

A General Framework for Anytime Approximation in Probabilistic Databases

Maarten Van den Heuvel and Floris Geerts

University of Antwerp, Belgium

Wolfgang Gatterbauer

Northeastern University, Boston, MA, USA

Martin Theobald

University of Luxembourg, Luxembourg

Abstract

Anytime approximation algorithms that compute the probabilities of queries over probabilistic databases can be of great use to statistical learning tasks. Those approaches have been based so far on either (i) sampling or (ii) branch-and-bound with model-based bounds. We present here a more general branch-and-bound framework that extends the possible bounds by using “dissociation,” which yields tighter bounds.

Introduction

Since calculating the exact probability of a query over probabilistic databases (PDBs) is #P-hard in general, probabilistic inference is a major bottleneck in several related applications, such as statistical relational learning (Raedt and Kersting 2017), and fast approximation methods are needed. *Anytime approximation methods* give approximate answers fast, yet allow the user to refine the answer by using additional time and resources. The current state of the art in anytime approximation for PDBs are either based on (i) sampling (Ré, Dalvi, and Suciu 2007) or on (ii) the model-based branch-and-bound approach by (Fink, Huang, and Olteanu 2013) implemented in the SPROUT system, where the latter outperforms the former. See (Van den Broeck and Suciu 2017) for a recent survey.

We propose here to improve the branch-and-bound approach by replacing model-based bounds with novel *dissociation-based bounds*, which were shown to dominate model-based bounds (Gatterbauer and Suciu 2014). The technique of dissociation is related to the variable splitting framework by (Choi and Darwiche 2010), and has so far only been applied at the first-order (query) level for upper bounds (Gatterbauer and Suciu 2017). One reason is that for conjunctive queries, these dissociation-based upper bounds are uniquely defined and proven to be better than any model-based bounds. In contrast, a whole spectrum of optimal oblivious lower bounds exists, which includes the model-based bounds as special cases. Important ingredients of our approach are strategies for quickly choosing good lower bounds, as well as a novel heuristic for the “branch” part of the algorithm based on *influence* or sensitivity (Kanagal, Li, and Deshpande 2011).

Background

The problem. We consider the evaluation of Boolean conjunctive queries (CQs) and illustrate our problem and approach with the following query Q over PDB D :

$$Q :- R(X), S(X, Y), T(Y)$$

	R	X	p	S	X	Y	p	T	Y	p
D	r_1	a	0.5	s_1	a	c	0.3	t_1	c	0.4
	r_2	b	0.6	s_2	a	d	0.4	t_2	d	0.8
				s_3	b	d	0.5			

Grounding Q over D leads to a propositional formula φ that is called the “lineage of Q over D ” in which each variable represents a tuple in the database. In our example, $\varphi = r_1(s_1t_1 \vee s_2t_2) \vee r_2s_3t_2$. Computing its probability $\mathbb{P}(\varphi)$ is equivalent to computing the probability $\mathbb{P}(Q)$, i.e. the probability that the query Q is true over PDB D . In general, calculating this probability is #P-hard in the number of variables in the lineage, and thus in the number of tuples in the database. We are interested in developing an approximation scheme that allows us to trade-off available time with required accuracy of approximation.

State of the art. The algorithm underlying SPROUT (Fink, Huang, and Olteanu 2013) approximates $\mathbb{P}(\varphi)$ by using lineage decompositions based on independence and determinism. When the smaller lineages obtained are read-once (i.e. have no multiple occurrences of the same variable) their probabilities are computed exactly in PTIME, otherwise they are approximated by *model-based bounds*: for each variable, all except one of its occurrences are set to true or 1 (resp. false or 0) while the remaining occurrence is assigned the original probability to get an upper (resp. lower) bound. SPROUT randomly selects these occurrences. The bounds are propagated back up, based on the decomposition, to obtain an approximation of $\mathbb{P}(\varphi)$. If this approximation is not accurate enough, Shannon expansion (SE) is applied on the decomposed lineages: a variable x is selected and, based on $\varphi \equiv (x \wedge \varphi|_{x=1}) \vee (\neg x \wedge \varphi|_{x=0})$, the decomposition and approximation process continues on $\varphi|_{x=1}$ and $\varphi|_{x=0}$ until the desired accuracy is obtained. SPROUT selects the most frequent variable (the one with the highest number of occurrences, ties broken randomly) for SE. Clearly, the challenge is to limit the number of SEs, as these can double the size of the formulas, resulting in higher computation cost.

Our approach

The branch-and-bound approach just described is generic: any technique for approximating the probability of lineages can be plugged in during decomposition, and similarly, any variable selection procedure for SE can be used.

We propose to use *dissociation* to obtain approximations. Dissociation replaces multiple occurrences of a variable with independent copies. Our example lineage φ becomes $\varphi' = r_1(s_1t_1 \vee s_2t_2) \vee r_2s_3t_2''$. By carefully assigning new probabilities p_U and p_L to both copies t_2' and t_2'' of t_2 , $\mathbb{P}_{p_U}(\varphi')$ is guaranteed to be an upper, and $\mathbb{P}_{p_L}(\varphi')$ to be a lower bound for $\mathbb{P}(\varphi)$. What makes dissociations particularly attractive for our framework is that dissociation bounds generalize model-based bounds, and that they can be calculated efficiently. As a consequence, we obtain better approximations in each step, hereby possibly reducing the number of SEs needed to obtain a desired accuracy. For SE, we propose a method based on influence. We next detail some aspects of our approach and highlight some of its advantages.

1. Better bounds. Following Theorem 4.8 in (Gatterbauer and Suciu 2014), the two green shaded areas in Fig. 1a show all possible assignments of probabilities to the occurrences t_2' and t_2'' in our lineage φ' that guarantee to upper or lower bound the true probability $\mathbb{P}(\varphi) = 0.384$. For upper bounds, the two model-based bounds provide an approximation of 0.419 and 0.441. By contrast, assigning the original probability $p(t_2) = 0.8$ to both t_2' and t_2'' results in the unique and optimal upper dissociation bound of 0.393.

For the lower bounds, the shaded area on the lower left shows all possible lower bound assignments, and all assignments on the curved border are possible optima. Here, the two model-based bounds are just a few of the possible options, achieving approximations of 0.201 and 0.286 respectively. The *symmetric* lower dissociation bound assigns both occurrences an equal share by setting $p_L(t_2) = p_L(t_2'') = 1 - \sqrt[2]{1 - p(t_2)}$ which gives 0.297 as lower bound. The best assignment lies slightly to the left and gives 0.302.

Part of our approach are *gradient-descent methods* that aim to find this optimal assignment. Computation of the gradient is closely related to the notion of *influence*, which can be computed in PTIME for read-once formulas (Kanagal, Li, and Deshpande 2011), as well as for dissociated formulas. While gradient-descent methods generally take longer to find the best lower bound than randomly assigning model-based bounds, the quality of the bound is often much better, thus justifying these additional computations.

2. Better variable selection. The state-of-the-art heuristic for choosing a variable for SE in SPROUT is selecting the *most frequent* variable. We instead select the variable for SE that has the highest *sum of influences* of each of its occurrences, as if they were independent. This ensures PTIME computation and turns out to be a better choice. Moreover, we can re-use the computed influences from our gradient-based optimization methods, thus avoiding re-computation.

3. Optimization trade-off. We have a “knob” to control how much time to spend on finding a good lower bound using gradient descent, before moving on to a next Shannon expansion. Recall that model-based bounds are fast to




Figure 1: a) Green shaded areas show all possible assignments for upper and lower dissociation bounds for our running example. Green values show resulting approximations. b) Our new methods (PBD and HB) achieve better accuracy-time tradeoffs than prior methods (MB and SD) on the Yago3 dataset.

compute; they are randomly selected, in one step. In contrast, descent methods perform multiple steps, and this extra computation pays off only when considerably better lower bounds are obtained, and further SEs can be avoided.

Experiments

We experimented with several approximation methods on the Yago3 dataset (Mahdisoltani, Biega, and Suchanek 2015). We obtained 380 lineages from 4 different queries by assigning different input probabilities to the data, using different join-orders to factorize the lineages, and by injecting different constants into the queries. Figure 1b shows the average approximation error over time for four different instantiations of our approach.

MB is the existing Model-based approach from SPROUT and performs slightly worse than SD, which uses Symmetric Dissociations for both upper and lower bounds. Both methods work best with the frequency heuristic for SE. Our new methods PGD (Projected Gradient Descent) and HB (Hybrid method) vastly outperform both approaches. Both methods use the optimal symmetric upper bounds and apply 10 gradient descent steps before using SE. But whereas PGD searches for the true best lower bounds, HB searches for the best possible model-based lower bound, by moving in a gradient direction. This makes the optimization faster, but produces slightly worse bounds than PGD. These methods work best with the influence heuristic for SE, but outperform the others even with the frequency heuristic.

Conclusion

We introduced an anytime approximation framework for probabilistic query evaluation. Our framework leverages novel dissociation bounds that generalize and improve upon model-based bounds. Our experimental results show notable improvements over the current state of the art, and we believe that the approach also has applications beyond PDBs in the broader area of statistical-relational learning (SRL).

Acknowledgements.

This work has been supported in part by NSF IIS-1762268 and FWO G042815N.

References

- Choi, A., and Darwiche, A. 2010. Relax, compensate and then recover. In *JSAI International Symposium on Artificial Intelligence*, volume 6797 of *Lecture Notes in Computer Science*, 167–180. Springer.
- Fink, R.; Huang, J.; and Olteanu, D. 2013. Anytime approximation in probabilistic databases. *VLDB J.* 22(6):823–848.
- Gatterbauer, W., and Suciu, D. 2014. Oblivious bounds on the probability of Boolean functions. *TODS* 39(1):1–34.
- Gatterbauer, W., and Suciu, D. 2017. Dissociation and propagation for approximate lifted inference with standard relational database management systems. *VLDB J.* 26(1):5–30.
- Kanagal, B.; Li, J.; and Deshpande, A. 2011. Sensitivity analysis and explanations for robust query evaluation in probabilistic databases. In *SIGMOD*, 841–852.
- Mahdisoltani, F.; Biega, J.; and Suchanek, F. M. 2015. YAGO3: A knowledge base from multilingual wikipedias. In *CIDR*.
- Raedt, L. D., and Kersting, K. 2017. Statistical relational learning. In *Encyclopedia of Machine Learning and Data Mining*. Springer. 1177–1187.
- Ré, C.; Dalvi, N.; and Suciu, D. 2007. Efficient top-k query evaluation on probabilistic data. In *ICDE*, 886–895.
- Van den Broeck, G., and Suciu, D. 2017. Query processing on probabilistic data: A survey. *Foundations and Trends in Databases* 7(3-4):197–341.