StaRAI: From a Probabilistic Propositional Model to a Highly Compressed Probabilistic Relational Model (Extended Abstract)

Marcel Gehrke[1], Malte Luttermann[2]

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

Abstract
This extended abstract summarises the content of our tutorial titled "StaRAI: From a Probabilistic Propositional Model to a Highly Compressed Probabilistic Relational Model" presented at the 18th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2025) held on September 23–26, 2025, in Hagen, Germany.

1. Introduction

Our surrounding world is inherently uncertain and relational. The field of Statistical Relational AI (StaRAI) has emerged to account for both uncertainty and relational modelling, for example in probabilistic graphical models [1]. In StaRAI, probabilistic models explicitly encode objects and relations, which enables algorithms to exploit repeated structures, i.e., subgraphs with matching associated probability functions as well as identical graph structure, for efficiency gains during inference [2, 3, 4]. To exploit repeated structures, they first have to be identified. While such repeated structures frequently occur in many practical applications, they are generally not explicitly represented in a learned model and thus cannot be exploited by inference algorithms. By identifying and compactifying repeated structures, the size of a model can be significantly reduced, resulting in inference algorithms no longer running in time exponential, but only polynomial, with respect to so-called domain sizes [5, 6, 7]. It is therefore crucial to efficiently identify and compactify repeated structures to enable efficient inference.

2. Compression of Probabilistic Relational Models

Our tutorial provides a look at recent advances in the field of computing a highly compressed probabilistic relational model from a given probabilistic propositional model. In the past, it has been shown that a colour passing procedure based on a belief propagation scheme can be used to identify repeated structures in a propositional probabilistic model [4, 8]. Identified repeated structures can then be exploited to compress the model, which results in a significant reduction in storage requirements and drastically increases the efficiency of inference algorithms while at the same time maintaining exact inference results [7]. The current state-of-the-art algorithm to identify and compress repeated structures in a propositional probabilistic model is the Advanced Colour Passing (ACP) algorithm [9], which builds on the CompressFactorGraph algorithm introduced by Kersting et al. [8]. ACP compresses repeated structures by introducing so-called logical variables such that the resulting compressed model allows for efficient inference independent of a specific inference algorithm, whereas the CompressFactorGraph algorithm is tied to the belief propagation inference algorithm. Additional major advancements of ACP compared to its predecessor are that ACP is able to handle (i) commutative functions (i.e., associated

The 18th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2025), September 23–26, 2025, Hagen, Germany
marcel.gehrke@uni-hamburg.de (M. Gehrke); malte.luttermann@dfki.de (M. Luttermann)
0000-0001-9056-7673 (M. Gehrke); 0009-0005-8591-6839 (M. Luttermann)

probability functions that map to a unique output value regardless of the order of some input values) and (ii) arbitrary orders of arguments in associated probability functions. By tackling both (i) and (ii), ACP solves significant practical limitations, as commutative functions frequently occur in practical applications and no specific order of arguments can be guaranteed in a learned model. Accounting for commutative functions and arbitrary orders of arguments requires additional offline computations, which have been efficiently implemented by using histograms that count the occurrences of specific input values for a (sub)set of a function's arguments [10, 11]. By deploying histograms within the ACP algorithm, ACP is able to efficiently identify and compress a given probabilistic propositional model while at the same time ensuring exact inference results when applying the resulting compressed model.

The ACP algorithm requires the associated probability functions to exactly match in order to identify and group repeated structures. However, in practical applications, associated probability functions slightly differ even if they should be considered equal due to inaccurate estimates, e.g., when learning the functions from data. To solve the problem of constructing a compressed representation from a given probabilistic propositional model taking inaccurate estimates of associated probability functions into account, the ε-Advanced Colour Passing (ε-ACP) algorithm has been introduced as a generalisation of the ACP algorithm [12]. The ε-ACP algorithm allows for a deviation between functions depending on a hyperparameter ε in order for functions to be considered equivalent. By doing so, ε-ACP efficiently identifies and compresses repeated structures that are not exactly identical. It has been proven that the change in query results induced by ε-ACP is strictly bounded and empirical experiments have shown that the change in query results is even close to zero in practice, validating that the compressed model returned by ε-ACP is an eligible approximation of the original model. Speller et al. [13] have further extended the ε-ACP algorithm to obtain a hierarchic version called Hierarchical Advanced Colour Passing (HACP) that efficiently computes a hierarchy of ε values, thereby yielding a hyperparameter-free approach with known error bounds on the change in query results. The HACP algorithm ensures that repeated structures that have been grouped for a specific choice of ε remain grouped for a larger value of ε as well, thus enhancing interpretability. Moreover, HACP allows to control the trade-off between compression and accuracy of query results by specifying an acceptable error bound on the change in query results, from which the corresponding value for ε is automatically derived (such that no value for ε has to be chosen manually).

To allow for the compression of a probabilistic propositional model containing functions whose function definitions are unknown, another extension of ACP has been proposed [14, 15]. By analysing the graph structure of the model, repeated structures are identified and known function definitions of associated probability functions are transferred to associated probability functions with the same surrounding graph structure but unknown function definitions. The equivalence of associated probability functions can also be determined by representing their function definitions as vectors and checking for the collinearity of these vectors, which allows to identify equivalent functions whose range (codomain) values lie on different scales without having to apply any normalisation step to the functions [16].

Obtaining a compressed probabilistic relational model is not only relevant to perform efficient probabilistic inference but also opens up new possibilities to efficiently perform causal inference efficiently. A compressed model computed by ACP that is enriched by causal knowledge can be used to perform efficient causal inference [17], even if only partial causal knowledge is available [18].

3. Conclusion

In the field of StaRAI, probabilistic models explicitly encode objects and relations, thereby allowing inference algorithms to exploit repeated structures for efficiency gains. Identifying and compressing these structures are thus fundamental problems that must be solved to enable efficient inference. Recently, there have been significant advancements in the field of computing a highly compressed probabilistic relational model from a given probabilistic propositional model. Specifically, the ACP algorithm and its variants have emerged to identify and compress repeated structures. This work provides an overview of recent advancements in the field of compressing probabilistic models.

Acknowledgements

This work was supported by the BMBF project AnoMed 16KISA057.

References

[1] D. Koller, N. Friedman, Probabilistic Graphical Models: Principles and Techniques, MIT Press, 2009.

[2] D. Poole, First-Order Probabilistic Inference, in: Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-2003), Morgan Kaufmann Publishers Inc., 2003, pp. 985–991.

[3] G. Van Den Broeck, N. Taghipour, W. Meert, J. Davis, L. De Raedt, Lifted Probabilistic Inference by First-Order Knowledge Compilation, in: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-2011), AAAI Press, 2011, pp. 2178–2185.

[4] B. Ahmadi, K. Kersting, M. Mladenov, S. Natarajan, Exploiting Symmetries for Scaling Loopy Belief Propagation and Relational Training, Machine Learning 92 (2013) 91–132.

[5] G. Van den Broeck, On the Completeness of First-Order Knowledge Compilation for Lifted Probabilistic Inference, in: Advances in Neural Information Processing Systems 24 (NIPS-2011), Curran Associates, Inc., 2011, pp. 1386–1394.

[6] N. Taghipour, D. Fierens, G. Van den Broeck, J. Davis, H. Blockeel, Completeness Results for Lifted Variable Elimination, in: Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS-2013), PMLR, 2013, pp. 572–580.

[7] M. Niepert, G. Van den Broeck, Tractability through Exchangeability: A New Perspective on Efficient Probabilistic Inference, in: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-2014), AAAI Press, 2014, pp. 2467–2475.

[8] K. Kersting, B. Ahmadi, S. Natarajan, Counting Belief Propagation, in: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI-2009), AUAI Press, 2009, pp. 277–284.

[9] M. Luttermann, T. Braun, R. Möller, M. Gehrke, Colour Passing Revisited: Lifted Model Construction with Commutative Factors, in: Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-2024), AAAI Press, 2024, pp. 20500–20507.

[10] M. Luttermann, J. Machemer, M. Gehrke, Efficient Detection of Exchangeable Factors in Factor Graphs, in: Proceedings of the Thirty-Seventh International Florida Artificial Intelligence Research Society Conference (FLAIRS-2024), Florida Online Journals, 2024.

[11] M. Luttermann, J. Machemer, M. Gehrke, Efficient Detection of Commutative Factors in Factor Graphs, in: Proceedings of the Twelfth International Conference on Probabilistic Graphical Models (PGM-2024), PMLR, 2024, pp. 38–56.

[12] M. Luttermann, J. Speller, M. Gehrke, T. Braun, R. Möller, M. Hartwig, Approximate Lifted Model Construction, in: Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-2025), IJCAI Organization, 2025.

[13] J. Speller, M. Luttermann, M. Gehrke, T. Braun, Compression versus Accuracy: A Hierarchy of Lifted Models, in: Proceedings of the Twenty-Eighth European Conference on Artificial Intelligence (ECAI-2025), IOS Press, 2025.

[14] M. Luttermann, R. Möller, M. Gehrke, Lifting Factor Graphs with Some Unknown Factors, in: Proceedings of the Seventeenth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2023), Springer, 2023, pp. 337–347.

[15] M. Luttermann, R. Möller, M. Gehrke, Lifting Factor Graphs with Some Unknown Factors for New Individuals, International Journal of Approximate Reasoning 179 (2025) 109371.

[16] M. Luttermann, R. Möller, M. Gehrke, Lifted Model Construction without Normalisation: A Vectorised Approach to Exploit Symmetries in Factor Graphs, in: Proceedings of the Third Learning on Graphs Conference (LoG-2024), PMLR, 2025, pp. 46:1–46:17.

[17] M. Luttermann, M. Hartwig, T. Braun, R. Möller, M. Gehrke, Lifted Causal Inference in Relational Domains, in: Proceedings of the Third Conference on Causal Learning and Reasoning (CLeaR-2024), PMLR, 2024, pp. 827–842.

[18] M. Luttermann, T. Braun, R. Möller, M. Gehrke, Estimating Causal Effects in Partially Directed Parametric Causal Factor Graphs, in: Proceedings of the Sixteenth International Conference on Scalable Uncertainty Management (SUM-2024), Springer, 2024, pp. 265–280.