COMPRESSION VS. ACCURACY:
COMPACT MODELS FOR EFFICIENCY AND INTERPRETATION

MALTE LUTTERMANN[1], JAN SPELLER[2], MARCEL GEHRKE[1], AND TANYA BRAUN[2]

[1]Institute of Humanities-Centered Artificial Intelligence, University of Hamburg
[2]Computer Science Department, University of Münster

AGENDA

1. Introduction and background [Tanya]
2. Compressing models [Malte]
3. Accuracy considerations [Jan]
4. Summary and future directions [Tanya]

PROBABILISTIC RELATIONAL MODELS & REASONING

- Combine objects and relations among them with a notion of uncertainty
  - Starting with uncertainty, adding objects and relations: Parfactor graphs
  - Starting with objects and relations, adding uncertainty: MLNs
  - Grounding semantics: Symmetries in factors / weights
- Reasoning algorithm for parfactor graphs, e.g., LVE
  - Sum out representative and reuse result for indistinguishable instances
  - Allows for tractable probabilistic inference

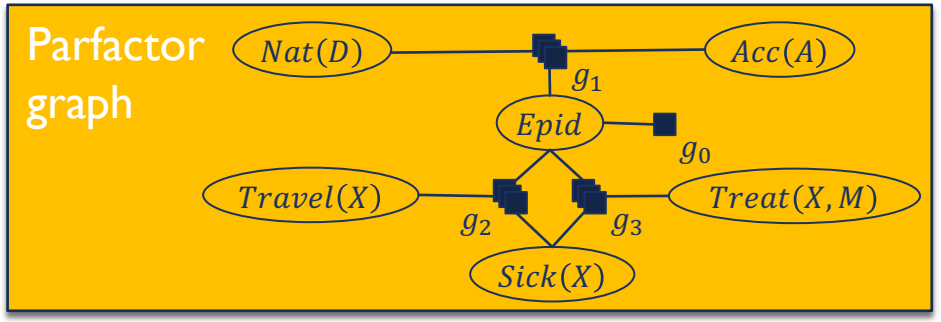

MLNs

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

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

$$\left(\sum_{t \in r(Treat(X,M))} g_3(Epid = e, Sick(X) = s, Treat(X, M) = t) \right)^{\#M|X}$$

COMPRESSION

- Advanced Colour Passing for (exact) lifted model construction
  - Efficient detection and compression of *commutative* factors
  - Efficient detection and grouping of *exchangeable* factors
  - Similar idea for learning an MLN by lifting a hypergraph representation
    [Kok and Domingos 2009]
- ε-ACP as a generalisation of ACP for approximate lifted model construction
  - Bound on the change in query results

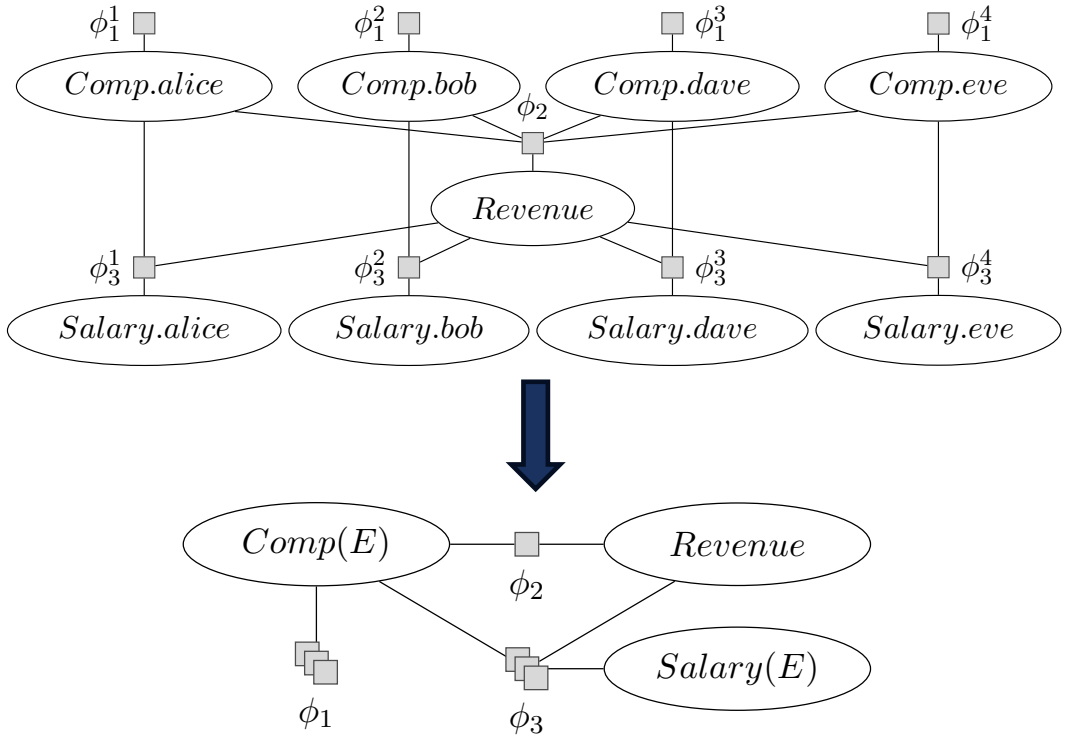

ACCURACY

- Goal: Relax strict equality between potentials while retaining interpretability and theoretical guarantees
  - Distance measure d_∞ determines degree of approximation
- HACP: hyperparameter-free hierarchy for interpretable model comparisons
  - Geometric interpretation: Close neighbourhood on the unit sphere → Representatives from their convex hull
- Both HACP and ε-ACP admit *asymptotic bounds* for every ε level / level of compression

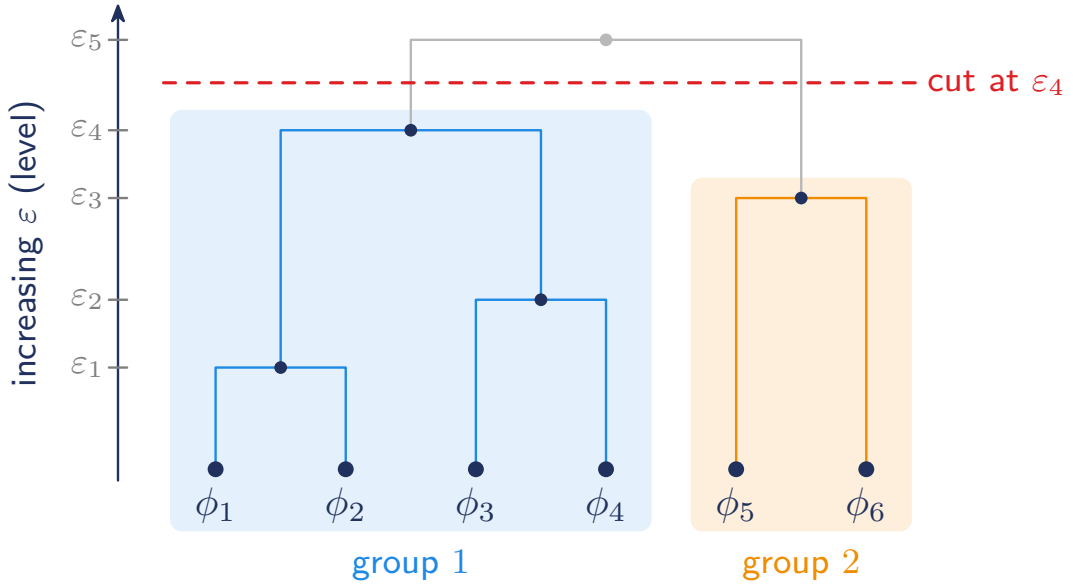

SOME THINGS WE DID NOT TALK ABOUT TODAY: MODELLING & REASONING

Probabilistic inference

- Assignment queries:
  - Lifted inference for MPE
    [de Salvo Braz et al. 2006, Apsel & Brafman 2012, Sarkhel et al. [2014], B & Möller 2019]
  - Lifted inference for MAP queries
    [Sharma et al. 2018, B 2020]
- Continuous inference
  [Choi et al. 2010, Hartwig et al. 2023]
- Lifted variational inference in hybrid models
  [Choi & Amir 2012]

Decision-theoretic modelling and inference

- Decision-theoretic ProbLog
  [Van den Broeck et al. 2010]
- Lifting in Markov decision process based formalisms
  [Boutelier et al. 2001, Sanner & Kersting 2010, B et al. 2022, Marwitz et al. 2026]
- Markov logic decision networks
  [Nath & Domingos 2009]
- (Temporal) decision parfactor models
  [Gehrke et al. 2019a, b]

SOME THINGS WE DID NOT TALK ABOUT TODAY: LEARNING PROBABILISTIC RELATIONAL MODELS DIRECTLY

- Learning MLNs
  - Log-linear model → allows for component-wise optimisation during weight learning

$$P(\omega) = \frac{1}{Z}\exp\left(\sum_{i=1}^{n} w_i n_i(\omega)\right) = \frac{1}{Z} weight(\omega) \rightarrow lweight(\omega) = \ln\exp\left(\sum_{i=1}^{n} w_i n_i(\omega)\right) = \sum_{i=1}^{n} w_i n_i(\omega)$$

- Focus on
  - Parameters (weights) w_i
    [Singla and Domingos 2005, Lowd and Dominogs 2007, Van Haaren et al. 2015]
  - Structure
    - Could consider learning logical statements, e.g., First-order Inductive Learning (FOIL) [Quinlan 1990], and then learn weights
    - Learn structure and weight at the same time, e.g., using ILP [Kok and Domingos 2005, Mihalkova and Mooney 2007], structural motifs [Kok and Domingos 2010], boosting [Khot et al. 2011]

Learning *relational dependency graphs* using boosting [Natarajan et al. 2012]

[selection of works, not complete; apologies to all we have missed]

THE LARGER SCOPE

Statistical Relational Learning & AI

- Study and design
  - intelligent agents
  - that reason about and
  - act in noisy worlds
  - composed of objects and relations among the objects

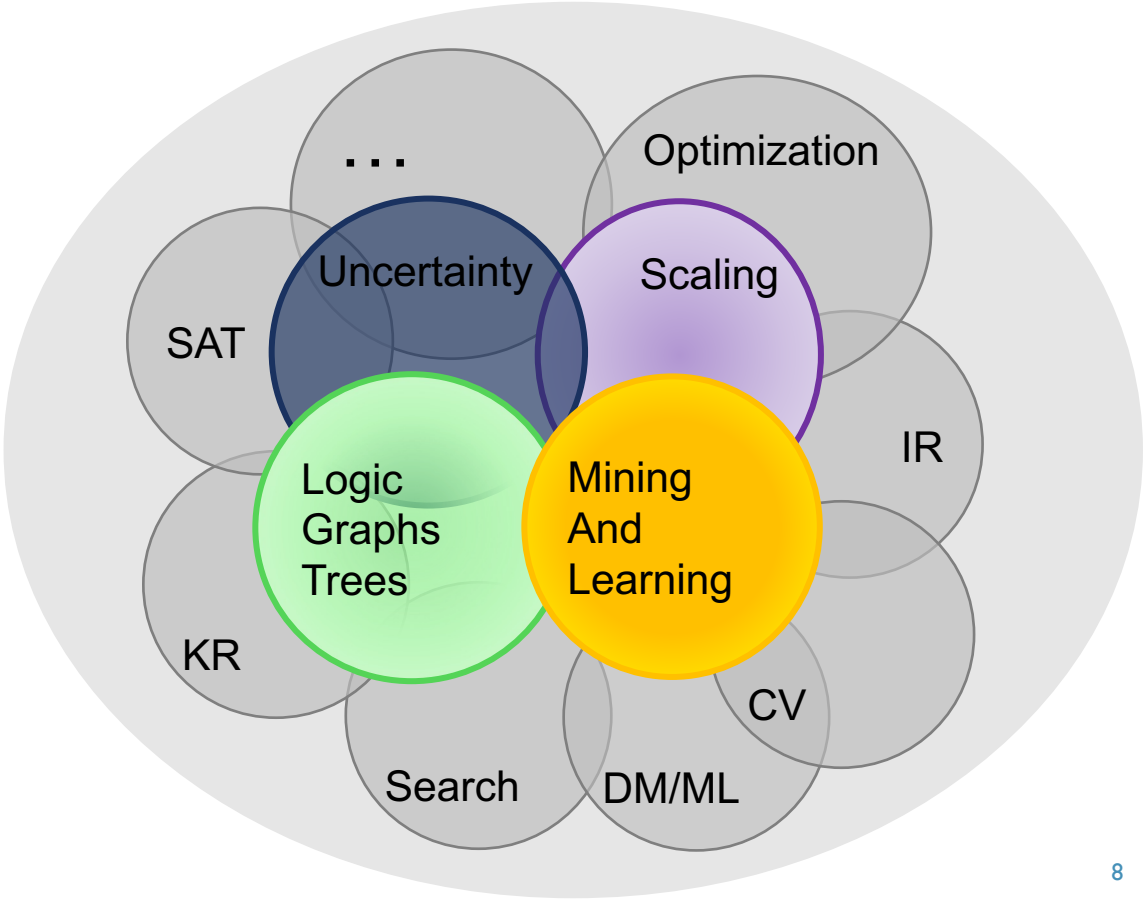

Kristian Kersting. Statistical Relational AI. Tutorial at KI 2018.

WHAT ELSE IS THERE TO DO?

- Further enhance lifting applicability
- Develop more robust direct learning algorithms
- Incorporate additional requirements such as
  - Privacy
  - Ethics
  - Explainability
  - Human-awareness
- And so much more…
- For slides, please go to (QR code goes to this address):
  https://www.uni-muenster.de/Informatik.AGBraun/en/research/tutorials/ijcai-26.html

Thank you!

BIBLIOGRAPHY

ORDERED ALPHABETICALLY

BIBLIOGRAPHY

- 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.
- B [2020]
  Tanya Braun: Rescued from a Sea of Queries: Exact Inference in Probabilistic Relational Models. *PhD thesis*, 2020.
- B 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.

BIBLIOGRAPHY

- 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.
- Gehrke et al. [2019a]
  Marcel Gehrke, Tanya Braun, Ralf Möller, Alexander Waschkau, Christoph Strumann, and Jost Steinhäuser. Lifted Maximum Expected Utility. In *Artificial Intelligence in Health*, 2019.

BIBLIOGRAPHY

- Gehrke et al. [2019b]
  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.
- Hartwig et al. [2023]
  Mattis Hartwig, Ralf Möller, and Tanya Braun: An Extended View on Lifting Gaussian Bayesian Networks. In *Artificial Intelligence Journal*, 2023.
- Khot et al. [2011]
  Tushar Khot, Sriraam Natarajan, Kristian Kersting, and Jude Shavlik: Learning Markov Logic Networks via Functional Gradient Boosting. In *ICDM-2011 Proceedings of the 11th IEEE International Conference on Data Mining*, 2011.
- Kok and Domingos [2005]
  Stanley Kok and Pedro Domingos: Learning the Structure of Markov Logic Networks. In *ICML-2005 Proceedings of the 22nd International Conference on Machine Learning*, 2005.

BIBLIOGRAPHY

- Kok and Domingos [2009]
  Stanley Kok and Pedro Domingos: Learning Markov Logic Network Structure via Hypergraph Lifting. In *ICML-2009 Proceedings of the 26th International Conference on Machine Learning*, 2009.
- Kok and Domingos [2010]
  Stanley Kok and Pedro Domingos: Learning Markov Logic Networks using Structural Motifs. In *ICML-2010 Proceedings of the 27th International Conference on Machine Learning*, 2010.
- Lowd and Domingos [2007]
  Daniel Lowd and Pedro Domingos: Efficient Weight Learning for Markov Logic Networks. In *PKDD-2007 Proceedings of the 11th European Conference on Principles and Practice of Knowledge Discovery in Databases*, 2007.
- Marwitz et al. [2026]
  Florian Andreas Marwitz, Tanya Braun, Ralf Möller, and Marcel Gehrke: Lifted Forward Planning in Relational Factored Markov Decision Processes with Concurrent Actions. In *AAMAS-26 Proc. of the 25th International Conference on Autonomous Agents and Multi-Agent Systems*, 2026.

BIBLIOGRAPHY

- Mihalkova and Mooney [2007]
  Lilyana Mihalkova and Raymond J. Mooney: Bottom-up Learning of Markov Logic Network Structure. In *ICML-2007 Proceedings of the 24th International Conference on Machine Learning*, 2007.
- Natarajan et al. [2012]
  Sriraam 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.
- 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.
- Quinlan [1990]
  John Ross Quinlan: Learning Logical Definitions from Relations. In *Machine Learning*, 1990.

BIBLIOGRAPHY

- 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.
- Sarkhel et al. [2014]
  Somdeb Sarkhel, Deepak Venugopal, Parag Singla, and Vibhav Gogate: Lifted MAP Inference for Markov Logic Networks. In *AISTATS-14 Proc. of the 17th International Conference on Artificial Intelligence and Statistics*, 2014.
- Sharma et al. [2018]
  Vishal Sharma, Noman Ahmed Sheikh, Happy Mittal, Vibhav Gogate, and Parag Singla: Lifted Marginal MAP Inference. In *UAI-18 Proc. of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018.
- Singla and Domingos [2005]
  Parag Singla and Pedro Domingos. Discriminative Training of Markov Logic Networks. In *AAAI2005 Proceedings of the 20th National Conference on Artificial Intelligence*, 2005.

BIBLIOGRAPHY

- 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 Haaren et al. [2015]
  Jan Van Haaren, Guy Van den Broeck, Wannes Meert, and Jesse Davis: Lifted generative learning of Markov logic networks. In *Machine Learning*, 2015.