Annals of Mathematics and Artificial Intelligence
https://doi.org/10.1007/s10472-024-09942-w

Knowledge compilation

Preface

Adnan Darwiche[1] · Pierre Marquis[2,3]

Accepted: 19 March 2024

Keywords Knowledge representation · Automated reasoning · Knowledge compilation

Mathematics Subject Classification (2010) 68T27 · 68T20 · 68T30 · 68T35

The area of knowledge compilation has been receiving an increased attention in AI and computer science more broadly. This area has a long tradition in AI and was originally motivated by the desire to compile knowledge into tractable representations to facilitate logical reasoning; see, e.g., [8] and [14, 15, 18, 19, 23, 25, 30]. A turning point in this area has been the work of [18] which presented a comprehensive theory based on compiling knowledge into *tractable Boolean circuits,* which are *deep* representations—in contrast to earlier efforts which focused on *flat* representations such as Conjunctive Normal Form (CNF) and Disjunctive Normal Form (DNF); see, e.g., [23, 25]. Tractable Boolean circuits are basically and-or circuits with inputs representing variables, their negations, or constants. These circuits are also said to be in Negation Normal Form (NNF) since negations are restricted to appear at inputs. NNF circuits are universal representations as they can represent any Boolean function, and they earn their tractability from additional properties they must satisfy. Among the most prominent properties are decomposability which leads to DNNF circuits [9]; determinism which leads to d-DNNF circuits [10], and decision/ordering which lead to OBDD circuits [7]. Further properties have also been identified which are based on either strengthening or weakening earlier properties, leading to other circuit types such as SDD circuits [13].

Beyond expanding the scope of target compilation languages, the work in [18] contributed a principled methodology for studies in knowledge compilation. In particular, it identified

✉ Adnan Darwiche
darwiche@cs.ucla.edu

✉ Pierre Marquis
marquis@cril.univ-artois.fr

1 Computer Science Department, UCLA, Los Angeles CA 90095, USA

2 CRIL, Univ. Artois & CNRS, rue Jean Souvraz, Lens F-62307, France

3 Institut Universitaire de France, Paris, France

two key dimensions for analyzing tractable representations: succinctness and tractability. The first dimension studies different tractable representations based on their relative sizes, and the second dimension compares them based on operations (queries and transformations) that they support in polytime. Queries compute answers to questions posed on the tractable representation, such as whether it implies a particular formula. Transformations manipulate or combine tractable representations while ensuring that the result remains in the same form. The work popularized the term *knowledge compilation map* for studies of this form, which have since then been conducted in broader contexts that transcend Boolean logic.

Even though knowledge compilation originated in the world of logical reasoning, the tractable forms it produced found applications in much broader areas. One early application was to probabilistic reasoning [11, 12], giving birth to an extensive line of work on *tractable arithmetic circuits;* see [16] for a recent survey. Knowledge compilation and tractable circuits have seen a number of additional and exciting developments over the years, which have broadened and extended their applications from reasoning, to learning and more recently to neuro-symbolic AI, by an area that is concerned with integrating neural and symbolic approaches to AI [5, 20, 24]. Other areas that benefited from knowledge compilation include database theory; see, e.g., [2, 22], function synthesis [1, 26], constrained sampling [27], and explainable AI [3, 4, 6, 17, 21, 28, 29]. It is not surprising then to see that the articles in this special issue span a significant spectrum as we discuss next.

This special issue of the Annals of Mathematics and Artificial Intelligence includes five papers concerned with various, yet distinct aspects of knowledge compilation.

Product configuration is a common task in industry used to develop product lines, where each product is composed from a set of reusable features (also called components). Typically, not every combination of selected features gives rise to a functional product. Hence, it is common to use feature models, which specify constraints that eliminate invalid combinations of features—such models can have thousands of features and hundreds of thousands of constraints. Feature-model analysis typically requires many queries over such constraints which makes it an ideal application of knowledge compilation. Practical questions that arise in this context include: how to implement various queries in feature-model analysis using classical queries studied in the knowledge compilation literature, and which tractable forms are most suitable for such an analysis? Such a study is contributed in the paper "On the benefits of knowledge compilation for feature-model analyses" by Chico Sundermann, Elias Kuiter, Tobias Heß, Heiko Raab, Sebastian Krieter, and Thomas Thüm, which also evaluates some existing knowledge compilers for feature-model analysis.

Boolean functional synthesis is concerned with generating specific Boolean functions that meet user-provided specifications. It was recently observed that if such specifications are compiled into a tractable form, then the synthesis problem is facilitated to the point where it can be conducted in polytime. This is covered in the paper "Tractable representations for Boolean functional synthesis" by S. Akshay, Supratik Chakraborty, and Shetal Shah, which provides a perspective on this area and then tackles two key questions: What tractable forms facilitate synthesis, and how can we ensure that a specification is compiled into that form? Both syntactic and semantic conditions are identified for NNF circuits to facilitate synthesis, and knowledge compilation algorithms are proposed for this purpose.

A switch-list representation (SLR) of a Boolean function can be viewed as a compressed version of its truth table, and can be exponentially smaller than the tabular representation. Yet, it is the closest to truth tables as far as its support for polytime queries and transformations compared to all prominent, tractable representations of Boolean functions known today. The SLR was introduced relatively recently and an in-depth study of this representation is contributed in the paper "A faster implementation of EQ and SE queries for switch-list representations" by Ondřej Čepek and James Weigle, with a particular focus on optimized, polytime algorithms for testing whether one SLR entails another and whether they are equivalent. These queries, particularly entailment, are known to be among the more difficult queries for other tractable forms that are commonly used in the literature.

Knowledge compilation can also be useful when reasoning about actions, particularly for belief tracking which is concerned with obtaining information about the current state of the world after a history of actions and observations. For example, knowledge compilation can be used to represent the information in a belief state that must be maintained throughout a history. A study of belief tracking from the perspective of knowledge compilation is contributed in the paper "A knowledge compilation perspective on queries and transformations for belief tracking" by Alexandre Niveau, Hector Palacios, Sergej Scheck, and Bruno Zanuttini. The authors define queries and transformations which pertain to belief tracking, and study them for propositional domains, considering a number of representations for belief states, actions, observations, and goals. Towards the definition of a knowledge compilation map for belief tracking, the authors also discuss the relative succinctness of representations.

Another area that benefited from knowledge compilation is preference reasoning. Given a set of combinatorial objects, also called alternatives, one specifies preferences over these objects which can be conditioned on their properties. Different languages for specifying preferences have been proposed which have different expressiveness and succinctness. These are specialized languages that are not directly related to Boolean logic as normally used in knowledge compilation, but that latter literature provides a methodology for studying preference languages. In the preferences domain, the input knowledge base corresponds to some statements in the used preference language which gives rise to a number of reasoning tasks (i.e., queries). Examples include whether the specified preference relation satisfies some properties (e.g., is antisymmetric), what preference relation is implied on two particular objects, which are the top preferred objects, which objects are equally preferred, and whether two knowledge bases encode the same preferences. One may also be interested in various transformations applied to preference statements, such as combining them conjunctively or disjunctively except that the semantics of these combinations are much more intricate than in the Boolean logic case. Still, one is interested in studying preference languages with respect to their expressiveness and succinctness, and their support for various polytime queries and transformations. Such a study is contributed in the paper "An extended knowledge compilation map for conditional preference statements-based and generalized additive utilities-based languages" by Hélène Fargier, Stefan Mengel, and Jérôme Mengin, which develops a corresponding knowledge compilation map.

Acknowledgements The support of the international research project MAKC "Modern Approaches to Knowledge Compilation" funded by CNRS is acknowledged by the invited editors.

Author Contributions A.D. and P.M. wrote the preface.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing interests The authors declare no competing interests.

References

1. Akshay, S., Arora, J., Chakraborty, S., Krishna, S.N., Raghunathan, D., Shah, S.: Knowledge compilation for Boolean functional synthesis. In: FMCAD. pp. 161–169. IEEE (2019)
2. Amarilli, A.: Provenance in databases and links to knowledge compilation. In: KOCOON workshop on knowledge compilation (2019), http://kocoon.gforge.inria.fr/slides/amarilli.pdf
3. Arenas, M., Bertossi, P.B.L., Monet, M.: The tractability of shap-score-based explanations over deterministic and decomposable Boolean circuits (2021)
4. Audemard, G., Koriche, F., Marquis, P.: On tractable XAI queries based on compiled representations. In: KR. pp. 838–849 (2020)
5. Besold, T.R., d'Avila Garcez, A., Bader, S., Bowman, H., Domingos, P., Hitzler, P., Kuehnberger, K.U., Lamb, L.C., Lowd, D., Lima, P.M.V., de Penning, L., Pinkas, G., Poon, H., Zaverucha, G.: Neural-symbolic learning and reasoning: A survey and interpretation (2017)
6. den Broeck, G.V., Lykov, A., Schleich, M., Suciu, D.: On the tractability of shap explanations (2021)
7. Bryant, R.E.: Graph-based algorithms for Boolean function manipulation. IEEE Transactions on Computers **C-35**, 677–691 (1986)
8. Cadoli, M., Donini, F.M.: A survey on knowledge compilation. AI Commun. **10**(3–4), 137–150 (1997)
9. Darwiche, A.: Decomposable negation normal form. J. ACM **48**(4), 608–647 (2001)
10. Darwiche, A.: On the tractable counting of theory models and its application to truth maintenance and belief revision. J. Appl. Non Class. Logics **11**(1–2), 11–34 (2001)
11. Darwiche, A.: A logical approach to factoring belief networks. In: Fensel, D., Giunchiglia, F., McGuinness, D.L., Williams, M. (eds.) Proceedings of the Eights International Conference on Principles and Knowledge Representation and Reasoning (KR-02), Toulouse, France, April 22-25, 2002. pp. 409–420. Morgan Kaufmann (2002)
12. Darwiche, A.: A differential approach to inference in Bayesian networks. J. ACM **50**(3), 280–305 (2003)
13. Darwiche, A.: SDD: A new canonical representation of propositional knowledge bases. In: IJCAI. pp. 819–826. IJCAI/AAAI (2011)
14. Darwiche, A.: Tractable knowledge representation formalisms. In: Tractability: Practical Approaches to Hard Problems, pp. 141–172. Cambridge University Press (2014)
15. Darwiche, A.: Three modern roles for logic in AI. In: PODS. pp. 229–243. ACM (2020)
16. Darwiche, A.: Tractable boolean and arithmetic circuits. In: Neuro-Symbolic Artificial Intelligence, Frontiers in Artificial Intelligence and Applications, vol. 342, pp. 146–172. IOS Press (2021)
17. Darwiche, A., Hirth, A.: On the reasons behind decisions. In: ECAI. Frontiers in Artificial Intelligence and Applications, vol. 325, pp. 712–720. IOS Press (2020)
18. Darwiche, A., Marquis, P.: A knowledge compilation map. J. Artif. Intell. Res. **17**, 229–264 (2002)
19. Darwiche, A., Marquis, P., Suciu, D., Szeider, S.: Recent trends in knowledge compilation (dagstuhl seminar 17381). Dagstuhl Reports **7**(9), 62–85 (2017)
20. Garcez, A., Lamb, L.: Neurosymbolic ai: The 3rd wave. ArXiv arXiv:2012.05876 (2020)
21. Huang, X., Izza, Y., Ignatiev, A., Cooper, M.C., Asher, N., Marques-Silva, J.: Efficient explanations for knowledge compilation languages (2021)
22. Jha, A.K., Suciu, D.: Knowledge compilation meets database theory: Compiling queries to decision diagrams. Theory Comput. Syst. **52**(3), 403–440 (2013)
23. Marquis, P.: Knowledge compilation using theory prime implicates. In: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI). pp. 837–845 (1995)
24. Raedt, L.D., Dumancic, S., Manhaeve, R., Marra, G.: From statistical relational to neuro-symbolic artificial intelligence. In: IJCAI. pp. 4943–4950. ijcai.org (2020)
25. Selman, B., Kautz, H.A.: Knowledge compilation and theory approximation. J. ACM **43**(2), 193–224 (1996)
26. Shah, P., Bansal, A., Akshay, S., Chakraborty, S.: A normal form characterization for efficient Boolean skolem function synthesis. In: LICS. pp. 1–13. IEEE (2021)
27. Sharma, S., Gupta, R., Roy, S., Meel, K.S.: Knowledge compilation meets uniform sampling. In: LPAR. EPiC Series in Computing, vol. 57, pp. 620–636. EasyChair (2018)

28. Shi, W., Shih, A., Darwiche, A., Choi, A.: On tractable representations of binary neural networks. In: KR. pp. 882–892 (2020)
29. Shih, A., Choi, A., Darwiche, A.: A symbolic approach to explaining Bayesian network classifiers. In: IJCAI. pp. 5103–5111. ijcai.org (2018)
30. Slivovsky, F.: An introduction to knowledge compilation. In: KOCOON workshop on knowledge compilation (2019), http://kocoon.gforge.inria.fr/slides/slivovsky.pdf