Estimating Causal Effects in Partially Directed Parametric Causal Factor Graphs

SUM 2024 – Palermo, Italy

Malte Luttermann[1], Tanya Braun[2], Ralf Möller[3], and Marcel Gehrke[3]

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

November 28, 2024

Motivation

- Goal: Make decisions under uncertainty
- Need to compute the effect of actions
- Need to apply the semantics of an intervention instead of conditioning
  - E.g., $P(Rev \mid do(Comp(alice) = \text{high}))$

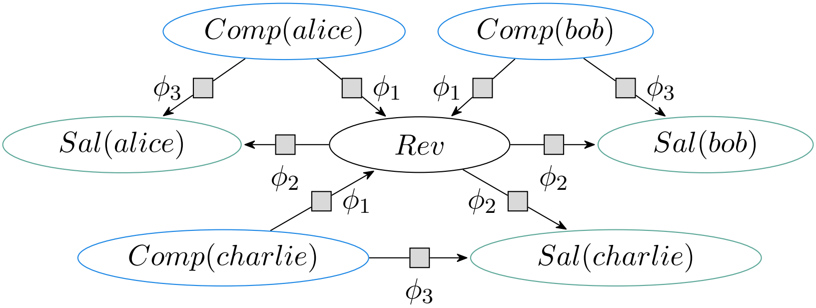

Motivation

- Ideally, we would like to have a first-order (lifted) representation
- Lifting uses a representative of indistinguishable individuals for computations
- Lifting exploits symmetries to speed up probabilistic inference

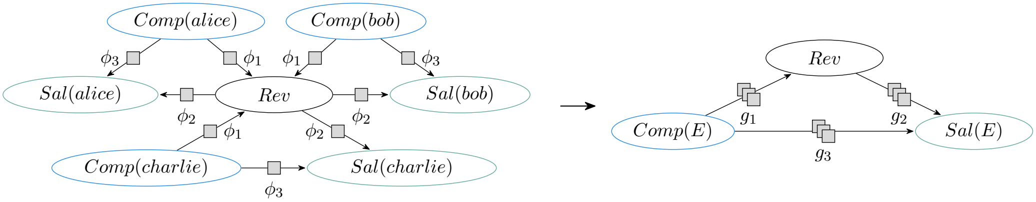

Motivation

- Ideally, we would like to have a first-order (lifted) representation
- Lifting uses a representative of indistinguishable individuals for computations
- Lifting exploits symmetries to speed up probabilistic inference

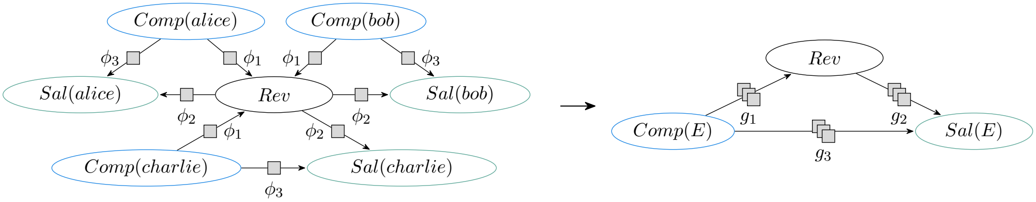

Problem

In general, we do not know all causal relationships.

Problem Setup

- Incorporate partial causal knowledge in a lifted representation
- Estimate causal effects in a partially directed lifted representation

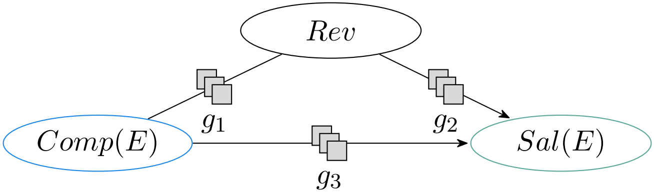

Problem Setup

- Incorporate partial causal knowledge in a lifted representation
- Estimate causal effects in a partially directed lifted representation

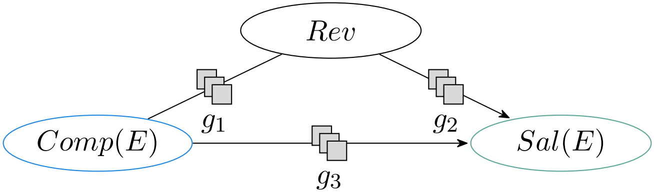

Main Contribution

A formalism to enable first-order decision making with partial causal knowledge.

Partially Directed Parametric Causal Factor Graphs (PPCFGs)

- Directed edges to represent known causal relationships
- Undirected edges for relationships with unknown causal directions
- Logical variables to represent groups of random variables

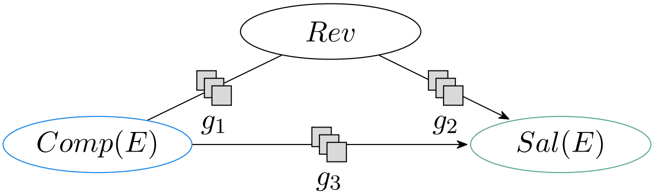

Partially Directed Parametric Causal Factor Graphs (PPCFGs)

- Directed edges to represent known causal relationships
- Undirected edges for relationships with unknown causal directions
- Logical variables to represent groups of random variables

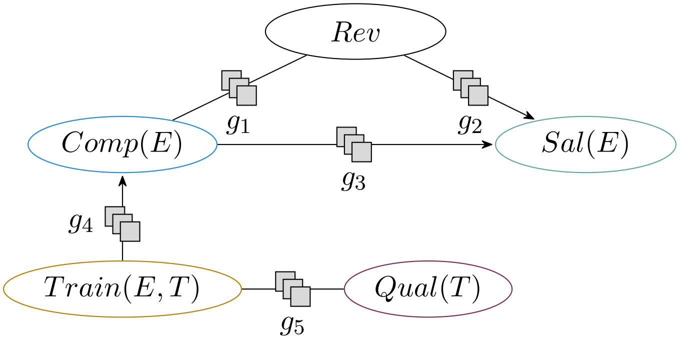

Partially Directed Parametric Causal Factor Graphs (PPCFGs)

- Full joint probability distribution encoded by a product over all ground factors:

$$P_M = \frac{1}{Z} \prod_{g \in \boldsymbol{G}} \prod_{\phi_k \in gr(g)} \phi_k(\mathcal{A}_k)$$

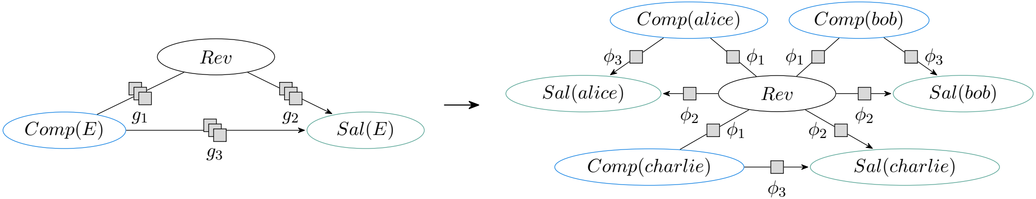

Partially Directed Parametric Causal Factor Graphs (PPCFGs)

- Full joint probability distribution encoded by a product over all ground factors:

$$P_M = \frac{1}{Z} \prod_{g \in \boldsymbol{G}} \prod_{\phi_k \in gr(g)} \phi_k(\mathcal{A}_k)$$

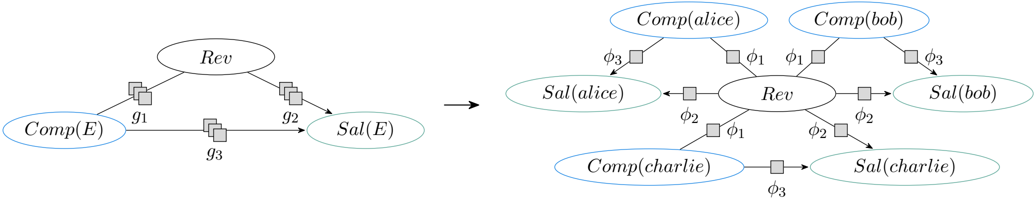

Note:
The definition of the full joint distribution is independent of causal relationships.

Interventions

- Recall: We want to compute the effect of actions
- Is it worth the costs to send an employee to a training course?
- What effect has sending all employees to a training course on the revenue?

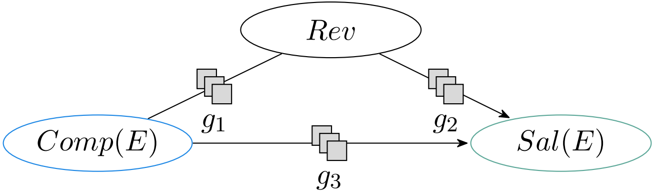

Interventions

- 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 $Comp(E)$

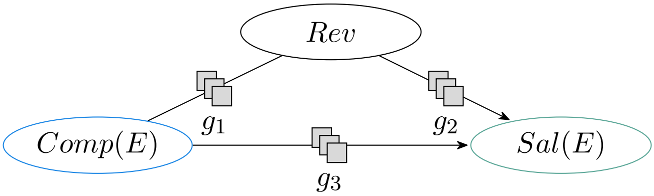

Interventions

- 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 $Comp(E)$

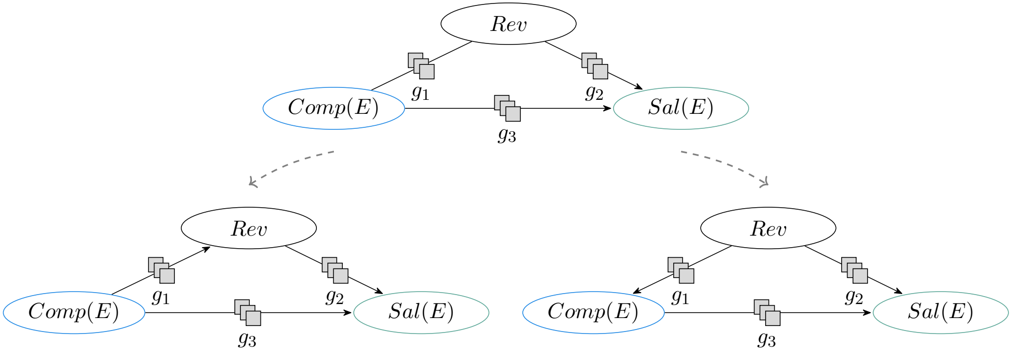

Interventions

- 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 $Comp(alice)$

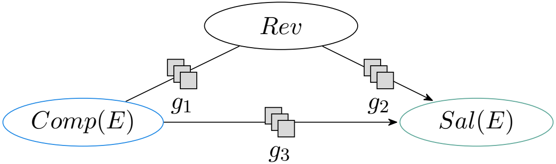

Interventions

- 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 $Comp(alice)$

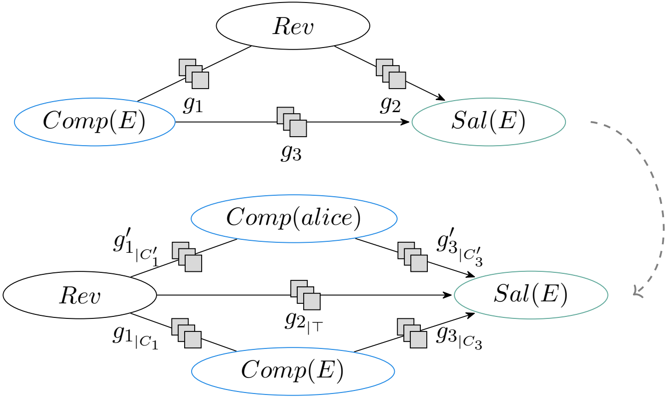

Algorithm to Estimate Causal Effects in PPCFGs

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

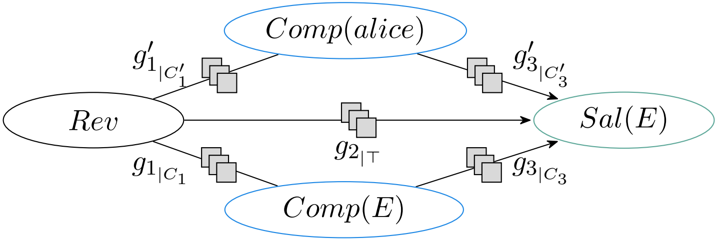

Note (1.)

Our algorithm only grounds necessary parts of the model.

Algorithm to Estimate Causal Effects in PPCFGs

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

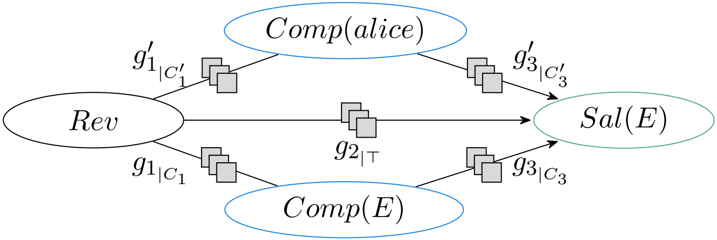

Theorem (2.)

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.

Conclusion

We enable decision making in first-order models with partial causal knowledge:

- PPCFGs as partially directed first-order representations
- Algorithm to estimate causal effects in PPCFGs on a lifted level