Learning Tractable Probabilistic Models from Inconsistent Local Estimates

Anonymous Author(s)
Affiliation
Address
email

Abstract

Tractable probabilistic models or probabilistic circuits which admit exact linear time posterior marginal inference are often preferred in practice over intractable models such as Bayesian and Markov networks. This is because although tractable models are slightly inferior to the intractable ones in terms of goodness-of-fit measures such as log-likelihood, they do not use approximate inference at prediction time and as a result exhibit superior predictive performance. In this paper, we consider the problem of improving a tractable model using a large number of local probability estimates, each defined over a small subset of variables (given observations) that are either available from experts or via an external process. Given a model learned from fully-observed data, the key idea in our approach is to update the parameters of the model via a gradient descent procedure that seeks to minimize a convex combination of two quantities: one that enforces closeness via KL divergence to the local estimates and another that enforces closeness to the given model. We show that although the gradients are NP-hard to compute on arbitrary graphical models, they can be efficiently computed over tractable models. We show via experiments that our approach yields tractable models that are significantly superior to the ones learned from data alone even when the local estimates have large error as well as robust to perturbations in the data model.

1 Introduction

A major issue with probabilistic graphical models (PGMs) [4, 9] such as Bayesian and Markov networks is that in most real-world domains, exact inference is computationally infeasible. As a result, approximate inference algorithms are often used in practice, but these algorithms can be inaccurate and may exhibit high variability. Thus, even if a PGM performs well on *goodness-of-fit* measures such as test set log-likelihood, its predictive performance, because of the use of approximate inference methods, can be quite poor.

To combat this problem, the field of tractable probabilistic models (TPMs) or probabilistic circuits (PCs) [10, 12, 22] seeks to learn probabilistic models that admit polynomial time exact inference algorithms. TPMs are more biased or less expressive than general probabilistic models (because they enforce tractability) and as a result may yield a slightly inferior (model) fit. But, the hope is that the slight drop in goodness-of-fit measures is offset by superior predictive performance because no approximations are used at prediction time. This hope is often realized in practice, and demonstrated in numerous empirical studies in literature (cf. [19, 22]).

In this paper, we focus on the following learning problem over TPMs. We assume that the learner has access to (1) a TPM learned from a small dataset having few or no missing variables and (2) a set containing *local estimates* where each estimate is a marginal probability distribution over a small subset of variables given observations. This learning problem is quite common in real-world

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

applications where the aforementioned local estimates are derived from some combination of noisy private data (e.g., in social networks [2, 7]), domain expertise and knowledge (e.g., in active learning [23]) and independent predictive models (e.g., in conditional TPMs [19, 24]).[1]

Contributions. We propose to express the above learning task as the following minimization problem: update the parameters of the given tractable model such that a linear combination of the following two KL distances is minimized: (1) the distance between the distribution represented by the original parameters and the one represented by the updated parameters; and (2) the sum of the distances between the given local estimates and the ones computed using the updated distribution. We derive a gradient-based method for solving this optimization task. Since the gradients require computation of marginal distributions over subsets of variables, they are NP-hard in general on arbitrary probabilistic models but can be computed efficiently on tractable models [25]. This shows the virtue of using tractable models in our learning settings.

An important aspect of our proposed method is that it works well even if there does not exist a valid joint probability distribution P such that the local estimates are marginal distributions of P. We call such estimates *inconsistent local estimates*. For example, consider two marginal probability distributions $P_A(x_1, x_2)$ and $P_B(x_2, x_3)$ over three random variables X_1, X_2, X_3 that are obtained by querying two experts (or local classifiers) A and B respectively. A joint probability distribution $P(x_1, x_2, x_3)$ over the three variables that is consistent with the two distributions exists iff $P_A(x_2) = P_B(x_2)$ where $P_A(x_2) = \sum_{x_1} P_A(x_1, x_2)$ and $P_B(x_2) = \sum_{x_3} P_B(x_2, x_3)$. Unfortunately, because of factors such as precision errors and user bias, in practice, $P_A(x_2)$ will often not equal $P_B(x_2)$. We show how our proposed linear combination of KL distances helps circumvent this issue.

We performed a controlled empirical evaluation of our proposed method using 20 popular datasets that have been used in numerous studies on tractable models [11]. Our results show that our approach that leverages local estimates yields significant improvements in both generative and predictive performance over the original model learned from data alone, even when the local estimates are inconsistent. Moreover, since the optimization problem is smooth, our procedure is guaranteed to reach a local optima under mild conditions.

Related Work. Vomlel [26, 27] studied the problem of integrating probabilistic knowledge bases where a joint probability distribution is constructed from low-dimensional probability distributions (local estimates). Vomlel used a classic optimization method called iterative proportional fitting procedure (IPFP) [8] and proposed a variant called the generalized expectation maximization algorithm (GEMA) for solving this problem. Vomlel provided convergence proofs for these methods; showing that IPFP converges when the local estimates are consistent and GEMA converges even if the local estimates are inconsistent. Unfortunately, Vomlel's approach has high computational complexity (is exponential in the treewidth of the graph defined over the local estimates) and is not practical. The method proposed in this paper does not have this limitation. Peng and Ding [14] proposed two polynomial time approximations for IPFP and applied them to Bayesian networks. However, their preliminary experimental study demonstrates that the error due to their approximations is quite high and convergence is not guaranteed if the local estimates are inconsistent. In contrast, our proposed method makes very few approximations and leverages tractable inference to yield a practical scheme.

2 Notation and Background

We use the following notation. Bold upper-case letters (e.g., $\boldsymbol{X}$, $\boldsymbol{Y}$, $\boldsymbol{U}$, etc.) are used to denote sets of discrete random variables while bold lower-case letters (e.g., $\boldsymbol{x}$, $\boldsymbol{y}$, etc.) denote an assignment of values to all variables in the corresponding set denoted by bold upper case letters (thus $\boldsymbol{x}$ and $\boldsymbol{y}$ denote an assignment to all variables in the set $\boldsymbol{X}$ and $\boldsymbol{Y}$ respectively). We use upper-case letters (e.g., X, U, V, etc.) to denote the variables. For simplicity of exposition, we assume that all variables are binary taking values from the set $\{0, 1\}$. A lower case letter (e.g., x, y, etc.) denotes an assignment of a value to the corresponding variable denoted by the upper case letter.

For simplicity of exposition, we present our algorithm and experimental results on a specific type of tractable model called cutset networks [18] which combines Bayesian networks and AND/OR graphs [6], noting that the algorithm in this paper can be easily extended to other tractable models such as sum-product networks [15], thin junction trees [1] and arithmetic circuits [3].

[1] We describe these applications in the supplementary material and focus on technical aspects in this paper.

2.1 Bayesian Networks

Bayesian networks (BNs) [4, 13] use directed acyclic graphs (DAGs) and conditional probability tables (CPTs) to compactly represent joint probability distributions over a set of random variables. Each node in the DAG denotes a random variable and is associated with a conditional probability distribution of the corresponding variable given its parents in the DAG. Formally, a Bayesian network is a triple $\langle \boldsymbol{X}, G, F \rangle$ where $\boldsymbol{X} = \{X_1, \ldots, X_n\}$ is a set of random variables; $F = \{F_1, \ldots, F_n\}$ is a set of conditional probability tables (CPTs); and $G(V, E)$ is a directed acyclic graph such that each vertex $V_i \in V$ is associated with the variable $X_i \in \boldsymbol{X}$ and E is the set of directed edges. Each node V_i is associated with the CPT $F_i = P_i(X_i|\boldsymbol{U}_i)$ that represents the conditional probability distribution of the variable X_i given its parents $\boldsymbol{U}_i \subseteq \boldsymbol{X} \setminus X_i$.

Let $\boldsymbol{x} = (x_1, \ldots, x_n)$ be an assignment of values to all variables in the set $\boldsymbol{X}$ and let $\boldsymbol{u}_i \sim \boldsymbol{x}$ denote that the value assignment to all variables in the set $\boldsymbol{U}_i$ is compatible with $\boldsymbol{x}$, namely $\boldsymbol{u}_i$ is the projection of $\boldsymbol{x}$ on $\boldsymbol{U}_i$. We will parameterize the Bayesian network using a set of parameters Θ where each $\theta_{x_i \boldsymbol{u}_i} \in \Theta$ is equal to the conditional probability $P(X_i = 1|\boldsymbol{U}_i = \boldsymbol{u}_i)$. Note that since $P(X_i = 1|\boldsymbol{U}_i = \boldsymbol{u}_i) + P(X_i = 0|\boldsymbol{U}_i = \boldsymbol{u}_i) = 1$, we do not need to have a separate parameter for $P(X_i = 0|\boldsymbol{U}_i = \boldsymbol{u}_i)$. Under this parameterization, the probability distribution represented by the Bayesian network can be written as:

$$\Pr(\boldsymbol{x}) = \prod_{\substack{i=1 \\ \boldsymbol{u}_i \sim \boldsymbol{x}}}^{n} (\theta_{x_i \boldsymbol{u}_i})^{x_i} (1 - \theta_{x_i \boldsymbol{u}_i})^{1 - x_i} \tag{1}$$

The two main reasoning tasks over Bayesian networks are most-probable explanation (MPE) and posterior marginal inference (MAR). The former seeks to find an assignment of values to all variables given evidence such that the probability is maximized. Formally, let $\boldsymbol{E} \subset \boldsymbol{X}$ be a set of evidence variables, let $\boldsymbol{Z} = \boldsymbol{X} \setminus \boldsymbol{E}$, and $\boldsymbol{e}$ be an assignment to all variables in $\boldsymbol{E}$, then the MPE task is given by $\max_{\boldsymbol{z}} \Pr(\boldsymbol{z}, \boldsymbol{e})$. The MAR task seeks to find the marginal probability distribution over each non-evidence variable given $\boldsymbol{e}$, namely compute $\Pr(z_j|\boldsymbol{e})$ for each variable $Z_j \in \boldsymbol{Z}$. Both tasks are known to be NP-hard in general but can be solved efficiently in practice on Bayesian networks having *small treewidth* using bucket (or variable) elimination and propagation algorithms [5]. Specifically, given a Bayesian network having treewidth w, the time and space complexity of these exact inference algorithms is bounded by $O(n2^{w+1})$ and $O(n2^w)$ respectively.

2.2 Cutset Networks

Cutset networks (CNs) [17] are tractable probabilistic models that combine two well-known classes of tractable models: AND/OR graphs [6] and tree Bayesian networks. Formally, a CN over a set of variables $\boldsymbol{X}$ ($\boldsymbol{X}$ may include latent variables) is defined recursively using the following three conditions: (1) A tree Bayesian network over $\boldsymbol{X}$ is a CN; (2) An OR node labeled by a variable $X_i \in \boldsymbol{X}$ such that $|\boldsymbol{X}| > 1$ with two child CNs, each defined over the set $\boldsymbol{X} \setminus \{X_i\}$ is a CN. We follow the convention that the left child of the OR node labeled by X_i represents conditioning over $X_i = 0$ and the right child represents conditioning over $X_i = 1$. The arcs from the OR node to its child nodes are labeled with probability values in $\mathbb{R}^+$ such that they sum to 1; and (3) Let $(\boldsymbol{X}_1, \boldsymbol{X}_2)$ be a partition of $\boldsymbol{X}$ such that $|\boldsymbol{X}| > 1$. Then, an AND node with two child CNs, one defined over $\boldsymbol{X}_1$ and the second defined over $\boldsymbol{X}_2$ is a CN.

Fig. 1 shows a CN defined over eight variables. In general, a full assignment $\boldsymbol{x}$ yields a rooted sub-graph in a CN.

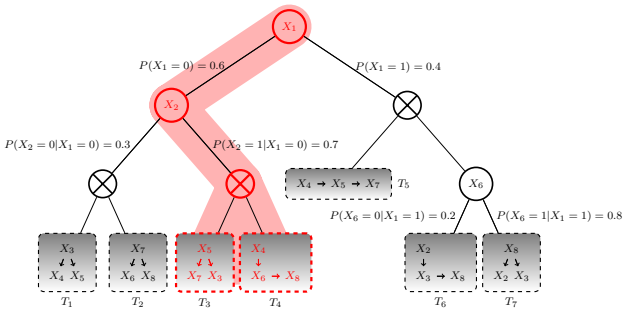

Figure 1: A cutset network defined over eight variables $\{X_1, \ldots, X_8\}$. OR nodes are denoted with variables in circle, AND nodes by cross mark in circle and leaf nodes (tree Bayesian networks) by shadowed dotted rectangles. Left and right arcs emanating from an OR node are labeled with conditional probabilities. The highlighted paths show the sub-tree consistent with the assignment $\{X_1 = 0, X_2 = 1, X_3 = \cdots = X_8 = 0\}$.

Following our notation for Bayesian networks, let Θ denote the set of parameters of the cutset network, such that $\theta_{x_i \boldsymbol{u}_i} \in \Theta$ is equal to the conditional probability $P(X_i = 1 | \boldsymbol{U}_i = \boldsymbol{u}_i)$. Given an assignment $\boldsymbol{x}$, when X_i is an OR node, $\boldsymbol{u}_i$ denotes the assignment along the path from the root of the CN to X_i. Alternatively, when X_i belongs to a tree Bayesian network in the CN, $\boldsymbol{u}_i$ denotes the assignment formed by composing the assignment along the path from the root of the CN to the tree Bayesian network with the assignment to the parent of X_i (if X_i has a parent in the network). Given this parameterization, the probability distribution associated with the cutset network is given by

$$\Pr(\boldsymbol{x}) = \prod_{\substack{i=1 \\ \boldsymbol{u}_i \sim \boldsymbol{x}}}^{n} (\theta_{x_i \boldsymbol{u}_i})^{x_i} (1 - \theta_{x_i \boldsymbol{u}_i})^{1 - x_i} \tag{2}$$

A distinguishing feature of CNs is that when they have no latent variables, both MAR and MPE inference can be performed over them in time that scales linearly with the size of the network. This can be accomplished by converting the CN in linear time to an AND/OR graph [6] or an arithmetic circuit [3] and performing two-passes over these circuits. When latent variables are present, CNs admit linear time MAR inference only while MPE inference is intractable in general.

3 Our Approach

In this section, we define the optimization problem for learning the parameters of a cutset network from local estimates and present a gradient-based algorithm for solving it.

3.1 The Learning Problem

The high-level goal of this paper is to improve the prediction quality and model fit of an existing model that is induced using global information or full data alone by leveraging *local information*. The latter is often available in practice from various sources such as prior knowledge, expert domain knowledge, external sensors or processes and (local) data over a small subset of features. A distinguishing property of local information is that it is typically available in large quantities and thus more robust as compared to global information/full data. However, because of various issues such as user bias (e.g., when information is obtained from different experts), locality and privacy constraints, local information is often noisy and inconsistent. Therefore, our specific goal is to learn cutset networks such that they optimally combine both local and global information, regularizing the two appropriately based on background knowledge.

We begin by describing required notation and simplifying assumptions. We assume that we have access to local information that can be summarized via pairwise distributions. Note that our algorithm can be easily extended to arbitrary (non-pairwise) local marginal distributions and we make the assumption for clarity of presentation only. Let D denote the KL divergence (distance) between two probability distributions defined over the same set of variables. Namely, $D(\mathcal{A}, \mathcal{B}) = \sum_{\boldsymbol{x}} \mathcal{A}(\boldsymbol{x}) \log \mathcal{A}(\boldsymbol{x}) - \sum_{\boldsymbol{x}} \mathcal{A}(\boldsymbol{x}) \log \mathcal{B}(\boldsymbol{x})$ where $\mathcal{A}$ and $\mathcal{B}$ are two distributions defined over $\boldsymbol{X}$. The first term equals the negative *entropy* of the distribution $\mathcal{A}$ and the second term equals the negative *cross-entropy* between $\mathcal{A}$ and $\mathcal{B}$.

Let $\mathcal{E}$ denote a subset of pairs of random variables, namely $\mathcal{E} \subseteq \{(X_j, X_k) | X_j, X_k \in \boldsymbol{X} \text{ and } j < k\}$.

Definition 1. *We say that a set of local estimates $\{\mathcal{P}_{jk}(X_j, X_k) | X_j, X_k \in \mathcal{E}\}$ is* ***consistent*** *iff there exists a probability distribution $\mathcal{P}$ such that each estimate $\mathcal{P}_{jk}(X_j, X_k)$ is a marginal probability distribution of $\mathcal{P}$. Otherwise, we say that the set is* ***inconsistent****.*

Given a set of local estimates and data, the learning problem is given by:

Given: A cutset network structure with parameters Θ representing a probability distribution $\mathcal{R}$, a set of pairwise local distributions $\mathcal{P}_{jk}(X_j, X_k)$ where $(X_j, X_k) \in \mathcal{E}$ and a fully observed dataset $\mathcal{X} = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}\}$.

To do: Find an assignment of values to all parameters in Θ such that the negative log-likelihood of $\mathcal{X}$ w.r.t. $\mathcal{R}_\Theta$ is minimized and the KL distance between $\mathcal{P}_{jk}(X_j, X_k)$ and $\mathcal{R}_\Theta(X_j, X_k)$ equals zero for all $X_j, X_k \in \mathcal{E}$.

Mathematically, we can express it as:

$$\underset{\boldsymbol{\Theta}}{\text{argmin}} \; - \sum_{\boldsymbol{x}^{(d)} \in \mathcal{X}} \log \mathcal{R}_\Theta(\boldsymbol{x}^{(d)}) \; s.t. \; \forall \; (X_j, X_k) \in \mathcal{E}, \; D(\mathcal{P}_{jk}(x_j, x_k), R_\Theta(x_j, x_k)) = 0 \tag{3}$$

Unfortunately, if the set $\{\mathcal{P}_{jk} | X_j, X_k \in \mathcal{E}\}$ is inconsistent (see Def. 1) then the above constrained optimization problem is infeasible.[2] To address this issue, we propose to use a linear combination of the objective and constraints (which is akin to lagrange relaxation). Mathematically,

$$\underset{\boldsymbol{\Theta}}{\text{argmin}} \; \lambda_1 \sum_{(X_j, X_k) \in \mathcal{E}} D(\mathcal{P}_{jk}(x_j, x_k), R_\Theta(x_j, x_k)) \; - \; \lambda_2 \left(\sum_{\boldsymbol{x}^{(d)} \in \mathcal{X}} \log \mathcal{R}_\Theta(\boldsymbol{x}^{(d)}) \right) \tag{4}$$

where $\lambda_1 \geq 0$ ad $\lambda_2 \geq 0$ are hyperparameters (technically, we only need one hyperparameter; we use two for convenience) that model the relative importance of the local and global statistics respectively.

3.2 Simplifying the Learning Problem

We simplify the optimization problem given in Eq. (4) by making the following observation. Since $P_{jk}(x_j, x_k) \log P_{jk}(x_j, x_k)$ is constant, we can remove it from the objective function and replace each term $D(\mathcal{P}_{jk}(x_j, x_k), R_\Theta(x_j, x_k))$ in Eq. (4) by $-\sum_{x_j, x_k} \mathcal{P}_{jk}(x_j, x_k) \log R_\Theta(x_j, x_k)$.

Negating the objective and making the substitutions described above yields the following maximization problem:

$$\underset{\boldsymbol{\Theta}}{\text{argmax}} \; \lambda_1 \sum_{(X_j, X_k) \in \mathcal{E}} \sum_{x_j, x_k} \mathcal{P}_{jk}(x_j, x_k) \log \mathcal{R}_\Theta(x_j, x_k) + \lambda_2 \left(\sum_{\boldsymbol{x}^{(d)} \in \mathcal{X}} \log \mathcal{R}_\Theta(\boldsymbol{x}^{(d)}) \right) \tag{5}$$

The optimization problem given in Eq. (5) is not concave in the parameters Θ but smooth and therefore can be solved using an iterative, gradient ascent algorithm. However, a drawback of this algorithm is that the gradient of the second term in Eq. (5) requires us to go over the whole data at each iteration yielding a slow algorithm. Although, stochastic gradient descent or mini-batch approaches can be used to address this issue, their convergence is quite slow in practice. Therefore, in order to reduce the computational complexity, we propose the following moment-matching approach.

Let $\mathcal{Q}$ denote the distribution associated with the cutset network having the same structure as $\mathcal{R}$. Thus, there is a one-to-one correspondence between the parameters of $\mathcal{Q}$ and $\mathcal{R}$. Let Π denote the set of parameters of $\mathcal{Q}$. Thus, given a parameter $\pi_{x_i, \boldsymbol{u}_i} \in \Pi$, there is a corresponding parameter $\theta_{x_i, \boldsymbol{u}_i}$ in Θ. Let the set of parameters Π be learned from data $\mathcal{X}$ by maximizing the log-likelihood. Since the parameters of cutset networks are conditional probability distributions, given $\mathcal{Q}$ learned from data, we can use negative cross-entropy between parameters of $\mathcal{Q}$ and $\mathcal{R}$ in lieu of the second term (log-likelihood) of Eq. (5). Mathematically,

$$\begin{aligned} &\underset{\boldsymbol{\Theta}}{\text{argmax}} \; \lambda_1 \sum_{(X_j, X_k) \in \mathcal{E}} \sum_{x_j, x_k} \mathcal{P}_{jk}(x_j, x_k) \log \mathcal{R}_\Theta(x_j, x_k) + \\ &\lambda_2 \sum_{\theta_{x_i, \boldsymbol{u}_i} \in \Theta} \pi_{x_i, \boldsymbol{u}_i} \log \theta_{x_i, \boldsymbol{u}_i} + (1 - \pi_{x_i, \boldsymbol{u}_i}) \log(1 - \theta_{x_i, \boldsymbol{u}_i}) \end{aligned} \tag{6}$$

3.3 Solving the Learning Problem via Gradient Ascent

We propose to solve the optimization task given in Eq. (6) using gradient ascent methods. To this end, in this section, we derive the gradients w.r.t. each parameter $\theta_{x_i, \boldsymbol{u}_i}$. The partial derivative of the second term w.r.t. $\theta_{x_i, \boldsymbol{u}_i}$ is straight-forward and given by

$$\lambda_2 \left(\frac{\pi_{x_i, \boldsymbol{u}_i}}{\theta_{x_i, \boldsymbol{u}_i}} - \frac{1 - \pi_{x_i, \boldsymbol{u}_i}}{1 - \theta_{x_i, \boldsymbol{u}_i}} \right) \tag{7}$$

The partial derivative of the first term of Eq. (6) is more involved and we summarize it in the following proposition. (Proofs are given in the appendix.)

[2] Note that the problem is infeasible even if the estimates are consistent and the cutset network is not an independency map (I-map) (cf. [9]) of $\mathcal{P}$

Algorithm 1: LCN-LIS ($\mathcal{X}$,$\mathcal{P}$, λ_1, λ_2, T)

Input : (1) Training examples $\mathcal{X}$ defined over a set of variables $\boldsymbol{X}$; (2) a set of inconsistent pairwise marginal statistics $\mathcal{P}_{jk}(X_j, X_k)$ where $(X_j, X_k) \in \mathcal{E}$ and $\mathcal{E} \subseteq \{X_j, X_k | X_j, X_k \in \boldsymbol{X} \text{ and } j < k\}$; (3) Two hyper parameters $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ and (4) A bound T on the number of iterations

Output : A Cutset Network $\mathcal{R}$

1 **begin**
2 Learn a Cutset Network $\mathcal{Q}$ from Data $\mathcal{X}$ using learning algorithms from literature (cf. [16–18]).
3 Initialize cutset network $\mathcal{R}$ to have the same structure as $\mathcal{Q}$. Let $\Theta = \{\theta_{x_i,\boldsymbol{u_i}}\}$ and $\Pi = \{\pi_{x_i,\boldsymbol{u_i}}\}$ denote the set of parameters of $\mathcal{R}$ and $\mathcal{Q}$ respectively.
4 Initialize: all parameters $\theta_{x_i,\boldsymbol{u_i}}$ of $\mathcal{R}$ to a random number between 0 and 1. Let $\theta^0_{x_i,\boldsymbol{u_i}}$ denote the initial value of $\theta_{x_i,\boldsymbol{u_i}}$
5 Initialize: $t = 0$
6 **repeat**
7 Initialize: $g_{x_i,\boldsymbol{u}_i} = 0$ for each $\theta_{x_i,\boldsymbol{u_i}} \in \Theta$
8 **for** *each possible value assignment* $X_j = x_j$ *and* $X_k = x_k$ *where* $(X_j, X_k) \in \mathcal{E}$ **do**
9 Set $X_j = x_j$ and $X_k = x_k$ as evidence in $\mathcal{R}$ and compute the conditional probabilities $\mathcal{R}_{\Theta^t}(\boldsymbol{u}_i, x_i | x_j, x_k)$ for each $\theta_{x_i,\boldsymbol{u}_i} \in \Theta$
10 **for** *each parameter* $\theta_{x_i,\boldsymbol{u_i}} \in \Theta$ **do**
11 Let:

$$\delta^+ = \frac{\mathcal{R}_{\Theta^t}(\boldsymbol{u}_i, X_i = 1 | x_j, x_k)}{\theta^t_{x_i,\boldsymbol{u}_i}} \text{ and } \delta^- = \frac{\mathcal{R}_{\Theta^t}(\boldsymbol{u}_i, X_i = 0 | x_j, x_k)}{(1 - \theta^t_{x_i,\boldsymbol{u}_i})}$$

12 Update:

$$g_{x_i,\boldsymbol{u}_i} = g_{x_i,\boldsymbol{u}_i} + \lambda_1 \mathcal{P}_{jk}(x_j, x_k)(\delta^+ - \delta^-) + \lambda_2 \left(\frac{\pi_{x_i,\boldsymbol{u}_i}}{\theta^t_{x_i,\boldsymbol{u}_i}} - \frac{1 - \pi_{x_i,\boldsymbol{u}_i}}{1 - \theta^t_{x_i,\boldsymbol{u}_i}} \right)$$

13 **end**
14 **end**
15 **for** *each parameter* $\theta_{x_i,\boldsymbol{u_i}} \in \Theta$ **do**
16 $\theta^{t+1}_{x_i,\boldsymbol{u}_i} = \theta^t_{x_i,\boldsymbol{u}_i} + \alpha \times g_{x_i,\boldsymbol{u}_i}$ // α: learning rate
17 **end**
18 $t = t + 1$
19 **until** *convergence or* $t \geq T$;
20 **return** $\mathcal{R}$ *with parameters* Θ^t
21 **end**

Proposition 2. *The partial derivative of*

$$\lambda_1 \sum_{(X_j,X_k) \in \mathcal{E}} \sum_{x_j,x_k} \mathcal{P}_{jk}(x_j, x_k) \log \mathcal{R}_\Theta(x_j, x_k)$$

w.r.t. $\theta_{x_i,\boldsymbol{u}_i}$ *is given by*

$$\lambda_1 \sum_{(X_j,X_k) \in \mathcal{E}} \sum_{x_j,x_k} \mathcal{P}_{jk}(x_j, x_k) \left(\frac{\mathcal{R}_\Theta(\boldsymbol{u}_i, X_i = 1 | x_j, x_k)}{\theta_{x_i,\boldsymbol{u}_i}} - \frac{\mathcal{R}_\Theta(\boldsymbol{u}_i, X_i = 0 | x_j, x_k)}{1 - \theta_{x_i,\boldsymbol{u}_i}} \right) \quad (8)$$

4 Formal Algorithm

Next, we formally present an algorithm that leverages the gradient equations given in Eqs. (7)–(8) (see Algorithm 1) for solving the learning problem given in Eq. (6). The algorithm, which we call learning cutset networks with local inconsistent statistics (LCN-LIS), takes as input training dataset $\mathcal{X}$, local statistics $P_{jk}(X_j, X_k)$, two hyperparameters λ_1 and λ_2 (real numbers) and an integer bound T on the number of iterations. It begins by learning a cutset network $\mathcal{Q}$ from the dataset $\mathcal{X}$ (step 2) and initializes $\mathcal{R}$ to have the same structure as $\mathcal{Q}$ (step 3). In steps 4–19, the algorithm runs the gradient ascent steps. The gradient ascent begins by setting all parameters to a random number between 0 and 1. Then, at each iteration t, for each pair $(X_j, X_k) \in \mathcal{E}$, it sets X_j and X_k as evidence and runs a two-pass inference algorithm over the cutset network [6] to compute the required conditional probabilities $\mathcal{R}_{\Theta^t}(\boldsymbol{u}_i, X_i = 1 | x_j, x_k)$ and $\mathcal{R}_{\Theta^t}(\boldsymbol{u}_i, X_i = 0 | x_j, x_k)$ (step 9).

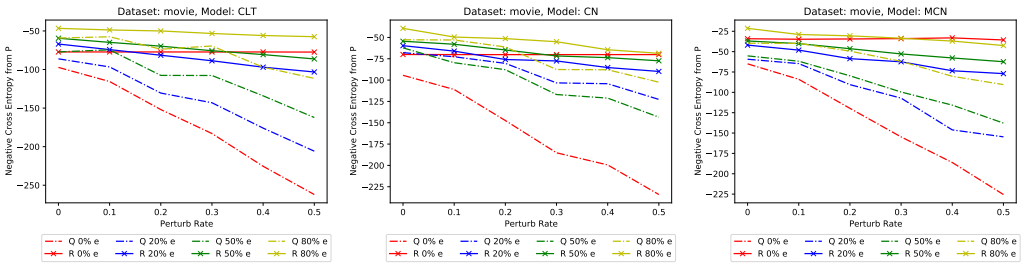

(a) movie; CLT (b) movie; CN (c) movie; MCN

Figure 2: Negative Cross Entropy between $\mathcal{P}$ and $\mathcal{Q}$, and between $\mathcal{P}$ and $\mathcal{R}$ with evidence of 0%, 20%, 50% and 80% on three different models: CLT, CN and MCN, as a function of perturb rate for a randomly chosen dataset: movie.

The algorithm then updates the gradient for each parameter $\theta_{x_i,\boldsymbol{u}_i}$ (steps 10–13; also see Eqs. (7)–(8)) given the assignment $X_j = x_j$ and $X_k = x_k$. In steps 15–17, the algorithm updates the parameters using the gradient estimates $g_{x_i,\boldsymbol{u}_i}$ and learning rate α. The algorithm terminates the gradient ascent on convergence or when the bound T on the number of iterations is reached. At termination, the algorithm returns $\mathcal{R}$ with parameters Θ^t.

The main virtue of our algorithm is that it has polynomial computational complexity. More specifically, the time (and space) complexity of Algorithm 1 is $O(|\mathcal{E}| \times |\Theta| \times T)$ where $|\mathcal{E}|$ denotes the number of pairwise (local) probability distributions provided to the algorithm as input, $|\Theta|$ the number of parameters and T the bound on the number of iterations.

Remarks. Note that a feasible solution to the optimization problem given in Eq. (6), and thus the one returned by Algorithm 1 *filters inconsistency* in the local estimates $\{\mathcal{P}_{jk}|X_j, X_k \in \mathcal{E}\}$ because it yields a globally consistent model $\mathcal{R}_\Theta$. Unlike previously proposed techniques for solving the optimization task in Eq. (3) such as the iterative proportional fitting procedure (IPFP) [8, 27] which will not converge when the local estimates are inconsistent [3], Algorithm 1 will converge to a local optima (because the objective is smooth).

To summarize, we derived and presented a gradient-based algorithm for learning the parameters of cutset networks in presence of local estimates (see Algorithm 1) and showed that the algorithm requires time and space that scales linearly with the number of given local statistics. Note that Theorem 2 can be easily extended to Bayesian and Markov networks. However, the problem is that computing the terms in the numerator in Eq. 8 will be NP-hard in general on Bayesian and Markov networks. This highlights another virtue of tractable models, local information, even if it is inconsistent can be efficiently integrated into a tractable model.

5 Experiments

We performed a detailed, controlled experimental study to evaluate the impact of using inconsistent local statistics on the quality of the learned model. Specifically, we used the following controls: (1) the accuracy of the cutset network $\mathcal{Q}$ learned from data $\mathcal{X}$ in Algorithm 1; (2) the strength/level of inconsistency in the local statistics $\mathcal{P}_{jk}(X_j, X_k)$; (3) the cutset network architecture; and (4) the number of evidence variables or observations available at test time (to test discriminative performance).

We used 20 benchmark datasets that have been widely used in previous studies [21, 22] to evaluate our new approach. The number of variables in these datasets vary from 16 to 1556, and all variables are binary. We ran Algorithm 1 for a maximum of 48 hours or 1000 iterations (namely $T = 1000$) or convergence,whichever was earlier. We used 5-fold cross-validation to select the values of the hyperparameters λ_1 and λ_2.

[3] IPFP and its extension called GEMA proposed by Vomlel [27] are computationally infeasible unless the treewidth of the primal graph obtained by combining the edges corresponding to $\mathcal{E}$ and the cutset network $\mathcal{Q}$ is bounded by a small constant and the cutset network is an I-map (cf. [9]) of a Markov network whose potentials are given by $\mathcal{E}$. The second requirement is particularly difficult to satisfy in practice.

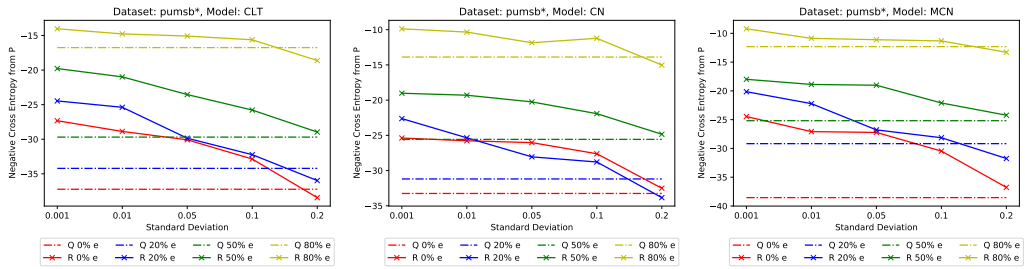

(a) pumsb*; CLT (b) pumsb*; CN (c) pumsb*; MCN

Figure 3: Negative Cross Entropy between $\mathcal{P}$ and $\mathcal{Q}$, and between $\mathcal{P}$ and $\mathcal{R}$, with evidence of 0%, 20%, 50% and 80% on three different models: CLT, CN and MCN, as a function of standard deviation σ (of Gaussian noise that is applied to the local statistics $\mathcal{P}_{jk}(X_j, X_k)$) on one randomly chosen dataset: pumsb*.

Table 1: Generative (0% evidence) performance measured using negative cross entropy between $\mathcal{P}$ and $\mathcal{Q}$, and between $\mathcal{P}$ and $\mathcal{R}$ on three different models: CLTs, CNs, MCNs.

		Negative cross-entropy with **0% Evidence**					
Datasets	**#var**	**CLTs**		**CNs**		**MCNs**	
		$\mathcal{Q}$	$\mathcal{R}$	$\mathcal{Q}$	$\mathcal{R}$	$\mathcal{Q}$	$\mathcal{R}$
nltcs	16	-7.62	**-6.85**	-6.86	**-6.17**	-7.32	**-6.03**
msnbc	17	-7.01	**-6.58**	-6.99	**-6.36**	-7.08	**-6.08**
kdd	64	-5.20	**-2.80**	-5.86	**-2.77**	-5.52	**-2.69**
plants	69	-18.64	**-17.23**	-17.70	**-15.55**	-17.04	**-15.05**
audio	100	-47.16	**-45.52**	-47.15	**-44.01**	-44.20	**-42.57**
jester	100	-60.95	**-59.62**	-61.08	**-57.45**	-57.92	**-56.70**
netflix	100	-62.98	**-61.46**	-63.85	**-60.93**	-58.31	**-57.73**
accidents	111	-38.00	**-34.21**	-37.06	**-33.14**	-35.72	**-32.43**
retail	135	-15.54	**-11.59**	-17.22	**-11.61**	-16.68	**-11.61**
pumsb*	163	-37.23	**-32.86**	-33.25	**-27.61**	-38.55	**-30.46**
dna	180	-103.18	**-95.89**	-124.06	**-99.41**	-98.53	**-90.58**
kosarek	190	-21.04	**-12.98**	-19.46	**-12.45**	-18.23	**-10.68**
msweb	294	-24.57	**-12.28**	-24.37	**-12.33**	-22.48	**-11.05**
book	500	-58.07	**-41.80**	-56.28	**-41.46**	-51.81	**-34.49**
movie	500	-97.27	**-77.21**	-94.39	**-70.00**	-64.99	**-34.14**
webkb	839	-207.12	**-172.65**	-199.76	**-162.30**	-181.05	**-126.06**
reuters	889	-152.53	**-113.42**	-146.47	**-109.21**	-149.49	**-111.29**
20newsg	910	-194.55	**-157.70**	-181.47	**-155.48**	-177.83	**-152.65**
bbc	1058	-325.51	**-270.74**	-265.46	**-210.53**	-243.87	**-223.60**
ad	1556	-112.01	**-47.27**	-93.96	**-41.34**	-96.79	**-53.70**
Total AVG		-79.81	**-64.03**	-75.14	**-59.01**	-69.67	**-55.43**

We used *negative cross-entropy* between the true $\mathcal{P}$ and the learned model (under various conditions) as our evaluation criteria. The higher the negative cross-entropy, better the model. For each dataset, we learned a mixtures of cutset networks [18] and used it as the true model $\mathcal{P}$. We used 10% of randomly chosen examples in the training set to learn $\mathcal{Q}$. Thus, the dataset used by Algorithm 1 has 90% fewer examples than the one used for learning $\mathcal{P}$. We did this to ensure that the true model differs significantly from $\mathcal{Q}$, which in turn will help us evaluate how local information improves the quality of an inferior starting point. We further controlled the quality of $\mathcal{Q}$ using a parameter h which we call the *perturb rate*. h lies between 0 and 100, and given a value for h, we replaced h% of the parameters in $\mathcal{Q}$ by a random number. We normalized $\mathcal{Q}$ to ensure that it is a valid probability distribution.

We used three types of cutset network structures: (1) cutset networks with depth 0 which are equivalent to Chow-Liu trees (CLTs); (2) cutset networks with no latent variables (CNs); and (3) mixtures of cutset networks (MCNs). The latter is a state-of-the-art model [20]. We learned both discriminative and generative cutset networks. In discriminative networks, we set L% of random variables as evidence $\boldsymbol{E}$ and learn a probability distribution over the variables $\boldsymbol{X} \setminus \boldsymbol{E}$ given an assignment $\boldsymbol{e}$ to the evidence variables. We used 4 values for L: 0, 20, 50 and 80. When $L = 0$, we get a generative model while the remaining models are discriminative.

We generated local statistics from $\mathcal{P}$ as follows. Since $\mathcal{P}$ is a tractable model, we can efficiently (in linear time) compute $\mathcal{P}_{jk}(X_j, X_k)$ for all $X_j, X_k \in \boldsymbol{X}$. To make them inconsistent, we added a value ϵ that is randomly sampled from a Gaussian distribution with 0 mean and standard deviation σ. We experimented with five values of σ: 0.001, 0.01, 0.05, 0.1, 0.2. Note that after adding ϵ, we have to normalize the distributions to ensure that they are valid.

5.1 Results

Figure 2 presents the negative cross-entropy scores achieved by CLTs, CNs and MCNs as a function of the perturb rate (h) for a randomly chosen dataset ('movie'), and for different values of L (the percentage of evidence variables). We present the plots for the remaining datasets in the supplementary material. These plots help us evaluate the impact that the quality of $\mathcal{Q}$ has (namely the impact of the starting point) on the model learned by Algorithm 1. We observe that as the perturb rate increases, there is a substantial drop in the negative cross-entropy between $\mathcal{P}$ and $\mathcal{Q}$. However, the negative cross entropy of $\mathcal{P}$ and $\mathcal{R}$ remains relatively flat. This shows that the use of local statistics significantly improves the model quality, especially when the model based on global data is inaccurate.

Figure 3 shows the negative cross-entropy scores achieved by CLTs, CNs and MCNs as a function of the standard deviation $\sigma \in \{0.001, 0.01, 0.05, 0.1, 0.2\}$ (of Gaussian noise that is applied to the local statistics $\mathcal{P}_{jk}(X_j, X_k)$) for dataset 'pumsb*', and for different values of L (the percentage of evidence variables). The results clearly show that when the local statistics contain more noise, the performance degrades as expected. When the σ is large ($\sigma > 0.2$), there is a non-zero chance that the final model $\mathcal{R}$ will be worse than $\mathcal{Q}$. Thus, based on our experimental results, our proposed algorithm is more likely to yield a better model than the one based on global information alone when the Gaussian noise has standard deviation ≤ 0.1.

Table 2: Discriminative (20% evidence) performance of $\mathcal{Q}$ and $\mathcal{R}$ on model CLTs, CNs, MCNs.

Datasets	#var	20% Evidence					
		CLTs		CNs		MCNs	
		$\mathcal{Q}$	$\mathcal{R}$	$\mathcal{Q}$	$\mathcal{R}$	$\mathcal{Q}$	$\mathcal{R}$
nltcs	16	-4.81	**-4.42**	-4.91	**-4.66**	-4.34	**-4.10**
msnbc	17	-5.78	**-5.61**	-6.23	**-6.10**	-5.25	**-4.89**
kdd	64	-5.10	**-3.55**	-5.34	**-3.06**	-4.81	**-3.38**
plants	69	-13.63	**-12.82**	-13.32	**-12.88**	-12.59	**-12.02**
audio	100	-42.35	**-40.19**	-40.58	**-37.15**	-37.86	**-35.27**
jester	100	-55.06	**-52.38**	-51.99	**-51.58**	-49.06	**-47.08**
netflix	100	-60.74	**-60.16**	-55.43	**-53.05**	-52.32	**-50.85**
accidents	111	-37.18	**-36.25**	-36.72	**-34.97**	-33.30	**-32.68**
retail	135	-15.12	**-12.42**	-15.05	**-13.02**	-14.14	**-11.65**
pumsb*	163	-34.21	**-32.23**	-31.19	**-28.78**	-29.17	**-28.14**
dna	180	-95.59	**-91.40**	-115.77	**-94.19**	-92.93	**-82.54**
kosarek	190	-16.98	**-12.71**	-13.32	**-12.01**	-11.68	**-9.04**
msweb	294	-13.47	**-11.46**	-14.04	**-11.98**	-12.14	**-11.37**
book	500	-26.51	**-23.10**	-42.92	**-21.25**	-32.43	**-20.19**
movie	500	-86.17	**-66.97**	-67.89	**-59.64**	-59.34	**-41.97**
webkb	839	-182.77	**-155.29**	-173.35	**-144.99**	-162.70	**-131.96**
reuters	889	-124.07	**-99.43**	-132.16	**-105.33**	-128.97	**-102.87**
20newsg	910	-182.35	**-151.92**	-167.36	**-152.8**	-163.80	**-142.50**
bbc	1056	-308.09	**-247.37**	-250.69	**-193.32**	-240.87	**-195.05**
ad	1558	-98.00	**-37.23**	-74.93	**-36.43**	-77.56	**-36.68**
Total AVG		-70.34	**-57.85**	-65.66	**-53.86**	-61.26	**-50.21**

Tables 1 and 2 show the generative (0% evidence) and discriminative (20% evidence) performance of $\mathcal{Q}$ and $\mathcal{R}$ respectively. Due to lack of space, we report the discriminative performances with 50% and 80% of the variables set as evidence in the supplementary material. We used $\sigma = 0.1$ and $h = 0$ to generate the experimental data given in the tables. Each number in each table is an average over 5 runs. To avoid clutter, we do not report the standard deviation because it was fairly small over all runs. These results help us analyze the impact of the cutset network architecture and number of evidence variables on our evaluation criteria. We observe that on an average, using local inconsistent statistics improves the negative cross entropy of each architecture by 17-23%. MCNs are the best performing model overall and Chow-Liu trees (CLTs) are significantly worse. There was no significant difference in the amount of improvement as we increased the number of evidence nodes. This suggests that our approach is equally useful for both discriminative and generative models.

6 Conclusion

In this paper, we presented a new method for learning the parameters of cutset networks in presence of noisy local estimates. Unlike conventional algorithms which use full i.i.d. data during the learning process, we proposed a novel approach that uses noisy local information to learn a more accurate and robust model. The key advantage of using local estimates is that they are often readily available as compared to full i.i.d. data. We also showed via experiments on benchmark datasets that our new algorithm greatly improves the quality of the initial model learned from i.i.d. data, even when the local estimates are inconsistent and noisy.

References

[1] F. R. Bach and M. I. Jordan. Thin Junction Trees. In *Advances in Neural Information Processing Systems*, pages 569–576. MIT Press, 2001.

[2] David L Banks and Nicholas Hengartner. Social networks. *Encyclopedia of Quantitative Risk Analysis and Assessment*, 4, 2008.

[3] A. Darwiche. A Differential Approach to Inference in Bayesian Networks. *Journal of the ACM*, 50:280–305, 2003.

[4] A. Darwiche. *Modeling and reasoning with Bayesian networks*. Cambridge university press, 2009.

[5] R. Dechter. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence*, 113:41–85, 1999.

[6] R. Dechter and R. Mateescu. AND/OR Search Spaces for Graphical Models. *Artificial Intelligence*, 171(2-3):73–106, 2007.

[7] Alain Degenne and Michel Forsé. *Introducing social networks*. Sage, 1999.

[8] W. E. Deming and F. F. Stephan. On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known. *The Annals of Mathematical Statistics*, 11(4): 427 – 444, 1940.

[9] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

[10] Y. Liang, J. Bekker, and Guy Van den Broeck. Learning the structure of probabilistic sentential decision diagrams. In *Uncertainty in Artificial Intelligence*, 2017.

[11] Daniel Lowd and Jesse Davis. Learning markov network structure with decision trees. In *ICDM*, pages 334–343. IEEE, 2010.

[12] Daniel Lowd and Pedro M. Domingos. Learning arithmetic circuits. In *Uncertainty in Artificial Intelligence*, pages 383–392, 2008.

[13] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

[14] Yun Peng and Zhongli Ding. Modifying bayesian networks by probability constraints. *CoRR*, abs/1207.1356, 2012.

[15] H. Poon and P. Domingos. Sum-Product Networks: A New Deep Architecture. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 337–346. AUAI Press, 2011.

[16] T. Rahman and V. Gogate. Learning ensembles of cutset networks. In *AAAI conference on Artificial Intelligence*, pages 3301–3307, 2016.

[17] T. Rahman and V. Gogate. Merging strategies for sum-product networks: From trees to graphs. In *Proceedings of the Thirty-Second Conference Conference on Uncertainty in Artificial Intelligence*, pages 617–626, 2016.

[18] T. Rahman, P. Kothalkar, and V. Gogate. Cutset networks: A simple, tractable, and scalable approach for improving the accuracy of chow-liu trees. In *Proceedings of the Twenty-Fifth European Conference on Machine Learning*, pages 630–645, 2014.

[19] T. Rahman, S. Jin, and V. Gogate. Cutset Bayesian Networks: A New Representation for Learning Rao-Blackwellised Graphical Models. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 5751–5757, 2019.

[20] T. Rahman, S. Jin, and V. Gogate. Look Ma, No Latent Variables: Accurate Cutset Networks via Compilation. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the Thirty-Sixth International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5311–5320. PMLR, 2019.

[21] A. Rooshenas and D. Lowd. Learning sum-product networks with direct and indirect variable interactions. In *Proceedings of the Thirty-First International Conference on Machine Learning*, pages 710–718, 2014.

[22] Amirmohammad Rooshenas and Daniel Lowd. Learning tractable graphical models using mixture of arithmetic circuits. In *AAAI*, 2013.

[23] B. Settles. *Active Learning*. Morgan & Claypool Publishers, 2012. ISBN 1608457257.

[24] Xiaoting Shao, Alejandro Molina, Antonio Vergari, Karl Stelzner, Robert Peharz, Thomas Liebig, and Kristian Kersting. Conditional sum-product networks: Modular probabilistic circuits via gate functions. *International Journal of Approximate Reasoning*, 140:298–313, 2022.

[25] Antonio Vergari, YooJung Choi, Anji Liu, Stefano Teso, and Guy Van den Broeck. A compositional atlas of tractable circuit operations for probabilistic inference. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, dec 2021.

[26] J. Vomlel. *Methods of Probabilistic Knowledge Integration*. Phd thesis, Czech Technical University, 1999.

[27] J. Vomlel. Integrating inconsistent data in a probabilistic model. *Journal of Applied Non-Classical Logics*, 14(3):367–386, 2004.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] In the supplement.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] In the supplementary material
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] In the Supplement
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In the section on Experiments and the Supplement
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A] Standard deviation was very small and so error bars were not provided.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In the supplement
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]

(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]