Towards Explainability of Approximate Lifted Model Construction

A Geometric Perspective

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

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

Perspectives on Humanities-Centered AI and Formal & Cognitive Reasoning Workshop – KI2025
September 16, 2025

Problem Setup

- Input: A factor graph $\psi(\boldsymbol{r}) = \prod_{j=1}^{m} \phi_j(\boldsymbol{r}_j) \to P_M(\boldsymbol{r}) = \frac{1}{Z}\psi(\boldsymbol{r})$
- Output: A parametric factor graph with approximately equivalent semantics as ψ
  - With a minimal approximation error
  - With theoretical guarantees for the change in query results

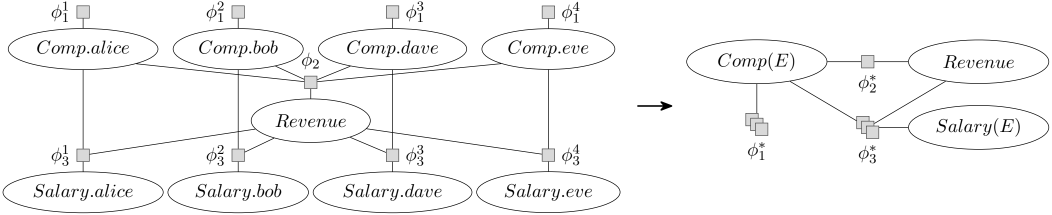

Problem Setup

- Input: A factor graph $\psi(\boldsymbol{r}) = \prod_{j=1}^{m} \phi_j(\boldsymbol{r}_j) \rightarrow P_M(\boldsymbol{r}) = \frac{1}{Z}\psi(\boldsymbol{r})$
- Output: A parametric factor graph with approximately equivalent semantics as ψ
  - With a minimal approximation error
  - With theoretical guarantees for the change in query results

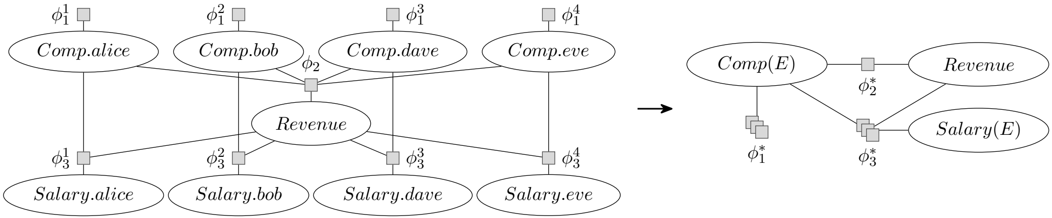

- Solutions (Previous Work):
  - Advanced Colour Passing Algorithm (ACP, Kersting et al., 2009; Ahmadi et al., 2013; Luttermann et al., 2024)
    - ε-ACP (Luttermann et al., 2025)
    - Hierarchical ACP (Speller et al., 2025)

Problem Setup

- Input: A factor graph $\psi(\boldsymbol{r}) = \prod_{j=1}^{m} \phi_j(\boldsymbol{r}_j) \rightarrow P_M(\boldsymbol{r}) = \frac{1}{Z}\psi(\boldsymbol{r})$
- Output: A parametric factor graph with approximately equivalent semantics as ψ
  - With a minimal approximation error
  - With theoretical guarantees for the change in query results

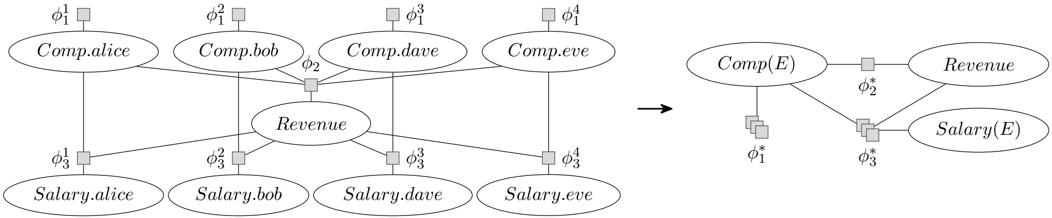

- Solutions (Previous Work):
  - Advanced Colour Passing Algorithm (ACP, Kersting et al., 2009; Ahmadi et al., 2013; Luttermann et al., 2024)
    - ε-ACP (Luttermann et al., 2025)
    - Hierarchical ACP (Speller et al., 2025)

$\rightarrow$ Better understanding of model choice

Approximate Lifted Model Construction

- Originally strict equality between potentials required
  - E.g., $\varphi_1 = \varphi_1$, $\varphi_2 = \varphi_2$, $\varphi_3 = \varphi_3$, $\varphi_4 = \varphi_4$

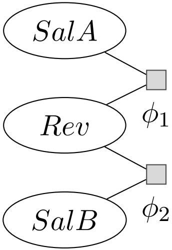

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

$SalB$	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

- Originally strict equality (after *normalisation*, $\alpha > 0$) between potentials required
  - E.g., $\varphi_1 = \alpha\varphi'_1$, $\varphi_2 = \alpha\varphi'_2$, $\varphi_3 = \alpha\varphi'_3$, $\varphi_4 = \alpha\varphi'_4$

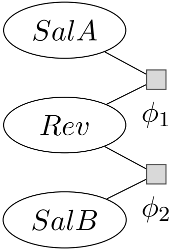

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

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

Approximate Lifted Model Construction

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

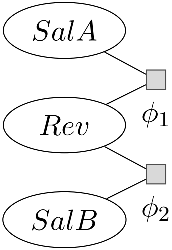

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

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

Approximate Lifted Model Construction

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

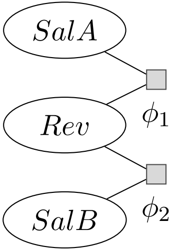

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

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

- How much deviation should be allowed and what is the impact on query results?

ε-Equivalence

- Potentials $\varphi_1 \in \mathbb{R}_{>0}$ and $\varphi_2 \in \mathbb{R}_{>0}$ are *ε-equivalent* if

$$\varphi_1 \in [\varphi_2 \cdot (1-\varepsilon), \varphi_2 \cdot (1+\varepsilon)] \text{ and}$$
$$\varphi_2 \in [\varphi_1 \cdot (1-\varepsilon), \varphi_1 \cdot (1+\varepsilon)] \text{ (Luttermann et al., 2025)}$$

- 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*; Notation: $\phi_1 =_\varepsilon \phi_2$

Example ($\varepsilon = 0.1$)

$\phi_1(SalA, Rev)$ 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

Grouping ε-Equivalent Factors

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

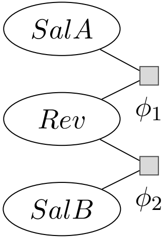

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

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

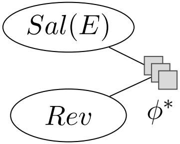

$Sal(E)$	Rev	$\phi^*(Sal(E), Rev)$
high	high	φ^*_1
high	low	φ^*_2
low	high	φ^*_3
low	low	φ^*_4

Grouping ε-Equivalent Factors

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

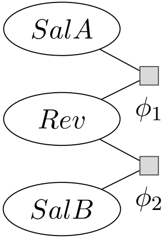

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

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

$\downarrow$

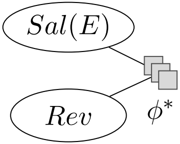

$Sal(E)$	Rev	$\phi^*(Sal(E), Rev)$
high	high	φ^*_1
high	low	φ^*_2
low	high	φ^*_3
low	low	φ^*_4

- How to choose φ^*_i?

Grouping ε-Equivalent Factors

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

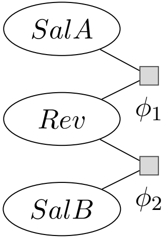

$SalA$	Rev	$\phi_1(SalA, Rev)$
high	high	φ_1
high	low	φ_2
low	high	φ_3
low	low	φ_4

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

$\downarrow$

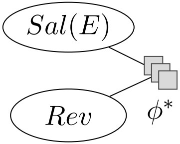

$Sal(E)$	Rev	$\phi^*(Sal(E), Rev)$
high	high	φ^*_1
high	low	φ^*_2
low	high	φ^*_3
low	low	φ^*_4

- How to choose φ^*_i? $\rightarrow$ $\varphi^*_i = \frac{1}{2}(\varphi_i + \varphi'_i)$, $\quad \phi^*(\boldsymbol{r}) = \frac{1}{m}\sum_{i=1}^{m} \phi_i(\boldsymbol{r})$ (currently)

One-dimensional ε-equivalence distance

- Potentials for convenience: $\phi_i(k) := \varphi_k$

Definition (Speller et al., 2025)

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

One-dimensional ε-equivalence distance

- Potentials for convenience: $\phi_i(k) := \varphi_k$

Definition (Speller et al., 2025)

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

Properties

(i) d_∞ is non-negative and symmetric.

(ii) It holds that $d_\infty(\phi_1, \phi_2) = 0$ if and only if $|\phi_1(k) - \phi_2(k)| = 0$ for all $k = 1, \dots, n$, which holds only if $\phi_1 = \phi_2$.

(iii) The triangle inequality $d_\infty(\phi_1, \phi_2) \leq d_\infty(\phi_1, \phi_3) + d_\infty(\phi_3, \phi_2)$ does *not* hold in general for three vectors ϕ_1, ϕ_2, ϕ_3.

(iv) 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.

One-dimensional ε-equivalence distance

- Potentials for convenience: $\phi_i(k) := \varphi_k$

Definition (Speller et al., 2025)

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

- Induces a distance measure for factors
- Enables pairwise comparisons and clustering
- $d_\infty(\phi_1, \phi_2)$ is equal to minimal ε, such that $\phi_1 =_\varepsilon \phi_2$
- Forms the backbone of hierarchical compression
  - Sorting of pairwise ε groups based on smallest d_∞ values
    (Agglomerative clustering algorithm based on d_∞ with complete linkage)

Towards Explainability: A Geometric Perspective in $\mathcal{L}^p$

For two ε-equivalent factors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$, it holds for $k = 1, \ldots, n$ that:

- Comparison of Lengths:
  - $\phi_{3-i}(k) \in \left[\phi_i(k) \cdot \frac{1}{1+\varepsilon}, \phi_i(k) \cdot (1+\varepsilon)\right]$ for $i = 1, 2$
  - $\frac{\phi_{3-i}(k)}{\phi_i(k)} \in \left[\frac{1}{1+\varepsilon}, 1+\varepsilon\right]$ for $i = 1, 2$

Towards Explainability: A Geometric Perspective in $\mathcal{L}^p$

For two ε-equivalent factors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$, it holds for $k = 1, \dots, n$ that:

- Comparison of Lengths:
  - $\phi_{3-i}(k) \in \left[\phi_i(k) \cdot \frac{1}{1+\varepsilon}, \phi_i(k) \cdot (1+\varepsilon)\right]$ for $i = 1, 2$
  - $\frac{\phi_{3-i}(k)}{\phi_i(k)} \in \left[\frac{1}{1+\varepsilon}, 1+\varepsilon\right]$ for $i = 1, 2$
- Distances in $\mathcal{L}^p$:
  - $\|\phi_1 - \phi_2\|_p \leq \varepsilon \|\phi_i\|_p$ holds for $i = 1, 2$
  - If $\|\phi_1 - \phi_2\|_p \leq \varepsilon$, then $\phi_1 =_{\varepsilon'} \phi_2$ holds with $\varepsilon' := \dfrac{\varepsilon}{\min\limits_{\substack{k=1,\dots,n,\\ i=1,2}} \{\phi_i(k)\}}$

Towards Explainability: A Geometric Perspective in $\mathcal{L}^p$

For two ε-equivalent factors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$, it holds for $k = 1, \ldots, n$ that:

- Comparison of Lengths:
  - $\phi_{3-i}(k) \in \left[\phi_i(k) \cdot \frac{1}{1+\varepsilon}, \phi_i(k) \cdot (1+\varepsilon)\right]$ for $i = 1, 2$
  - $\frac{\phi_{3-i}(k)}{\phi_i(k)} \in \left[\frac{1}{1+\varepsilon}, 1+\varepsilon\right]$ for $i = 1, 2$
- Distances in $\mathcal{L}^p$:
  - $\|\phi_1 - \phi_2\|_p \leq \varepsilon \|\phi_i\|_p$ holds for $i = 1, 2$
  - If $\|\phi_1 - \phi_2\|_p \leq \varepsilon$, then $\phi_1 =_{\varepsilon'} \phi_2$ holds with $\varepsilon' := \dfrac{\varepsilon}{\min\limits_{\substack{k=1,\ldots,n,\\ i=1,2}} \{\phi_i(k)\}}$

Example ($\varepsilon = 0.01$)

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

- $\|\phi_1 - \phi_2\|_p^p = \sum_{k=1}^2 |\phi_1(k) - \phi_2(k)|^p = 0.01^p + 0^p = 0.01^p = \varepsilon^p$
- $\varepsilon' = \dfrac{\varepsilon}{\min\limits_{\substack{k=1,\ldots,n,\\ i=1,2}} \{\phi_i(k)\}} = \frac{0.01}{0.01} = 1$

Towards Explainability: A Geometric Perspective in $\mathcal{L}^p$

For two ε-equivalent factors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$, it holds for $k = 1, \ldots, n$ that:

- Comparison of Lengths:
  - $\phi_{3-i}(k) \in \left[\phi_i(k) \cdot \frac{1}{1+\varepsilon}, \phi_i(k) \cdot (1+\varepsilon)\right]$ for $i = 1, 2$
  - $\frac{\phi_{3-i}(k)}{\phi_i(k)} \in \left[\frac{1}{1+\varepsilon}, 1+\varepsilon\right]$ for $i = 1, 2$
- Distances in $\mathcal{L}^p$:
  - $\|\phi_1 - \phi_2\|_p \leq \varepsilon \|\phi_i\|_p$ holds for $i = 1, 2$
  - If $\|\phi_1 - \phi_2\|_p \leq \varepsilon$, then $\phi_1 =_{\varepsilon'} \phi_2$ holds with $\varepsilon' := \dfrac{\varepsilon}{\min\limits_{\substack{k=1,\ldots,n,\\ i=1,2}} \{\phi_i(k)\}}$

$+$

$$\phi_1, \phi_2 \in S_p^{n-1} := \{\phi \in \mathbb{R}^n : \|\phi\|_p = 1\}$$

Towards Explainability: A Geometric Perspective in $\mathcal{L}^p$

For two ε-equivalent factors $\phi_1, \phi_2 \in \mathbb{R}^n_{>0}$, it holds for $k = 1, \dots, n$ that:

- Comparison of Lengths:
  - $\phi_{3-i}(k) \in \left[\phi_i(k) \cdot \frac{1}{1+\varepsilon}, \phi_i(k) \cdot (1+\varepsilon)\right]$ for $i = 1, 2$
  - $\frac{\phi_{3-i}(k)}{\phi_i(k)} \in \left[\frac{1}{1+\varepsilon}, 1+\varepsilon\right]$ for $i = 1, 2$
- Distances in $\mathcal{L}^p$:
  - $\|\phi_1 - \phi_2\|_p \leq \varepsilon \|\phi_i\|_p$ holds for $i = 1, 2$
  - If $\|\phi_1 - \phi_2\|_p \leq \varepsilon$, then $\phi_1 =_{\varepsilon'} \phi_2$ holds with $\varepsilon' := \dfrac{\varepsilon}{\min\limits_{\substack{k=1,\dots,n,\\ i=1,2}} \{\phi_i(k)\}}$

$+$

$$\phi_1, \phi_2 \in S_p^{n-1} := \{\phi \in \mathbb{R}^n : \|\phi\|_p = 1\}$$

  - $\|\phi_1 - \phi_2\|_p \leq \varepsilon$
  - $\phi_{3-i} \in \overline{B^p_\varepsilon}(\phi_i) \cap S_p^{n-1}$ for $i = 1, 2$ with $\overline{B^p_\varepsilon}(\phi_i) := \{\phi \in \mathbb{R}^n : \|\phi - \phi_i\|_p \leq \varepsilon\}$

ε-Equivalent Groups

Theorem (Speller et al., 2025)

Let $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ *denote a group of pairwise* ε*-equivalent factors and let* $\phi_\omega(\boldsymbol{r}) = \sum_{i=1}^{m} \omega_i \phi_i(\boldsymbol{r})$ *be the weighted mean with weights* $\omega_i \geq 0$ *and* $\sum_{i=1}^{m} \omega_i = 1$ *for all assignments* $\boldsymbol{r}$*. Then,* $\boldsymbol{G}^* = \{\phi_1, \ldots, \phi_m, \phi_\omega\}$ *is a group of pairwise* ε*-equivalent factors.*

ε-Equivalent Groups

Theorem (Speller et al., 2025)

Let $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ denote a group of pairwise ε-equivalent factors and let $\phi_\omega(\boldsymbol{r}) = \sum_{i=1}^m \omega_i \phi_i(\boldsymbol{r})$ be the weighted mean with weights $\omega_i \geq 0$ and $\sum_{i=1}^m \omega_i = 1$ for all assignments $\boldsymbol{r}$. Then, $\boldsymbol{G}^* = \{\phi_1, \ldots, \phi_m, \phi_\omega\}$ is a group of pairwise ε-equivalent factors.

$$\downarrow$$

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

ε-Equivalent Groups

Theorem (Speller et al., 2025)

Let $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ denote a group of pairwise ε-equivalent factors and let $\phi_\omega(\boldsymbol{r}) = \sum_{i=1}^m \omega_i \phi_i(\boldsymbol{r})$ be the weighted mean with weights $\omega_i \geq 0$ and $\sum_{i=1}^m \omega_i = 1$ for all assignments $\boldsymbol{r}$. Then, $\boldsymbol{G}^* = \{\phi_1, \ldots, \phi_m, \phi_\omega\}$ is a group of pairwise ε-equivalent factors.

$$\downarrow$$

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

- $\phi^*(\boldsymbol{r}) = \frac{1}{m} \sum_{i=1}^m \phi_i(\boldsymbol{r})$
- $\phi_{gm}(\boldsymbol{r}) := \sqrt[m]{\prod_{i=1}^m \phi_i(\boldsymbol{r})}$

ε-Equivalent Groups

Theorem (Speller et al., 2025)

Let $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ denote a group of pairwise ε-equivalent factors and let $\phi_\omega(\boldsymbol{r}) = \sum_{i=1}^m \omega_i \phi_i(\boldsymbol{r})$ be the weighted mean with weights $\omega_i \geq 0$ and $\sum_{i=1}^m \omega_i = 1$ for all assignments $\boldsymbol{r}$. Then, $\boldsymbol{G}^* = \{\phi_1, \ldots, \phi_m, \phi_\omega\}$ is a group of pairwise ε-equivalent factors.

$$\downarrow$$

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

- $\phi^*(\boldsymbol{r}) = \frac{1}{m} \sum_{i=1}^m \phi_i(\boldsymbol{r})$
- $\phi_{gm}(\boldsymbol{r}) := \sqrt[m]{\prod_{i=1}^m \phi_i(\boldsymbol{r})}$

$+$ $\phi_i \in S_{n-1}^p$

ε-Equivalent Groups

Theorem (Speller et al., 2025)

Let $\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ denote a group of pairwise ε-equivalent factors and let $\phi_\omega(\boldsymbol{r}) = \sum_{i=1}^m \omega_i \phi_i(\boldsymbol{r})$ be the weighted mean with weights $\omega_i \geq 0$ and $\sum_{i=1}^m \omega_i = 1$ for all assignments $\boldsymbol{r}$. Then, $\boldsymbol{G}^* = \{\phi_1, \ldots, \phi_m, \phi_\omega\}$ is a group of pairwise ε-equivalent factors.

$$\downarrow$$

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

- $\phi^*(\boldsymbol{r}) = \frac{1}{m} \sum_{i=1}^m \phi_i(\boldsymbol{r})$
- $\phi_{gm}(\boldsymbol{r}) := \sqrt[m]{\prod_{i=1}^m \phi_i(\boldsymbol{r})}$

$$+ \qquad \boxed{\phi_i \in S^p_{n-1}}$$

- $\|\phi_\omega\|_p \in \left[\frac{1}{1+\varepsilon}, 1\right]$ or $\phi_\omega \in \overline{B^p_1}(0) \setminus B^p_{1/(1+\varepsilon)}(0) = \left\{\phi \in \mathbb{R}^n : \frac{1}{1+\varepsilon} \leq \|\phi\|_p \leq 1\right\}$

Convex Extensions

$\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ pairwise ε-equivalent

- Convex Supersets
  - All pairwise ε-equivalent convex supersets of G, contain conv$(\phi_1, \ldots, \phi_m)$
  - Non-unique extension of this set in every dimension ($=_\varepsilon$ not transitive)
  - For hierarchical ACP $\rightarrow$ next compression level is build on maximal deviation d_∞ to all ϕ_i

Convex Extensions

$\boldsymbol{G} = \{\phi_1, \ldots, \phi_m\}$ pairwise ε-equivalent

- Convex Supersets
  - All pairwise ε-equivalent convex supersets of G, contain $\mathrm{conv}(\phi_1, \ldots, \phi_m)$
  - Non-unique extension of this set in every dimension ($=_\varepsilon$ not transitive)
  - For hierarchical ACP $\rightarrow$ next compression level is build on maximal deviation d_∞ to all ϕ_i

- Finitely Generated Convex Cone?
  - Scale all factors $\rightarrow$ Check ε-equivalence $\rightarrow$ Choose representative $\rightarrow \quad \phi_\omega \in \mathrm{conv}(\phi_1/\|\phi_1\|_p, \ldots, \phi_m/\|\phi_m\|_p)$
    - ε-Equivalence is *not* closed under rescaling onto S^p_{n-1} (For given group $G \subset S^p_{n-1}$ of ε-equivalent factors $\rightarrow \frac{\phi_\omega}{\|\phi_\omega\|_p} \in S^p_{n-1}$ is *not* necessarily pairwise ε-equivalent to G)
    - Sequence of operations plays major role for sorting *ε-equivalent* groups

Special Case: Euclidean Perspective

Definition (Cosine distance)

$$D_{\cos}\colon \mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0} \to [0,1], \quad D_{\cos}(\phi_1, \phi_2) := 1 - \frac{\phi_1 \cdot \phi_2}{||\phi_1||_2 \cdot ||\phi_2||_2}$$

which is equal to $1 - \cos(\theta)$ for one $\theta \in [0, \pi]$.

Special Case: Euclidean Perspective

Definition (Cosine distance)

$$D_{\cos}\colon \mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0} \to [0,1], \quad D_{\cos}(\phi_1,\phi_2) := 1 - \frac{\phi_1 \cdot \phi_2}{||\phi_1||_2 \cdot ||\phi_2||_2}$$

which is equal to $1-\cos(\theta)$ for one $\theta \in [0,\pi]$.

- If ϕ_1 and ϕ_2 are *exchangeable*, then it holds $D_{\cos}(\phi_1,\phi_2) = 0$.

Special Case: Euclidean Perspective

Definition (Cosine distance)

$$D_{\cos} \colon \mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0} \to [0,1], \quad D_{\cos}(\phi_1, \phi_2) := 1 - \frac{\phi_1 \cdot \phi_2}{||\phi_1||_2 \cdot ||\phi_2||_2}$$

which is equal to $1 - \cos(\theta)$ for one $\theta \in [0, \pi]$.

- If ϕ_1 and ϕ_2 are *exchangeable*, then it holds $D_{\cos}(\phi_1, \phi_2) = 0$.

$\oplus$ $\quad$ $\phi_i \in S^p_{n-1}$

Special Case: Euclidean Perspective

Definition (Cosine distance)

$$D_{\cos}\colon \mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0} \to [0,1], \quad D_{\cos}(\phi_1, \phi_2) := 1 - \frac{\phi_1 \cdot \phi_2}{||\phi_1||_2 \cdot ||\phi_2||_2}$$

which is equal to $1 - \cos(\theta)$ for one $\theta \in [0, \pi]$.

- If ϕ_1 and ϕ_2 are *exchangeable*, then it holds $D_{\cos}(\phi_1, \phi_2) = 0$.

$$\oplus \quad \boxed{\phi_i \in S^p_{n-1}}$$

- $2D_{\cos}(\phi_1, \phi_2) = ||\phi_1 - \phi_2||_2$
- $\phi_1 =_\varepsilon \phi_2 \Rightarrow$ the angle $\theta(\varepsilon) \leq \begin{cases} \arccos(1 - \frac{\varepsilon}{2}) & \text{for } \varepsilon \leq 2 \\ \frac{\pi}{2} & \text{for } \varepsilon > 2. \end{cases}$

Summary

- Possible generalisation via choice of representative $\phi^* \in \mathrm{conv}(\phi_1, \ldots, \phi_m)$
  - Theoretical guarantees for the change in query results still apply
  - Use of different loss functions to find optimal ϕ^*

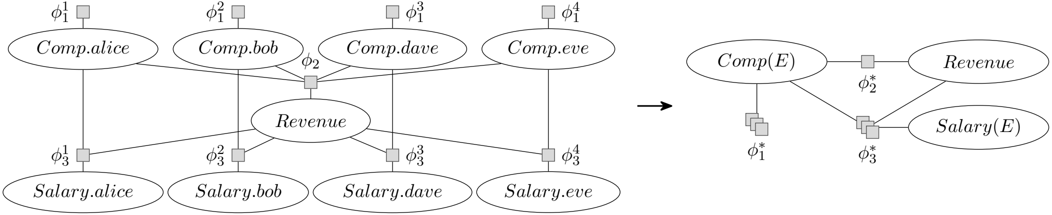

Summary

- Possible generalisation via choice of representative $\phi^* \in \text{conv}(\phi_1, \ldots, \phi_m)$
  - Theoretical guarantees for the change in query results still apply
  - Use of different loss functions to find optimal ϕ^*
- Choice of ε controls...
  - ...the trade-off between exactness and compactness
  - ...the distance in $\mathcal{L}^p$
  - ...the angle in $\mathcal{L}^2$ (controlling a convex cone)
  - ...a hierarchy of convex supersets containing each other

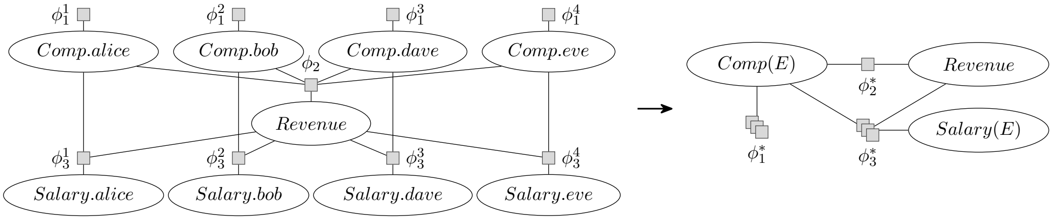

References

- Babak Ahmadi et al. (2013). »Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training«. *Machine Learning* 92.1, pp. 91–132.
- Kristian Kersting et al. (2009). »Counting Belief Propagation«. *UAI-09 Proc. of the 25th Conference on Uncertainty in Artificial Intelligence*. AUAI Press, pp. 277–284.
- Malte Luttermann et al. (2024). »Colour Passing Revisited: Lifted Model Construction with Commutative Factors«. *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-2024)*. AAAI Press, pp. 20500–20507.
- Malte Luttermann et al. (2025). »Approximate Lifted Model Construction«. *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-2025)*. https://arxiv.org/abs/2504.20784.
- Jan Speller et al. (2025). »Compression versus Accuracy: A Hierarchy of Lifted Models«. https://arxiv.org/abs/2505.22288.