StaRAI: From a Probabilistic Propositional Model to a Highly Compressed Probabilistic Relational Model

ECSQARU 2025 – Hagen, Germany

Marcel Gehrke[1], Malte Luttermann[2]

[1]Institute for Humanities-Centered Artificial Intelligence, University of Hamburg, Germany
[2]German Research Center for Artificial Intelligence (DFKI), Lübeck, Germany

September 23, 2025

Agenda

1. Introduction to relational models [Marcel]
2. Compressing probabilistic relational models [Malte]
3. Application: Lifted causal inference [Malte]
   - Lifted computation of causal effects
   - Lifted computation of causal effects with partial causal knowledge
4. Summary [Marcel]

An Ice Cream Example

- Correlation $\neq$ causation

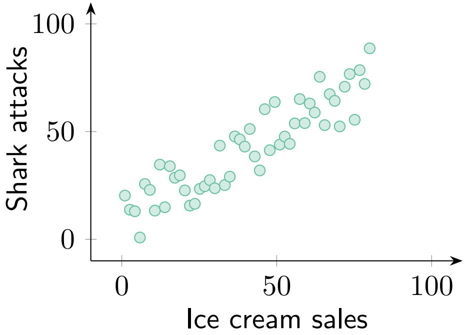

An Ice Cream Example

- Correlation $\neq$ causation

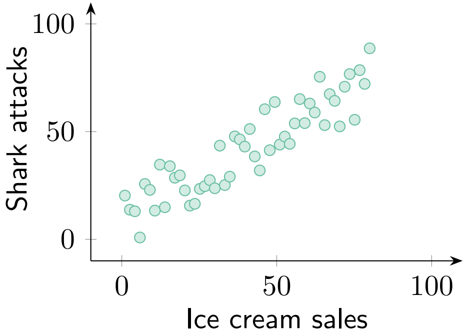

- Possible causal explanations:

1. $I \longrightarrow A$

An Ice Cream Example

- Correlation $\neq$ causation

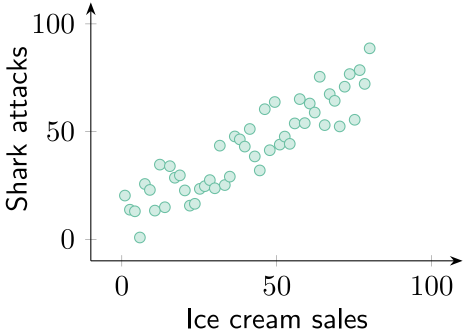

- Possible causal explanations:

1. $I \rightarrow A$
2. $I \leftarrow A$

An Ice Cream Example

- Correlation $\neq$ causation

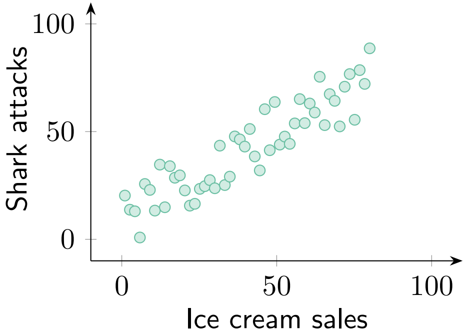

- Possible causal explanations:

1. 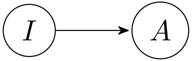

2. 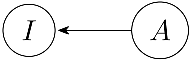

3. 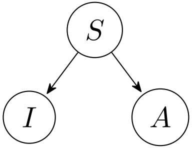

Explanation of the Ice Cream Example Data

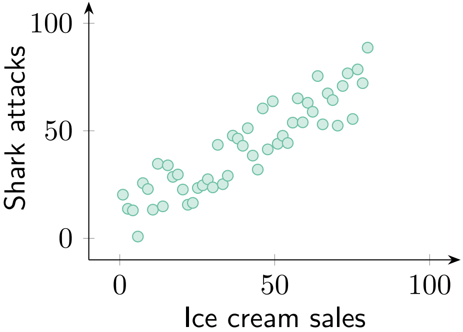

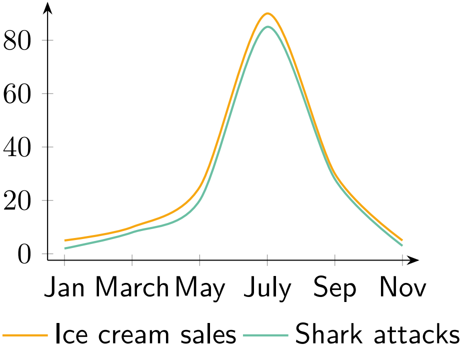

Learnings from the Ice Cream Example

- For *prediction*, correlation is sufficient

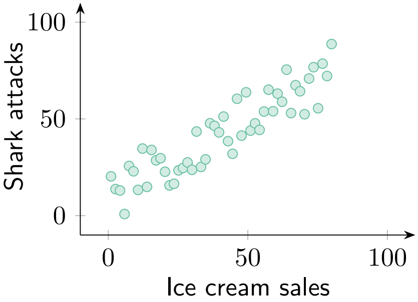

Learnings from the Ice Cream Example

- For *prediction*, correlation is sufficient
  - E.g., knowing ice cream sales suffices to predict shark attacks

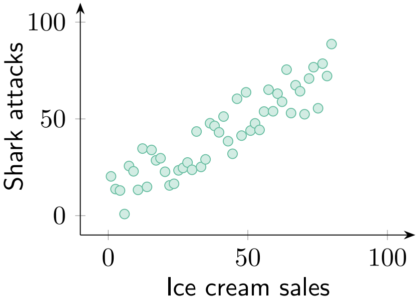

Learnings from the Ice Cream Example

- For *prediction*, correlation is sufficient
  - E.g., knowing ice cream sales suffices to predict shark attacks
- For *decision making (acting)*, causal information is required

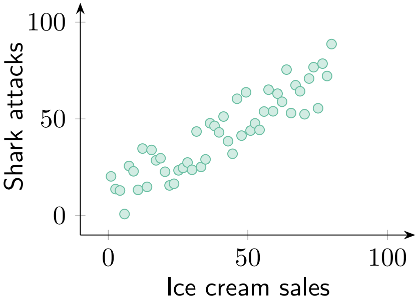

Learnings from the Ice Cream Example

- For *prediction*, correlation is sufficient
  - E.g., knowing ice cream sales suffices to predict shark attacks
- For *decision making (acting)*, causal information is required
  - E.g., Reducing ice cream sales will *not* reduce shark attacks

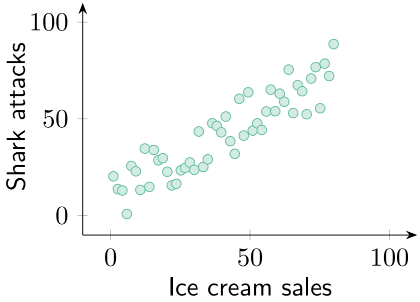

Causal Models

A *causal model* consists of

1. a causal graph G, and
2. a probability distribution P.

A	$\phi_1(A)$
$\vdots$	$\vdots$

A	B	$\phi_2(A, B)$
$\vdots$	$\vdots$	$\vdots$
$\vdots$	$\vdots$	$\vdots$

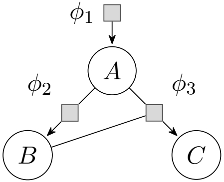

A	B	C	$\phi_3(A, B, C)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$

Brendan J. Frey (2003). »Extending Factor Graphs so as to Unify Directed and Undirected Graphical Models«. *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence (UAI-2003)*. Morgan Kaufmann Publishers Inc., pp. 257–264.

Parametric Causal Factor Graphs

- Parametric causal factor graphs (PCFGs) use logical variables to represent groups of random variables
- Full joint probability distribution encoded as a product over all ground factors

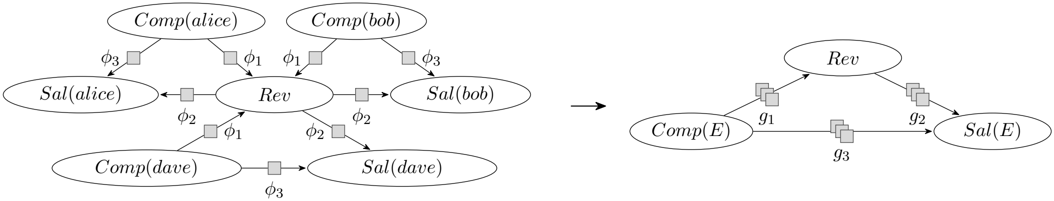

Malte Luttermann, Mattis Hartwig, Tanya Braun, Ralf Möller, and Marcel Gehrke (2024). »Lifted Causal Inference in Relational Domains«. *Proceedings of the Third Conference on Causal Learning and Reasoning (CLeaR-2024)*. PMLR, pp. 827–842.

Lifted Causal Inference in PCFGs I

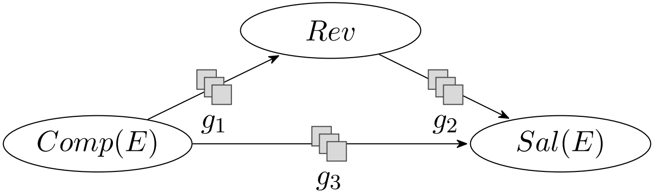

- Is it worth the costs to send an employee to a training course?

$$P(Rev \mid do(Comp(alice) = \mathrm{high})) - P(Rev \mid do(Comp(alice) = \mathrm{low})) = \,?$$

Lifted Causal Inference in PCFGs I

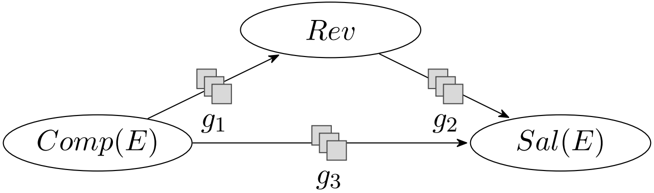

- Is it worth the costs to send an employee to a training course?

$$P(Rev \mid do(Comp(alice) = \text{high})) - P(Rev \mid do(Comp(alice) = \text{low})) = \,?$$

- What effect has sending all employees to a training course on the revenue?

$$P(Rev \mid do(Comp(E) = \text{high})) - P(Rev \mid do(Comp(E) = \text{low})) = \,?$$

Lifted Causal Inference in PCFGs II

- E.g., $P(Rev \mid do(Comp(E) = \text{high}))$
  - Sets fixed value $Comp(E) = \text{high}$
  - Removes incoming influences from $Comp(E)$

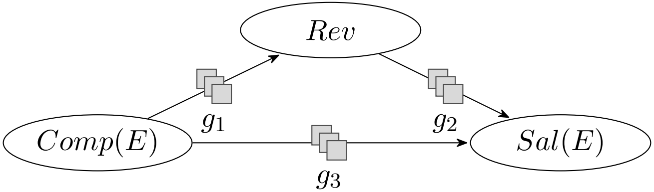

Lifted Causal Inference in PCFGs II

- E.g., $P(Rev \mid do(Comp(E) = \text{high}))$
  - Sets fixed value $Comp(E) = \text{high}$
  - Removes incoming influences from $Comp(E)$

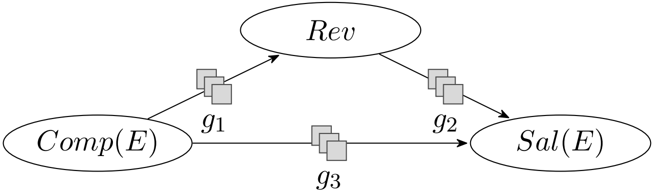

- $do(Comp(E) = \text{high})$ is shorthand for $do(Comp(e_1) = \text{high}, \ldots, Comp(e_k) = \text{high})$, where $\text{dom}(E) = \{e_1, \ldots, e_k\}$
- In non-lifted model, every $e_i \in \text{dom}(E)$ has to be considered separately

Lifted Causal Inference in PCFGs III

- An intervention on a propositional random variable requires splitting of nodes
- E.g., $P(Rev \mid do(Comp(alice) = high))$
  - Removes $alice$ from $Comp(E)$
  - Adds an additional node $Comp(alice)$

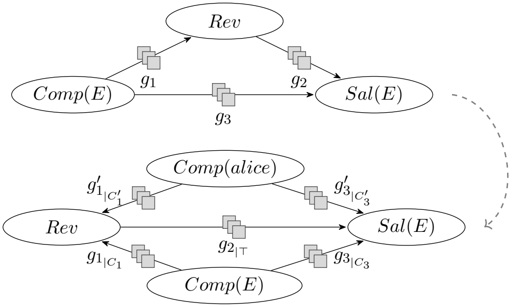

Partially Directed Parametric Causal Factor Graphs

- Often not all causal relationships are known
- Directed edges to represent known causal relationships
- Undirected edges for relationships with unknown causal directions

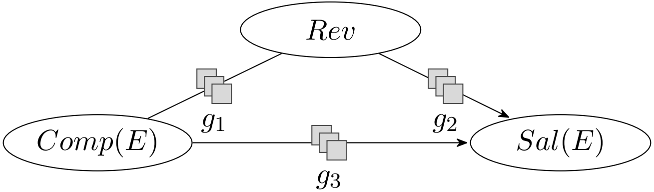

Malte Luttermann, Tanya Braun, Ralf Möller, and Marcel Gehrke (2024). »Estimating Causal Effects in Partially Directed Parametric Causal Factor Graphs«. *Proceedings of the Sixteenth International Conference on Scalable Uncertainty Management (SUM-2024)*. Springer, pp. 265–280.

Lifted Causal Inference in Partially Directed PCFGs I

- An intervention is defined on a fully directed graph
- E.g., $P(Rev \mid do(Comp(E) = \text{high}))$
  - Sets fixed value $Comp(E) = \text{high}$
  - Removes incoming influences from $Comp(E)$

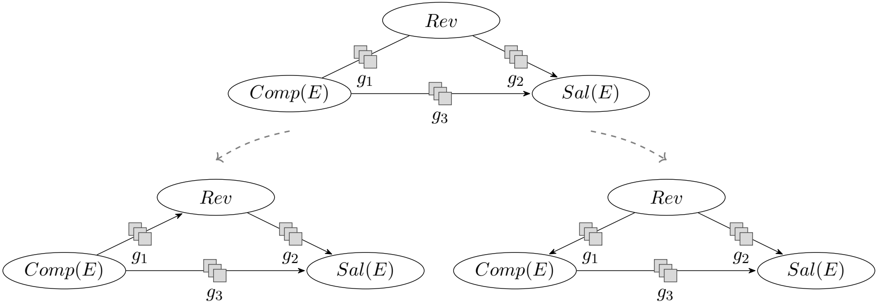

Lifted Causal Inference in Partially Directed PCFGs II

General algorithm:

1. Split nodes of interventional variables (avoid grounding as much as possible)
2. Enumerate relevant edge directions to compute the effect of an action

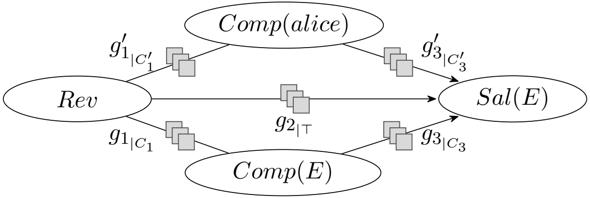

Lifted Causal Inference in Partially Directed PCFGs II

General algorithm:

1. Split nodes of interventional variables (avoid grounding as much as possible)
2. Enumerate relevant edge directions to compute the effect of an action

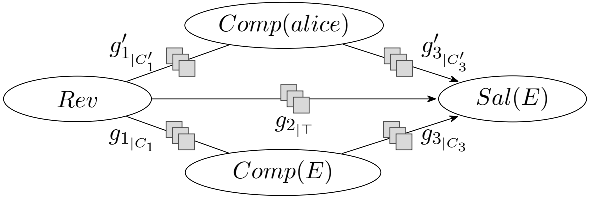

Theorem

To compute the effect of an intervention, it is sufficient to consider the directions of the undirected edges that are connected to the random variables on which we intervene.

Lifted Causal Inference – Summary

- (Partially directed) PCFGs enable lifted causal inference
- Grounding is avoided whenever possible
- Only relevant edge directions of undirected edges are considered for causal inference

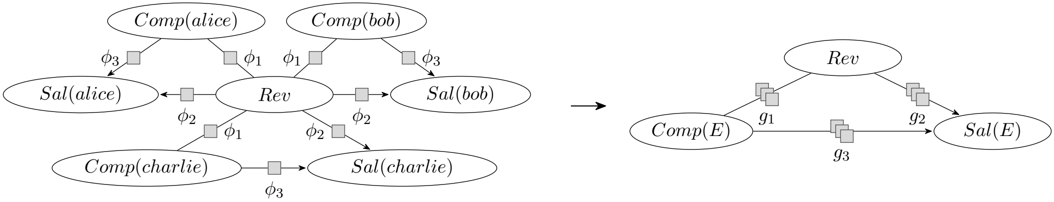

Agenda

1. Introduction to relational models [Marcel]
2. Compressing probabilistic relational models [Malte]
3. Application: Lifted causal inference [Malte]
4. Summary [Marcel]

Bibliography

- Brendan J. Frey (2003). »Extending Factor Graphs so as to Unify Directed and Undirected Graphical Models«. *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence (UAI-2003)*. Morgan Kaufmann Publishers Inc., pp. 257–264.
- Malte Luttermann, Tanya Braun, Ralf Möller, and Marcel Gehrke (2024). »Estimating Causal Effects in Partially Directed Parametric Causal Factor Graphs«. *Proceedings of the Sixteenth International Conference on Scalable Uncertainty Management (SUM-2024)*. Springer, pp. 265–280.
- Malte Luttermann, Mattis Hartwig, Tanya Braun, Ralf Möller, and Marcel Gehrke (2024). »Lifted Causal Inference in Relational Domains«. *Proceedings of the Third Conference on Causal Learning and Reasoning (CLeaR-2024)*. PMLR, pp. 827–842.