

Towards Privacy-Aware Bayesian Networks: A Credal Approach

Niccolò Rocchi ^{a,b}, Fabio Stella ^{a,*} and Cassio de Campos ^c

^aUniversity of Milano-Bicocca, Italy

^bFondazione IRCCS Istituto Nazionale dei Tumori, Italy

^cEindhoven University of Technology, the Netherlands

Abstract. *Bayesian networks* (BN) are versatile probabilistic graphical models that enable efficient knowledge representation and inference. These models have proven effective across diverse domains, including healthcare, bioinformatics, economics, law, and image processing. The structure and parameters of a BN can be obtained by domain experts or directly learned from available data. However, as privacy concerns escalate, it becomes increasingly critical for publicly released models to safeguard sensitive information in training data. Typically, released models do not prioritize privacy by design, and the issue equally affects BNs. In particular, *tracing attacks* from adversaries can combine the released BN with auxiliary data to determine whether specific individuals belong to the data from which the BN was learned. The current approach to addressing this privacy issue involves introducing noise into the learned parameters. While this method offers robust protection against tracing attacks, it also significantly impacts the model's utility, in terms of both the significance and accuracy of the resulting inferences. Hence, high privacy may be attained, but at the cost of releasing a possibly ineffective model. This paper introduces *credal networks* (CN) as a novel and practical solution for balancing the model's privacy and utility. Specifically, after adapting the notion of tracing attacks, we demonstrate that a CN enables the masking of the learned BN, thereby reducing the probability of successful tracing attacks. As CNs are obfuscated but not noisy versions of BNs, they can achieve meaningful inferences while safeguarding the privacy of the released model. Moreover, we identify key learning information that must be concealed to prevent attackers from recovering the BN underlying the released CN. Finally, we conduct a set of numerical experiments to analyze how privacy gains can be modulated by tuning the CN hyperparameters. Our results confirm that CNs provide a principled, practical, and effective approach towards the development of privacy-aware probabilistic graphical models.

1 Introduction

Experts across various domains rely on *artificial intelligence* (AI) systems to support their daily work and critical decision-making. Unlike opaque AI methods, transparent models such as *Bayesian networks* (BN) [18] inspire greater trust and encourage broader adoption, providing rich insights into the data generation process. A BN consists of (i) a graph whose nodes are associated with random variables and whose edges encode conditional (in)dependencies among

the variables, and (ii) a compact set of parameters that allow for representing the variables' joint probability distribution efficiently. When interpreted in *causal* terms [23], a BN becomes a powerful tool to model cause-effect relationships, essential for understanding how decisions propagate towards different outcomes [4, 15]. These reasons have fueled the widespread adoption of BNs in fields where statistical and causal reasoning are central, including oncology [16], neuroscience [26], and genetics [29]. Learning a BN involves two main tasks: recovering its graphical structure and estimating its parameters [27]. These tasks can be accomplished using data alone, by leveraging the expertise of domain specialists, or through a combination of both methods [8]. Even more, multi-center studies are beneficial for overcoming the limitations of individual data centers, such as data scarcity and biases [37]. However, strict privacy regulations often prohibit the direct sharing of individual data across institutions. As a result, *federated learning* [38] approaches emerged. A typical scenario involves all centers learning their own BNs locally and then sharing them with a global aggregator to develop a *global* model [30].

Releasing the learned BN can still leave sensitive information vulnerable [19]. Adversaries who combine the model with auxiliary data (for example, public population registries) can develop *tracing attacks* [7], which determine whether specific individual records were included in the training data set, i.e., the data used to learn the BN. Existing privacy-preserving strategies for BNs typically rely on adding random noise to the network parameters. In this case, the obtained model is said to be *sanitized*. Typical choices for the noise include the Laplacian random perturbation [39], which has been proven to ensure *differential privacy* [13]. Although differential privacy provides robust protection against tracing attacks, perturbation techniques can lead to inconsistencies in the BN marginal distributions when not properly managed [3]. Furthermore, the magnitude of the injected noise can be substantial, yet unnoticeable in the sanitized BN. Consequently, the released model may not accurately represent the underlying data, leading to potentially unreliable inferences [5]. Murakonda et al. [19] established an approximation of the effectiveness of tracing attacks on BNs. They demonstrated that, for any given error rate, the attacker's power is dependent on the network complexity and the sample size. Nevertheless, finding a practical way to *balance* the model's privacy and utility remains a challenging, short-blanket problem that is not adequately addressed in the BN literature [24, 28].

In this work, we build on this gap to introduce a novel approach to the privacy-utility balance problem in BNs, inspired by the relation-

* Corresponding author. Email: fabio.stella@unimib.it

Table 1. Notation adopted throughout the paper; “distrib.” stands for distribution, while “param.” for parameters.

| Symbol | Meaning |
|--|--|
| $ S $ | The cardinality of a set S . |
| \mathbf{X} | Discrete random variables having support $\text{supp}(\mathbf{X})$. |
| $P(\mathbf{X} \mid \theta_0)$ | Multinomial distrib. over $\text{supp}(\mathbf{X})$, parametrized by θ_0 . |
| $p(\cdot \mid \theta_0)$ | Probability mass function, $p(\mathbf{x} \mid \theta_0) = P(\mathbf{X} = \mathbf{x} \mid \theta_0)$. |
| (\mathcal{G}, θ_0) | BN underlying $P(\mathbf{X} \mid \theta_0)$. |
| Θ | Space of param. that factorize according to \mathcal{G} , $\theta_0 \in \Theta$. |
| \mathcal{P} | General population over variables \mathbf{X} . |
| \mathcal{R}, \mathcal{T} | Reference and target populations, $\mathcal{R}, \mathcal{T} \subset \mathcal{P}$. |
| $\hat{K}_{\mathcal{T}}(\mathbf{X})$ | Estimated credal set from data \mathcal{T} , $\hat{K}_{\mathcal{T}}(\mathbf{X}) \subset \Theta$. |
| $\hat{\theta}_{\mathcal{T}}$ | MLE of θ_0 from data \mathcal{T} over space Θ . |
| $\hat{\theta}_{\mathcal{T}}^{\mathcal{X}}$ | MLE of θ_0 from data \mathcal{T} over space $\hat{K}_{\mathcal{T}}(\mathbf{X}) \cap \Theta$. |
| $L(\mathbf{x}, \theta)$ | Log-likelihood ratio (LLR) function over $\text{supp}(\mathbf{X}) \times \Theta$. |
| $L(\mathbf{x})$ | LLR function where 2nd argument is $\hat{\theta}_{\mathcal{T}}$, i.e. $L(\mathbf{x}, \hat{\theta}_{\mathcal{T}})$. |

ship between privacy and Walley’s *imprecise probability* theory [32] pointed out in the recent research [2, 21]. In particular, we leverage *credal networks* (CN) [11], an enhanced version of BNs that incorporates epistemic uncertainty into their parameters. We theoretically and empirically show that releasing a CN as the *masked* version of a BN provides a robust defense against tracing attacks. The parameters of a CN are represented as intervals, or more generally as sets, which enclose the point estimates of the ground-truth BN parameters. Therefore, they can yield *guaranteed* inferential bounds using diverse (exact or approximate) algorithmic approaches [1, 12, 22]. While the inferential bounds may sometimes be quite broad, potentially leading to “I don’t know” conclusions, they still ensure correct responses to queries as opposed to noisy BNs. Consequently, CNs are valuable for offering a more usable and explainable form of privacy [6]. Moreover, by controlling the level of uncertainty injected into the BN, a CN can practically balance the privacy level and the uncertainty of inferences.

The paper first introduces the concepts and notation of BNs and CNs (Section 2), followed by a formal definition of a tracing attack against a BN, along with known theoretical results (Section 3). The notion of tracing attack is then extended to CNs while proving its consistency (Section 4.1). Section 4.2 proves, in the sample size limit, that releasing a CN ensures equal or higher privacy than releasing a BN (Theorem 3). Finally, Section 5 presents and discusses an extensive set of numerical experiments that simulate tracing attacks on different networks. Results suggest that the privacy gained by a CN may even exceed the theoretical bound established in Theorem 3.

2 Probabilistic Graphical Models

Probabilistic graphical models (PGMs) [18] are practical tools for representing and reasoning with multidimensional probability distributions. This article focuses on two PGMs: BNs and CNs. The notation adopted throughout the paper is summarized in Table 1.

BNs and CNs both rely on the structure provided by a graph. A graph \mathcal{G} is a pair $\mathcal{G} = (\mathbf{X}, \mathcal{E})$, where \mathbf{X} is the set of *nodes* and \mathcal{E} is a set of edges between pairs of nodes. If $X \rightarrow Y$ in \mathcal{G} then X is said to be a *parent* of Y ; Π_Y denotes the set of parents of Y . A *directed and acyclic graph* (DAG) is a graph where all edges are directed ($X \rightarrow Y$) and there is no sequence of nodes $X \rightarrow \dots \rightarrow Y$ so that $X = Y$. In the following, we consider DAGs \mathcal{G} whose nodes are associated with random variables. Namely, $X \in \mathbf{X}$ represents both a node in \mathcal{G} and a random variable with given marginal distribution $P(X)$. Assuming to know $P(X)$, we denote by $\text{supp}(X) = \{x : P(X = x) > 0\}$ its support. Similarly, $\text{supp}(\mathbf{X})$ is the Cartesian

product $\times_{X \in \mathbf{X}} \text{supp}(X)$. All supports are assumed to be fixed and known. Our work focuses on *categorical* random variables, i.e., a set \mathbf{X} so that $|\text{supp}(\mathbf{X})| < \infty$. In this case, the vector of \mathbf{X} has a joint probability distribution $P(\mathbf{X} \mid \theta_0)$ being a Multinomial(θ_0) over $\text{supp}(\mathbf{X})$, for some set of parameters θ_0 .

2.1 Bayesian Networks

BNs [18] efficiently represent high-dimensional probability distributions, facilitating effective knowledge representation and inference.

Definition 1 (BN). A BN is a pair (\mathcal{G}, θ_0) where \mathcal{G} is a DAG over random variables \mathbf{X} and θ_0 a set of parameters that factorizes accordingly to \mathcal{G} , namely:

$$P(\mathbf{X} \mid \mathcal{G}, \theta_0) = \prod_{X \in \mathbf{X}} P(X \mid \Pi_X, \theta_X), \quad (1)$$

where θ_X , which is part of θ_0 , are parameters defining the local distribution of X .

As we restrict ourselves to categorical random variables, $p(x \mid \Pi_X = \pi_X, \theta_X)$ follows a multinomial for any $\pi_X \in \text{supp}(\Pi_X)$. Hence, the *complexity* of a BN is determined by the total number of configurations of \mathbf{X} :

$$C(\mathcal{G}) = \sum_{X \in \mathbf{X}} |\text{supp}(\Pi_X)| \cdot (|\text{supp}(X)| - 1). \quad (2)$$

A BN can be learned from data alone, and we focus on this setting. However, domain experts can also build BNs by hand [34], but this process may be highly time-consuming, especially as the number of nodes and parameters increases. Hybrid approaches leverage both data and experts’ knowledge to mitigate human-level and data-level biases [8]. Those approaches proved effective in various scenarios, especially when causal relationships are sought [36].

Learning a BN involves two main tasks: learning the DAG (the *structure learning* task), and estimating its parameters (the *parameter learning* task). Denoting the available data by \mathcal{D} , we can express the whole learning task in a Bayesian fashion:

$$\underbrace{P(\mathcal{G}, \theta \mid \mathcal{D})}_{\text{BN learning}} = \underbrace{P(\mathcal{G} \mid \mathcal{D})}_{\text{Structure learning}} \cdot \underbrace{P(\theta \mid \mathcal{G}, \mathcal{D})}_{\text{Parameter learning}}, \quad (3)$$

where we seek the optimum of the left-hand-side term by exploring the spaces of DAGs and parameters.

There are two primary approaches for the structure learning task [17, 27]. The first, known as the *score-based* approach, aims to optimize $P(\mathcal{G} \mid \mathcal{D})$ by maximizing the graph likelihood $P(\mathcal{D} \mid \mathcal{G})$, which is treated as the graph *score*. The second approach, called *constraint-based*, iteratively narrows down the space of all DAGs that are consistent with \mathcal{D} .

In this paper, we assume that the structure learning task has been accomplished, thus \mathcal{G} is known. Consequently, we omit the condition $(\cdot \mid \mathcal{G})$ in all subsequent formulas for convenience. The parameter learning task can be reformulated as follows:

$$\underbrace{P(\theta \mid \mathcal{D})}_{\text{Posterior}} \propto \underbrace{P(\theta)}_{\text{Prior}} \cdot \underbrace{P(\mathcal{D} \mid \theta)}_{\text{Likelihood}}. \quad (4)$$

Consider Θ as the space of all possible sets of parameters that factorize accordingly to \mathcal{G} , namely those for which it holds Eq. (1). There are two common approaches to parameter learning. The first one is the *Maximum likelihood estimation* (MLE), which assumes a

uniform prior and aims at maximizing the likelihood of the parameters, namely at finding $\hat{\theta}_{\mathcal{D}} = \operatorname{argmax}_{\theta \in \Theta} P(\mathcal{D} \mid \theta)$. The second, more robust approach is known as *Bayesian posterior estimation* [18]. The prior $P(\theta)$ is assumed to follow a Dirichlet distribution, which is the conjugate prior of a multinomial likelihood. Hence, the posterior follows a Dirichlet distribution, which can then be maximized (more details are provided in Section 2.2).

2.2 Credal Networks

CNs [11] are models based on Walley’s imprecise probability theory [32] and introduce epistemic uncertainty in the BNs’ parameters.

Definition 2 (Credal set). *A credal set for a random variable X is a closed and convex set of probability distributions over $\operatorname{supp}(X)$, denoted by $K(X)$. The joint credal set over \mathbf{X} is defined similarly over $\operatorname{supp}(\mathbf{X})$.*¹

Being X categorical, $K(X)$ can be thought of as a convex hull in $[0, 1]^{|\operatorname{supp}(X)|-1}$, which is the whole parameter space for X . As such, it can be defined either by its extreme points or through linear constraints [35]. In the following (with a slight abuse of notation), we write $\theta^K \in K(X)$ to refer to the element $P(X \mid \theta^K) \in K(X)$.

A *conditional credal set* $K(X \mid Y)$ is usually defined elementwise, that is, obtained from $K(X, Y)$ by applying the Bayes rule to each of its elements [10] (and possibly taking the convex hull). A collection of conditional credal sets $K(X \mid Y)$ is said to be *separately specified* when $K(X \mid Y = y_1)$ and $K(X \mid Y = y_2)$ are unrelated for $y_1 \neq y_2$, i.e., not constrained by each other.

Definition 3 (Locally specified CNs). *A locally specified CN consists of a DAG \mathcal{G} where each node X is associated with a local credal set $K(X \mid \Pi_X)$.*

As is common in the literature, we consider $K(X \mid \Pi_X)$ to be separately specified. It becomes now visible how a locally and separately specified CN is akin to a BN. For any point parameter describing $P(X \mid \Pi_X = \pi_X)$ in the BN, there is now a closed set $K(X \mid \Pi_X = \pi_X)$ in the CN. What is missing, however, is a way to *combine* the local credal sets, so that a global-to-local relation as the one in Eq. (1) holds. While there is no single way to do this, a practical choice is offered by the CN *strong extension* [10].

Definition 4 (Strong extension). *The strong extension of a locally and separately specified CN is a pair $(\mathcal{G}, K(\mathbf{X}))$, where $K(\mathbf{X})$ is the convex hull of:*

$$\left\{ \prod_{X \in \mathbf{X}} P(X \mid \Pi_X, \theta_X) : P(X \mid \pi_X, \theta_X) \in K(X \mid \pi_X) \right\}. \quad (5)$$

In a strong extension, each distribution within $K(\mathbf{X})$ factorizes accordingly to \mathcal{G} and Eq. (1). As such, the extension mimics the same conditional independence properties of a standard BN for all its extreme distributions [11].

Remark 1. The following discussion introduces two approaches to defining the local sets $K(X \mid \Pi_X)$. The resulting CNs are both locally and separately specified, and we consider their strong extensions $(\mathcal{G}, K(\mathbf{X}))$ throughout the paper. Other types of CNs are out of the scope of the present article and can be found in Corani and de Campos [9]. Our main results hold regardless of how the CN was obtained, as long as they are locally and separately specified.

The first approach to defining $K(X \mid \Pi_X)$ is based on the *imprecise Dirichlet model* (IDM) [31]. The second approach consists of perturbing a given BN. Let X be a random variable, $k = |\operatorname{supp}(X)|$, and assume we want to estimate the parameters $\theta_X = \{\theta_1, \dots, \theta_k\}$ where $\theta_i = p(x_i \mid \theta_X)$. In the *precise* context of the Bayesian posterior estimator, the parameters are assumed to follow a Dirichlet distribution $Dir(\mathbf{c})$. Here, $\mathbf{c} = \{c_1, \dots, c_k\}$ are constrained to $c_i > 0 \forall i$ and $\sum_i c_i = S$, where S is a positive constant [31]. Then, θ_i can be estimated from data as:

$$\theta_i = \frac{n_i + c_i}{N + S}, \quad (6)$$

where n_i is the count of data items for which $X = x_i$, and N is the data sample size. Consider now an *uncertain* version of the Bayesian posterior estimator. The *local* application of IDM [9] allows each c_i to vary within the interval $[0, S]$, leading to a set of priors for each θ_i . The maximum a posteriori estimate for θ_i becomes now the interval:

$$p(x_i) = \left[\frac{n_i}{N + S}, \frac{n_i + S}{N + S} \right]. \quad (7)$$

Parameters of the conditional distributions can be derived similarly:

$$p(x_i \mid \pi_j) = \left[\frac{n_{ij}}{n_j + S}, \frac{n_{ij} + S}{n_j + S} \right], \quad (8)$$

where n_{ij} is the count of data items for which $(X = x_i, \Pi_X = \pi_j)$.

Another approach to define $K(X \mid \Pi_X)$ is known as ε -contamination [11], and consists in injecting uncertainty into a given BN. The BN may have been learned from data or provided by experts. Definition 5 describes the concept for marginal distributions, but it extends naturally to each local conditional probability.

Definition 5 (ε -contamination). *Let X be a discrete random variable with probability mass function $p(\cdot \mid \theta_X)$. Given $\varepsilon \in (0, 1)$, the ε -contamination of $p(\cdot \mid \theta_X)$ is defined as the set of all probability mass functions $r(\cdot)$ so that:*

$$r(x) \geq (1 - \varepsilon)p(x \mid \theta_X), \quad \forall x \in \operatorname{supp}(X). \quad (9)$$

3 Problem Statement

Consider a *general population* \mathcal{P} of individuals, where \mathbf{X} represents the measured variables for each of them, and let (\mathcal{G}, θ_0) be the unknown BN that generated \mathcal{P} . Let also $(\mathcal{G}, \hat{\theta}_{\mathcal{T}})$ be the BN learned from a subsample of individuals $\mathcal{T} \subset \mathcal{P}$, which we call the *target population*. Data in \mathcal{T} are *private*, meaning that no one has direct access to them, apart from the data owner and the learning algorithm. However, the BN may be released for research purposes. Let the attacker have access to the released BN and auxiliary data $\mathcal{R} \subset \mathcal{P}$, referred to as the *reference population*. In this setting, a *tracing attack* [7], also known as *membership inference attack*, consists of detecting whether a specific individual $\mathbf{x} \in \mathcal{P}$ belongs to \mathcal{T} . In a medical field, \mathcal{P} could represent all patients suffering from a particular disease, \mathcal{T} a hospital database that handles some patients belonging to \mathcal{P} , and \mathcal{R} a publicly available clinical registry. Another example involves a competing insurance company that aims to attract more clients from a rival firm. In this context, \mathcal{P} represents a vulnerable category of individuals, while \mathcal{T} and \mathcal{R} denote subsets of people associated with the competing firm and the insurance company, respectively.

The attacker information, namely (\mathcal{G}, θ_0) and \mathcal{R} , is exploited to build a *decision rule* $\phi(\cdot)$ so that, given $\mathbf{x} \in \mathcal{P}$, it yields either $\phi(\mathbf{x}) = 1$, hence concluding $\mathbf{x} \in \mathcal{T}$, or $\phi(\mathbf{x}) = 0$, concluding

¹ Support here means the set of states that *might* have positive probability.

$\mathbf{x} \notin \mathcal{T}$. A possible choice for the rule $\phi(\cdot)$ involves conducting the *log-likelihood ratio* (LLR) test [20]. In our case, the test is expressed by the following competing hypotheses. We assign $\phi(\mathbf{x}) = 1$ whenever we reject H_0 in favor of H_1 , and $\phi(\mathbf{x}) = 0$ otherwise.

$$\begin{cases} H_0 : \mathbf{x} \notin \mathcal{T} \\ H_1 : \mathbf{x} \in \mathcal{T}. \end{cases} \quad (10)$$

Surely, this test has only limited qualities in exposing elements of \mathcal{T} , but these are typically considered too much to be allowed [14].

The definition of the test statistic depends on the values of $\hat{\theta}_{\mathcal{T}}$ and θ_0 [20] (see Table 1 for notation). However, θ_0 is the vector of *true* BN parameters, which are unknown. When \mathcal{R} is large enough, the attacker can assume that the MLE estimate $\hat{\theta}_{\mathcal{R}}$ is a reasonable approximation of θ_0 . Thus, the LLR statistic is given by:

$$L : \text{supp}(\mathbf{X}) \rightarrow \mathbb{R}, \quad L(\mathbf{x}) = \log \left[\frac{P(\mathbf{x} | \hat{\theta}_{\mathcal{T}})}{P(\mathbf{x} | \hat{\theta}_{\mathcal{R}})} \right]. \quad (11)$$

The test rejects H_0 whenever the LLR statistic is *large enough*. Specifically, the test operates as follows. Let α be the test Type I error, namely $\alpha = P(\phi(\mathbf{x}) = 1 | \mathbf{x} \notin \mathcal{T})$ represents the error incurred by the attacker when concluding that $\mathbf{x} \in \mathcal{T}$ when the opposite holds instead. The attacker first sets an *error level* α that can be tolerated, and estimates the distribution of $L(\cdot)$ under the null hypothesis H_0 . Next, a threshold $\tau(\alpha)$ is chosen, so that $P(L(\mathbf{x}) > \tau(\alpha) | \mathbf{x} \notin \mathcal{T}) = \alpha$.² Finally, the null hypothesis $H_0 : \mathbf{x} \notin \mathcal{T}$ is rejected in favor of H_1 whenever $L(\mathbf{x}) > \tau(\alpha)$. Otherwise, if $L(\mathbf{x}) \leq \tau(\alpha)$, there is insufficient evidence to conclude that $\mathbf{x} \in \mathcal{T}$. The *attack power* is defined as $\beta = P(\phi(\mathbf{x}) = 1 | \mathbf{x} \in \mathcal{T})$. It has been proved that the LLR test achieves the highest power among all other tests, for any given error level α [20].

Murakonda et al. [19] proved a relationship between α and β that is independent of the parameters $\hat{\theta}_{\mathcal{T}}$. Instead, this relationship can be deduced from the graphical structure and the sample size of the target population \mathcal{T} .

Theorem 1 ([19]). *For any error level α , it holds:*

$$z_\alpha + z_{1-\beta} \approx \sqrt{\frac{C(\mathcal{G})}{|\mathcal{T}|}}, \quad (12)$$

where z_s , $0 < s < 1$, is the quantile at level $1 - s$ of the Standard Normal distribution $\mathcal{N}(0, 1)$.

4 Theoretical Analysis

In this section, we adapt the notion of tracing attacks to CNs. Then, we address and provide answers to the following research questions:

- RQ1** Does releasing a CN lead to a lower tracing attack power than releasing a BN?
- RQ2** Are there some conditions under which the attacker can infer the true BN by observing the released CN?
- RQ3** Does the attack power depend on the CN uncertainty?

4.1 Attacking Credal Networks

When a BN is released, the attacker has access to $(\mathcal{G}, \hat{\theta}_{\mathcal{T}})$ and also to the reference population \mathcal{R} . Therefore, the MLE $\hat{\theta}_{\mathcal{R}}$ can be obtained

² If $F(\cdot)$ is the cumulative distribution function (CDF) of $L(\cdot)$ under H_0 , then $\tau(\alpha) = 1 - F^{-1}(\alpha)$.

from \mathcal{R} and the knowledge about the released DAG \mathcal{G} . On the contrary, when a CN is released, the attacker only has access to a *neighborhood* of $\hat{\theta}_{\mathcal{T}}$, namely the credal set $\hat{K}_{\mathcal{T}}(\mathbf{X}) \subset \Theta$ estimated from \mathcal{T} . Therefore, a new *uncertain* version of the tracing attack must be designed.

While the test hypotheses are the same as in Eq. (10), the LLR statistic now depends on the choice of $\theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})$. By using the properties of logarithms, we rewrite Eq. (11) and adapt it to the *uncertain case* as follows:

$$\begin{aligned} L &: \text{supp}(\mathbf{X}) \times \hat{K}_{\mathcal{T}}(\mathbf{X}) \rightarrow \mathbb{R}, \\ L(\mathbf{x}, \theta^K) &= \log[P(\mathbf{x} | \theta^K)] - \log[P(\mathbf{x} | \hat{\theta}_{\mathcal{R}})]. \end{aligned} \quad (13)$$

The decision rules related to the uncertain version of the test are then $\phi(\cdot, \theta^K)$, one for each $\theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})$. By setting $\theta^K = \hat{\theta}_{\mathcal{T}}$ in Eq. (13), we obtain the *precise* version of $L(\cdot)$ and $\phi(\cdot)$ as described in Section 3.

From the attacker perspective, a point $\theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})$ has to be selected for $\phi(\cdot, \cdot)$ so that, whenever $\phi(\mathbf{x}, \theta^K) = 1$ then $\phi(\mathbf{x}, \hat{\theta}_{\mathcal{T}}) = 1$. If this is the case, we say the test defined by $\phi(\cdot, \theta^K)$ is *consistent* with $\phi(\cdot, \hat{\theta}_{\mathcal{T}})$. If a random point is chosen within $\hat{K}_{\mathcal{T}}(\mathbf{X})$, it may coincide with $\hat{\theta}_{\mathcal{T}}$ by chance. However, any other point would represent the MLE estimate on some other subset $\mathcal{H} \subset \mathcal{P}$, where \mathcal{H} is unknown. In the worst case, $\mathcal{H} \cap \mathcal{T} = \emptyset$. Definition 6 provides the notion of a tracing attack against CNs.

Definition 6 (Tracing attack against CNs). *Let $\mathcal{T}, \mathcal{R} \subset \mathcal{P}$ the target and reference population respectively, and $(\mathcal{G}, \hat{K}_{\mathcal{T}}(\mathbf{X}))$ the CN estimated from \mathcal{T} . Let also $\hat{\theta}_{\mathcal{R}}^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})$ be the MLE estimate on \mathcal{R} , namely:*

$$\hat{\theta}_{\mathcal{R}}^K = \underset{\theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})}{\text{argmax}} \sum_{\mathbf{x} \in \mathcal{R}} P(\mathbf{x} | \theta^K). \quad (14)$$

We define the tracing attack against $(\mathcal{G}, \hat{K}_{\mathcal{T}}(\mathbf{X}))$ as the attack given by $L(\cdot, \hat{\theta}_{\mathcal{R}}^K)$ and $\phi(\cdot, \hat{\theta}_{\mathcal{R}}^K)$.

Theorem 2 proves the test's consistency in expectation. Lemmas 1 and 2 lead it.

Remark 2. The following results hold in the sample size limit, leveraging the asymptotic distribution of the log-likelihood function and the LLR statistic [20].

Lemma 1. $\mathbb{E} \left[L(\mathbf{x}, \hat{\theta}_{\mathcal{R}}^K) | \mathbf{x} \in \mathcal{T} \right] \leq \mathbb{E} \left[L(\mathbf{x}, \hat{\theta}_{\mathcal{T}}) | \mathbf{x} \in \mathcal{T} \right]$.

Proof. The second term of $L(\mathbf{x}, \theta^K)$, namely $\log P(\mathbf{x} | \hat{\theta}_{\mathcal{R}})$, is constant for any $\theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})$. Recalling $\hat{\theta}_{\mathcal{T}}$ is the MLE estimate over \mathcal{T} , we have:

$$\sum_{\mathbf{x} \in \mathcal{T}} L(\mathbf{x}, \theta^K) \leq \sum_{\mathbf{x} \in \mathcal{T}} L(\mathbf{x}, \hat{\theta}_{\mathcal{T}}), \quad \forall \theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X}). \quad (15)$$

The conclusion follows by exploiting Remark 2. \square

Notice that the inequality in Eq. (15) holds for any $\theta \in \Theta$ in the left-hand side term.

Lemma 2. *For any $\alpha \in [0, 1]$, it holds $\tau(\alpha, \hat{\theta}_{\mathcal{T}}) \leq \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K)$.*

Proof. Recalling $\hat{\theta}_{\mathcal{R}}^K$ is the MLE estimate over \mathcal{R} constrained to $\hat{K}_{\mathcal{T}}(\mathbf{X})$, we obtain:

$$\sum_{\mathbf{x} \in \mathcal{R}} L(\mathbf{x}, \hat{\theta}_{\mathcal{T}}) \leq \sum_{\mathbf{x} \in \mathcal{R}} L(\mathbf{x}, \hat{\theta}_{\mathcal{R}}^K) \leq \sum_{\mathbf{x} \in \mathcal{R}} L(\mathbf{x}, \hat{\theta}_{\mathcal{R}}) = 0, \quad (16)$$

where the first inequality exploits the fact that $\hat{\theta}_{\mathcal{T}} = \hat{\theta}_{\mathcal{T}}^K$. Assume two tracing attacks are conducted: one against the BN and the other against the CN, both with the same error level α . The former uses $L(\cdot, \hat{\theta}_{\mathcal{T}})$ as the test statistic, while the latter uses $L(\cdot, \hat{\theta}_{\mathcal{R}}^K)$. Under the null hypothesis $H_0 : x \notin \mathcal{T}$, the related thresholds are so that:

$$P(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{T}}) > \tau(\alpha, \hat{\theta}_{\mathcal{T}}) + l) = \alpha, \quad (17)$$

$$P(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{R}}^K) > \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K) + l) = \alpha, \quad (18)$$

where $l = \log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{R}})$ is a constant. Due to Eq. (16) and Remark 2, the two left-hand side terms of Eq. (17) and Eq. (18) are nested iff $\tau(\alpha, \hat{\theta}_{\mathcal{T}}) \leq \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K)$. \square

From Lemma 1 and Lemma 2 we obtain the consistency (in expectation) of Definition 6, summarized in Theorem 2. Practically speaking, whenever the attacker rejects the null hypothesis H_0 and concludes that $\mathbf{x} \in \mathcal{T}$ by using the released CN, the same conclusion would likely be reached if the attacker had instead utilized the BN. Note that this holds regardless of how the CN was learned.

Theorem 2. *In expectation over $\mathbf{x} \in \mathcal{T}$, it holds:*

$$L(\mathbf{x}, \hat{\theta}_{\mathcal{R}}^K) \geq \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K) \Rightarrow L(\mathbf{x}, \hat{\theta}_{\mathcal{T}}) \geq \tau(\alpha, \hat{\theta}_{\mathcal{T}}). \quad (19)$$

If the attacker picks randomly $\theta^K \in \hat{K}_{\mathcal{T}}(\mathbf{X})$, the consistency of the test (Definition 6) could not be proved. In that case, the inequalities in Eq. (16) may not hold.

4.2 Privacy Guarantees

This section investigates how privacy varies when releasing a CN instead of a BN. In particular, the following theorem answers the research question **RQ1**.

Theorem 3. *For any given value of the Type I error α , the power of the attack against the CN is no higher than that achieved against the BN. Specifically, for any $\alpha \in [0, 1]$, it holds $\beta(\hat{\theta}_{\mathcal{R}}^K) \leq \beta(\hat{\theta}_{\mathcal{T}})$.*

Proof. From the definition of the test power, it holds the following:

$$\begin{aligned} \beta(\hat{\theta}_{\mathcal{R}}^K) &= P(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{R}}^K) > \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K) + l \mid \mathbf{x} \in \mathcal{T}) = \\ &= \frac{1}{|\mathcal{T}|} \sum_{\mathbf{x} \in \mathcal{T}} \mathbb{I}(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{R}}^K) > \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K) + l) \leq \\ &\leq \frac{1}{|\mathcal{T}|} \sum_{\mathbf{x} \in \mathcal{T}} \mathbb{I}(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{T}}) > \tau(\alpha, \hat{\theta}_{\mathcal{R}}^K) + l) \leq \\ &\leq \frac{1}{|\mathcal{T}|} \sum_{\mathbf{x} \in \mathcal{T}} \mathbb{I}(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{T}}) > \tau(\alpha, \hat{\theta}_{\mathcal{T}}) + l) = \\ &= P(\log P(\mathbf{x} \mid \hat{\theta}_{\mathcal{T}}) > \tau(\alpha, \hat{\theta}_{\mathcal{T}}) + l \mid \mathbf{x} \in \mathcal{T}) = \\ &= \beta(\hat{\theta}_{\mathcal{T}}), \quad (20) \end{aligned}$$

where $\mathbb{I}(\cdot) \in \{0, 1\}$ is the *indicator function* of its argument, and the inequalities exploit Remark 2. \square

For the final part of this section, assume the attacker knows how the released CN was learned, for instance, whether the local application of IDM or the ε -contamination was employed. In this scenario, the privacy gained by using CNs may be compromised, allowing $\hat{\theta}_{\mathcal{T}}$ to be derived. Specifically, Lemmas 3 and 4 address the research question **RQ2**.

Lemma 3. *Consider a CN learned by the local application of IDM. The true BN can be recovered whenever the attacker knows S .*

Each parameter interval is of width $\frac{S}{N+S}$, whose value can be obtained by the difference of the interval extremes. Not only can $\hat{\theta}_{\mathcal{T}}$ be obtained, but the sample size $|\mathcal{T}|$ can also be inferred, providing the attacker with absolute frequencies. In the light of Lemma 3, knowledge about S and N must be kept secret while releasing the CN learned with IDM. Different S_X could be used as an additional protection technique, e.g., one for each random variable $X \in \mathbf{X}$.

Lemma 4. *Consider an ε -contaminated CN defined in Section 2.2. Then, the attacker can recover the underlying BN without additional information.*

By exploiting Eq. (9), each parameter interval is uniquely defined by the original parameter value and ε , which is the interval width. Trade-offs between the uncertainty included in the network and the privacy to be preserved depend on the specific case. Intuitively, larger parameter intervals lead to greater privacy, but also result in a decrease in the model's utility. Depending on S , or ε , we can mark two extreme situations:

1. $\hat{K}_{\mathcal{T}}(\mathbf{X}) \equiv \{\hat{\theta}_{\mathcal{T}}\}$. In this case, we release a single BN, and the power of a tracing attack is approximated by Theorem 1 [19]. Hence, we release an unmasked model, thereby obtaining the minimum privacy.
2. $\hat{K}_{\mathcal{T}}(\mathbf{X}) \equiv \Theta$. In this case, each parameter interval is $[0, 1]$, thus only the DAG is released. Under Definition 6, the LLR statistic is constantly zero. Hence, we release the least information while ensuring the highest privacy.³

5 Experiments

We conduct experiments to answer research question **RQ3**. In particular, we aim to empirically estimate how enhanced privacy, i.e., the power β of the test, varies based on: (i) the complexity of the DAG, and (ii) the uncertainty introduced in the CN. This section investigates the local application of IDM only, and not the ε -contaminated class of CNs, due to the pointed concerns regarding privacy breaches (Section 4.2). Hence, the uncertainty is represented by a global constant S (Section 2.2), and the attacker does not necessarily know how the CN was learned. We note that, even for large networks, the experiments remain computationally feasible. We make our source code available for reproducibility and research purposes [25].

5.1 Experimental Setting

We consider all the DAGs derived from the combinations of: number of nodes $M \in \{10, 20, 50, 100\}$; edge density $E \in \{1, 2, 4\}$. The latter indicates the average number of ingoing edges per node. The consideration of large networks stems from practical frameworks, such as omics data modeling [29]. The following presents the experimental pipeline for a given DAG \mathcal{G} with M nodes and $M \cdot E$ edges.

Data generation A BN (\mathcal{G}, θ_0) over binary variables is randomly generated and taken as the ground truth for the data generation process. Then, we sample 10 000 individuals from (\mathcal{G}, θ_0) and consider them as the general population \mathcal{P} . The reference and target population \mathcal{R}, \mathcal{T} are randomly sampled from \mathcal{P} with sample size 5 000 and 500, respectively. Subsampling from \mathcal{P} is performed 20 times to ensure variability in the sampling.

³ The DAG may still reveal private information [33], but it is secure against tracing attacks.

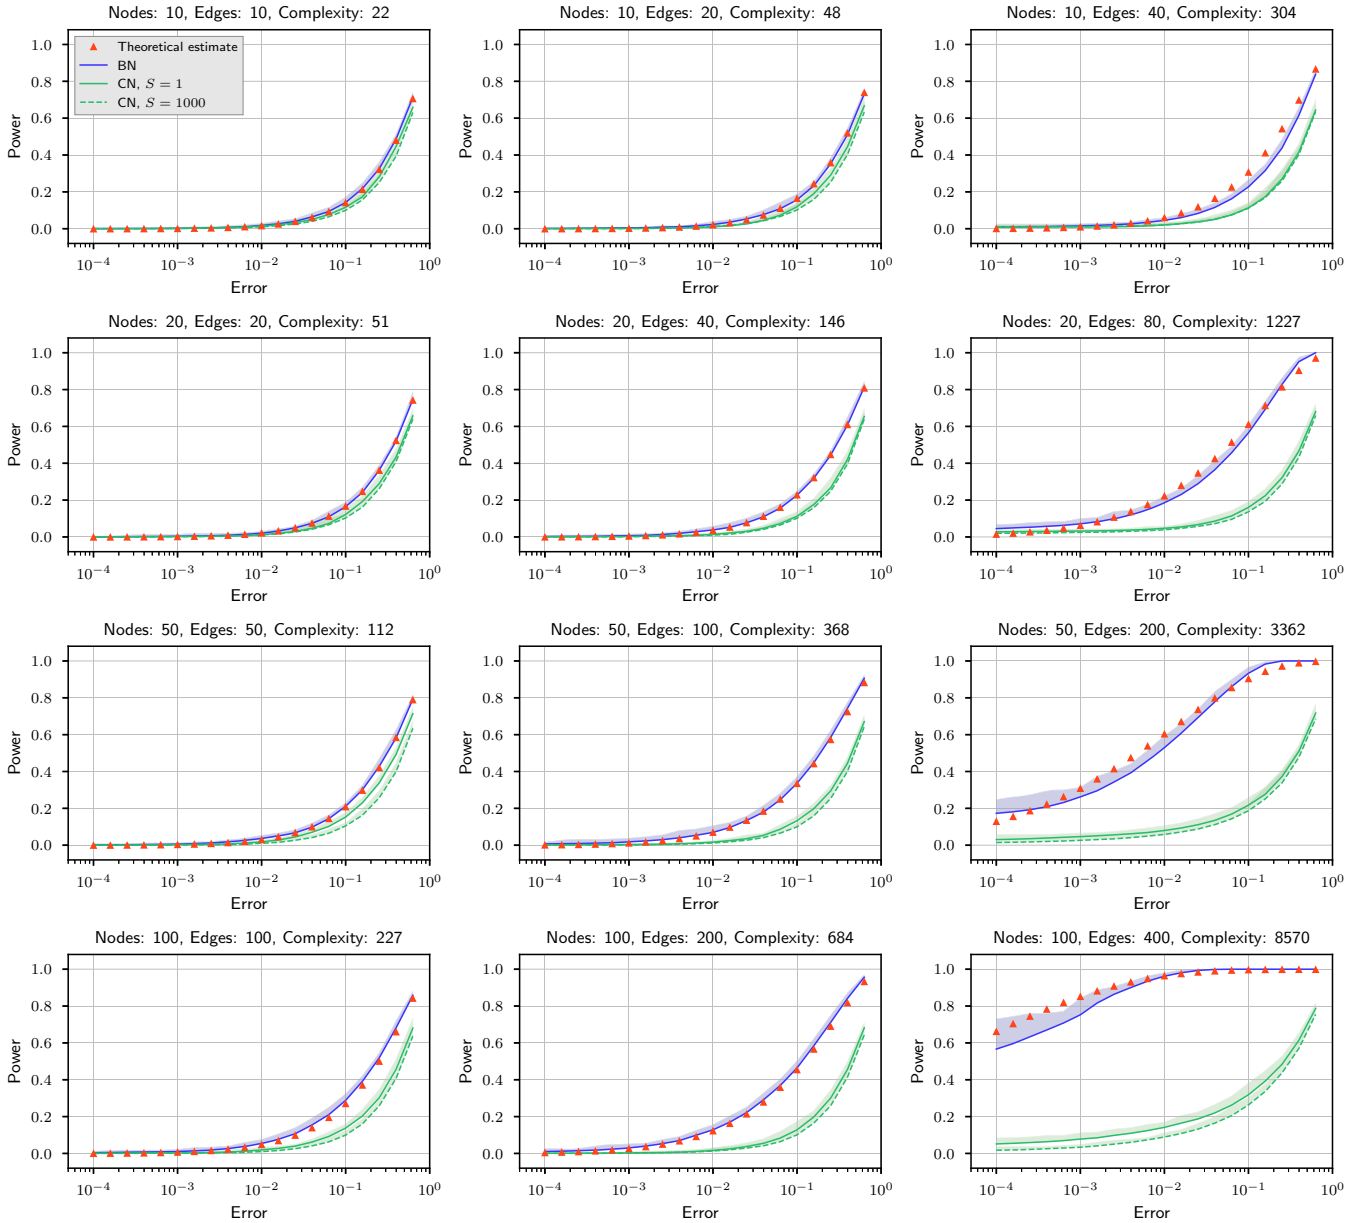


Figure 1. Experimental results show error vs. power rates for tracing attacks across various node (rows) and edge (columns) configurations. For each setup, \mathcal{T} and \mathcal{R} are sampled 20 times from \mathcal{P} . Solid lines represent average power, while shaded areas indicate the maximum power achieved.

Tracing attack against BN Given a sequence of Type I errors $\alpha_1, \dots, \alpha_k \in [10e^{-4}, 0.631]$, $k = 20$, the aim is to compute the corresponding power of a tracing attack, i.e. β_1, \dots, β_k . First, for each α_i , we compute the theoretical estimate of β_i as given by Theorem 1, by letting $C(\mathcal{G})$ be the complexity of \mathcal{G} and $|\mathcal{T}| = 500$ the target sample size. To conduct the attack against the BN, we perform the MLE estimation of $(\mathcal{G}, \hat{\theta}_{\mathcal{R}})$ and $(\mathcal{G}, \hat{\theta}_{\mathcal{T}})$ from \mathcal{R} and \mathcal{T} , respectively. Recall that $\hat{\theta}_{\mathcal{T}}$ is released by the modeler and is sensitive information, while $\hat{\theta}_{\mathcal{R}}$ is always accessible to the attacker. Then, for each α_i , we perform the tracing attack by using the statistics $L(\cdot, \hat{\theta}_{\mathcal{T}})$. In particular, for a fixed α_i , the test considers any data point $\mathbf{x} \in \mathcal{P}$ and assign either $\mathbf{x} \notin \mathcal{T}$ or $\mathbf{x} \in \mathcal{T}$. The true positive rate β_i is computed for any error level α_i .

Tracing attack against CN The tracing attack against the CN is performed similarly to the one described in the previous paragraph. The only difference consists in estimating $\hat{K}_{\mathcal{T}}(\mathbf{X})$, released by the modeler instead of $\hat{\theta}_{\mathcal{T}}$. The credal set is estimated from the target population for a fixed S . The attacker, having access to $\hat{K}_{\mathcal{T}}(\mathbf{X})$, needs to estimate $\hat{\theta}_{\mathcal{T}}^K$ to perform the attack as in Definition 6. For ease of computation, we estimate $\hat{\theta}_{\mathcal{R}}^K$ by sampling 500 points inside the credal set and taking the one giving the highest likelihood on \mathcal{R} . This procedure should be accurate enough, as we work with intervals over binary variables. Then, for each error level α_i , the tracing attack is performed on the whole \mathcal{P} by employing the statistics $L(\cdot, \hat{\theta}_{\mathcal{R}}^K)$, and the true positive rate β_i is computed.

5.2 Results & Discussion

Results are shown in Figure 1. We report the results on two CNs: those learned by the local application of IDM with $S \in \{1, 1000\}$. We also consider other configurations, namely $S \in \{10, 50, 100\}$. However, the corresponding attack powers are not significantly different from the power obtained with $S = 1$, because the sample size is relatively large in such a situation. Therefore, we choose not to plot them.

From the attacker's perspective, the *optimal curve* would be L-shaped towards the top-left corner; in other words, the attack that achieves the maximum power value $\beta = 1$ when the error level is set to its minimum value $\alpha = 0$. From the privacy-preserving perspective, instead, a graphical model ensures complete privacy when its curve is L-shaped towards the bottom-right corner of the plot; in other words, a power $\beta = 0$ for any Type I error $\alpha < 1$ and a power $\beta = 1$ only for $\alpha = 1$.

Our results demonstrate that CNs are not less private than BNs, as the green curves (representing CNs) are consistently below, or coincide with, the blue curves (representing BNs). This statement is further supported by the fact that the maximum power achieved with CNs does not intersect with the BN curves for all complexities above 146 and $\alpha > 10^{-2}$. Still, CNs are more private than BNs in most experiments. The area between the blue and green curves indicates the privacy gained by exposing the CN instead of the BN. We observe that this area increases as the network grows in complexity. Moreover, the desired L-shaped behavior, from the privacy-preserving perspective, appears to be achieved in very complex networks. There is, in fact, a significant gap between the BN and CN curves in the scenario with $M = 100$ nodes and $M \cdot E = 400$ edges. For error levels of $\alpha = 10^{-4}, 10^{-3}$, the attack power against the BN is approximately $\beta = 0.6, 0.8$, respectively, whereas the power against the CNs is around $\beta = 0, 0.1$, respectively.

The two considered CNs (those with $S = 1$ and $S = 1000$) exhibit similar behaviors, particularly at low complexities. However, we notice that choosing $S = 1000$ has a positive impact on privacy, albeit to a small extent. This suggests that smaller credal sets, which lead to less uncertain inferences, can be utilized while preserving similar privacy levels as more uncertain models. In other words, by introducing a small amount of uncertainty into the CN, we can achieve precise inferences (indicating high model's utility) while maintaining poor attack power (indicating high privacy). Thus, one could see a CN learned through the local IDM with $S = 1$ already as a reasonable balance between privacy and the model's utility.

Interestingly, for any value of the Type I error α , the power β of the attack against a CN appears consistent across all complexities: the power curve seems to share the same convex function across experiments, which is slightly elevated only for very complex networks. In contrast, the BN curves adhere to the theoretical estimate as expected, indicating that the model's complexity has a minimal impact on privacy in CNs.

6 Conclusions

This work introduces a promising and ready-to-use approach to the issue of privacy-aware parameter learning in BNs. In particular, we adopt the imprecise probability theory to obfuscate the BNs' parameters before their release; the degree of introduced uncertainty effectively hinders tracing attacks despite not neutralizing them completely. Contrary to a noisy BN, the released CN retains significance regarding the actual BN parameters. It thus keeps its utility for further

inferences, as one can use the CN to issue provably correct inferential bounds, which are impossible after adding noise to a BN.

Our framework is directly relevant to healthcare applications, where data privacy is paramount and data sets are often small and fragmented across institutions. It also aligns well with federated learning scenarios, where individual data sources, such as hospitals, collaboratively train AI models. While most of the existing federated learning research focuses on the structure learning task, little attention has been given to the parameter learning task, especially under population heterogeneity. We hope this work can begin a research direction to close that gap.

7 Future Work

We believe we can achieve sharper results than Theorem 2 and stronger guarantees than Theorem 3 by investigating the non-asymptotic behavior of the log-likelihood and LLR functions; the experimental results also support this argument. Nevertheless, the paper's insights are already applicable to practice and pave the way for a new line of research involving CNs and privacy. Moreover, we aim to identify other sensitive information than that provided in Section 4.2, to prevent the reconstruction of the masked BN. Motivated by our findings, a fruitful research direction involves providing a theoretical approximation of the attack effectiveness against CNs, similar to the one discussed in Murakonda et al. [19].

On the empirical side, we recognize the need for more ablation studies to evaluate the sensitivity of results across different networks, sample sizes, and uncertainty levels. Moreover, while CNs achieve higher utility than noisy BNs in providing correct inferential bounds, this is primarily a qualitative property. An empirical evaluation of the accuracy and calibration of both models is an ongoing work that we are addressing. Future work should also investigate the applicability of our framework to BNs with non-binary variables and alternative classes of CNs. Nevertheless, as supported by the literature, we firmly believe that the binary case of CNs has already revealed most of the interesting properties. Finally, we will evaluate and fine-tune our methodology using real-world data as soon as suitable and relevant scenarios become available.

Acknowledgements

This work was supported by the MUR under the grant "Dipartimenti di Eccellenza 2023-2027" of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca, Milan, Italy, and by the National Plan for NRRP Complementary Investments (Project n. PNC0000003 - AdvanCed Technologies for Human-centrEd Medicine (ANTHEM)). Niccolò Rocchi is funded by NextGeneration EU and Fondazione IRCCS Istituto Nazionale dei Tumori, Milan, Italy, in the Horizon Europe project IDEA4RC framework. Cassio de Campos thanks the support from the Eindhoven Artificial Intelligence Systems Institute, EU European Defence Fund via project KOIOS (EDF-2021-DIGIT-REFL-KOIOS), and the Dutch Research Council (NWO) via project NGF.1609.242.024.

References

- [1] A. Antonucci, C. P. de Campos, D. Huber, and M. Zaffalon. Approximate credal network updating by linear programming with applications to decision making. *Int. J. Approx. Reason.*, 58:25–38, 2015. ISSN 0888-613X. doi: 10.1016/j.ijar.2014.10.003.

- [2] J. Bailie and R. Gong. General inferential limits under differential and pufferfish privacy. *Int. J. Approx. Reason.*, 172:109242, 2024. ISSN 0888-613X. doi: 10.1016/j.ijar.2024.109242.
- [3] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, SIGMOD/PODS07*, pages 273–282. ACM, 2007. doi: 10.1145/1265530.1265569.
- [4] A. Bernasconi, A. Zanga, P. J. F. Lucas, M. Scutari, S. Di Cosimo, M. C. De Santis, E. La Rocca, P. Baili, I. Cavallo, P. Verderio, C. M. Ciniselli, S. Pizzamiglio, A. Blanda, P. Perego, P. Vallerio, F. Stella, A. Trama, and T. A. W. Group. From real-world data to causally interpretable models: A bayesian network to predict cardiovascular diseases in adolescents and young adults with breast cancer. *Cancers*, 16(21), 2024. doi: 10.3390/cancers16213643. URL <https://www.mdpi.com/2072-6694/16/21/3643>.
- [5] R. Binkyte, C. A. Pinzón, S. Lestyán, K. Jung, H. H. Arcolezi, and C. Palamidessi. Causal discovery under local privacy. In F. Locatello and V. Didelez, editors, *Proceedings of the Third Conference on Causal Learning and Reasoning*, volume 236 of *Proceedings of Machine Learning Research*, pages 325–383. PMLR, 2024. URL <https://proceedings.mlr.press/v236/binkyte24a.html>.
- [6] B. Bullek, S. Garboski, D. J. Mir, and E. M. Peck. Towards understanding differential privacy: When do people trust randomized response technique? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17*, pages 3833–3837. ACM, 2017. doi: 10.1145/3025453.3025698.
- [7] N. Carlini, S. Chien, M. Nasr, S. Song, A. Terzis, and F. Tramèr. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1897–1914, 2022. doi: 10.1109/SP46214.2022.9833649.
- [8] A. C. Constantinou, Z. Guo, and N. K. Kitson. The impact of prior knowledge on causal structure learning. *Knowl. Inf. Syst.*, 65(8):3385–3434, 2023. ISSN 0219-115-023-01858-x.
- [9] G. Corani and C. de Campos. A tree augmented classifier based on extreme imprecise dirichlet model. *Int. J. Approx. Reason.*, 51(9):1053–1068, 2010. ISSN 0888-613X. doi: 10.1016/j.ijar.2010.08.007. URL <https://www.sciencedirect.com/science/article/pii/S0888613X10001027>.
- [10] F. Cozman and C. de Campos. *Local computation in credal networks*, pages 5–11. ECAI, 2004. URL <https://research.tue.nl/en/publications/local-computation-in-credal-networks>.
- [11] F. G. Cozman. Credal networks. *Artif. Intell.*, 120(2):199–233, 2000. ISSN 0004-3702. doi: 10.1016/s0004-3702(00)00029-1.
- [12] C. P. de Campos and F. Cozman. Inference in credal networks through integer programming. In *Proceedings of the 5th International Symposium on Imprecise Probability: Theories and Applications*, pages 145–154, 2007.
- [13] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. *Journal of Privacy and Confidentiality*, 7(3):17–51, 2017. ISSN 2575-8527. doi: 10.29012/jpc.v7i3.405.
- [14] C. Dwork, A. Smith, T. Steinke, and J. Ullman. Exposed! a survey of attacks on private data. *Annual Review of Statistics and Its Application*, 4(1):61–84, 2017. ISSN 2326-831X. doi: 10.1146/annurev-statistics-060116-054123.
- [15] S. Feuerriegel, D. Frauen, V. Melnychuk, J. Schweisthal, K. Hess, A. Curth, S. Bauer, N. Kilbertus, I. S. Kohane, and M. van der Schaar. Causal Machine Learning for Predicting Treatment Outcomes. *Nat. Med.*, 30(4):958–968, 2024.
- [16] M. Grube, C. Reijnen, P. J. F. Lucas, F. Kommos, F. K. F. Kommos, S. Y. Brucker, C. B. Walter, E. Oberlechner, B. Krämer, J. Andress, F. Neis, A. Staebler, J. M. A. Pijnenborg, and S. Kommos. Improved preoperative risk stratification in endometrial carcinoma patients: external validation of the endorisk bayesian network model in a large population-based case series. *J. Cancer Res. Clin. Oncol.*, 149(7):3361–3369, Aug. 2022. doi: 10.1007/s00432-022-04218-4.
- [17] N. K. Kitson, A. C. Constantinou, Z. Guo, Y. Liu, and K. Chobtham. A survey of bayesian network structure learning. *Artif. Intell. Rev.*, 56(8):8721–8814, 2023. ISSN 1573-7462. doi: 10.1007/s10462-022-10351-w.
- [18] D. Koller and N. Friedman. *Probabilistic graphical models. Adaptive computation and machine learning*. MIT Press, Cambridge, Mass. [u.a.], [nachdr.] edition, 2010. ISBN 9780262013192.
- [19] S. Kumar Murakonda, R. Shokri, and G. Theodorakopoulos. Quantifying the privacy risks of learning high-dimensional graphical models. In A. Banerjee and K. Fukumizu, editors, *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 2287–2295. PMLR, 2021. URL <https://proceedings.mlr.press/v130/kumar-murakonda21a.html>.
- [20] E. Lehmann and J. P. Romano. *Testing Statistical Hypotheses*. Springer International Publishing, 2022. ISBN 9783030705787. doi: 10.1007/978-3-030-70578-7.
- [21] Q. Li, C. Zhou, B. Qin, and Z. Xu. Local differential privacy for belief functions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9):10025–10033, 2022. ISSN 2159-5399. doi: 10.1609/aaai.v36i9.21241.
- [22] D. D. Mauá and F. G. Cozman. Thirty years of credal networks: Specification, algorithms and complexity. *Int. J. Approx. Reason.*, 126:133–157, 2020. ISSN 0888-613X. doi: 10.1016/j.ijar.2020.08.009.
- [23] J. Pearl. *From Bayesian Networks to Causal Networks*, pages 157–182. Springer US, 1995. ISBN 9781489914248. doi: 10.1007/978-1-4899-1424-8_9.
- [24] M. Rigaki and S. Garcia. A survey of privacy attacks in machine learning. *ACM Computing Surveys*, 56(4):1–34, 2023. ISSN 1557-7341. doi: 10.1145/3624010.
- [25] N. Rocchi. Niccolo-rocchi/bn-privacy: Initial release, Aug. 2025. URL <https://doi.org/10.5281/zenodo.16693199>.
- [26] S. Roy, R. K. W. Wong, and Y. Ni. Directed cyclic graph for causal discovery from multivariate functional data. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 42762–42774. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/8549ab0f323b841955e70ca383b27d1-Paper-Conference.pdf.
- [27] M. Scanagatta, A. Salmerón, and F. Stella. A survey on bayesian network structure learning from data. *Lect. Notes. Artif. Int.*, 8(4):425–439, 2019. ISSN 2192-6360. doi: 10.1007/s13748-019-00194-y.
- [28] J. Shao, Z. Li, W. Sun, T. Zhou, Y. Sun, L. Liu, Z. Lin, Y. Mao, and J. Zhang. A survey of what to share in federated learning: Perspectives on model utility, privacy leakage, and communication efficiency, 2023.
- [29] P. Suter, E. Dazert, J. Kuipers, C. K. Y. Ng, T. Boldanova, M. N. Hall, M. H. Heim, and N. Beerenwinkel. Multi-omics subtyping of hepatocellular carcinoma patients using a bayesian network mixture model. *PLOS Comput. Biol.*, 18(9):e1009767, 2022. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1009767.
- [30] P. Torrijos, J. A. Gámez, and J. M. Puerta. *FedGES: A Federated Learning Approach for Bayesian Network Structure Learning*, pages 83–98. Springer Nature Switzerland, 2025. ISBN 9783031789809. doi: 10.1007/978-3-031-78980-9_6.
- [31] P. Walley. Inferences from multinomial data: Learning about a bag of marbles. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):3–34, 1996. ISSN 1467-9868. doi: 10.1111/j.2517-6161.1996.tb02065.x.
- [32] P. Walley. Towards a unified theory of imprecise probability. *Int. J. Approx. Reason.*, 24(2):125–148, 2000. ISSN 0888-613X. doi: 10.1016/S0888-613X(00)00031-1.
- [33] L. Wang, Q. Pang, and D. Song. Towards practical differentially private causal graph discovery. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5516–5526. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/3b13b1eb44b05f57735764786fab9c2c-Paper.pdf.
- [34] H. Xiao-xuan, W. Hui, and W. Shuo. Using expert’s knowledge to build bayesian networks. In *2007 International Conference on Computational Intelligence and Security Workshops (CISW 2007)*, pages 220–223, 2007. doi: 10.1109/CISW.2007.4425484.
- [35] M. Zaffalon. The naive credal classifier. *J. Stat. Plan. Infer.*, 105(1):5–21, 2002. ISSN 0378-3758. doi: 10.1016/s0378-3758(01)00201-4.
- [36] A. Zanga, E. Ozkirimli, and F. Stella. A survey on causal discovery: Theory and practice. *Int. J. Approx. Reason.*, 151:101–129, 2022. ISSN 0888-613X. doi: 10.1016/j.ijar.2022.09.004.
- [37] A. Zanga, A. Bernasconi, P. J. F. Lucas, H. Pijnenborg, C. Reijnen, M. Scutari, and F. Stella. Causal discovery with missing data in a multicentric clinical study. In J. M. Juares, M. Marcos, G. Stiglic, and A. Tucker, editors, *Artificial Intelligence in Medicine*, pages 40–44, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-34344-5.
- [38] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021. ISSN 0950-7051. doi: 10.1016/j.knosys.2021.106775.
- [39] J. Zhang, G. Cormode, C. M. Procopius, D. Srivastava, and X. Xiao. Privabases: Private data release via bayesian networks. *ACM Trans. Database Syst.*, 42(4):1–41, 2017. doi: 10.1145/3134428.