Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CHAI
Humanities-Centered AI

Universität Münster

Compression vs. Accuracy: Compact Models for Efficiency and Interpretation

IJCAI-ECAI 2026 – Bremen, Germany

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

[1]Institute for Humanities-Centered Artificial Intelligence, University of Hamburg, Germany
[2]Data Science Group, University of Münster, Germany

August 15, 2026

Agenda

1. Introduction and background [Tanya]
   - Motivation, use cases
   - Probabilistic (relational) models
   - Lifted probabilistic inference
2. Compressing models [Malte]
   - The basics of colour passing
   - Recent advances in colour passing
3. Accuracy considerations [Jan]
   - Hierarchical colour passing
   - Geometric interpretation
4. Summary and future directions [Tanya]

Approximate Lifted Model Construction I

- So far: Strict equality between potentials required
  - E.g., $\varphi_1 = \varphi'_1$, $\varphi_2 = \varphi'_2$, $\varphi_3 = \varphi'_3$, $\varphi_4 = \varphi'_4$

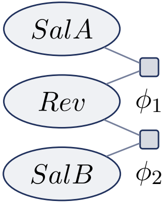

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ'_1
high	low	φ'_2
low	high	φ'_3
low	low	φ'_4

Approximate Lifted Model Construction II

- Next: Allow for a small deviation between potentials for practical applicability
  - E.g., $\varphi_1 \approx \varphi'_1$, $\varphi_2 \approx \varphi'_2$, $\varphi_3 \approx \varphi'_3$, $\varphi_4 \approx \varphi'_4$

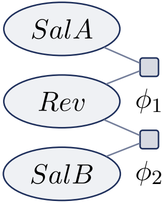

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ'_1
high	low	φ'_2
low	high	φ'_3
low	low	φ'_4

Malte Luttermann, Jan Speller, Marcel Gehrke, Tanya Braun, Ralf Möller, and Mattis Hartwig (2025). »Approximate Lifted Model Construction«. *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-2025)*. IJCAI Organization, pp. 9077–9085.

Approximate Lifted Model Construction III

- Next: Allow for a small deviation between potentials for practical applicability
  - E.g., $\varphi_1 \approx \varphi'_1$, $\varphi_2 \approx \varphi'_2$, $\varphi_3 \approx \varphi'_3$, $\varphi_4 \approx \varphi'_4$

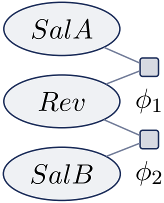

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ'_1
high	low	φ'_2
low	high	φ'_3
low	low	φ'_4

- How can factors that are not strictly equivalent be grouped?
- How much deviation should be allowed and what is the impact on query results?

ε-Equivalence

- Potentials $\varphi \in \mathbb{R}^+$ and $\varphi' \in \mathbb{R}^+$ are ε-*equivalent* if

$$\varphi \in [\varphi' \cdot (1 - \varepsilon), \varphi' \cdot (1 + \varepsilon)] \text{ and}$$
$$\varphi' \in [\varphi \cdot (1 - \varepsilon), \varphi \cdot (1 + \varepsilon)]$$

ε-Equivalence

- Potentials $\varphi \in \mathbb{R}^+$ and $\varphi' \in \mathbb{R}^+$ are ε-*equivalent* if

$$\varphi \in [\varphi' \cdot (1-\varepsilon), \varphi' \cdot (1+\varepsilon)] \text{ and}$$
$$\varphi' \in [\varphi \cdot (1-\varepsilon), \varphi \cdot (1+\varepsilon)]$$

- Factors $\phi_1(R_1, \ldots, R_n)$ and $\phi_2(R'_1, \ldots, R'_n)$ are ε-equivalent if all potentials in their potential tables are ε-equivalent

ε-Equivalence

- Potentials $\varphi \in \mathbb{R}^+$ and $\varphi' \in \mathbb{R}^+$ are ε-*equivalent* if

$$\varphi \in [\varphi' \cdot (1-\varepsilon), \varphi' \cdot (1+\varepsilon)] \text{ and}$$
$$\varphi' \in [\varphi \cdot (1-\varepsilon), \varphi \cdot (1+\varepsilon)]$$

- Factors $\phi_1(R_1, \ldots, R_n)$ and $\phi_2(R'_1, \ldots, R'_n)$ are ε-equivalent if all potentials in their potential tables are ε-equivalent
- E.g., for $\varepsilon = 0.1$, $\phi_1(SalA, Rev)$ and $\phi_2(SalB, Rev)$ are ε-equivalent:

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	0.81
high	low	0.32
low	high	0.51
low	low	0.21

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	0.84
high	low	0.31
low	high	0.51
low	low	0.20

Non-Transitivity of ε-Equivalence I

- ε-equivalence is *not transitive*: $\varphi =_\varepsilon \varphi'$ and $\varphi' =_\varepsilon \varphi''$ do *not* imply $\varphi =_\varepsilon \varphi''$

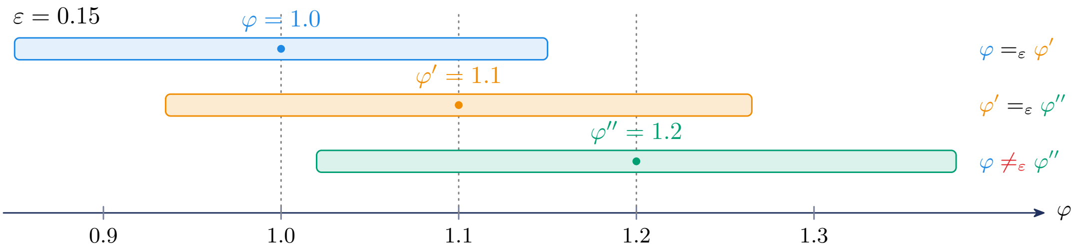

- Each band is an ε-window $[\varphi(1-\varepsilon), \varphi(1+\varepsilon)]$ (here $\varepsilon = 0.15$); a centre lies in another band iff the two are ε-equivalent

Non-Transitivity of ε-Equivalence II

- Grouping naively would cause cascading errors

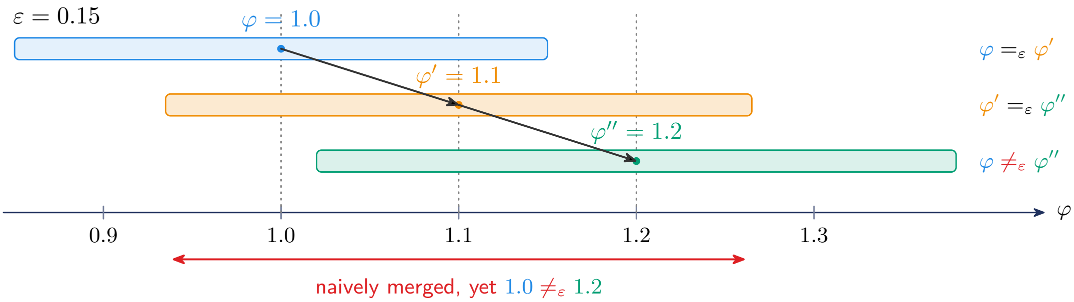

$\Rightarrow$ Ensure pairwise ε-equivalence of all factors in a group

Grouping ε-Equivalent Factors

- A representative potential table is required to construct a lifted representation

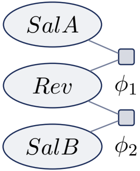

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ'_1
high	low	φ'_2
low	high	φ'_3
low	low	φ'_4

$\downarrow$

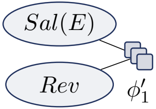

$Sal(E)$	Rev	$\phi'_1(Sal(E), Rev)$
high	high	φ^*_1
high	low	φ^*_2
low	high	φ^*_3
low	low	φ^*_4

Grouping ε-Equivalent Factors

- A representative potential table is required to construct a lifted representation

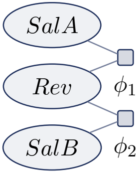

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ'_1
high	low	φ'_2
low	high	φ'_3
low	low	φ'_4

$\downarrow$

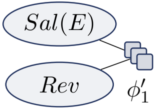

$Sal(E)$	Rev	$\phi'_1(Sal(E), Rev)$
high	high	φ^*_1
high	low	φ^*_2
low	high	φ^*_3
low	low	φ^*_4

- How to choose φ^*_i?

Minimisation of the Approximation Error I

For a group of pairwise ε-equivalent factors $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$, we want to set

$$\phi_1 = \phi^*, \ldots, \phi_m = \phi^*$$

Minimisation of the Approximation Error I

For a group of pairwise ε-equivalent factors $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$, we want to set

$$\phi_1 = \phi^*, \ldots, \phi_m = \phi^*$$

such that

$$\phi^* = \arg\min_{\phi_j} \sum_{\phi_i \in \boldsymbol{G}} Err(\phi_i, \phi_j),$$

Minimisation of the Approximation Error I

For a group of pairwise ε-equivalent factors $\boldsymbol{G} = \{\phi_1, \dots, \phi_m\}$, we want to set

$$\phi_1 = \phi^*, \dots, \phi_m = \phi^*$$

such that

$$\phi^* = \underset{\phi_j}{\arg\min} \sum_{\phi_i \in \boldsymbol{G}} Err(\phi_i, \phi_j),$$

where $Err(\phi_i, \phi_j)$ is the sum of squared deviations between ϕ_i's and ϕ_j's potentials:

$$Err(\phi_i, \phi_j) = \sum_{(r_1, \dots, r_n)} \Big(\phi_i(r_1, \dots, r_n) - \phi_j(r_1, \dots, r_n)\Big)^2$$

Minimisation of the Approximation Error II

Theorem

The optimal choice for ϕ^* is the arithmetic mean over $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$

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

and the arithmetic mean ϕ^* is ε-equivalent to ϕ_i for $i = 1, \ldots, m$.

Minimisation of the Approximation Error III

- Choose φ_i^* as the row-wise arithmetic mean of the potentials

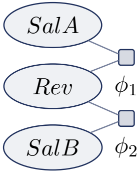

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ'_1
high	low	φ'_2
low	high	φ'_3
low	low	φ'_4

$\downarrow$

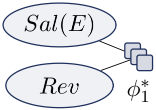

$Sal(E)$	Rev	$\phi_1^*(Sal(E), Rev)$
high	high	$\varphi_1^* = (\varphi_1 + \varphi'_1) \;/\; 2$
high	low	$\varphi_2^* = (\varphi_2 + \varphi'_2) \;/\; 2$
low	high	$\varphi_3^* = (\varphi_3 + \varphi'_3) \;/\; 2$
low	low	$\varphi_4^* = (\varphi_4 + \varphi'_4) \;/\; 2$

- Check for ε-equivalence instead of strict equivalence between factors

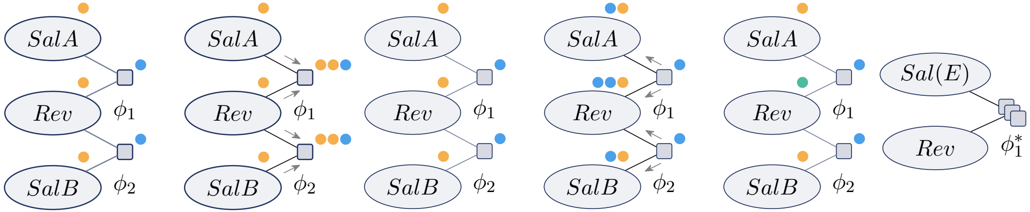

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	Rev	$\phi_2(SalB, Rev)$
high	high	φ_1'
high	low	φ_2'
low	high	φ_3'
low	low	φ_4'

$Sal(E)$	Rev	$\phi_1^*(Sal(E), Rev)$
high	high	$(\varphi_1 + \varphi_1')\ /\ 2$
high	low	$(\varphi_2 + \varphi_2')\ /\ 2$
low	high	$(\varphi_3 + \varphi_3')\ /\ 2$
low	low	$(\varphi_4 + \varphi_4')\ /\ 2$

Bounding the Change in Query Results I

Theorem (Bound on the Change in Query Results)

Let $d = \ln\,(1+\varepsilon)^{2m}$ with $m = |\boldsymbol{\Phi}|$ (number of factors). Then, after running ε-ACP, the change of every query result $p = P_M(r \mid \boldsymbol{\xi})$ is bounded by

$$\frac{p\,e^{-d}}{p\,(e^{-d}-1)+1} \;\leq\; P_{M'}(r \mid \boldsymbol{\xi}) \;\leq\; \frac{p\,e^{d}}{p\,(e^{d}-1)+1}.$$

- Bounds apply to arbitrary queries and factor graphs
- Slightly improved bound for arithmetic mean is sharp

Hei Chan and Adnan Darwiche (2005). »A Distance Measure for Bounding Probabilistic Belief Change«. *International Journal of Approximate Reasoning* 38, pp. 149–174.

Bounding the Change in Query Results II

- Bounds for $m = 10$ (left), $m = 100$ (middle), and $m = 1000$ (right) factors
- Dashed line: $\varepsilon = 0.01$, solid line: $\varepsilon = 0.001$
- x-axes depict the original probability p, y-axes reflect the bound on the change in p

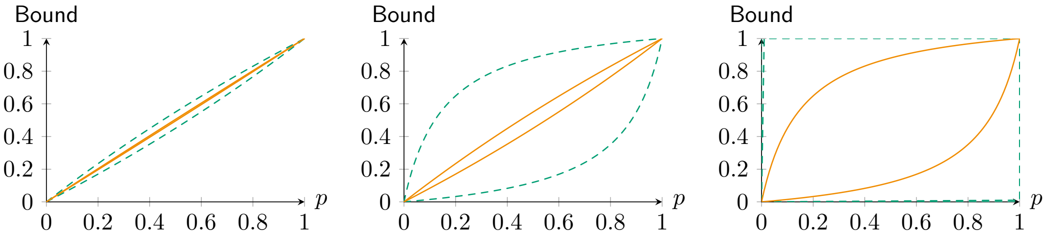

- Bounds apply to arbitrary queries and factor graphs

The ε-Advanced Colour Passing Algorithm – Experiments

- ε-ACP is a generalisation of ACP that assigns identical colours to groups of ε-equivalent factors (instead of equivalent factors)
- Left: Comparison of run times for lifted inference
- Right: Quotients of query results p' in the modified factor graph and p in the original factor graph

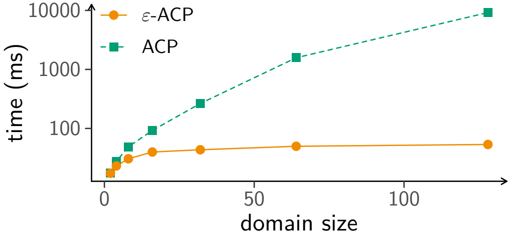

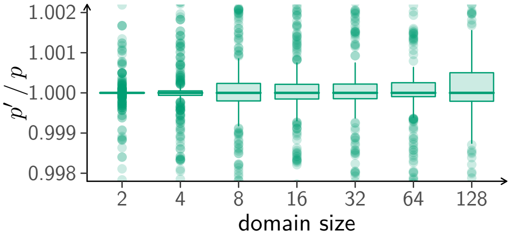

Approximate Lifted Model Construction – Summary

- Ensure the practical applicability of lifted model construction
- Hyperparameter ε to control the trade-off between exactness and compactness
- Theoretical guarantees for the change in query results

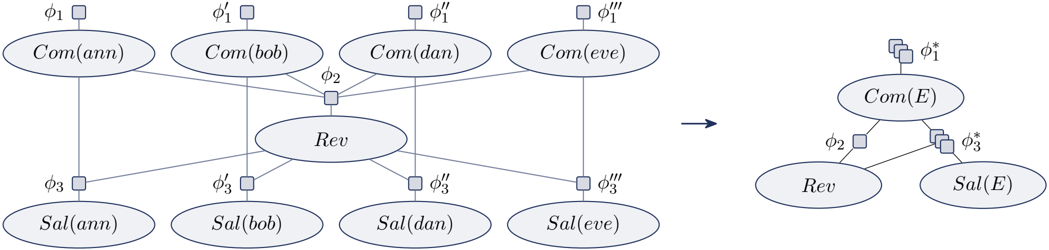

One-Dimensional ε-Equivalence Distance

- Treat a factor's potential table as a tuple $\phi \in \mathbb{R}^n_{>0}$
- 1DEED reduces pairwise ε-equivalence to a single number:

$$d_\infty(\phi_1, \phi_2) = \max_{k=1,\ldots,n} \frac{|\phi_1(k) - \phi_2(k)|}{\min\{\phi_1(k), \phi_2(k)\}}$$

Theorem (1DEED Characterises ε-Equivalence)

For all $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$ it holds that

$$\phi_1 =_\varepsilon \phi_2 \iff d_\infty(\phi_1, \phi_2) \leq \varepsilon.$$

$\Rightarrow$ d_∞ is the *smallest* ε for which ϕ_1 and ϕ_2 are still ε-equivalent (distance measure)

One-Dimensional ε-Equivalence Distance – Example

- Let $\phi_1 = (0.81, 0.32, 0.51, 0.21)$ and $\phi_2 = (0.84, 0.31, 0.51, 0.20)$

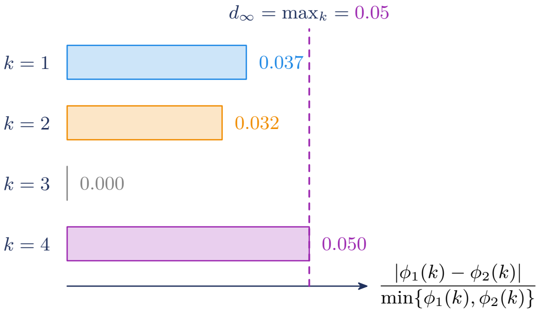

- One scalar per factor pair enables efficient comparison, indexing, and ordering

- 1DEED captures relative deviations well, but is *not* a metric

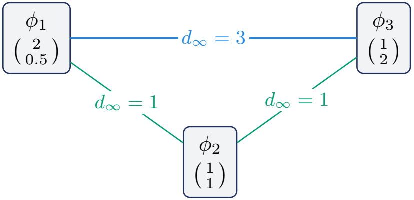

$$\underbrace{d_\infty(\phi_1, \phi_2) + d_\infty(\phi_2, \phi_3) = 1 + 1 = 2}_{\text{detour}} \quad < \quad \underbrace{3 = d_\infty(\phi_1, \phi_3)}_{\text{direct}}$$

Proposition

Triangle inequality does *not* hold for d_∞.

Hierarchical Lifted Model Construction I

- Recall: ε-ACP groups factors whose potentials differ by at most a factor of $(1 \pm \varepsilon)$
- ε is a hyperparameter – finding a suitable value may require many ε-ACP runs
- Worse, different ε can yield inconsistent models: a larger ε may *split* a group instead of merging groups

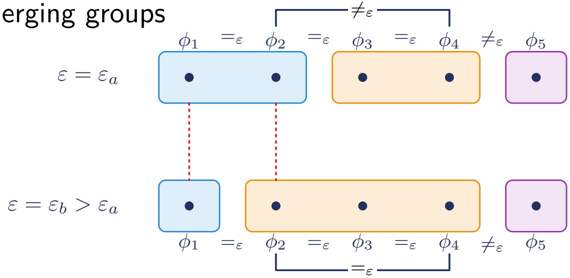

$\{\phi_1, \phi_2\}$ grouped at ε_a,
but split at $\varepsilon_b \Rightarrow$ *not nested*

Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025a). »Compression versus Accuracy: A Hierarchy of Lifted Models«. *Proceedings of the Twenty-Eighth European Conference on Artificial Intelligence (ECAI-2025)*. IOS Press, pp. 5051–5058.

Hierarchical Lifted Model Construction II

- Idea: a hyperparameter-free construction yielding a *hierarchy* of models
- *Structure preservation*: Once grouped at some ε, factors stay grouped for every larger ε
- Larger ε only *merges* smaller-ε groups $\Rightarrow$ nested, interpretable models

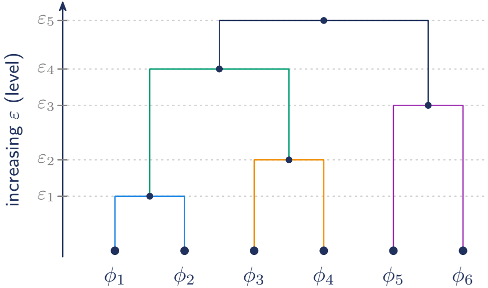

Hierarchical Ordering of ε-Equivalent Factors

- Compute all pairwise distances $\Lambda_{ij} = d_\infty(\phi_i, \phi_j)$
- Merge greedily by *smallest* distance
- On each merge, keep the *maximum* distance per column (complete linkage)
  - Ensures pairwise ε-equivalence per group
- Agglomerative clustering in $O(m^3)$, computed *once*

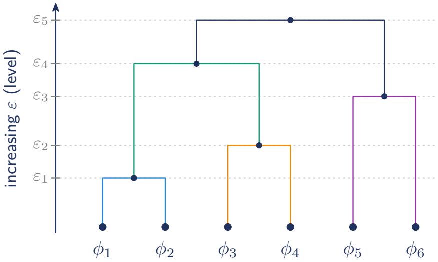

$\Rightarrow$ Ordered thresholds $\varepsilon_1 < \varepsilon_2 < \cdots < \varepsilon_{m-1}$ with a consistently nested grouping

Maintaining the Distance Matrix Λ

- Upper-triangular matrix $\Lambda_{ij} = d_\infty(\phi_i, \phi_j)$; each step merges the pair with the *smallest* active entry

Λ	ϕ_1	ϕ_2	ϕ_3	ϕ_4	ϕ_5	ϕ_6
ϕ_1	0	$\boxed{\varepsilon_{12}}$	ε_{13}	ε_{14}	ε_{15}	ε_{16}
ϕ_2	0	0	ε_{23}	ε_{24}	ε_{25}	ε_{26}
ϕ_3	0	0	0	ε_{34}	ε_{35}	ε_{36}
ϕ_4	0	0	0	0	ε_{45}	ε_{46}
ϕ_5	0	0	0	0	0	ε_{56}
ϕ_6	0	0	0	0	0	0

- Smallest entry is $\varepsilon_{12} \Rightarrow$ record $\varepsilon_1 = \varepsilon_{12}$ and merge ϕ_1 and ϕ_2

Maintaining the Distance Matrix Λ

- Upper-triangular matrix $\Lambda_{ij} = d_\infty(\phi_i, \phi_j)$; each step merges the pair with the *smallest* active entry

$$
\begin{array}{c|cccccc}
\Lambda & \phi_1 & \phi_2 & \phi_3 & \phi_4 & \phi_5 & \phi_6 \\
\hline
\phi_1 & 0 & \boxed{\varepsilon_{12}} & \varepsilon^*_{13} & \varepsilon^*_{14} & \varepsilon^*_{15} & \varepsilon^*_{16} \\
\phi_2 & 0 & 0 & \cancel{\varepsilon_{23}} & \cancel{\varepsilon_{24}} & \cancel{\varepsilon_{25}} & \cancel{\varepsilon_{26}} \\
\phi_3 & 0 & 0 & 0 & \varepsilon_{34} & \varepsilon_{35} & \varepsilon_{36} \\
\phi_4 & 0 & 0 & 0 & 0 & \varepsilon_{45} & \varepsilon_{46} \\
\phi_5 & 0 & 0 & 0 & 0 & 0 & \varepsilon_{56} \\
\phi_6 & 0 & 0 & 0 & 0 & 0 & 0
\end{array}
$$

- Merge in place: $\Lambda_{1k} \leftarrow \max\{\Lambda_{1k}, \Lambda_{2k}\}$ (keeps groups pairwise ε-equivalent), then deactivate row 2

The Hierarchical Advanced Colour Passing Algorithm (HACP)

- Pick a level i (i.e., ε_4) – cut the hierarchy
- Three phases:
  1. Load the groups at level i
  2. One colour per group, run standard ACP
  3. Set potentials to the group mean
- All levels follow from a *single* ordering – no ε to tune

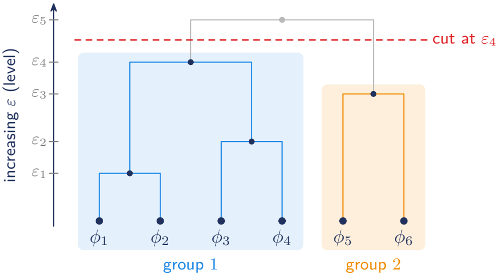

Hierarchy of Error Bounds

Proposition (Bounds Carry Over To HACP)

HACP inherits the same *sharp* error bounds as ε-ACP. The maximal absolute change of any query result $p = P_M(r \mid \boldsymbol{\xi})$ to $p' = P_{M'}(r \mid \boldsymbol{\xi})$ is

$$p_{\max\Delta} = \max_{\text{for any } r|\boldsymbol{\xi}} |p - p'| \leq \frac{\sqrt{e^d} - 1}{\sqrt{e^d} + 1}, \qquad d = \ln\,(1+\varepsilon)^{2m}.$$

- Increasing ε (higher levels) grows the bound $\Rightarrow$ A *hierarchy of error bounds*
- Computed *before* running HACP – no need to inspect the full factor graph

Note

A $10\times$ larger ε affects the bound roughly like $10\times$ more factors m.

Choosing a Level for a Target Accuracy

- Every level ε_i comes with a guaranteed bound on $p_{\max\Delta}$
- Reverse view: Pick the largest ε (most compression) that still meets a target accuracy $p_{\max\Delta} \leq p^*_\Delta$

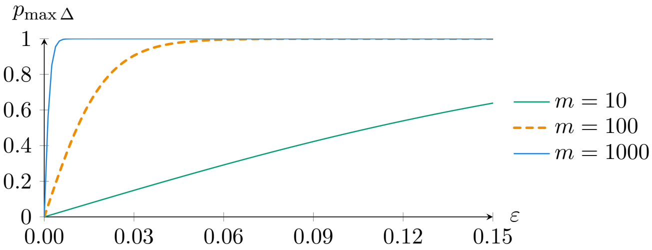

- Larger ε or more factors m loosen the bound

Target Accuracy To Determine an Admissible ε

- Invert the bound: A target accuracy p^*_Δ yields the largest ε we may use

Theorem (Guaranteed Accuracy)

Any $\varepsilon \leq \varepsilon_{\max}(p^*_\Delta)$ guarantees $p_{\max\Delta} \leq p^*_\Delta$, with $\varepsilon_{\max}(p^*_\Delta) = \tanh\left(\frac{1}{m} \ln \frac{1+p^*_\Delta}{1-p^*_\Delta}\right)$.

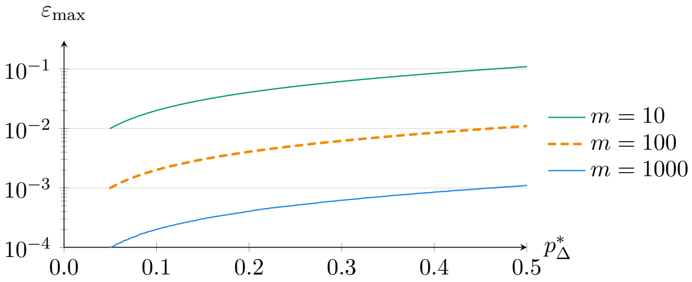

Hierarchical Lifted Model Construction – Summary

- Hyperparameter-free *hierarchy* of lifted models that is structure-preserving
- 1DEED: A single scalar governs ε-equivalence, the ordering, and the bounds
- Explicit compression-versus-accuracy trade-off with guarantees and interpretability

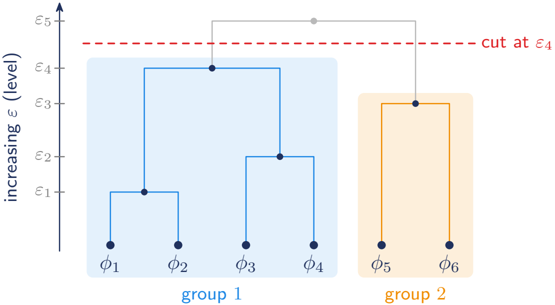

Pairs of ε-Equivalent Vectors I

- $\phi_1 =_\varepsilon \phi_2 \iff d_\infty(\phi_1, \phi_2) = \max_k \frac{|\phi_1(k) - \phi_2(k)|}{\min\{\phi_1(k), \phi_2(k)\}} \leq \varepsilon$

Theorem (Distances in $\mathcal{L}^p$)

- $\phi_1 =_\varepsilon \phi_2 \Rightarrow \|\phi_1 - \phi_2\|_p \leq \varepsilon \|\phi_i\|_p$
- $\|\phi_1 - \phi_2\|_p \leq \varepsilon \Rightarrow \phi_1 =_{\varepsilon'} \phi_2$

  holds with $\varepsilon' := \dfrac{\varepsilon}{\min\limits_{k,i}\{\phi_i(k)\}}$

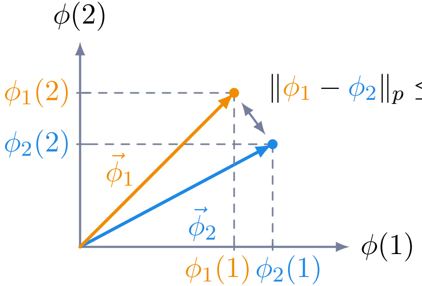

Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025b). »Towards Explainability of Approximate Lifted Model Construction: A Geometric Perspective«. *Proceedings of the Eleventh Workshop on Formal and Cognitive Reasoning (FCR-2025)*. CEUR, pp. 41–56.

Pairs of ε-Equivalent Vectors II

- $\phi_1 =_\varepsilon \phi_2 \quad \Longleftrightarrow \quad d_\infty(\phi_1, \phi_2) = \max_k \frac{|\phi_1(k) - \phi_2(k)|}{\min\{\phi_1(k), \phi_2(k)\}} \leq \varepsilon$

Example ($p = 1$):

$$\widehat{\phi}_1 = \frac{\phi_1}{\|\phi_1\|_1} = \begin{pmatrix} 0.50 \\ 0.50 \end{pmatrix}, \quad \widehat{\phi}_2 = \frac{\phi_2}{\|\phi_2\|_1} = \begin{pmatrix} 0.65 \\ 0.35 \end{pmatrix}$$

- $\phi_1 =_{0.43} \phi_2 \quad [0.5 \approx (1 + 0.43) \cdot 0.35]$
- $\|\phi_1 - \phi_2\|_1 = \sum_{k=1}^{2} |\phi_1(k) - \phi_2(k)| = 0.3 \leq \varepsilon \|\phi_i\|_1$

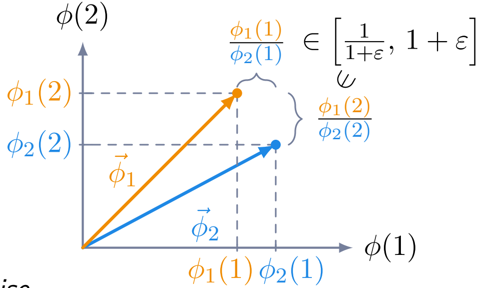

- Relational deviation is evaluated *dimension-wise*
- As with the ℓ_∞-norm, *every deviation must be small*

Pairs of ε-Equivalent Vectors III

- $\phi_1 =_\varepsilon \phi_2 \quad \iff \quad d_\infty(\phi_1, \phi_2) = \max_k \frac{|\phi_1(k) - \phi_2(k)|}{\min\{\phi_1(k), \phi_2(k)\}} \leq \varepsilon$

Example ($p \geq 1$)**:**

$$\phi_1 = \begin{pmatrix} 0.01 \\ 1.0 \end{pmatrix}, \quad \phi_2 = \begin{pmatrix} 0.02 \\ 1.0 \end{pmatrix}$$

- $\|\phi_1 - \phi_2\|_p^p = \sum_{k=1}^{2} |\phi_1(k) - \phi_2(k)|^p = 0.01^p = \varepsilon^p$
- $\varepsilon' = \frac{\varepsilon}{\min\limits_{\substack{k=1,\ldots,n,\\ i=1,2}} \{\phi_i(k)\}} = \frac{0.01}{0.01} = 1 \quad \Rightarrow \quad \phi_1 =_{\varepsilon'} \phi_2$

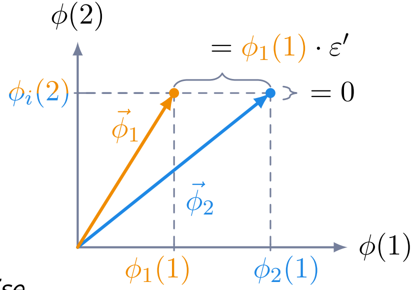

- Relational deviation is evaluated *dimension-wise*
- As with the ℓ_∞-norm, *every deviation must be small*

From Pairs to ε-Equivalent Groups I

- Let $G = \{\phi_1, \phi_2, \phi_3, \phi_4\}$ be pairwise ε-equivalent ($\phi_i =_\varepsilon \phi_j$)

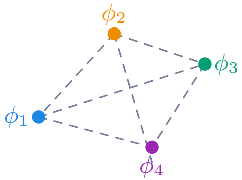

ε-equivalent

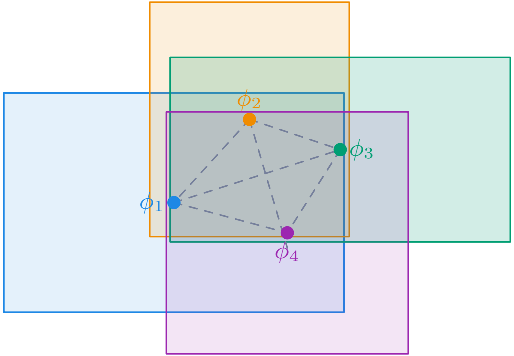

Per-factor ε-equivalent boxes

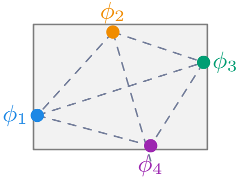

Intersection of boxes

From Pairs to ε-Equivalent Groups II

Theorem (Convex Combinations)

Let $G = \{\phi_1, \ldots, \phi_m\}$ be pairwise ε-equivalent. Then every convex combination

$$\phi_\omega = \sum_{i=1}^{m} \omega_i \phi_i \text{ with } \sum_{i=1}^{m} \omega_i = 1, \omega_i \geq 0$$

is pairwise ε-equivalent to all factors in G.

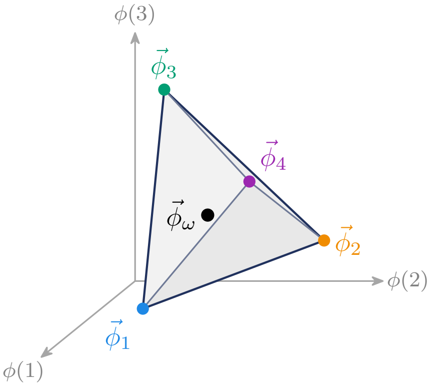

Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025b). »Towards Explainability of Approximate Lifted Model Construction: A Geometric Perspective«. *Proceedings of the Eleventh Workshop on Formal and Cognitive Reasoning (FCR-2025)*. CEUR, pp. 41–56.

Convex Hull

- The convex hull is the *smallest* convex set containing the group

$$\mathrm{conv}(\phi_1, \ldots, \phi_m) = \left\{ \sum_{i=1}^{m} \omega_i \phi_i \;\middle|\; \omega_i \geq 0,\ \sum_{i=1}^{m} \omega_i = 1 \right\}$$

- Choosing a representative in conv:
  - $\phi_{\text{mean}} = \frac{1}{m} \sum_{i=1}^{m} \phi_i$
  - $\phi_{\text{gm}} = \sqrt[m]{\prod_{i=1}^{m} \phi_i}$

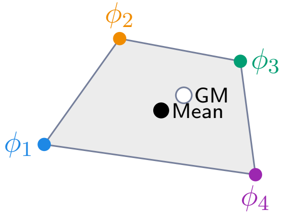

Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025b). »Towards Explainability of Approximate Lifted Model Construction: A Geometric Perspective«. *Proceedings of the Eleventh Workshop on Formal and Cognitive Reasoning (FCR-2025)*. CEUR, pp. 41–56.

Scaling Leaves the Semantics Unchanged

- Exchangeable factors: Identical up to a scalar $\alpha \in \mathbb{R}^+$ (and permutation)
- Scaling any factor by $\alpha \in \mathbb{R}^+$ leaves the semantics unchanged

$$P_M = \frac{1}{Z} \prod_{j=1}^{m} \phi_j \quad \xrightarrow{\phi_k \to \alpha \cdot \phi_k} \quad \frac{1}{\alpha \cdot Z} \cdot \alpha \prod_{j=1}^{m} \phi_j = \frac{1}{Z} \prod_{j=1}^{m} \phi_j$$

- The normalisation constant Z absorbs α
- Thus: Scaling several factors by *different* scalars leaves semantics unchanged

Malte Luttermann, Ralf Möller, and Marcel Gehrke (2025). »Lifted Model Construction without Normalisation: A Vectorised Approach to Exploit Symmetries in Factor Graphs«. *Proceedings of the Third Learning on Graphs Conference (LoG-2024)*. PMLR, 46:1–46:17.

The Role of Normalisation

- Exchangeable factors may lie on different scales
  - E.g., when learning potentials by counting occurrences in different data sources
- Not the potentials themselves, but their *ratios* determine the semantics of the model

A	B	$\phi(A, B)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

A	B	$\phi'(A, B)$
high	high	$\alpha \cdot \varphi_1$
high	low	$\alpha \cdot \varphi_2$
low	high	$\alpha \cdot \varphi_3$
low	low	$\alpha \cdot \varphi_4$

Malte Luttermann, Ralf Möller, and Marcel Gehrke (2025). »Lifted Model Construction without Normalisation: A Vectorised Approach to Exploit Symmetries in Factor Graphs«. *Proceedings of the Third Learning on Graphs Conference (LoG-2024)*. PMLR, 46:1–46:17.

Normalising Potentials

- Divide every potential by the sum of its table Σ
- After normalisation (fair comparison)
  - $\Rightarrow$ Exchangeable factors satisfy $\Sigma' = 1$
  - $\Rightarrow$ ACP applies unchanged
  - $\Rightarrow$ Standardized initialization for ε-ACP and HACP

A	$\phi_1(A)$
high	8
low	2

B	$\phi_2(B)$
high	4
low	1

$$\xrightarrow[\div \Sigma]{\text{normalise}}$$

	$\phi_1' = \phi_2'$
high	0.8
low	0.2

Normalisation via Vector Representation

- Represent a factor's potential table as a vector $\vec{\phi}$
- Exchangeable factors have *collinear* vectors (in $\mathcal{L}^2$ cosine distance $D_{\cos} = 0$)

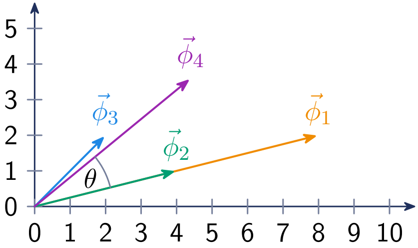

$$D_{\cos}(\phi_1, \phi_4) = 1 - \frac{\vec{\phi}_1 \cdot \vec{\phi}_4}{\|\vec{\phi}_1\|_2 \, \|\vec{\phi}_4\|_2} = \frac{1}{2}\|\vec{\phi}_1 - \vec{\phi}_4\|_2 = 1 - \cos(\theta)$$

for one $\theta \in [0, \pi]$

E.g., for $\vec{\phi}_1 = (8, 2)$ and $\vec{\phi}_2 = (4, 1)$ $\Rightarrow$ $\vec{\phi}_1 = 2 \cdot \vec{\phi}_2$:

$$D_{\cos}(\phi_1, \phi_2) = 1 - \frac{8 \cdot 4 + 2 \cdot 1}{\sqrt{68} \cdot \sqrt{17}} = 1 - \frac{34}{2 \cdot 17} = 0$$

Marcel Gehrke, Ralf Möller, and Tanya Braun (2020). »Taming Reasoning in Temporal Probabilistic Relational Models«. *Proceedings of the Twenty-Fourth European Conference on Artificial Intelligence (ECAI-2020)*. IOS Press, pp. 2592–2599.

Generalised Normalisation in $\mathcal{L}^p$

- Scaling leaves model's semantics unchanged
- *Geometrically*, we normalise away the scale and keep only the direction

- $S_p^{n-1} := \{\phi \in \mathbb{R}^n : \|\phi\|_p = 1\}$ unit p-sphere
- $\widehat{\phi}_i := \frac{\phi_i}{\|\phi_i\|_p} \in \mathbb{R}_+^n \cap S_p^{n-1}$
- Exchangeable factors $\alpha' \cdot \phi_1' = \alpha \cdot \phi_1$ collapse onto the *same* point $\widehat{\vec{\phi}}_1$ after normalisation

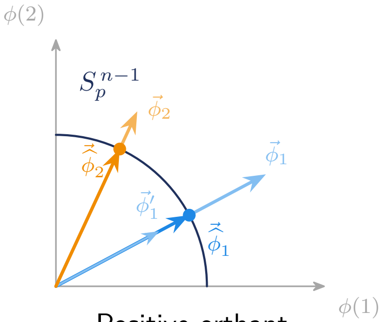

Positive orthant

ε-Equivalence as a Neighbourhood Under the Sphere

Theorem (ε-Equivalence on the Sphere)

$\widehat{\phi}_1, \widehat{\phi}_2 \in S_p^{n-1}$: $\widehat{\phi}_1 =_\varepsilon \widehat{\phi}_2 \;\Rightarrow\; \|\widehat{\phi}_1 - \widehat{\phi}_2\|_p \leq \varepsilon$,

i.e., $\widehat{\phi}_2 \in \overline{B_\varepsilon^p}(\widehat{\phi}_1) \cap S_p^{n-1}$

Theorem (Bounded Length Under the Sphere)

$\phi_\omega = \sum_{i=1}^m \omega_i \widehat{\phi}_i \in \operatorname{conv}(\widehat{\phi}_1, \ldots, \widehat{\phi}_m)$ and $\widehat{\phi}_i =_\varepsilon \widehat{\phi}_j$:

$$\frac{1}{1+\varepsilon} \leq \|\phi_\omega\|_p \leq 1$$

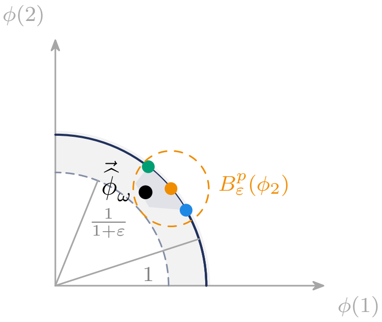

Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025b). »Towards Explainability of Approximate Lifted Model Construction: A Geometric Perspective«. *Proceedings of the Eleventh Workshop on Formal and Cognitive Reasoning (FCR-2025)*. CEUR, pp. 41–56.

Hierarchical Order of Operations

- Unambiguous process
  1. Normalise all factors ϕ_i onto S_p^{n-1}
  2. Check pairwise ε-equivalence of $\widehat{\phi}_i =_\varepsilon \widehat{\phi}_j$
  3. Choose $\phi_\omega \in \mathrm{conv}(\widehat{G})$ for all groups G
- All pairwise ε-equivalent convex supersets of G contain $\mathrm{conv}(\widehat{G})$ as possible representatives

  $\Rightarrow$ Larger ε leads to larger $G' \supset G$:

  $$\Rightarrow \quad \mathrm{conv}(\widehat{G}) \subset \mathrm{conv}(\widehat{G}')$$

  $\Rightarrow$ New weighted $\phi'_\omega \in \mathrm{conv}(\widehat{G}')$

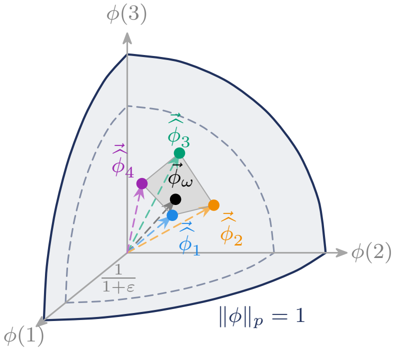

Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025b). »Towards Explainability of Approximate Lifted Model Construction: A Geometric Perspective«. *Proceedings of the Eleventh Workshop on Formal and Cognitive Reasoning (FCR-2025)*. CEUR, pp. 41–56.

Why Not Simplify to a Convex Cone?

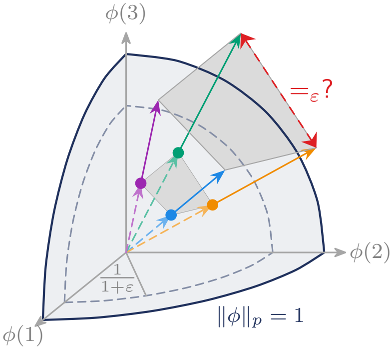

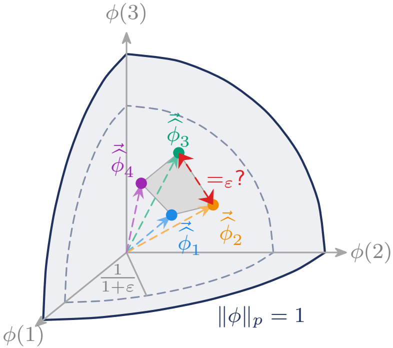

- Testing of ε-equivalence before Normalising: $\phi_2 =_\varepsilon \phi_3 \not\Rightarrow \widehat{\phi}_2 =_\varepsilon \widehat{\phi}_3$
- Normalising before testing of ε-equivalence: $\phi_2 \neq_\varepsilon \phi_3 \not\Rightarrow \widehat{\phi}_2 \neq_\varepsilon \widehat{\phi}_3$

Geometric Interpretability – Takeaways

- Scaling leaves model's semantics unchanged
- ε-equivalence describes a compact geometric region in $\mathcal{L}^p$
- In $\mathcal{L}^p$: ε-equivalence means bounded *shell radius* (in $\mathcal{L}^2$ bounded *angle*)
- Every point inside the convex hull of pairwise ε-equivalent factors remains a valid representative (e.g. arithmetic mean)
- ε-ACP combines ε-equivalent groups of factors (re-organising for different ε)
- HACP Each hierarchy level determines ε-equivalent convex areas under the unit p-sphere, where higher levels are merging overlapping sets
- Order of operations in HACP is meaningful and necessary

ε-Commutativity

- ε-commutativity: Outputs invariant under argument permutations *up to ε-equivalence* (here $\{ComA, ComB\}$ if $\varphi_3 =_\varepsilon \varphi'_3$, $\varphi_4 =_\varepsilon \varphi'_4$)

$ComA$	$ComB$	Rev	$\phi(ComA, ComB, Rev)$
high	high	high	φ_1
high	high	low	φ_2
high	low	high	φ_3
high	low	low	φ_4
low	high	high	φ'_3
low	high	low	φ'_4
low	low	high	φ_5
low	low	low	φ_6

Malte Luttermann, Jan Speller, Tanya Braun, Marcel Gehrke, and Ralf Möller (2026). »Lifted Model Construction under Approximate Commutativity«. *Proceedings of the Seventeenth International Conference on Scalable Uncertainty Management (SUM-2026)*. Forthcoming.

ε-Commutativity – Compression

- Replace each symmetry set of ε-equivalent potentials by one representative φ_i^* (the arithmetic mean minimises the deviation)

$ComA$	$ComB$	Rev	$\phi(ComA, ComB, Rev)$
high	high	high	φ_1
high	high	low	φ_2
high	low	high	φ_3
high	low	low	φ_4
low	high	high	φ_3'
low	high	low	φ_4'
low	low	high	φ_5
low	low	low	φ_6

$\longrightarrow$

$\#_E[Com(E)]$	Rev	$\phi(\#_E[Com(E)], Rev)$
$[2, 0]$	high	φ_1^*
$[2, 0]$	low	φ_2^*
$[1, 1]$	high	φ_3^*
$[1, 1]$	low	φ_4^*
$[0, 2]$	high	φ_5^*
$[0, 2]$	low	φ_6^*

ε-Commutativity – Structural Properties

Theorem (ε-Commutativity Sets)

- *Downward closure*: If C is ε-commutative, so is every subset of C
- *Not closed under unions*: Pairwise ε-commutativity does *not* imply the whole set

- Detecting ε-commutative arguments is harder than in the exact case

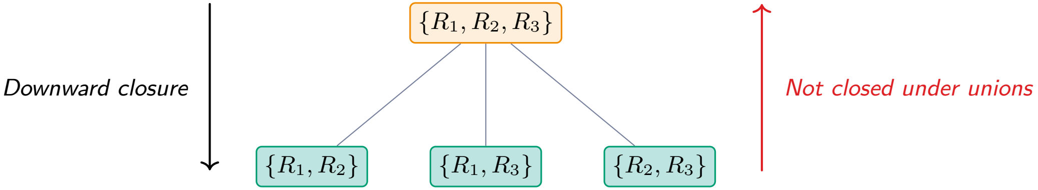

Bounding the Change in Query Results I

Theorem (Bound on the Change in Query Results)

Let $d = \ln\,(1+\varepsilon)^{2m}$, where m is the number of ε-commutative factors. Then, after replacing each symmetry set of ε-equivalent potentials by their arithmetic mean, the change of every query result $p = P_M(r \mid \boldsymbol{\xi})$ is bounded by

$$\frac{p\,e^{-d}}{p(e^{-d}-1)+1} \leq P_{M'}(r \mid \boldsymbol{\xi}) \leq \frac{p\,e^{d}}{p(e^{d}-1)+1}.$$

- Bounds apply to arbitrary queries and factor graphs

Hei Chan and Adnan Darwiche (2005). »A Distance Measure for Bounding Probabilistic Belief Change«. *International Journal of Approximate Reasoning* 38, pp. 149–174.

Bounding the Change in Query Results II

- Same bounds as before apply (and can even be tightened)
- Bounds for $m = 10$ (left), $m = 100$ (middle), and $m = 1000$ (right) factors
- Dashed line: $\varepsilon = 0.01$, solid line: $\varepsilon = 0.001$
- x-axes depict the original probability p, y-axes reflect the bound on the change in p

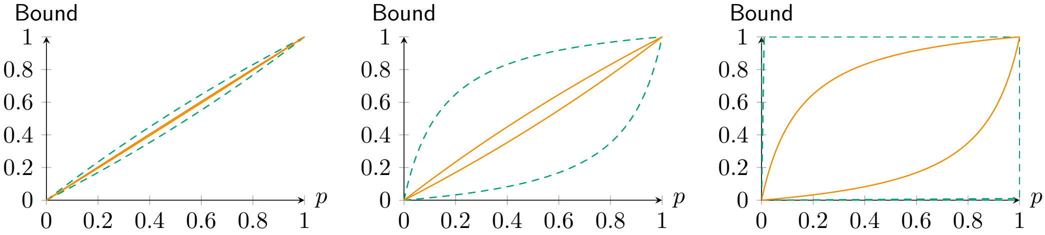

- Bounds apply to arbitrary queries and factor graphs

Agenda

1. Introduction and background [Tanya]
   - Motivation, use cases
   - Probabilistic (relational) models
   - Lifted probabilistic inference
2. Compressing models [Malte]
   - The basics of colour passing
   - Recent advances in colour passing
3. Accuracy considerations [Jan]
   - Hierarchical colour passing
   - Geometric interpretation
4. Summary and future directions [Tanya]

Bibliography I

- Hei Chan and Adnan Darwiche (2005). »A Distance Measure for Bounding Probabilistic Belief Change«. *International Journal of Approximate Reasoning* 38, pp. 149–174.
- Marcel Gehrke, Ralf Möller, and Tanya Braun (2020). »Taming Reasoning in Temporal Probabilistic Relational Models«. *Proceedings of the Twenty-Fourth European Conference on Artificial Intelligence (ECAI-2020)*. IOS Press, pp. 2592–2599.
- Malte Luttermann, Ralf Möller, and Marcel Gehrke (2025). »Lifted Model Construction without Normalisation: A Vectorised Approach to Exploit Symmetries in Factor Graphs«. *Proceedings of the Third Learning on Graphs Conference (LoG-2024)*. PMLR, 46:1–46:17.

Bibliography II

- Malte Luttermann, Jan Speller, Tanya Braun, Marcel Gehrke, and Ralf Möller (2026). »Lifted Model Construction under Approximate Commutativity«. *Proceedings of the Seventeenth International Conference on Scalable Uncertainty Management (SUM-2026)*. Forthcoming.
- Malte Luttermann, Jan Speller, Marcel Gehrke, Tanya Braun, Ralf Möller, and Mattis Hartwig (2025). »Approximate Lifted Model Construction«. *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-2025)*. IJCAI Organization, pp. 9077–9085.
- Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025a). »Compression versus Accuracy: A Hierarchy of Lifted Models«. *Proceedings of the Twenty-Eighth European Conference on Artificial Intelligence (ECAI-2025)*. IOS Press, pp. 5051–5058.

Bibliography III

- Jan Speller, Malte Luttermann, Marcel Gehrke, and Tanya Braun (2025b). »Towards Explainability of Approximate Lifted Model Construction: A Geometric Perspective«. *Proceedings of the Eleventh Workshop on Formal and Cognitive Reasoning (FCR-2025)*. CEUR, pp. 41–56.