

Predictive Uncertainty Estimation for Tractable Deep Probabilistic Models

Julissa Villanueva Llerena

Institute of Mathematics and Statistics University of São Paulo, Brazil
jgville@ime.usp.br

Abstract

Tractable Deep Probabilistic Models (TPMs) are generative models based on arithmetic circuits that allow for exact marginal inference in linear time. These models have obtained promising results in several machine learning tasks. Like many other models, TPMs can produce over-confident incorrect inferences, especially on regions with small statistical support. In this work, we will develop efficient estimators of the predictive uncertainty that are robust to data scarcity and outliers. We investigate two approaches. The first approach measures the variability of the output to perturbations of the model weights. The second approach captures the variability of the prediction to changes in the model architecture. We will evaluate the approaches on challenging tasks such as image completion and multilabel classification.

1 Introduction

Despite the enormous success of deep models in several machine learning tasks, there exists a concern about their robustness, reliability and interpretability [Caruana *et al.*, 2015]. Several efforts have been made to deal with these issues, such as estimating the uncertainty about predictions to measure how much the model “knows what it knows” [Lakshminarayanan *et al.*, 2017]. A growing consensus in that sense holds that the model “may simply not be the best judge of its trustworthiness” [Jiang *et al.*, 2018].

Sum-Product Networks (SPNs) are a popular type of TPMs [Poon and Domingos, 2011] that can be viewed as a specific type of compositional deep neural networks. An SPN is an acyclic computation graph with an input layer that computes univariate probability distributions and internal neurons computing either weighted sums or products. It is important to note that each neuron of an SPN represents a joint probability distribution over the variables connected directly or indirectly to it. Selective SPNs are a subclass of SPNs that additionally allows for efficient Maximum-A-Posteriori (MAP) inference and parameter learning [Peharz *et al.*, 2014]. Probabilistic Sentential Decision Diagrams (PSDDs) further constrain selective SPNs to allow for the representation of logical constraints on the domain [Kisa *et al.*, 2014].

Despite the vantages of the probabilistic semantic of SPNs against general deep models, such the efficient architecture learning from data, SPNs can also generalize poorly on regions with insufficient statistical support or conflicting data, leading to unreliable, overconfident and prior-dependent inferences. It is thus relevant to ascertain the predictive uncertainty of inferences, that is, to provide a confidence measure for each prediction.

Predictive uncertainty arises from two different sources: *aleatory uncertainty* and *model uncertainty*. Aleatory uncertainty is the irreducible uncertainty which derives from the natural complexity of the data. Model uncertainty is the (ideally) reducible uncertainty in selecting a model (parameters and structure) that reasonably explain the data [Gal, 2016].

This project aims to investigate new approaches to estimate predictive uncertainty for TPMs. To build scalable estimators, we will focus on efficient estimators that can be obtained in time proportional to the computation of a prediction (often linear in the size of the model), excluding many of the current Bayesian inference approaches, making the problem non-trivial.

2 Research Proposal

In order to increase the robustness of TPMs, we propose two approaches to obtain robust predictive uncertainty estimators. The first approach estimates the uncertainty introduced by the misfit of model parameters by the sensitivity of the predictions to perturbations in the parameters [Mauá *et al.*, 2018; Antonucci *et al.*, 2019]. Initially, we consider measuring the sensitivity while varying all parameters independently. Since many of the parameters are usually correlated (e.g. they are estimated using overlapping portions of the dataset), we will also consider connecting the perturbations of correlated weights. As this is a more challenging task than the previous one it may involve using approximate solutions to ensure efficiency.

The second approach estimates the uncertainty introduced by a mis-specification of the architecture (or structure). This consists in estimating the variability of the predictions as the network structure changes by using an ensemble of SPNs (not limited to the highest score models). One approach to obtaining a diverse set of architectures is by sampling structures. For that purpose, we can sample uniformly or sample

the high-score structures while ensuring diversity. The algorithms developed will be evaluated by their performance on real-world tasks such as image completion, missing data imputation and multilabel classification.

3 Contribution

As our first contribution, we devised a polynomial-time algorithm to decide whether a given MAP configuration is *robust* with respect to changes in multiples parameters of a Selective SPN. We evaluated our algorithms in data imputation and multilabel classification tasks. The results showed that our approach can discriminate easy- and hard-to-classify instances, often more accurately than criteria based on the probabilities induced by the model [Villanueva and Mauá, 2019]. We are currently carrying out more extensive empirical evaluation, and adapting some of the algorithms to the formalism of PSDDs.

4 Related Work

There have been several proposals for performing Bayesian inference with SPNs, although not with the aim of estimating predictive uncertainty. Computing the posterior distribution is intractable for SPNs [Rashwan *et al.*, 2016], and one usually resorts to approximate techniques such as moment matching [Rashwan *et al.*, 2016] and variational inference [Zhao *et al.*, 2016] and Gibbs sampling [Vergari *et al.*, 2019]. Even though they are approximate, prior dependent and very inefficient compared to the computational cost of a typical SPN inference.

Recently, there has been many proposals for achieving predictive uncertainty estimation for general neural networks. Drop-out is a popular approach, that efficiently approximates Bayesian inference [Gal and Ghahramani, 2016]; note that drop-out violates the structural assumptions of TPMs. Another development consists in approximating Bayesian inference using an ensemble of neural networks [Lakshminarayanan *et al.*, 2017; Hu *et al.*, 2020], similar to what we propose for TPMs.

There is recent work on uniform sampling of PSDD structures [Mattei *et al.*, 2019]; this can be used for generating a diverse ensemble of selective SPNs that can be used to estimate predictive uncertainty.

Acknowledgments

This study was financed by the Brazilian Agency *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES) - Finance Code 001 and CNPq grants no. 303920/2016-5 and 420669/2016-7.

References

[Antonucci *et al.*, 2019] Alessandro Antonucci, Alessandro Facchini, and Lilith Mattei. Credal sentential decision diagrams. In *ISIPTA*, 2019.

[Caruana *et al.*, 2015] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *ACM KDD*, 2015.

[Gal and Ghahramani, 2016] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.

[Gal, 2016] Yarin Gal. *Uncertainty in deep learning*. PhD thesis, University of Cambridge, 2016.

[Hu *et al.*, 2020] Shi Hu, Nicola Pezzotti, Dimitrios Mavroedidis, and Max Welling. Simple and accurate uncertainty quantification from bias-variance decomposition. *CoRR*, abs/2002.05582, 2020.

[Jiang *et al.*, 2018] Heinrich Jiang, Been Kim, Melody Guan, and Maya Gupta. To trust or not to trust a classifier. In *NeurIPS*, 2018.

[Kisa *et al.*, 2014] Doga Kisa, Guy Van den Broeck, Arthur Choi, and Adnan Darwiche. Probabilistic sentential decision diagrams. In *KR*, 2014.

[Lakshminarayanan *et al.*, 2017] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NeurIPS*, 2017.

[Mattei *et al.*, 2019] Lilith Mattei, Decio Soares, Alessandro Antonucci, Denis Mauá, and Alessandro Facchini. Exploring the space of probabilistic sentential decision diagrams. In *Workshop of TPM*, 2019.

[Mauá *et al.*, 2018] Denis Deratani Mauá, Diarmaid Conaty, Fabio Gagliardi Cozman, Katja Poppenhaeager, and Cassio Polpo de Campos. Robustifying sum-product networks. *International Journal of Approximate Reasoning*, 2018.

[Peharz *et al.*, 2014] Robert Peharz, Robert Gens, and Pedro Domingos. Learning selective sum-product networks. In *Workshop of TPM*, 2014.

[Poon and Domingos, 2011] Hoifung Poon and Pedro Domingos. Sum-product networks: A new deep architecture. In *UAI*, 2011.

[Rashwan *et al.*, 2016] Abdullah Rashwan, Han Zhao, and Pascal Poupart. Online and distributed bayesian moment matching for parameter learning in sum-product networks. In *Artificial Intelligence and Statistics*, 2016.

[Vergari *et al.*, 2019] Antonio Vergari, Alejandro Molina, Robert Peharz, Zoubin Ghahramani, Kristian Kersting, and Isabel Valera. Automatic bayesian density analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.

[Villanueva and Mauá, 2019] Julissa Villanueva and Denis Mauá. Robust analysis of map inference in selective sum-product networks. In *ISIPTA*, 2019.

[Zhao *et al.*, 2016] Han Zhao, Tameem Adel, Geoff Gordon, and Brandon Amos. Collapsed variational inference for sum-product networks. In *ICML*, 2016.