z)$ iff

$P(X|Z,Y)= P(X|Z)$ for all z,y with $P(z,y)>0$.

Other notations: $D_P(X;Y|Z) =\text{def}= \neg\, I_P(X;Y|Z)$

Contextual independence: for not all z.

The independence model of a distribution

The independence map (model) M of a distribution P is the set of the valid independence triplets:

$M_P = \{I_{P,1}(X_1;Y_1|Z_1), \ldots, I_{P,K}(X_K;Y_K|Z_K)\}$

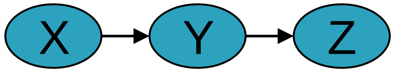

If P(X,Y,Z) is a Markov chain, then

$M_P = \{D(X;Y), D(Y;Z), I(X;Z|Y)\}$

Normally/almost always: D(X;Z)

Exceptionally: I(X;Z)

The independence map of a N-BN

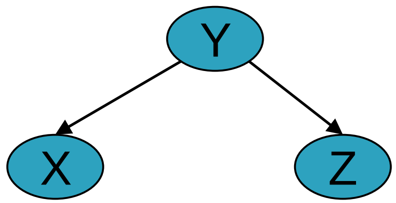

If P(Y,X,Z) is a naive Bayesian network, then
M_P={D(X;Y), D(Y;Z), I(X;Z|Y)}
Normally/almost always: D(X;Z)
Exceptionally: I(X;Z)

Bayesian networks

- Directed acyclic graph (DAG)
  - nodes – random variables/domain entities
  - edges – direct probabilistic dependencies (edges– causal relations
- Local models – $P(X_i|Pa(X_i))$
- Three interpretations:

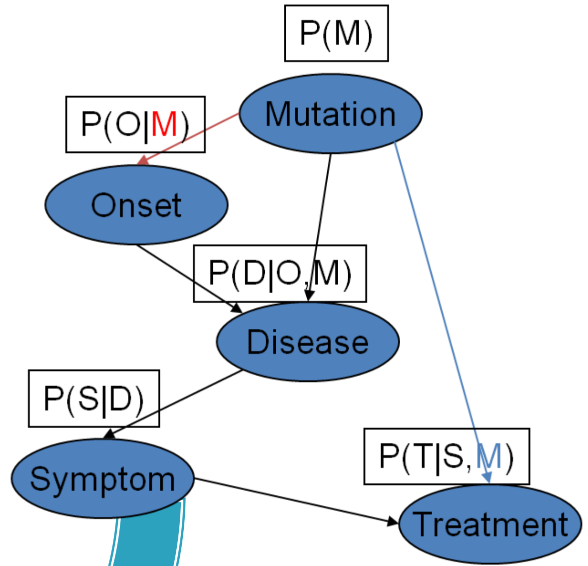

1. Causal model

3. Concise representation of joint distributions

$$P(M,O,D,S,T) =$$
$$P(M)P(O|M)P(D|O,M)P(S|D)P(T|S,M)$$

$M_P=\{I_{P,1}(X_1;Y_1|Z_1),\ldots\}$

2. Graphical representation of (in)dependencies

Inferring independencies from structure: d-separation

$I_G(X;Y|Z)$ denotes that X is d-separated (directed separated) from Y by Z in directed graph G.

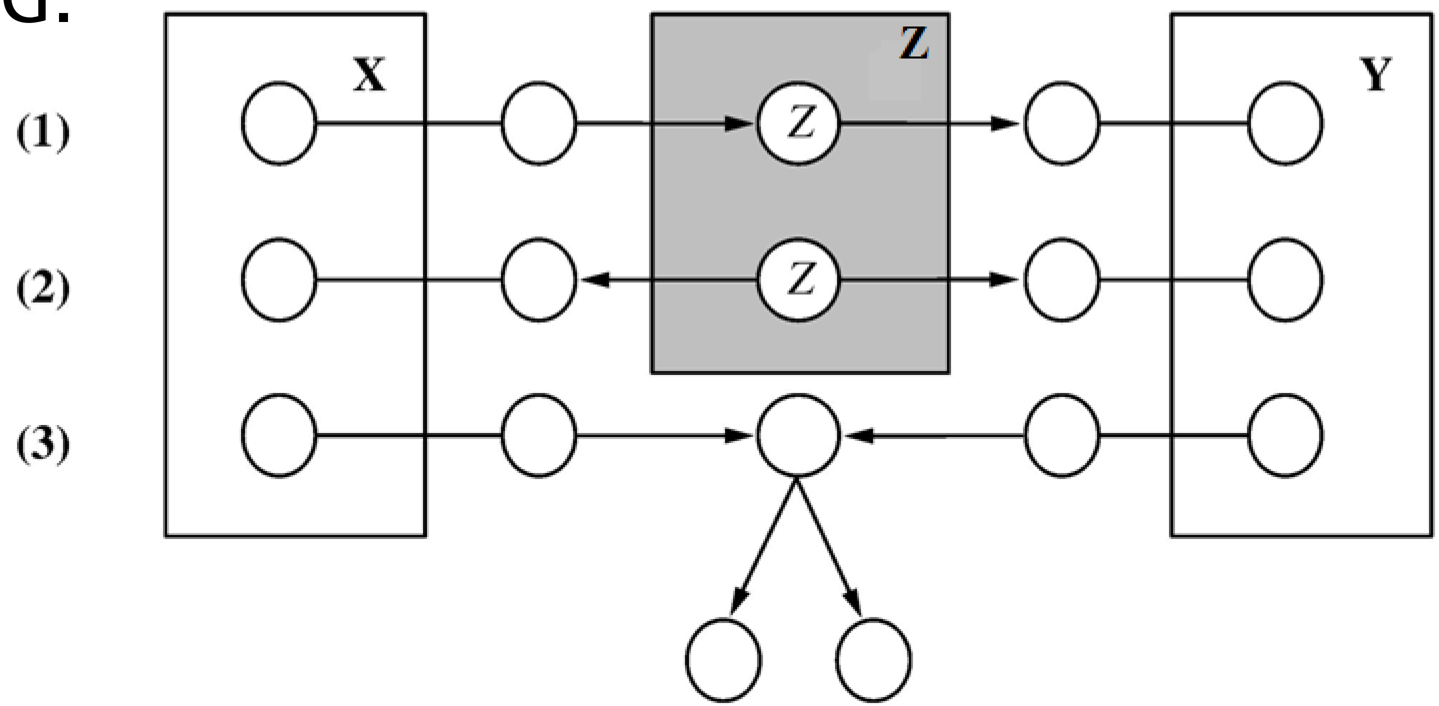

d-separation and the global Markov condition

Definition 7 *A distribution $P(X_1, \ldots, X_n)$ obeys the global Markov condition w.r.t. DAG G, if*

$$\forall\, X, Y, Z \subseteq U \; (X \perp\!\!\!\perp Y | Z)_G \Rightarrow (X \perp\!\!\!\perp Y | Z)_P, \tag{9}$$

where $(X \perp\!\!\!\perp Y | Z)_G$ denotes that X and Y are d-separated by Z, that is if every path p between a node in X and a node in Y is blocked by Z as follows

1. *either path p contains a node n in Z with non-converging arrows (i.e. $\rightarrow n \rightarrow$ or $\leftarrow n \rightarrow$),*
2. *or path p contains a node n not in Z with converging arrows (i.e. $\rightarrow n \leftarrow$) and none of its descendants of n is in Z.*

Representation of independencies

D-separation provides a sound and complete, computationally efficient algorithm to read off an (in)dependency model consisting the independencies that are valid in all distributions Markov relative to G, that is $\forall\ X, Y, Z \subseteq V$

$$(X \perp\!\!\!\perp Y|Z)_G \Leftrightarrow ((X \perp\!\!\!\perp Y|Z)_P \text{ in all P Markov relative to G}). \quad (10)$$

For certain distributions exact representation is not possible by Bayesian networks, e.g.:

1. Intransitive Markov chain: X➔Y➔Z
2. Pure multivariate cause: {X,Z}➔Y
3. Diamond structure:

P(X,Y,Z,V) with M_P={D(X;Z), D(X;Y), D(V;X), D(V;Z), I(V;Y|{X,Z}), I(X;Z|{V,Y}).. }.

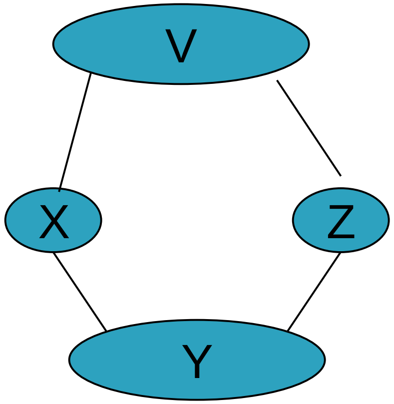

Association vs. Causation: Markov chain

Causal models:

Markov chain

$P(X_1,...)$

$M_P=\{I(X_{i+1};X_{i-1}|X_i)\}$

„first order Markov property”

Flow of time?

The building block of causality: v-structure (arrow of time)

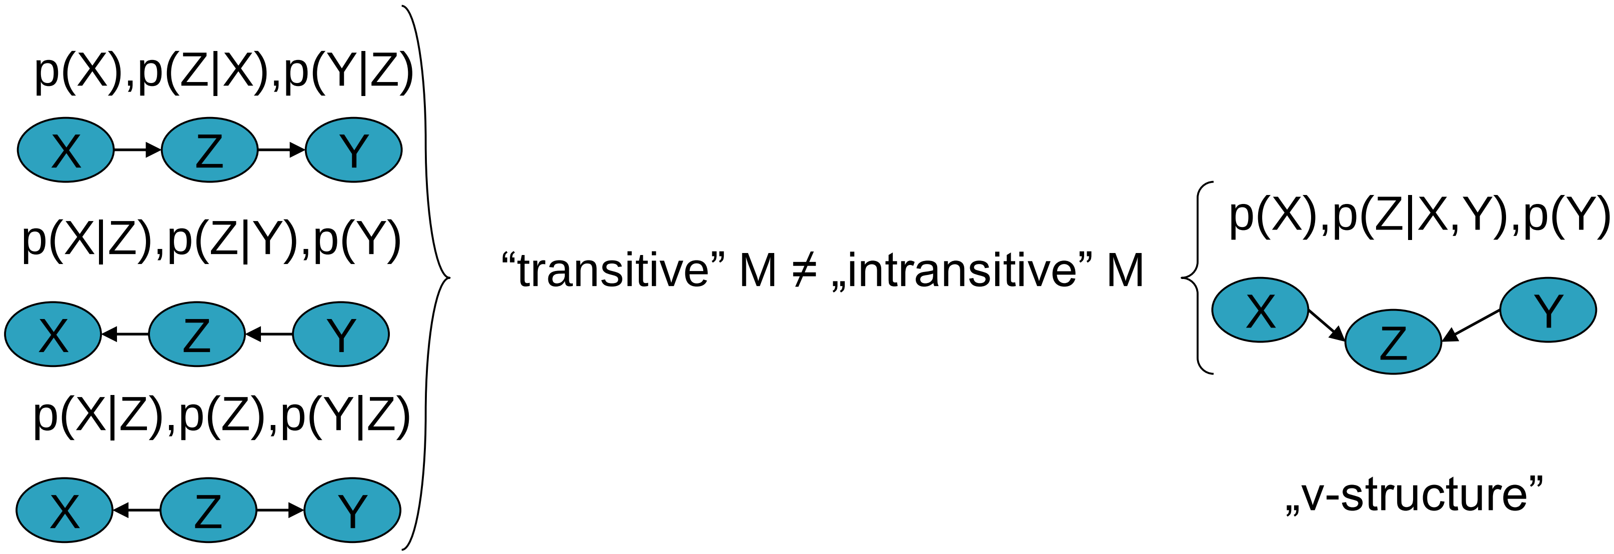

M_P={D(X;Z), D(Z;Y), D(X,Y), I(X;Y|Z)}

M_P={D(X;Z), D(Y;Z), I(X;Y), D(X;Y|Z) }

Often: present knowledge renders future states conditionally independent. (confounding)

Ever(?): present knowledge renders past states conditionally independent. (backward/atemporal confounding)

Interventional inference in causal Bayesian networks

- (Passive, observational) inference
  - P(Query|Observations)
- **Interventionist inference**
  - **P(Query|Observations, Interventions)**
- Counterfactual inference
  - P(Query| Observations, Counterfactual conditionals)

Interventions and graph surgery

If G is a causal model, then compute p(Y|do(X=x)) by

1. deleting the incoming edges to X
2. setting X=x
3. performing standard Bayesian network inference.

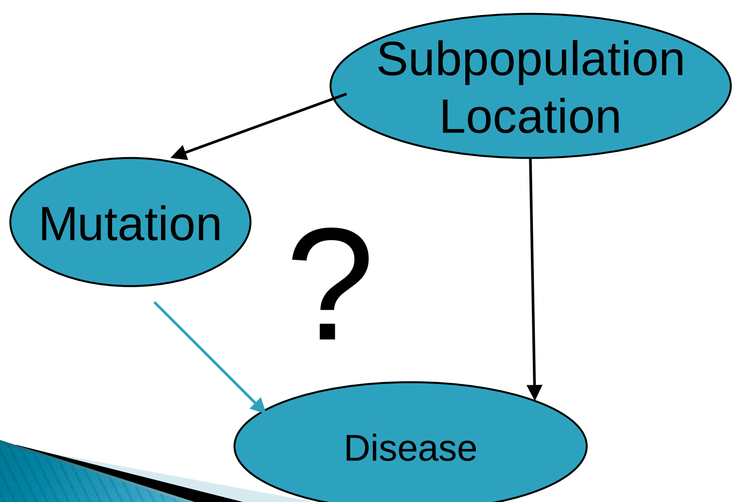

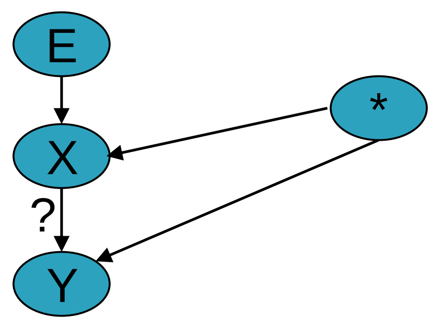

Summary

- Can we represent exactly (in)dependencies by a BN?
  - ***almost always***
- Can we interpret
  - edges as causal relations
    - with no hidden variables?
      - ***compelled edges as a filter***
    - in the presence of hidden variables?
      - ***Sometimes, e.g. confounding can be excluded in certain cases***
    - in local models as autonomous mechanisms?
      - ***a priori knowledge, e.g. Causal Markov Assumption***
- Can we infer the effect of interventions in a causal model?
  - ***Graph surgery with standard inference in BNs***
- Optimal study design to infer the effect of interventions?
  - ***With no hidden variables: yes, in a non-Bayesian framework***
  - ***In the presence of hidden variables: open issue***
- Suggested reading
  - J. Pearl: Causal inference in statistics, 2009