



StaRAI: From a Probabilistic Propositional Model to a Highly Compressed Probabilistic Relational Model

[Marcel Gehrke](#)¹, [Malte Luttermann](#)²




¹Institute of Humanities-Centered Artificial Intelligence, University of Hamburg
²German Research Center for Artificial Intelligence (DFKI)



AGENDA

1. Introduction to relational models [Marcel]
2. Compressing probabilistic relational models [Malte]
3. Application: Lifted causal inference [Malte]
4. Summary [Marcel]

STATISTICAL RELATIONAL ARTIFICIAL INTELLIGENCE (STARAI)




EXPLOITING SYMMETRIES

10 $Presents(X, P, C) \Rightarrow Attends(X, C)$

- Exchangeable random variables in the full joint probability distribution
 - Inference using representatives
 - Tractability in terms of domain sizes
 - Also speeds up causal inference

3.75 $Publishes(X, C) \wedge FarAway(C) \Rightarrow Attends(X, C)$



OBTAINING A HIGHLY COMPRESSED PROBABILISTIC RELATIONAL MODEL

- Identifying symmetries within factors to compress probabilistic relational models even further
 - by identifying exchangeable factors
 - by identifying commutative factors
- To identify exchangeable and commutative factors use so-called buckets to drastically reduce the search space

A	B	$\phi_1(A, B)$	b
true	true	φ_1	[2, 0]
true	false	φ_2	[1, 1]
false	true	φ_3	[1, 1]
false	false	φ_4	[0, 2]

b	$\phi_1^>(b)$	$\phi_2^>(b)$
[2, 0]	$\langle \varphi_1 \rangle$	$\langle \varphi_1 \rangle$
[1, 1]	$\langle \varphi_2, \varphi_3 \rangle$	$\langle \varphi_3, \varphi_2 \rangle$
[0, 2]	$\langle \varphi_4 \rangle$	$\langle \varphi_4 \rangle$

Possible rearrangements:
 [2, 0]: $B \rightarrow \{1, 2\}, C \rightarrow \{1, 2\}$
 [1, 1]: $B \rightarrow \{2\}, C \rightarrow \{1\}$
 [0, 2]: $B \rightarrow \{1, 2\}, C \rightarrow \{1, 2\}$
 Intersection: $B \rightarrow \{2\}, C \rightarrow \{1\}$


A	B	R	$\phi(A, B, R)$	b
true	true	true	φ_1	[3, 0]
true	true	false	φ_4	[2, 1]
true	false	true	φ_2	[2, 1]
true	false	false	φ_5	[1, 2]
false	true	true	φ_2	[2, 1]
false	true	false	φ_5	[1, 2]
false	false	true	φ_3	[1, 2]
false	false	false	φ_6	[0, 3]

b	$\phi(b)$
[3, 0]	$\langle \varphi_1 \rangle$
[2, 1]	$\langle \varphi_4, \varphi_2, \varphi_2 \rangle$
[1, 2]	$\langle \varphi_5, \varphi_5, \varphi_3 \rangle$
[0, 3]	$\langle \varphi_6 \rangle$

Initial candidates: $\{\{A, B, R\}\}$
 [3, 0] ($|\phi(b)| < 2$): $\{\{A, B, R\}\}$
 [2, 1]: $\{\{A, B\}\}$
 [1, 2]: $\{\{A, B\}\}$
 [0, 3] ($|\phi(b)| < 2$): $\{\{A, B\}\}$

APPROXIMATE LIFTED MODEL CONSTRUCTION

- Allow for small deviations in potentials, while compressing models
 - Make it more practically
 - Group factors in case they are ϵ -equivalent
 - Known bounded error based on ϵ
 - Hierarchical extension exists
- Resolve unknown factors while compressing models
 - Find best fitting group based on graph structure and known factors having a similar graph structure



SOME THINGS WE DID NOT TALK ABOUT TODAY: MORE ON LIFTED INFERENCE

- Lifted evidence [Van den Broeck & Davis 12]
- Lifted queries [Braun & Möller 18]
- Assignment queries: Lifted inference for MPE [de Salvo Braz et al. 06, Apsel & Brafman 12, Braun & Möller 19] and MAP queries [Braun 20]
- Continuous inference [Choi et al. 10, Hartwig et al. 23]
- Lifted variational inference in hybrid models [Choi & Amir 12]

Plus, learning of relational models:
e.g., boosted relational learning
[Natarajan et al. 12]


SOME THINGS WE DID NOT TALK ABOUT TODAY: DECISION MAKING

- Models:
 - Decision-theoretic ProbLog [Van den Broeck et al. 10]
 - First-order (partially observable) Markov decision processes (FO (PO)MDPs) [Boutelier et al. 01]
 - Markov logic decision networks [Nath & Domingos 09]
 - (Temporal) decision parfactor models [G et al. 19b, c]
 - Lifting the agent set in decentralised POMDPs (multi-agent setting) [Braun et al. 22]
- Solution methods:
 - Symbolic dynamic programming for FO POMDPs [Sanner & Kersting 10]
 - L(D)JT for decision parfactor models [G et al. 19b, c]

WHAT ELSE IS THERE TO DO?

- Develop more robust learning algorithms
- Incorporate additional requirements such as
 - Privacy [G et al. 24]
 - Ethics
 - Explainability
 - Human-awareness
- And so much more...

Thank you!





BIBLIOGRAPHY

ORDERED ALPHABETICALLY

BIBLIOGRAPHY

- **Ahmadi et al (2013)**
Babak Ahmadi, Kristian Kersting, Martin Mladenov, and Sriraam Natarajan: Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training. In Machine Learning, 2013.
- **Apsel and Brafman (2012)**
Udi Apsel and Ronen I. Brafman. Exploiting Uniform Assignments in First-Order MPE. In Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence, 2012.
- **Boutilier et al. (2010)**
Craig Boutilier, Ray Reiter, and Bob Price: Symbolic Dynamic Programming for First-order MDPs. In IJCAI-01 Proceedings of the 17th International Joint Conference on Artificial Intelligence, 2001.
- **Braun & Möller (2018)**
Tanya Braun and Ralf Möller: Parameterised Queries and Lifted Query Answering. In IJCAI-18 Proceedings of the 27th International Joint Conference on Artificial Intelligence, 2018.
- **Braun (2020)**
Tanya Braun: Rescued from a Sea of Queries: Exact Inference in Probabilistic Relational Models. PhD thesis, 2020.

BIBLIOGRAPHY

- **Braun et al. (2022)**
Tanya Braun, Marcel Gehrke, Florian Lau, and Ralf Möller: Lifting in Multi-agent Systems under Uncertainty. In UAI-22 International Conference on Uncertainty in Artificial Intelligence, 2022.
- **Choi et al. (2010)**
Jaesik Choi, Eyal Amir, and David J. Hill: Lifted Inference for Relational Continuous Models. In UAI-10 Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence, 2010.
- **Choi & Amir (2012)**
Jaesik Choi and Eyal Amir: Lifted Relational Variational Inference. In UAI-12 Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence, 2012.
- **de Salvo Braz et al. (2006)**
Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth: MPE and Partial Inversion in Lifted Probabilistic Variable Elimination. In AAAI-06 Proceedings of the 21st Conference on Artificial Intelligence, 2006.
- **G et al. (2019b)**
Marcel Gehrke, Tanya Braun, Ralf Möller, Alexander Waschkau, Christoph Strumann, and Jost Steinhäuser. Lifted Maximum Expected Utility. In Artificial Intelligence in Health, 2019.

BIBLIOGRAPHY

- **G et al. (2019c)**
Marcel Gehrke, Tanya Braun, and Ralf Möller. Lifted Temporal Maximum Expected Utility. In CanadianAI-19 Proceedings of the 32nd Canadian Conference on Artificial Intelligence, 2019.
- **G et al. (2024)**
Marcel Gehrke, Johannes Liebenow, Esfandiar Mohammadi, and Tanya Braun. Lifting in Support of Privacy-Preserving Probabilistic Inference. In KI - Künstliche Intelligenz, 2024.
- **Gogate & Domingos (2011)**
Vibhav Gogate and Pedro Domingos: Probabilistic Theorem Proving. In UAI-11 Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence, 2011.
- **Gogate et al. (2012)**
Vibhav Gogate, Abhay Jha, and Deepak Venugopal: Advances in Lifted Importance Sampling. In AAAI-12 Proceedings of the 26th AAAI Conference on Artificial Intelligence, 2012.
- **Hartwig et al. (2023)**
Mattis Hartwig, Ralf Möller, and Tanya Braun: An Extended View on Lifting Gaussian Bayesian Networks. In Artificial Intelligence Journal, 2023.
- **Natarajan et al. (2012)**
Srijaam Natarajan, Tushar Khot, Kristian Kersting, Bernd Gutmann, and Jude Shavlik: Gradient-based Boosting for Statistical Relational Learning: The Relational Dependency Network Case. In Machine Learning, 2012.

BIBLIOGRAPHY

*PRMs are a true backbone of AI, and this tutorial emphasized only some central topics. We definitely did not cite all publications relevant to the whole field of PRMs here. We would like to thank all our colleagues for making their slides available (see some of the references to papers for respective credits). Slides or parts of it are almost always modified.

- **Nath & Domingos (2009)**
Aniruddh Nath and Pedro Domingos. A Language for Relational Decision Theory. In Proceedings of the 6th International Workshop on Statistical Relational Learning, 2009.
- **Niepert (2012)**
Mathias Niepert: Markov Chains on Orbits of Permutation Groups. In UAI-12 Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence, 2012
- **Sanner & Kersting (2010)**
Scott Sanner and Kristian Kersting. Symbolic Dynamic Programming for First-order POMDPs. In AAAI-10 Proceedings of the 24th AAAI Conference on Artificial Intelligence, 2010.
- **Van den Broeck et al. (2010)**
Guy Van den Broeck, Ingo Thon, Martijn van Otterlo, and Luc De Raedt: DTProbLog: A Decision-theoretic Probabilistic Prolog. In AAAI-10 Proceedings of the 24th AAAI Conference on Artificial Intelligence, 2010.
- **Van den Broeck & Davis (2012)**
Guy Van den Broeck and Jesse Davis: Conditioning in First-Order Knowledge Compilation and Lifted Probabilistic Inference. In AAAI-12 Proceedings of the 26th AAAI Conference on Artificial Intelligence, 2012.