Inference in Statistical Relational AI

Tanya Braun[0000-0003-0282-4284] and Marcel Gehrke[0000-0001-9056-7673]

Institute of Information Systems, Universität zu Lübeck, Lübeck
{braun,gehrke}@ifis.uni-luebeck.de

Abstract. In recent years, a need for efficient inference algorithms on compact representations of large relational databases became apparent, e.g., machine learning or decision making. This need has lead to advances in probabilistic relational modelling for artificial intelligence, also known as statistical relational AI (StarAI). This tutorial provides a detailed introduction into exact inference in StarAI.

Keywords: StarAI · Exact inference · Uncertainty.

1 Introduction

Our real world consists of many individuals or objects connected to each other, and a whole lot of uncertainty. Probabilistic relational models (PRMs) allow for modelling individuals/objects and relations between them as well as incorporating uncertainties. Random variables are used to set individuals into relations to form PRMs. Within such a model, probabilities are used to represent uncertainty. Uncertainty may range from uncertain properties of individuals over uncertain relationships among individuals to uncertain identities or even existence.

PRMs exist in the form of Markov logic networks [8] or parameterised factor models [7] among others. They allow for compactly representing a world filled with many objects and recurring patterns. Inference in such models includes answering queries, e.g., for a probability of an event or a most probable explanation of a current state of a world. Query answering algorithms aim at answering such queries in an efficient way, leveraging relational structures as much as possible.

To leverage relational structures, individuals are treated identically as long as nothing is known about them [6]. Lifting, first introduced by Poole [7], uses exchangeability of random variables and their dependencies to speed up runtimes by avoiding repeated calculations in variable elimination (VE). Variable elimination [11] is one of the standard algorithms to perform query answering in probabilistic models. Since Poole's first paper, researchers have taken up lifting and applied it to various well-understood algorithms that work on propositional models such as knowledge compilation (KC) on the basis of weighted model counting (WMC) [2] to answer queries.

2 Overview

This tutorial provides an overview of PRMs, inference problems in them, and inference algorithms to solve the problems with a focus on lifted inference algorithms. It provides a deeper understanding of exact inference, i.e., methods that solve an inference problem without any approximations during calculations. In this main part of the tutorial, we present algorithms that are on the one hand rooted in VE and on the other hand rooted in KC. Next to VE and KC, we delve into the specifics of

- lifted VE (LVE) [7, 9],
- the junction tree algorithm (JT) [4],
- the lifted junction tree algorithm (LJT) incorporating LVE and JT [1],
- the interface algorithm (IA) for temporal models [5], and
- the lifted dynamic junction tree algorithm, which combines LJT and IA [3] as well as
- weighted first-order model counting and first-order KC [10].

Goal: At the end of the tutorial, each participant should have an understanding of what inference algorithms are able to accomplish and how lifting is able to alleviate inference in a world full of objects and repeated structures.

References

1. Braun, T., Möller, R.: Lifted Junction Tree Algorithm. In: Proceedings of KI 2016: Advances in Artificial Intelligence. pp. 30–42. Springer (2016)
2. Darwiche, A., Marquis, P.: A Knowledge Compilation Map. Journal of Artificial Intelligence Research **17(1)**, 229–264 (2002)
3. Gehrke, M., Braun, T., Möller, R.: Lifted Dynamic Junction Tree Algorithm. In: Proceedings of the International Conference on Conceptual Structures. pp. 55–69. Springer (2018)
4. Lauritzen, S.L., Spiegelhalter, D.J.: Local Computations with Probabilities on Graphical Structures and Their Application to Expert Systems. Journal of the Royal Statistical Society. Series B: Methodological **50**, 157–224 (1988)
5. Murphy, K.P.: Dynamic Bayesian Networks: Representation, Inference and Learning. Ph.D. thesis, University of California, Berkeley (2002)
6. Niepert, M., Van den Broeck, G.: Tractability through Exchangeability: A New Perspective on Efficient Probabilistic Inference. In: AAAI-14 Proceedings of the 28th AAAI Conference on Artificial Intelligence. pp. 2467–2475. AAAI Press (2014)
7. Poole, D.: First-order Probabilistic Inference. In: IJCAI-03 Proceedings of the 18th International Joint Conference on Artificial Intelligence. pp. 985–991. IJCAI Organization (2003)
8. Richardson, M., Domingos, P.: Markov Logic Networks. Machine Learning **62(1-2)**, 107–136 (2006)
9. Taghipour, N., Fierens, D., Davis, J., Blockeel, H.: Lifted Variable Elimination: Decoupling the Operators from the Constraint Language. Journal of Artificial Intelligence Research **47(1)**, 393–439 (2013)
10. Van den Broeck, G., Taghipour, N., Meert, W., Davis, J., De Raedt, L.: Lifted Probabilistic Inference by First-order Knowledge Compilation. In: IJCAI-11 Proceedings of the 22nd International Joint Conference on Artificial Intelligence. pp. 2178–2185. IJCAI Organization (2011)
11. Zhang, N.L., Poole, D.: A Simple Approach to Bayesian Network Computations. In: Proceedings of the 10th Canadian Conference on Artificial Intelligence. pp. 171–178. Springer (1994)