Compression versus Accuracy: A Hierarchy of Lifted Models

Jan Speller[1], *Malte Luttermann*[2,3], *Marcel Gehrke*[3] *and Tanya Braun*[1]

1. Motivation and Problem Setup

- Factor graphs compactly encode a probability distribution
- Semantics of a factor graph G over a set of factors $\boldsymbol{\Phi}$:

$$P_G = \frac{1}{Z} \prod_{\phi \in \boldsymbol{\Phi}} \phi$$

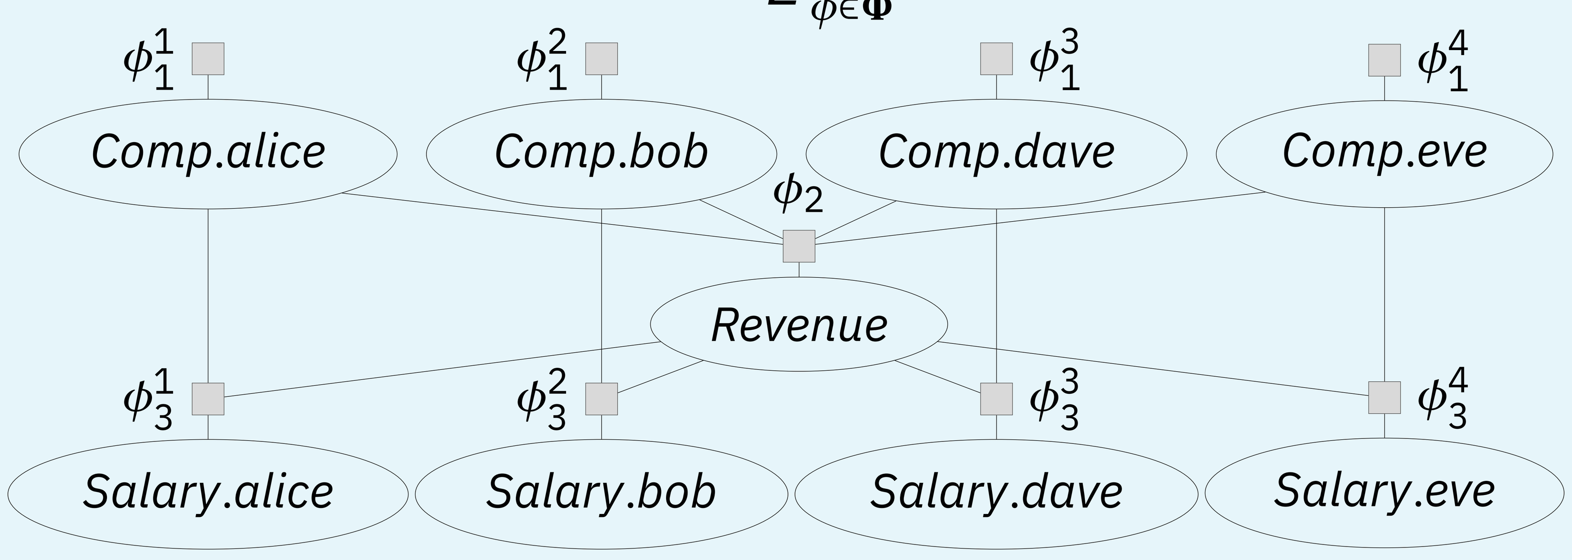

- Enable lifted inference
- Represent groups of random variables by logical variables

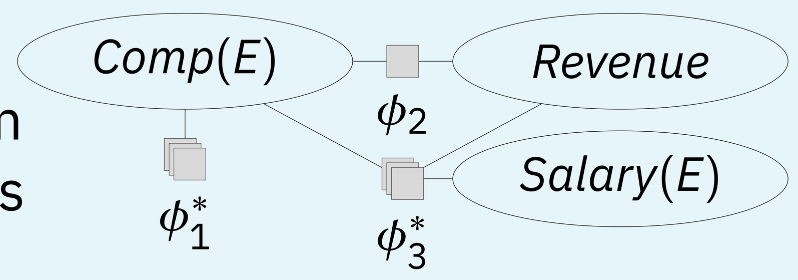

Problem Setup

Input: A factor graph G and level(s) of compression

Output: A hierarchy of parametric factor graph(s) entailing approximately equivalent semantics to G

- With hierarchical grouping structure for different levels of compression
- With hierarchical order of error bounds
- With theoretical guarantees for query results

2. Previous Work: ε-Advanced Colour Passing (ε-ACP)

- Factors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$ are ε-equivalent :⇔ for all potentials $\phi_1(k), \phi_2(k) \in \mathbb{R}_{>0}$ in their potential tables it holds that
  $\phi_1(k) \in [\phi_2(k) \cdot (1-\varepsilon), \phi_2(k) \cdot (1+\varepsilon)]$ and
  $\phi_2(k) \in [\phi_1(k) \cdot (1-\varepsilon), \phi_1(k) \cdot (1+\varepsilon)]$
- Assign colours to random variables according to their ranges and evidence
- Assign colours to factors according to their potential tables
- Pass colours to detect ε-equivalent symmetries in the graph

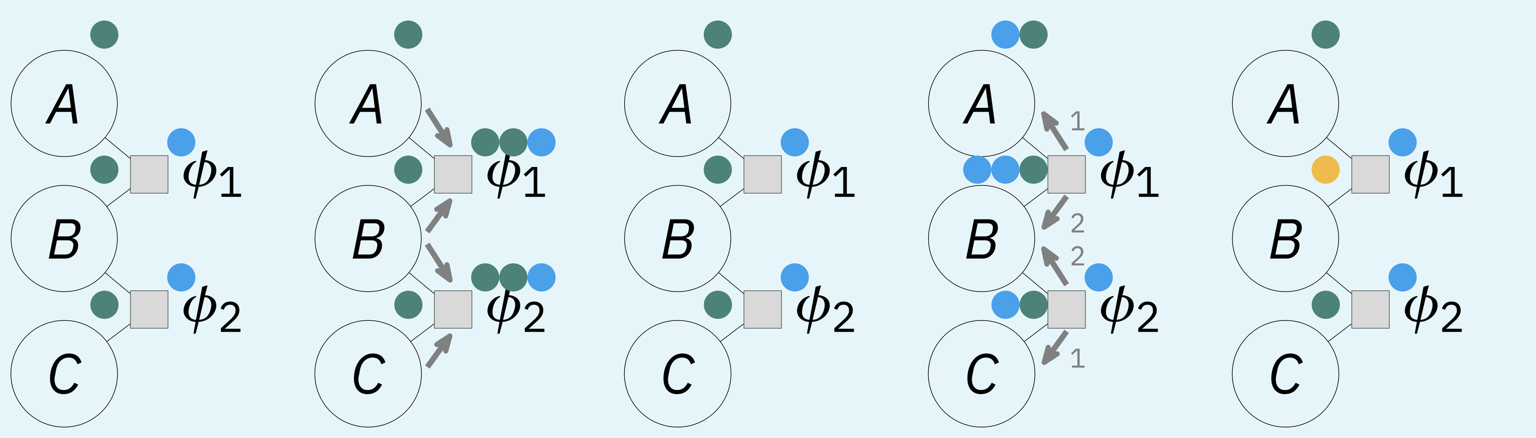

- Limitations: No guaranteed consistency of ε-equivalent groupings for different ε values
  - No informed choice of ε
- Solution: Hierarchical groupings for increasing ε

3. One-dimensional-ε-equivalence-distance (1DEED)

One-dimensional ε-equivalence distance (1DEED) is defined as:

$$d_\infty : \mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0} \to \mathbb{R}_{\geq 0} \quad d_\infty(\phi_1, \phi_2) := \max_{k=1,\ldots,n} \left\{ \frac{|\phi_1(k) - \phi_2(k)|}{\min\{|\phi_1(k)|, |\phi_2(k)|\}} \right\}$$

Properties:

- d_∞ is non-negative and symmetric
- $d_\infty(\phi_1, \phi_2) = 0 \Leftrightarrow |\phi_1(k) - \phi_2(k)| = 0$ for $k = 1, \ldots, n \Leftrightarrow \phi_1 = \phi_2$
- Triangle inequality does *not* hold in general
- Theorem: Two vectors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$ are *ε-equivalent* if and only if $d_\infty(\phi_1, \phi_2) \leq \varepsilon$ holds.

4. Hierarchical Advanced Colour Passing (HACP)

Algorithm 1: Determine ordered ε-vector and nested list of factors

(i) Compute pairwise 1DEED for factors (upper triangular matrix)

(ii) Run agglomerative clustering algorithm based on 1DEED with complete linkage within maximal deviation

- Choose level(s) of compression within ε-vector ($\varepsilon_i < \varepsilon_{i+1}$)

HACP: Use ε-ACP (generalisation of ACP) proceeding as follows:

(i) Pass groups of pairwise ε-equivalent factors based on nested list

(ii) Assign colours to factors according to the provided groups and run the colour passing procedure from ε-ACP

- Ensures identical potentials in resulting groups of factors
- Goal: Apply smallest possible change to potential tables
- Minimise sum of squared deviations between potentials:

$$\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\} \overset{\text{replaced by}}{\Longrightarrow} \boldsymbol{G}^* = \{\phi^*, \ldots, \phi^*\}$$

pairwise ε-equivalent $\qquad \phi^*(r_1, \ldots, r_n) = \frac{1}{m} \sum_{i=1}^{m} \phi_i(r_1, \ldots, r_n)$

- Corollary: If $\varepsilon = 0$, HACP is equivalent to ACP and ε-ACP.
- HACP preserves structural consistency and comparability

5. Compression versus Accuracy

Guaranteed bounds of change in query results

- Theorem: The maximal absolute deviation between any initial probability $p = P_M(r \mid \boldsymbol{e})$ of r given $\boldsymbol{e}$ in model M and the probability $p' = P_{M'}(r \mid \boldsymbol{e})$ in the modified model M' resulting from running HACP can be bounded by

$$p_{\max\Delta} := \max_{\text{for any } r \mid \boldsymbol{e}} |p - p'| \leq \frac{\sqrt{e^d} - 1}{\sqrt{e^d} + 1} \text{ with } d = \ln \left(\frac{\left(1 + \frac{m-1}{m}\varepsilon\right)(1+\varepsilon)}{1 + \frac{1}{m}\varepsilon} \right)^m$$

- Theorem: For any given $p^*_\Delta \in (0, \frac{1}{2}]$, the output of HACP guarantees for any $\varepsilon \in (0,1)$, which is smaller or equal to

$$\varepsilon_1 = -\frac{1 + \frac{m-1}{m} - \frac{1}{m}\sqrt[m]{e^d}}{2\frac{m-1}{m}} + \sqrt{\left(-\frac{1 + \frac{m-1}{m} - \frac{1}{m}\sqrt[m]{e^d}}{2\frac{m-1}{m}}\right)^2 - \frac{1 - \sqrt[m]{e^d}}{\frac{m-1}{m}}}$$

with $d = \ln\left(\frac{p^*_\Delta + 1}{1 - p^*_\Delta}\right)^2$ the bound $p_{\max\Delta} \leq p^*_\Delta$.

- Graphical illustration of theorems controlling the bound

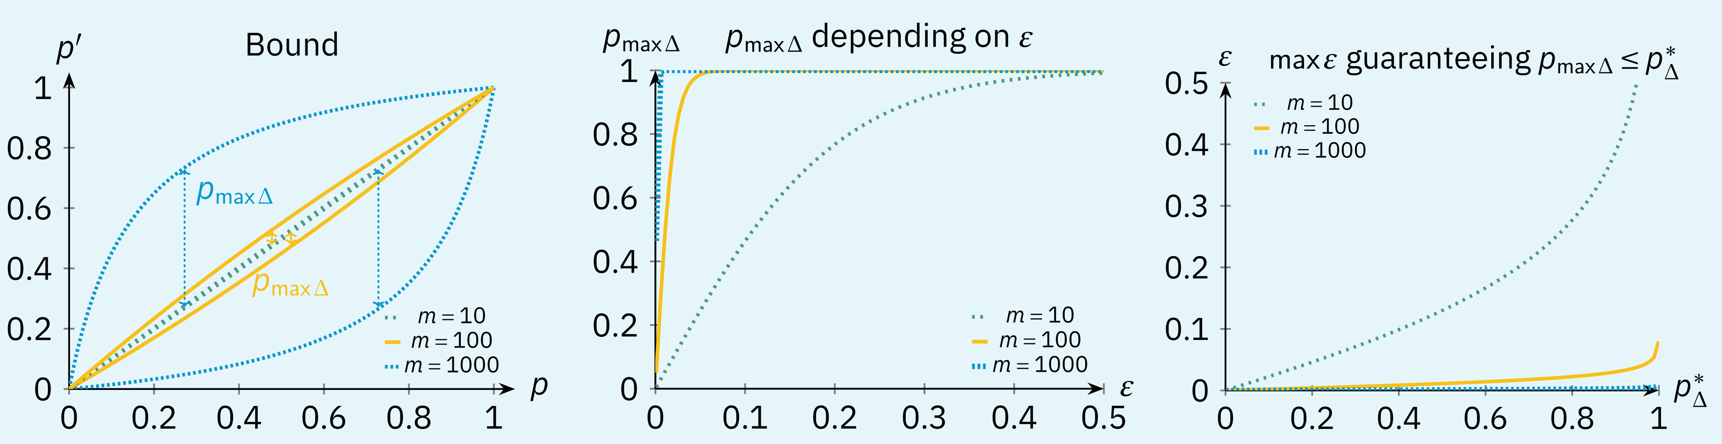

- Left: $\varepsilon = 0.001$; All: Dashed (blue) line: $m = 1000$, solid (yellow) line: $m = 100$, loosely dashed (green) line: $m = 10$
- Monotonic dependency of $p_{\max\Delta}$, ε, and m
- Bounds apply to arbitrary queries and factor graphs
- Pre-specification of maximal permissible ε or $p_{\max\Delta}$ values

6. Summary

- Novel framework for hierarchical lifting and model reconciliation
- Introduction of 1DEED as a practical tool for ε-equivalence
- Hierarchical trade-off between compression and accuracy
- Enables preanalysis of theoretical error bounds

University of Münster[1], *German Research Center for Artificial Intelligence (Lübeck)*[2], *University of Hamburg*[3] *ECAI 2025 – Bologna, Italy*