Two Optimal Strategies for Active Learning of Causal Models from Interventions

Alain Hauser Peter Bühlmann

Seminar für Statistik, ETH Zürich

PGM 2012, Granada

Causal model: example

Random variables:
X_1: taxis honking
X_2: Jonas awake
X_3: Alain awake
X_4: watermelons eaten

Causal model: example

Random variables:
X_1: taxis honking
X_2: Jonas awake
X_3: Alain awake
X_4: watermelons eaten

Directed acyclic graph (**DAG**)
of causal dependencies:

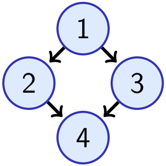

Causal model: example

Random variables:

X_1: taxis honking
X_2: Jonas awake
X_3: Alain awake
X_4: watermelons eaten

Directed acyclic graph (**DAG**) of causal dependencies:

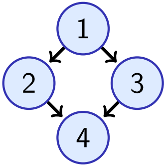

Factorization of density:

$$f(x) = f(x_1)f(x_2|x_1)f(x_3|x_1)f(x_4|x_2, x_3)$$

f has **Markov property** of D

Intervention: example

Random variables:
X_1: taxis honking
X_2: Jonas awake
X_3: Alain awake
X_4: watermelons eaten

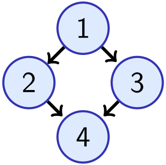

True DAG D

Observational density: $f(x) = f(x_1)f(x_2|x_1)f(x_3|x_1)f(x_4|x_2, x_3)$

Intervention: example

Random variables:

X_1: taxis honking

X_2: Jonas awake

X_3: Alain awake

X_4: watermelons eaten

Intervention at X_2: waking Jonas

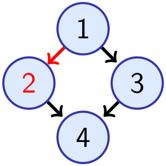

Observational density: $f(x) = f(x_1) f(x_2|x_1) f(x_3|x_1) f(x_4|x_2, x_3)$

Intervention: example

Random variables:
X_1: taxis honking
X_2: Jonas awake
X_3: Alain awake
X_4: watermelons eaten

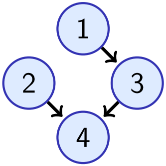

Intervention DAG $D^{(\{2\})}$

Observational density: $f(x) = f(x_1)f(x_2|x_1)f(x_3|x_1)f(x_4|x_2,x_3)$

Interventional density: $f(x|\text{do}(X_2 = U)) = f(x_1)\tilde{f}(x_2)f(x_3|x_1)f(x_4|x_2,x_3)$

Markov equivalence

A probability density in general obeys the Markov properties of **several** DAGs; those DAGs are called **Markov equivalent**
$\rightsquigarrow$ **limited identifiability** under observational data

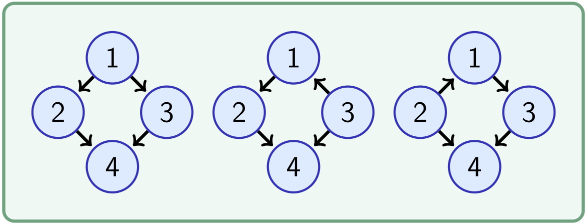

D D_1 D_2

Markov equivalence

A probability density in general obeys the Markov properties of **several** DAGs; those DAGs are called **Markov equivalent**
$\rightsquigarrow$ **limited identifiability** under observational data

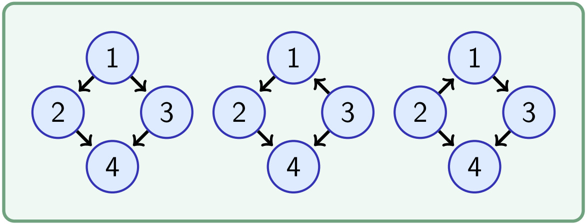

D D_1 D_2

On the other hand, intervention effects **do** depend on the DAG
$\rightsquigarrow$ **improved identifiability** of causal models under interventional data

Interventional Markov equivalence

Assume experiment in which **different** interventions at targets $I_1, I_2, \ldots$ are performed, summarized as **family of targets** $\mathcal{I} = \{I_1, I_2, \ldots\}$.
Note: observational case corresponds to special family $\mathcal{I} = \{\emptyset\}$

Definition ($\mathcal{I}$-Markov equivalence; Hauser and Bühlmann, 2012)

Given a family of targets $\mathcal{I}$, two DAGs D_1 and D_2 are called **$\mathcal{I}$-Markov equivalent** if they produce the same class of tuples of interventional densities.

Interventional Markov equivalence

Assume experiment in which **different** interventions at targets $I_1, I_2, \ldots$ are performed, summarized as **family of targets** $\mathcal{I} = \{I_1, I_2, \ldots\}$.
Note: observational case corresponds to special family $\mathcal{I} = \{\emptyset\}$

Definition ($\mathcal{I}$-Markov equivalence; Hauser and Bühlmann, 2012)

Given a family of targets $\mathcal{I}$, two DAGs D_1 and D_2 are called **$\mathcal{I}$-Markov equivalent** if they produce the same class of tuples of interventional densities.

In words: two DAGs D_1 and D_2 are $\mathcal{I}$-Markov equivalent if they are **statistically indistinguishable** from data produced from interventions at the targets in $\mathcal{I}$.

Interventional essential graph

Definition

Let $\mathcal{I}$ be a family of targets. The **$\mathcal{I}$-essential graph** of some DAG D is defined as $\mathcal{E}_{\mathcal{I}}(D) := \bigcup_{D' \sim_{\mathcal{I}} D} D'$.

In words: $\mathcal{E}_{\mathcal{I}}(D)$ is a partially directed graph

- having the same skeleton as D
- with a **directed edge** where the corresponding arrows of all DAGs $\mathcal{I}$-equivalent to D have the same orientation
- with an **undirected edge** where the orientation of the corresponding arrow is *not* common to all DAGs $\mathcal{I}$-equivalent to D

Interventional essential graph

Definition

Let $\mathcal{I}$ be a family of targets. The **$\mathcal{I}$-essential graph** of some DAG D is defined as $\mathcal{E}_{\mathcal{I}}(D) := \bigcup_{D' \sim_{\mathcal{I}} D} D'$.

In words: $\mathcal{E}_{\mathcal{I}}(D)$ is a partially directed graph

- having the same skeleton as D
- with a **directed edge** where the corresponding arrows of all DAGs $\mathcal{I}$-equivalent to D have the same orientation
- with an **undirected edge** where the orientation of the corresponding arrow is *not* common to all DAGs $\mathcal{I}$-equivalent to D

Properties:

- unique representation of an $\mathcal{I}$-Markov equivalence class
- chain graph with chordal chain components (Hauser and Bühlmann, 2012)

Interventional Markov equivalence: example

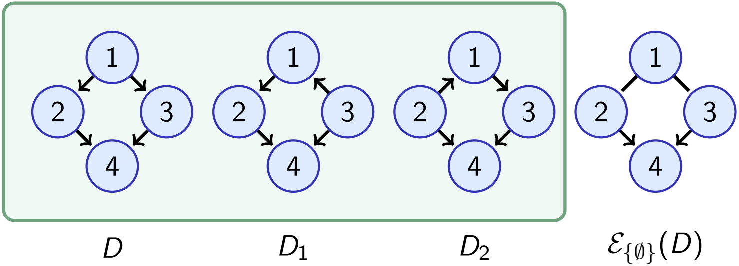

Observational Markov equivalence class of D with corresponding essential graph

Interventional Markov equivalence: example

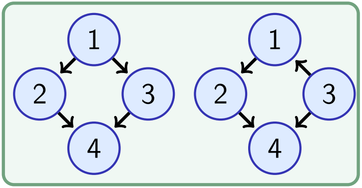

D $\quad D_1$

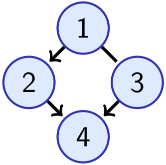

$\mathcal{E}_{\{\emptyset,\{2\}\}}(D)$

Interventional Markov equivalence class of D for family of targets $\mathcal{I} = \{\emptyset, \{2\}\}$. Corresponds to an experiment which measures

- observational data ($I = \emptyset$)
- interventional data from an intervention at X_2 ($I = \{2\}$)

Active learning: overview

Up to now: given list of interventions; characterization of identifiability via interventional essential graphs

Active learning: overview

Up to now: given list of interventions; characterization of identifiability via interventional essential graphs

Problem

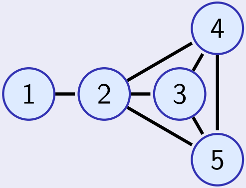

Given list of interventions performed so far and corresponding interventional essential graph, find "optimal" intervention target for maximal improvement of identifiability of causal models

Active learning: overview

Up to now: given list of interventions; characterization of identifiability via interventional essential graphs

Problem

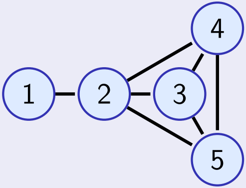

Given list of interventions performed so far and corresponding interventional essential graph, find "optimal" intervention target for maximal improvement of identifiability of causal models

Objectives: assessing identifiability

- Number of edges orientable after one (single-vertex) intervention $\rightsquigarrow$ OPTSINGLE

Active learning: overview

Up to now: given list of interventions; characterization of identifiability via interventional essential graphs

Problem

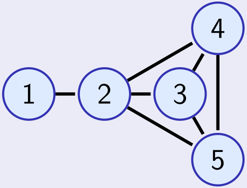

Given list of interventions performed so far and corresponding interventional essential graph, find "optimal" intervention target for maximal improvement of identifiability of causal models

Objectives: assessing identifiability

- Number of edges orientable after one (single-vertex) intervention $\rightsquigarrow$ OptSingle
- Number of interventions (at arbitrary targets) needed for full identifiability $\rightsquigarrow$ OptUnb

OptSingle: overview

- Yields **single-vertex intervention** that **maximizes number of orientable edges** in worst case
- Implementation: local algorithm that finds optimal intervention target in "local" fashion, only considering neighborhood of candidate vertices
- Complexity: in worst case exponential, depending on clique number of $\mathcal{I}$-essential graph

OptSingle: worst case example

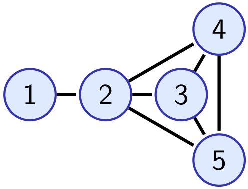

OptSingle: worst case example

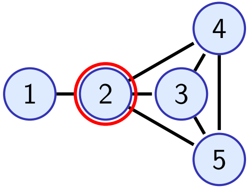

OptSingle: Find vertex that **guarantees** orientability of a maximum of edges after intervention

OptSingle: worst case example

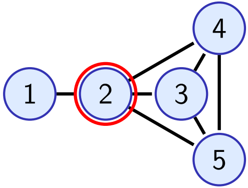

Intervention at vertex 2

OptSingle: worst case example

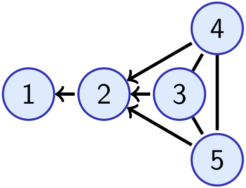

OptSingle: worst case example

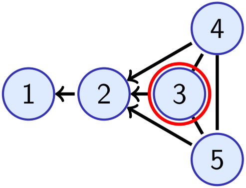

OptSingle: Find vertex that **guarantees** orientability of a maximum of edges after intervention

OptSingle: worst case example

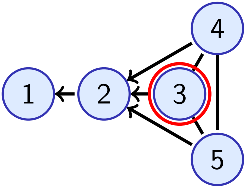

Intervention at vertex 3

OptSingle: worst case example

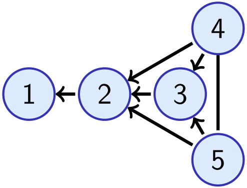

OptSingle: worst case example

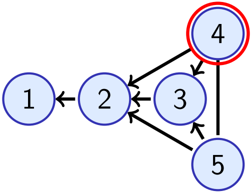

OptSingle: Find vertex that **guarantees** orientability of a maximum of edges after intervention

OptSingle: worst case example

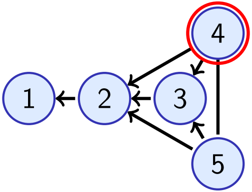

Intervention at vertex 4

OptUnb: overview

- Yields an intervention target (of arbitrary size) that **maximally reduces clique number** of interventional essential graph
- Iterative application of OptUnb yields **minimum** set of intervention targets that guarantees full identifiability for all causal models in interventional Markov equivalence class
- Implementation: based on LexBFS (Rose, 1970) and greedy coloring; exploits chordality of chain components
- Complexity: linear-time algorithm
- Proof of optimality proves conjecture of Eberhardt (2008) concerning number of interventions necessary and sufficient for full identifiability

OptUnb: worst case example

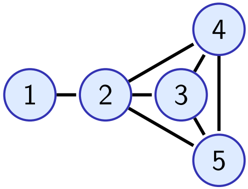

OptUnb: worst case example

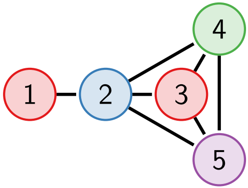

Optimal coloring using greedy coloring on LexBFS-ordering (Rose, 1970)

OptUnb: worst case example

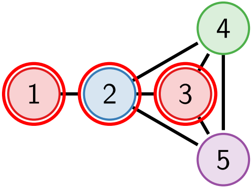

Intervention at lower half of colors

OptUnb: worst case example

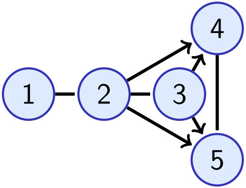

OptUnb: worst case example

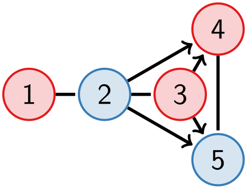

Optimal coloring using greedy coloring on LexBFS-ordering (Rose, 1970)

OptUnb: worst case example

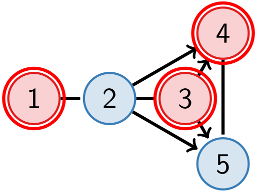

Intervention at lower half of colors

Evaluating active learning algorithms: simulation results

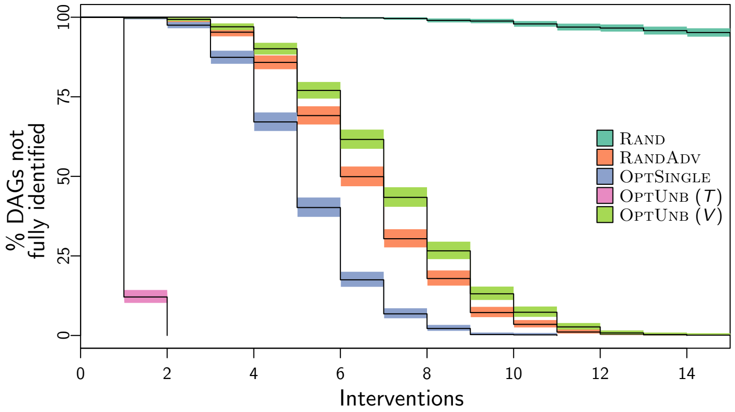

Number of intervention steps needed for full identifiability of DAGs ($p = 40$), measured in targets (T) or intervened variables (V). Thin lines: Kaplan-Meier estimates; colored bands: 95% confidence region.

Conclusions

- Causal models not fully identifiable from observational data
- Interventional data improves identifiability; information gain depends on intervention target

Conclusions

- Causal models not fully identifiable from observational data
- Interventional data improves identifiability; information gain depends on intervention target
- Two new active learning algorithms for proposing valuable intervention targets:
  - OptSingle: finds single-vertex target that maximizes number of orientable edges after intervention
  - OptUnb: finds target that maximally reduces clique size of interventional essential graph; iterative application guarantees full identifiability with a minimum number of interventions

Conclusions

- Causal models not fully identifiable from observational data
- Interventional data improves identifiability; information gain depends on intervention target
- Two new active learning algorithms for proposing valuable intervention targets:
  - OptSingle: finds single-vertex target that maximizes number of orientable edges after intervention
  - OptUnb: finds target that maximally reduces clique size of interventional essential graph; iterative application guarantees full identifiability with a minimum number of interventions
- Both strategies lead to significantly faster identification of causal models than randomly chosen interventions

References

F. Eberhardt. Almost optimal intervention sets for causal discovery. In *UAI*, pages 161–168, 2008.

A. Hauser and P. Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *JMLR*, 13:2409–2464, 2012.

D.J. Rose. Triangulated graphs and the elimination process. *Journal of Mathematical Analysis and Applications*, 32(3):597–609, 1970.

¡Muchas gracias!