Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

Universität Münster

UNIVERSITY OF LÜBECK
INSTITUTE OF INFORMATION SYSTEMS

Lifted Causal Inference in Relational Domains

Malte Luttermann, Mattis Hartwig, Tanya Braun, Ralf Möller, and Marcel Gehrke

1. Motivation

Problem Setup

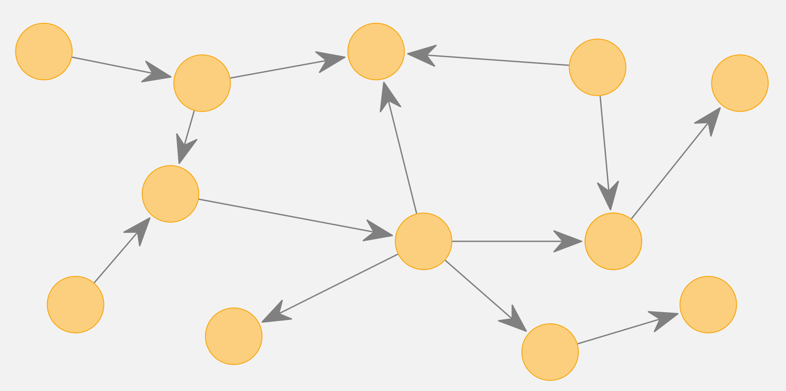

- Our world is inherently relational, i.e., data is often not independent and identically distributed (i.i.d)
  - → Need to represent various objects
  - → Need to represent relationships between objects
- Run time for causal inference increases drastically when considering each object individually

Our Contributions

- Parametric causal factor graphs to efficiently encode causal relationships in relational domains
- Lifted causal inference algorithm to reason about the effects of actions in relational domains

2. Background

- Parametric factor graphs (PFGs) compactly represent a probability distribution over a relational domain
- PFGs are undirected graphical models that use logical variables to represent groups of random variables

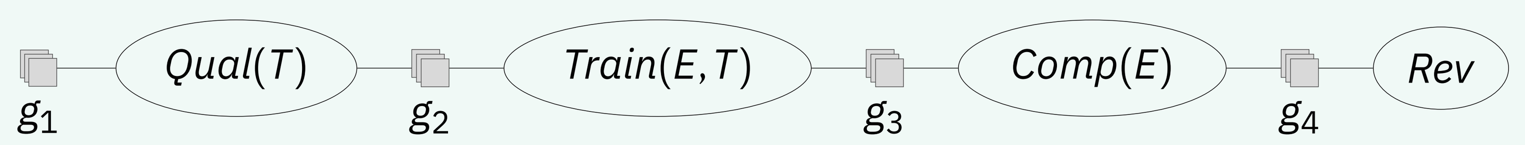

- Semantics of a PFG G over a set of parametric factors $\boldsymbol{G} = \{g_1, \dots, g_m\}$ is defined by grounding the model:

$$P_G = \frac{1}{Z} \prod_{g \in \boldsymbol{G}} \prod_{\phi \in gr(g)} \phi$$

- The ground model for $\mathscr{D}(E) = \{alice, bob, dave, eve\}$ and $\mathscr{D}(T) = \{t_1, t_2\}$ is given by:

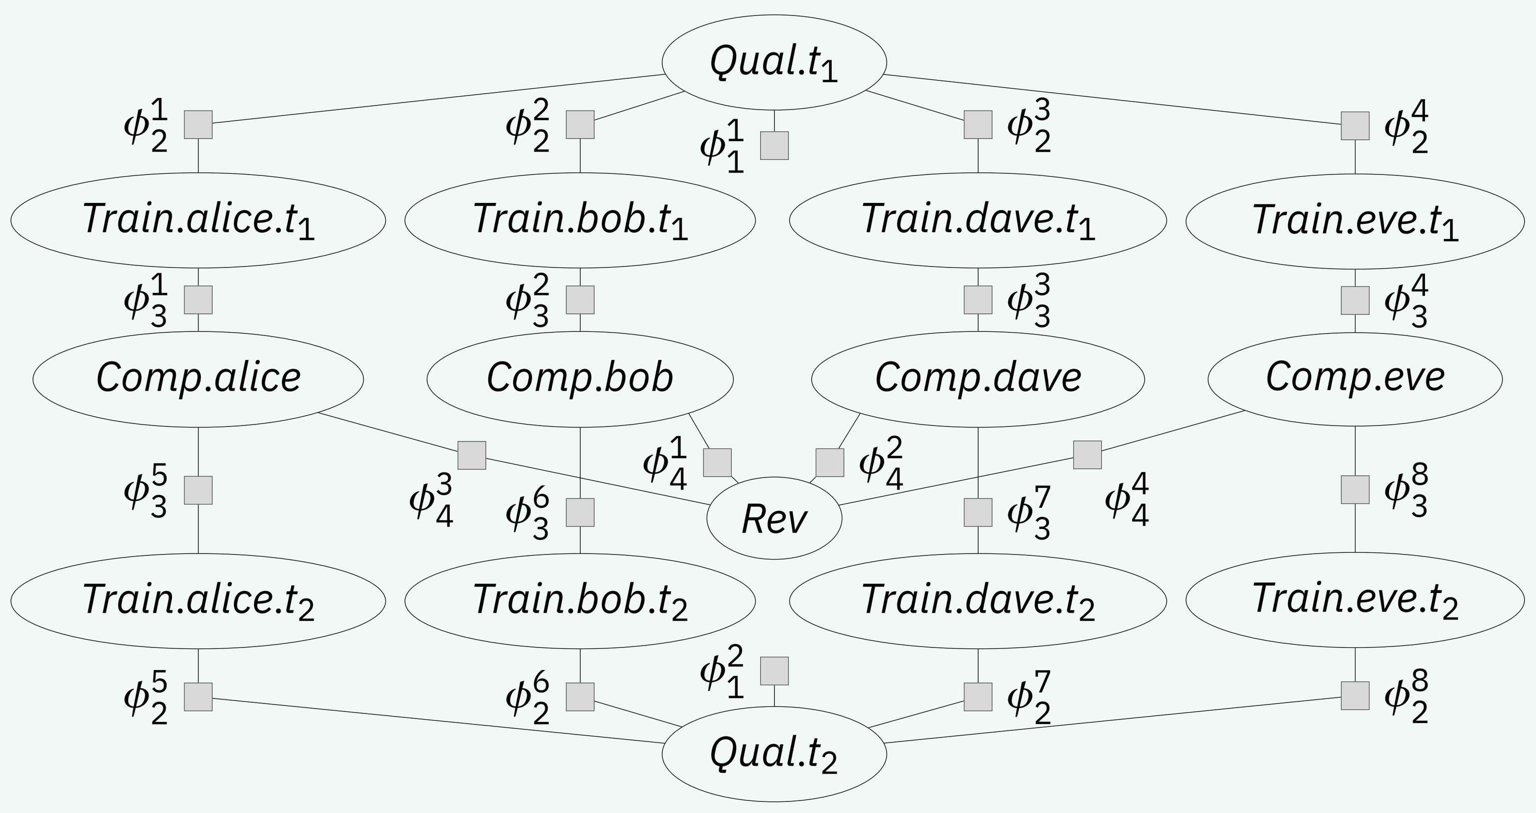

Limitations of Existing Models

(L1) PFGs do not capture causal relationships
(L2) Existing causal models do not support lifted inference

3. Parametric Causal Factor Graphs (PCFGs)

- Goal: Incorporate causal knowledge into PFGs (L1)
- Directed parametric factors represent causal relationships
- Definition of d-separation analogous to Bayesian networks

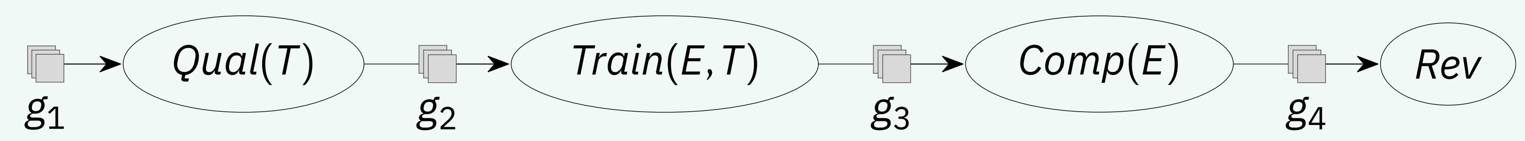

- Semantics is defined as a product over all factors in the ground model (as for PFGs)

4. The Lifted Causal Inference (LCI) Algorithm

- Goal: Lifted computation of causal effects in PCFGs (L2)
- Main steps of LCI:
  - (1) Split directed parametric factors based on the intervention variables
  - (2) Set potentials in parent factors of intervention variables according to the intervention
  - (3) Call a lifted inference algorithm on the modified graph from steps (1)+(2)
- For example, query $P(Rev \mid do(Train(bob, t_1)) = true)$ yields:

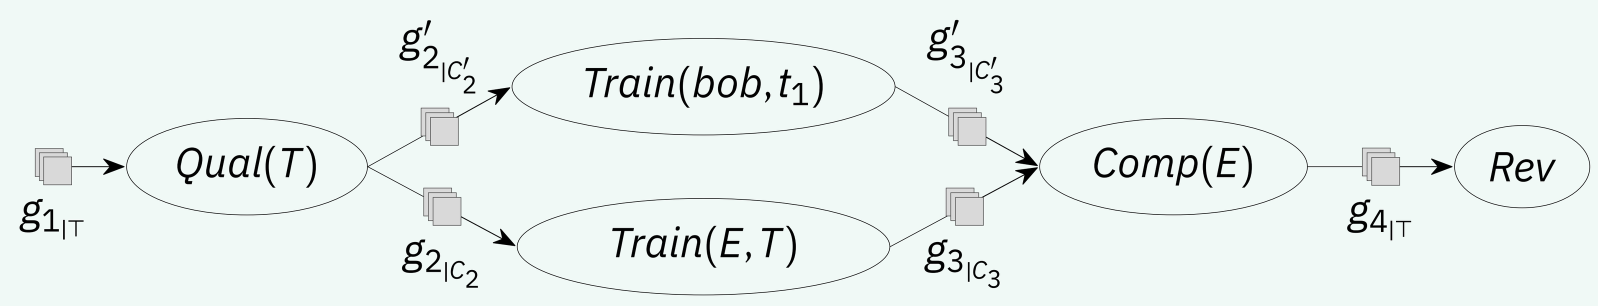

- g'_2 and g'_3 restrict E to bob and T to t_1 while g_2 and g_3 are restricted to the remaining combinations of E and T
- g'_2 maps each assignment of its arguments with $Train(bob, t_1) = true$ to 1 and all other assignments to 0
- In general, LCI can deal with arbitrary queries of the form $P(R_1, \dots, R_\ell \mid do(R'_1 = r'_1, \dots, R'_k = r'_k))$
- LCI also efficiently handles queries on parameterised random variables, such as $P(Rev \mid do(Train(E, t_1)) = true)$

5. Experiments

- Comparison of run times to compute the interventional distribution for a single intervention in the modified graph
- We compare
  - lifted inference on a PCFG,
  - propositional inference on a Bayesian network (BN), and
  - propositional inference on a factor graph (FG)