Parameteric Bayesian Networks: Finding the Right Probabilities.

Abstract for SynCoP 2023

Bahare Salmani and Joost-Pieter Katoen

February 2023

Bayesian networks (BNs) [1] are probabilistic graphical models that are widely used to reason under uncertainty. The key is that they reflect (in)dependencies, e.g., the causality information between the entities. This enables (i) introducing the knowledge from the application domain into the model and (ii) compactly representing the joint probability distributions over the set of random variables. The (in)dependencies between the variables are stored as a directed acyclic graph. The probabilities are stored in conditional probability tables (CPTs): each CPT row denotes a local probability distribution. In the normal setting, the CPT entries are constant probabilities. However, *not all the local probability distributions are always known a priori.*

Parametric Bayesian networks. We consider parametric Bayesian networks (pBNs) that have gained attention in many studies [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]: the CPT entries are no longer constant probabilities, yet multivariate polynomials with rational coefficients over a set of unknown parameters e.g., p and q. We consider two key tasks on pBNs: *sensitivity analysis* and *parameter tuning*. Sensitivity analysis starts with *computing sensitivity functions*: rational functions that are the solution to a given pBN query (e.g., a conditional probability) in terms of the network parameters. Parameter tuning takes a further step and aims at synthesizing the unknown parameters with respect to a given constraint. Parameters instantiations are joint values for the unknown parameters, e.g., $p = 0.3$ and $q = 0.55$. We address two synthesis problems. Let the hypothesis H, the evidence E, and the threshold λ be given. Is there an instantiation for the parameters that make the constraint $\phi : \Pr(H|E) \leq \lambda$ (or $\geq \lambda$) hold? The problem is referred to as *feasibility checking*. What is (a subset of) all instantiations that satisfy (or reject) the constraint ϕ? The latter problem is related to *parameter space partitioning*.

Motivation. Parameter tuning for pBNs is in general computationally hard [16]. The existing tools and techniques in the literature only address limited subclasses of pBNs: (i) the pBN only contains a single (or a couple of) unknown parameters at a time [11, 13, 17], (ii) the parameters only occur in a single CPT [12] or in the same clique of the junction tree [10], and (iii) a parameter always only occurs in a single probability distribution, i.e., no parameter dependencies between multiple distributions are allowed. Such restrictions yield sensitivity functions that are linear in each parameter and make the computation less expensive, yet such over-simplification is often problematical, e.g., for a pBN with parameters all in a single CPT, the constraint of interest may be unsatisfiable, but allowing the parameters in several CPTs yields satisfying instantiations for the same constraint. Moreover, for BNs often we are interested in the problem *minimal distance parameter tuning*: Let B be a BN that is parameterized to the pBN $\mathcal{B}$. For a given distance measure function, what are the instantiations that satisfy the constraint ϕ while minimizing the distance from the original values of the parameters in the original BN B? A pBN with parameters in multiple CPTs often yields a smaller overall distance compared to when the parameters are only in a single CPT [12].

Main objective. *We aimed to ease the existing restrictions in the literature and extend the spectrum of sensitivity analysis and parameter tuning on pBNs.*

Parameter synthesis for Markov chains. Synthesizing the parameters for Markov models has been significantly developed over the last two decades [18, 19, 20, 21, 22, 23, 24, 25]. The techniques, e.g., range over the gradient-based methods [26], convex optimization [23, 27], and region verification techniques. The parameter lifting algorithm (PLA) [22] e.g., offers a simple technique to verify the *regions* (e.g., $0.2 \leq p \leq 0.4$ and $0.35 \leq q \leq 0.6$) with respect to the given constraint ϕ and also allows handling models with parameter dependencies. We built upon those existing techniques from the probabilistic model checking (PMC) [28, 29, 30] community to analyze pBNs.

Method. We proposed a translation from (p)BNs to (p)MCs that relates (p)BN queries to reachability probabilities in (p)MCs. We proved the correctness of our translation and implemented it as a prototypical tool on top of the probabilistic model checker, *Storm* [31]. This enables (a) computing pBN sensitivity functions by the existing algorithms in Storm such as state elimination, (b) tuning the parameters using the existing feasibility checking algorithms such as Quadratically-Constrained Quadratic Program (QCQP) and Gradient-Descent (GD), (c) partitioning the parameter space to satisfying and rejecting subregions using parameter lifting algorithm (PLA). We performed an extensive empirical study to evaluate the effectiveness and scalability of our techniques and compared them to the pBN baseline tools: *Bayes server* and *SamIam*.

Research questions. We focused our study on the following research questions.

- Can we perform sensitivity analysis on parametric BNs fast and with arbitrary parameter dependencies?
- What are the decisive factors in the computation time of pBN sensitivity functions using PMC techniques?
- How effective are the parametric MC feasibility checking techniques for analyzing parametric BNs?
- How do the number of parameters influence the feasibility analysis time?
- What is the scalability of parameter space partitioning on pBNs?
- How does the *coverage factor* affect the partitioning time?

The main findings. In a nutshell, our main findings are as follows.

- Inference queries on (p)BNs correspond to reachability probabilities in (p)MCs.
- Temporal logic yields a flexible technique for formalizing inference queries.
- Current techniques for computing solution functions in PMC scale much better than existing techniques for parametric BNs: pBNs with hundreds of parameters can be analyzed in a reasonable time.
- Parameter synthesis technique from PMC can deal with multiple parameters within a BN, possibly occurring in different CPTs.
- Gradient descent is the favorable technique for finding a suitable parameter instantiation and can deal with parametric BNs with up to hundreds of parameters.
- Parameter space partitioning is very effective for parameter tuning for parametric BNs with up to ten parameters.

Our results extend the conference paper [32]. The extended paper including the new results is under publication in the Journal of Artificial Intelligence Research (JAIR).

References

[1] Adnan Darwiche. *Modeling and Reasoning with Bayesian Networks.* Cambridge University Press, 2009.

[2] Veerle MH Coupé and Linda C van der Gaag. *Practicable sensitivity analysis of Bayesian belief networks*, volume 1998. Utrecht University: Information and Computing Sciences, 1998.

[3] Veerle MH Coupé, Linda C van der Gaag, and J Dik F Habbema. Sensitivity analysis: an aid for belief-network quantification. *The Knowledge Engineering Review*, 15(3):215–232, 2000.

[4] Marek J. Druzdzel and Linda C. van der Gaag. Building probabilistic networks:" where do the numbers come from?". *IEEE Trans. Knowl. Data Eng.*, 12(4):481–486, 2000.

[5] Finn Verner Jensen. Gradient descent training of Bayesian networks. In *ESCQARU*, volume 1638 of *Lecture Notes in Computer Science*, pages 190–200. Springer, 1999.

[6] Kathryn Blackmond Laskey. Sensitivity analysis for probability assessments in Bayesian networks. *IEEE Trans. Syst. Man Cybern.*, 25(6):901–909, 1995.

[7] Enrique F. Castillo, José Manuel Gutiérrez, and Ali S. Hadi. Parametric structure of probabilities in Bayesian networks. In *ECSQARU*, volume 946 of *Lecture Notes in Computer Science*, pages 89–98. Springer, 1995.

[8] Enrique F. Castillo, José Manuel Gutiérrez, and Ali S. Hadi. Goal oriented symbolic propagation in Bayesian networks. In *AAAI/IAAI, Vol. 2*, pages 1263–1268. AAAI Press / The MIT Press, 1996.

[9] Enrique F. Castillo, José Manuel Gutiérrez, and Ali S. Hadi. Sensitivity analysis in discrete Bayesian networks. *IEEE Trans. Syst. Man Cybern. Part A*, 27(4):412–423, 1997.

[10] Uffe Kjærulff and Linda C. van der Gaag. Making sensitivity analysis computationally efficient. In *UAI*, pages 317–325. Morgan Kaufmann, 2000.

[11] Hei Chan and Adnan Darwiche. When do numbers really matter? *J. Artif. Intell. Res.*, 17:265–287, 2002.

[12] Hei Chan and Adnan Darwiche. Sensitivity analysis in Bayesian networks: From single to multiple parameters. In *UAI*, pages 67–75. AUAI Press, 2004.

[13] Veerle M. H. Coupé and Linda C. van der Gaag. Properties of sensitivity analysis of Bayesian belief networks. *Ann. Math. Artif. Intell.*, 36(4):323–356, 2002.

[14] Silja Renooij. Co-variation for sensitivity analysis in Bayesian networks: Properties, consequences and alternatives. *Int. J. Approx. Reason.*, 55(4):1022–1042, 2014.

[15] Janneke H. Bolt and Linda C. van der Gaag. Balanced tuning of multi-dimensional Bayesian network classifiers. In *ECSQARU*, volume 9161 of *Lecture Notes in Computer Science*, pages 210–220. Springer, 2015.

[16] Johan Kwisthout and Linda C. van der Gaag. The computational complexity of sensitivity analysis and parameter tuning. In *UAI*, pages 349–356. AUAI Press, 2008.

[17] Linda C van der Gaag, Silja Renooij, and Veerle MH Coupé. Sensitivity analysis of probabilistic networks. In *Advances in Probabilistic Graphical Models*, pages 103–124. Springer, 2007.

[18] Conrado Daws. Symbolic and parametric model checking of discrete-time Markov chains. In *ICTAC*, volume 3407 of *Lecture Notes in Computer Science*, pages 280–294. Springer, 2004.

[19] Ruggero Lanotte, Andrea Maggiolo-Schettini, and Angelo Troina. Parametric probabilistic transition systems for system design and analysis. *Formal Aspects Comput.*, 19(1):93–109, 2007.

[20] Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. Probabilistic reachability for parametric Markov models. *Int. J. Softw. Tools Technol. Transf.*, 13(1):3–19, 2011.

[21] Christian Dehnert, Sebastian Junges, Nils Jansen, Florian Corzilius, Matthias Volk, Harold Bruintjes, Joost-Pieter Katoen, and Erika Ábrahám. Prophesy: A probabilistic parameter synthesis tool. In *CAV (1)*, volume 9206 of *Lecture Notes in Computer Science*, pages 214–231. Springer, 2015.

[22] Tim Quatmann, Christian Dehnert, Nils Jansen, Sebastian Junges, and Joost-Pieter Katoen. Parameter synthesis for Markov models: Faster than ever. In *ATVA*, volume 9938 of *Lecture Notes in Computer Science*, pages 50–67, 2016.

[23] Murat Cubuktepe, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, and Ufuk Topcu. Synthesis in pMDPs: A tale of 1001 parameters. In *ATVA*, volume 11138 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 2018.

[24] Paul Gainer, Ernst Moritz Hahn, and Sven Schewe. Accelerated model checking of parametric Markov chains. In *ATVA*, volume 11138 of *Lecture Notes in Computer Science*, pages 300–316. Springer, 2018.

[25] Xinwei Fang, Radu Calinescu, Simos Gerasimou, and Faisal Alhwikem. Fast parametric model checking through model fragmentation. In *ICSE*, pages 835–846. IEEE, 2021.

[26] Linus Heck, Jip Spel, Sebastian Junges, Joshua Moerman, and Joost-Pieter Katoen. Gradient-descent for randomized controllers under partial observability. In *VMCAI*, volume 13182 of *Lecture Notes in Computer Science*, pages 127–150. Springer, 2022.

[27] Murat Cubuktepe, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, and Ufuk Topcu. Convex optimization for parameter synthesis in mdps. *IEEE Trans. Autom. Control.*, 67(12):6333–6348, 2022.

[28] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking.* MIT Press, 2008.

[29] Joost-Pieter Katoen. The probabilistic model checking landscape. In *LICS*, pages 31–45. ACM, 2016.

[30] Christel Baier, Luca de Alfaro, Vojtech Forejt, and Marta Kwiatkowska. Model checking probabilistic systems. In *Handbook of Model Checking*, pages 963–999. Springer, 2018.

[31] Christian Hensel, Sebastian Junges, Joost-Pieter Katoen, Tim Quatmann, and Matthias Volk. The probabilistic model checker storm. *Int. J. Softw. Tools Technol. Transf.*, 24(4):589–610, 2022.

[32] Bahare Salmani and Joost-Pieter Katoen. Fine-tuning the odds in Bayesian networks. In *ECSQARU*, volume 12897 of *Lecture Notes in Computer Science*, pages 268–283. Springer, 2021.