Empirical Bayesian Approaches for Robust Constraint-based Causal Discovery under Insufficient Data

Zijun Cui[1] , **Naiyu Yin**[1] , **Yuru Wang**[*2] and **Qiang Ji**[1]

[1]Rensselaer Polytechnic Institute
[2]Northeast Normal University

Abstract

Causal discovery is to learn cause-effect relationships among variables given observational data and is important for many applications. Existing causal discovery methods assume data sufficiency, which may not be the case in many real world datasets. As a result, many existing causal discovery methods can fail under limited data. In this work, we propose Bayesian-augmented frequentist independence tests to improve the performance of constraint-based causal discovery methods under insufficient data: 1) We firstly introduce a Bayesian method to estimate mutual information (MI), based on which we propose a robust MI based independence test; 2) Secondly, we consider the Bayesian estimation of hypothesis likelihood and incorporate it into a well-defined statistical test, resulting in a robust statistical testing based independence test. We apply proposed independence tests to constraint-based causal discovery methods and evaluate the performance on benchmark datasets with insufficient samples. Experiments show significant performance improvement in terms of both accuracy and efficiency over SOTA methods.

1 Introduction

Learning causal relations has been a fundamental and widely-investigated topic. The causal relations are captured by a directed acyclic graph (DAG), and a directed link in DAG captures cause-effect relation between two variables connected by the link [Glymour *et al.*, 2019]. Specifically, a directed link from node X to node Y indicates the cause-effect relation between cause variable X and effect variable Y. Causal discovery aims at learning a DAG capturing causal-effect relationships among a set of random variables from observational data, and one of the dominant approaches for causal discovery is through the structure causal model (SCM). Existing causal discovery methods focus on learning a DAG with high confidence from sufficient data samples [Yu *et al.*, 2019]. Not much attention, however, has been paid to performance improvement of causal discovery under limited data. Such

---
*Corresponding author

work is important, as even in the era of big data, there are still domains in which the availability of data is very limited. For example, in biological or clinical disciplines, data can be severely insufficient either because of high cost or lack of cases from which data is collected [Mukherjee and Speed, 2007]. Furthermore, even for applications with a vast amount of data, the data may not adequately cover all possible states of the nodes, leading to insufficient data for certain states. For example, the observed data under the absence of earthquake is adequate, while the observed data under the occurrence of earthquake is limited, due to the fact that earthquake rarely happens in nature.

In this paper, we employ constraint-based methods to learn a DAG through independence tests from observational data. Constraint-based causal discovery methods can be performed globally or locally. Global approaches aim at learning cause-effect relationships among all random variables, such as PC-stable [Colombo and Maathuis, 2014], and Sepset consistent PC (SC-PC) [Li *et al.*, 2019]. Global causal discovery methods discussed above learn DAGs that are in the same markov equivalent class of ground truth DAG. Further tests under certain assumptions about the graph or data distribution are needed to resolve the causal ambiguity [Glymour *et al.*, 2019]. In this paper, we focus on learning markov equivalent DAGs. In contrast to global approaches, local approaches identify the direct causes and effects of a target variable, represented by a causal Markov Blanket [Gao and Ji, 2015; Yang *et al.*, 2021]. A causal Markov Blanket captures local relationships of a target variable by identifying its parents, children, and spouses. For both global and local approaches, the main challenge of constraint-based causal discovery methods is that their performance highly depends on the accuracy of the independence test. Independence test error, even one mistake in independence decision, can propagate throughout the graph, causing a sequence of errors and resulting in an erroneous DAG with incorrect orientations [Spirtes, 2010]. Hence, to perform a robust constraint-based causal discovery, it is crucial to improve the robustness of the independence test.

To improve the causal discovery performance under insufficient data, we propose to introduce Bayesian approaches to independence tests for accurate and efficient constraint-based causal discovery. Specifically, two Bayesian-augmented frequentist independence tests are proposed, whereby we

use Bayesian approach to reliably estimate, under low data regime, independence test statistics used by frequentist independence tests. For both MI estimation (Sec.3.1) and hypothesis likelihood estimation (Sec.3.2), we employ Bayesian inference to calculate statistics by considering the entire parameter space instead of using a point estimate one. Given the estimated Bayesian statistics, we follow the standard frequentist framework to perform independence test. The proposed Bayesian-augmented independence tests are then applied to improve the constraint-based causal structure learning. We evaluate both local and global causal discovery performance with proposed independence tests on benchmark datasets and compare them to state-of-the-art methods. We empirically demonstrate the effectiveness of the proposed Bayesian approaches in improving both the accuracy and efficiency of the local and global causal discovery under insufficient data.

2 Related Work

To handle causal discovery under insufficient data, some methods downsize the problem domain to sub-domains. Rohekar *et al.*, [2020] approximated the structure by performing independence tests with a small fixed size of the condition set. The structure was then refined by iteratively increasing the condition set. A similar idea was explored in the Recursive Autonomy Identification (RAI) method [Yehezkel and Lerner, 2009]. Related works along this line always assume that there exist sufficient data for sub-domains. Besides, Claassen and Heskes [2012] estimated the posterior distribution of the independence hypothesis between two variables, based on which reliability was quantified. The causal discovery was then processed in decreasing order of reliability. Rohekar *et al.*, [2018] estimated posterior distribution of DAG through bootstrap samples. The negative effect from independence tests error was minimized through model averaging.

Some causal discovery methods address the limited data issue by directly improving the independence test [Marx and Vreeken, 2018]. A Bayesian-augmented frequntist independence test based on Bayes Factor (BF) was proposed [Natori *et al.*, 2017] whereby Bayesian parameter estimate is employed in computing BF while the value of BF is then applied to a frequentist independence test. The proposed independence test is incorporated into RAI, achieving competitive DAG learning performance. However, a threshold is required in [Natori *et al.*, 2017] and the selection of threshold can be heuristic. Instead, we propose to formulate the Bayes Factor into a well-defined statistical independence test without requiring threshold tuning.

In addition, different approaches have been proposed for robust independence tests under insufficient data. These methods, however, are not aimed at improving the causal discovery performance. Seok and Seon Kang [2015] improved the estimation of mutual information (MI) by partitioning the whole sample space into sub-regions. For better MI estimation under limited data, Bayesian approaches have been widely considered [Hutter, 2002; Archer *et al.*, 2013]. Besides, shrinkage estimators have also been employed [Sechidis *et al.*, 2019; Hausser and Strimmer, 2009]. Another category of recent independence test techniques are

focused on developing non-parametric methods to improve efficiency, such as CCIT [Sen *et al.*, 2017] and RCIT [Strobl *et al.*, 2019]. These works assume the availability of sufficient data and are mainly focused on continuous variables, while we are focused on discrete ones.

3 Proposed Methods

We consider two types of independence test: MI based and statistical testing based independence tests. We introduce Bayesian approaches to improve both types of independence tests through a Bayesian-augmented frequentist framework. Particularly, for MI based approach, we employ empirical Bayesian approach for better MI estimation under limited data. For statistical testing based approach, we consider the empirical Bayesian estimation of hypothesis likelihood and formulate it into χ^2 statistical independence test, providing an accurate p-value under limited data.

3.1 Bayesian Approach for Mutual Information Based Independence Test

The mutual information (MI) of two discrete random variables X and Y is defined as $MI(X;Y) = \sum_{i=1}^{K_x} \sum_{j=1}^{K_y} P(x_i, y_j) \log \frac{P(x_i,y_j)}{P(x_i)P(y_j)}$, where K_x and K_y denote the total number of possible states of X and Y respectively. $P(x_i, y_j)$, $P(x_i)$, and $P(y_j)$ represent the joint probability of (X, Y), and the marginal probabilities of X and Y respectively. By definition, $MI(X;Y) = 0$ if and only if X and Y are independent. In practice, the true MI is unknown, and the estimated MI is always larger than zero. In the following, we denote the probability distribution parameters as $\boldsymbol{\theta}$, i.e., $P(x_i) = \theta_i, P(y_j) = \theta_j$ and $P(x_i, y_j) = \theta_{ij}$. Conventionally, MLE is employed to estimate $\boldsymbol{\theta}$ from data as $\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} p(D|\boldsymbol{\theta})$, where $P(D|\boldsymbol{\theta})$ is the likelihood of parameter $\boldsymbol{\theta}$ given the data D. MI is then estimated as $MI = MI(X;Y|\hat{\boldsymbol{\theta}})$. When data is insufficient, MLE is not reliable [Geweke and Singleton, 1980] and MI tends to be overestimated. Instead, the full Bayesian MI is estimated from data over the entire parameter and hyper-parameter space, i.e.,

$$
\begin{aligned}
\hat{MI}^{fB} &= MI(X;Y|D) \\
&= \int \int MI(X;Y|\boldsymbol{\theta}, \alpha) p(\boldsymbol{\theta}, \alpha|D) d\boldsymbol{\theta} d\alpha \\
&= \int \int MI(X;Y|\boldsymbol{\theta}) p(\boldsymbol{\theta}|\alpha, D) p(\alpha|D) d\boldsymbol{\theta} d\alpha
\end{aligned}
\tag{1}
$$

where α is the hyper-parameter for symmetric Dirichlet prior of $\boldsymbol{\theta}$[1]. The full Bayesian MI is the expected MI over the joint posterior distribution of the parameters and hyper-parameter, i.e., $p(\boldsymbol{\theta}, \alpha|D)$. The integration over hyper-parameter α can be computationally challenging [Archer *et al.*, 2013]. Instead of marginalizing out α, we propose to maximize it out. Particularly, we approximate the integration over the hyper-parameter space by its mode α^* that maximizes a posterior

---

[1]As we have no prior preference on the elements of the Dirichlet distribution, we assume symmetric Dirichlet distribution, i.e., each entry in $\boldsymbol{\alpha}$ shares the same value and we denote it as α.

(MAP) of α, i.e., $\alpha^* = \arg\max_\alpha p(\alpha|D)$. By assuming uniform distribution of $p(\alpha)$, we have $\alpha^* = \arg\max_\alpha p(D|\alpha)$. The likelihood $p(D|\alpha)$ can be computed as (See Appx.A for details),

$$p(D|\alpha) = N! \frac{\Gamma(K\alpha)}{\Gamma(K\alpha + N)} \prod_{i=1}^{K} \frac{\Gamma(\alpha + n_i)}{\Gamma(\alpha)n_i!} \quad (2)$$

where K is the number of states for the random variable, n_i is the number of samples for state i, and $N = \sum_i^K n_i$. $P(D|\alpha)$ follows Polya distribution and $\Gamma(x)$ is the gamma function. We solve for α^* with a fixed-point update [Minka, 2000].

Given α^*, the full Bayesian method is converted to the empirical Bayesian method, and we have the proposed empirical Bayesian MI $\hat{MI}^{eB}$ defined as,

$$\hat{MI}^{eB} = \int MI(X;Y|\boldsymbol{\theta})p(\boldsymbol{\theta}|D,\alpha^*)d\boldsymbol{\theta} \quad (3)$$

with a closed-form solution (See Appx.A for details):

$$\hat{MI}^{eB} = \psi(N + \alpha^*K + 1) - \sum_{ij} \frac{n_{ij} + \alpha^*}{N + \alpha^*K}[\psi(n_i + \alpha^*K_y + 1)$$
$$+ \psi(n_j + \alpha^*K_x + 1) - \psi(n_{ij} + \alpha^* + 1)$$
$$(4)$$

where $\psi(x)$ is the digamma function. n_i and n_j are the number of samples for $X = i$ and $Y = j$ respectively, and n_{ij} is the number of samples for $(X,Y) = (i,j)$. The closed-form solution to empirical Bayesian estimation of MI was firstly presented in [Hutter, 2002], based on which we propose our approach. Our contribution lies in automatically estimating α^* by maximizing $p(\alpha|D)$ instead of manually selecting α as in Hutter's method. Given the estimated MI, we compare it against a pre-defined threshold for independence test. If MI is smaller than the threshold, two random variables will be declared to be independent, and dependent otherwise.

3.2 Bayesian Approach for Statistical Testing Based Independence Test

We now introduce our proposed Bayesian approach to improve the statistical testing based independence test. We firstly consider a standard independence test, G test [McDonald, 2009], which is a likelihood ratio test with null hypothesis assuming two random variables are independent. G test is a widely used statistical test. As the same with other statistical tests, G test doesn't require threshold tuning and the significance level is set to be 5% by default. The formula for the statistic G reads as $G = -2\sum_{i=1}^{K_x}\sum_{j=1}^{K_y} n_{ij} \ln \frac{\hat{\theta}_i \hat{\theta}_j}{\hat{\theta}_{ij}}$ where $\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} P(D|\boldsymbol{\theta})$. Samples $D = \{D_n\}_{n=1}^N$ are $i.i.d$ given parameter $\hat{\boldsymbol{\theta}}$ and the statistic G follows asymptotic $\chi^2_{df=(K_x-1)(K_y-1)}$ distribution, based on which a statistical test can be performed (See Appx.B for detailed derivations). As MLE parameter estimates are not reliable under insufficient data, leading to inaccurate estimation of the likelihood of hypothesis, we instead consider the empirical Bayesian estimation. Specifically, we employ the Bayes Factor (BF) [Kass and Raftery, 1995] which defines the ratio of

expected likelihoods of null hypothesis(H_0) and that of the alternative hypothesis(H_1) over all possible parameter settings with the posterior distributions of parameters under null and alternate hypothesis respectively,

$$BF = \frac{P(D|H_0, \alpha^0)}{P(D|H_1, \alpha^1)} = \frac{\int P(D|\boldsymbol{\theta})P(\boldsymbol{\theta}|H_0, \alpha^0)d\boldsymbol{\theta}}{\int P(D|\boldsymbol{\theta})P(\boldsymbol{\theta}|H_1, \alpha^1)d\boldsymbol{\theta}} \quad (5)$$

where α^0 and α^1 are the hyper-parameters for the symmetric Dirichlet prior under null and alternate hypothesis respectively. Both hypothesis likelihoods $P(D|H_0, \alpha^0)$ and $P(D|H_1, \alpha^1)$ can be analytically solved, and BF can be computed (See Appx.C for detailed derivations). However, BF can't be directly applied to a statistical test because samples $D = \{D_n\}_{n=1}^N$ are not $i.i.d$ given hyper-parameter α and BF no longer follows the χ^2 distribution under the null hypothesis. Detailed discussions on this are in Appx.C. Instead, we propose to approximate $P(D|\alpha)$ by a multinomial distribution and calculate modified parameters of multinomial distribution with α taken into account, as both capture the distributions for integer random variables, i.e.,

$$P(D|\alpha) \approx P(D|\tilde{\boldsymbol{\theta}}) = \frac{N!}{\prod_{i=1}^K n_i!} \prod_{i=1}^K \tilde{\theta}_i^{n_i} \quad (6)$$

where K is the total number of states, and $N = \sum_{i=1}^K n_i$ is the total number samples with n_i being the number of samples for state i. $\tilde{\theta}_i$ are the modified parameters of the multinomial distribution. $\tilde{\theta}_i = \frac{g(n_i,\alpha)}{g(N,K\alpha)}$ with $g(n_i,\alpha) = an_i + b\alpha$ where $\Lambda = \begin{pmatrix} a \\ b \end{pmatrix}$ are unknown coefficients. In summary, the motivations for the proposed approximation in Eq. 6 are two folds: 1) we formulate BF into a statistical test whereby threshold can be automatically decided; 2) different from G statistic using MLE parameter estimates, we use modified parameter estimates with prior α incorporated. By plugging the $P(D|\alpha)$ (defined in Eq. 2) into Eq. 6, it is clear that to satisfy Eq. 6, we must have $n_i \ln g(n_i,\alpha) = \ln\Gamma(n_i + \alpha) - \ln\Gamma(\alpha)$. Given $\{n_i\}_{i=1}^K$ and α, we can construct a system of K such equations through which we can solve for Λ^*, i.e.,

$$\Lambda^* = \arg\min_\Lambda ||M\Lambda - T||_2^2 = (M^tM)^{-1}M^tT \quad (7)$$

with $M = \begin{pmatrix} n_1, \alpha \\ n_2, \alpha \\ ... \\ n_K, \alpha \end{pmatrix}$, $T = \begin{pmatrix} t(n_1, \alpha) \\ t(n_2, \alpha) \\ ... \\ t(n_K, \alpha) \end{pmatrix}$, and $t(n_i, \alpha) = \exp(\frac{1}{n_i}(\ln\Gamma(n_i + \alpha) - \ln\Gamma(\alpha)))$. Given Λ^*, we have $\tilde{\theta}_i$ as $\tilde{\theta}_i = \frac{g(n_i,\alpha)}{g(N,K\alpha)} = \frac{a^*n_i + b^*\alpha}{a^*N + b^*K\alpha}$. $P(D|\tilde{\boldsymbol{\theta}})$ can well approximate $P(D|\alpha)$ (Empirical justifications can be found in Appx.D). Our proposed approximation is different from the method provided in [Minka, 2000] where the Polya distribution $P(D|\alpha)$ is interpreted as a multinomial distribution with modified counts $\tilde{n}_i$. In addition, our proposed estimation can better approximate the Polya distribution given the symmetric Dirichlet prior (Detailed derivations and empirical evaluation are in Appx.D) compared to [Minka, 2000]. We then approximate the hypothesis likelihood under null and alternative hypothesis respectively and obtain a modified Bayes Factor $\tilde{BF}$

$$\tilde{BF} = \frac{P(D|\tilde{\boldsymbol{\theta}}, H_0)}{P(D|\tilde{\boldsymbol{\theta}}, H_1)} = \frac{\prod_{i=1}^{K_x} \tilde{\theta}_i^{n_i} \prod_{j=1}^{K_y} \tilde{\theta}_j^{n_j}}{\prod_{i=1, j=1}^{K_x K_y} \tilde{\theta}_{ij}^{n_{ij}}} \quad (8)$$

We obtain the statistic BF_{chi2} for the statistical test as,

$$BF_{chi2} = -2\ln \tilde{BF} = -2 \sum_{i=1}^{K_x} \sum_{j=1}^{K_y} n_{ij} \ln \frac{\tilde{\theta}_i \tilde{\theta}_j}{\tilde{\theta}_{ij}} \quad (9)$$

The statistic BF_{chi2} asymptomatically follows the $\chi^2_{df=(K_x-1)(K_y-1)}$ distribution (Details are in Appx.D). If p-value is smaller than the significance level, we reject the null hypothesis and accept the alternative hypothesis. It is worth noting that BF can be directly applied for a frequentist independence test where a pre-defined threshold η is required [Natori *et al.*, 2017]. The value of the threshold is unconstrained [Kass and Raftery, 1995] making it hard to be properly selected. Instead, our approach only requires a significant level for independence test which is usually set to be 5% by default. Like conventional likelihood ratio test, our method indeed requires the asymptotic assumption. But with the use of Bayesian estimation, our method is less reliant on asymptotic assumption as demonstrated by experiments. As conventional likelihood ratio test is sound, our method should also be sound.

It is generally believed that Bayesian approach for parameter estimation is better than MLE under insufficient data [Kruschke, 2013], which motivates our Bayesian-augmented frequentist approaches. We theoretically show that Bayesian estimation is always better than MLE for parameter estimation with smaller estimation variance (Details are in Appx.E). Through exhaustive experiments, we further empirically demonstrate the robustness of Bayesian approaches under limited data through improved performance on both independence test and causal discovery. While our discussion focuses on marginal independence test, our methods can be straightforwardly applied to conditional independence test as they share the same mechanism. In fact, when applied to causal discovery, our methods are applied to primarily perform conditional independence tests.

4 Experiments

We evaluate both the local and global constraint-based causal discovery performance on benchmark datasets. Our work is to improve independence tests, so as to improve causal discovery under insufficient data. We thus focus our evaluations on constraint-based methods. Through exhaustive experiments, we show that our approaches can significantly improve causal discovery performance in terms of both accuracy and efficiency over state-of-the-art methods. Besides, we compare proposed independence tests to state-of-the-art independence tests to further show the effectiveness of the proposed methods.

Experiment Settings. We employ six benchmark datasets[2] that are widely used for causal discovery evaluation: CHILD, INSURANCE, ALARM, HAILFINDER, CHILD3

---

[2]https://www.bnlearn.com/bnrepository/.

and CHILD5. Statistical information of datasets are in Appx. F. The causal discovery performance is evaluated in terms of both accuracy and efficiency. For accuracy, we employ the structural hamming distance (SHD) [Tsamardinos *et al.*, 2006a]. SHD computes the number of extra and incorrect (missing and reverse) edges in the learned causal structure compared to the ground truth one. For efficiency, we consider the number of conducted independence test. We perform evaluation on a number of small sized datasets. These small sample sizes are chosen to mimic insufficient data scenario through significantly small number of samples per configuration. For each sample size, we repeat 10 runs and report the averaged performance over 10 runs. In addition, we report standard derivation of SHD. All the experiments are performed on a laptop with a 2.3 GHz 8-Core Intel Core i9 processor using CPU only (Specific running time can be found in Appx. F).

4.1 Local Constraint-based Causal Discovery

For the local causal discovery, we employ Causal Markov Blanket (CMB) [Gao and Ji, 2015], which is the state-of-the-art method. CMB employs constraint-based approach, which performs conditional independence test using MI to identify the CMB of a target node. We incorporate the proposed independence tests into CMB and compare to the original CMB. We denote cI^{eB} as the CMB with empirical Bayesian MI estimation and cBF_{chi2} as CMB with BF_{chi2} independence test. SHD is 0 if learned CMB is identical to the ground truth CMB. Details on algorithm settings (e.g., hyper-parameters) are in Appx. F.

From Table 1, we can see that both cI^{eB} and cBF_{chi2} outperform the CMB on all datasets in terms of both accuracy and efficiency under insufficient data. The number of performed independence test reduces dramatically. On ALARM dataset, cI_{eB} only performs 2210 independence tests on average, while CMB requires 17907 tests on average. The proposed methods improve the accuracy significantly. On INSURANCE dataset, cBF_{chi2} improves the averaged SHD by 3.92 compared to CMB. From the results we can see that, by introducing Bayesian approaches, both the accuracy and the efficiency can be improved. Comparing the performance between the two proposed methods, cBF_{chi2} achieves overall better accuracy, and cI^{eB} is more efficient with the fewest number of independence test on all datasets.

It is worth noting that the number of independence test increases with reduced samples in CMB, but decreases with the proposed methods. The reason is that under insufficient data, MLE will lead to an overestimated MI. Hence, conventional MI based independence test is likely to declare dependence when data size is small, resulting in a large number of independence test. As the sample size increases, the incorrect dependency declarations will be corrected and the number of independence tests will decrease. On the other hand, our methods are more accurate and show a preference of independence under insufficient data, resulting a small number of performed independence test.

Dataset	Size	SHD			#Independence Test		
		cI^{eB}	cBF_{chi2}	CMB	cI^{eB}	cBF_{chi2}	CMB
CHILD	100	2.90±0.28	2.65±0.40	5.94±0.65	1008	1154	16869
	300	2.61±0.26	2.64±0.59	6.95±0.63	1709	1926	14578
	500	2.29±0.31	2.24±0.84	4.52±0.58	2524	4751	13873
	MEAN	2.60	**2.51**	5.80	**1747**	2610	15107
INSURANCE	100	3.89±0.34	3.98±0.39	7.18±0.66	1261	1363	22168
	300	3.47±0.21	3.24±0.12	7.59±0.57	1541	2977	18043
	500	3.11±0.21	2.98±0.13	7.20±0.67	1477	3949	14881
	MEAN	3.49	**3.40**	7.32	**1426**	2763	18364
ALARM	100	2.69±0.07	2.39±0.19	5.20±0.71	1424	1109	27492
	300	2.50±0.19	2.27±0.15	4.36±0.83	2398	3885	14900
	500	2.40±0.11	2.26±0.19	3.53±0.62	2807	4766	11328
	MEAN	2.53	**2.31**	4.36	**2210**	3253	17907
HAILFINDER	500	3.33±0.02	4.22±0.04	7.90±0.11	676	1923	183350
	800	3.56±0.01	4.49±0.13	7.12±0.09	1098	2145	169705
	1000	3.56±0.09	4.45±0.08	7.10±0.11	1924	2621	119815
	MEAN	**3.48**	4.39	7.37	**1233**	2229	157620
CHILD3	500	2.46±0.23	2.53±0.18	4.72±0.28	7168	7417	14789
	800	3.01±0.13	2.67±0.11	3.57±0.21	6720	7802	9765
	1000	2.90±0.07	2.57±0.23	3.09±0.19	8424	8285	9516
	MEAN	2.79	**2.59**	3.79	**7437**	7835	11357
CHILD5	500	2.87±0.05	2.62±0.19	5.00±0.15	5234	11126	16819
	800	2.66±0.21	3.02±0.13	5.75±0.32	8236	11424	51967
	1000	2.82±0.23	2.99±0.07	4.34±0.19	13384	9956	36888
	MEAN	**2.78**	2.88	5.03	**8951**	10835	26322

Table 1: Local causal discovery performance under insufficient data

4.2 Global Constraint-based Causal Discovery

Majority of global causal discovery algorithms are under causal sufficiency assumption, whereby all random variables are observed in data and there is no latent variable. However, causal sufficiency assumption can be violated since the real data may fail to capture the values for all the variables, leaving some variables to be latent. To address this issue, several recent causal discovery methods [Ramsey *et al.*, 2012; Colombo *et al.*, 2012] have been developed to identify latent common confounders of the observed variables. In our evaluations, we mainly focus on standard algorithms that are under causal sufficiency assumptions. We firstly employ RAI [Yehezkel and Lerner, 2009] as our baseline and compare to two state-of-the-art methods. Then, to demonstrate that our proposed methods can consistently improve causal discovery performance, we consider well-known DAG learning algorithms: PC [Spirtes *et al.*, 2000] and MMHC [Tsamardinos *et al.*, 2006b] as two additional baselines. In the end, we consider the algorithms without causal sufficiency assumption to demonstrate that our proposed methods can be applied to different causal discovery methods, independent of the existence of latent confounders.

Global causal discovery with causal sufficiency assumption. We employ RAI as our baseline algorithm and incorporate the proposed independence tests. We denote rI^{eB} as the RAI with empirical Bayesian MI estimation, and rBF_{chi2} as RAI with BF_{chi2} independence test. We compare our approaches to two state-of-the-art methods: RAI-

BF method [Natori *et al.*, 2017] and PC-stable [Colombo and Maathuis, 2014]. SC-PC[3] can't be performed under insufficient data smoothly, and thus we exclude this method for comparison. SHD is 0 if the learned DAG and the ground truth DAG belong to the same equivalence class. Details on algorithm settings (e.g., hyper-parameters) are in Appx. F.

From Table 2, we can see that rBF_{chi2} outperforms RAI-BF and PC-stable on almost all datasets in terms of both accuracy and efficiency. rI^{eB} also achieves overall better accuracy and significantly improves efficiency. For example, on CHILD3, rBF_{chi2} improves the SHD by 21.7 and 27.6 compared to RAI-BF and PC-stable. In terms of efficiency, on HAILFINDER, rI^{eB} only performs 2215 independence tests in average, while RAI-BF requires 10212 tests in average. Comparing between the two proposed methods, rBF_{chi2} achieves better accuracy and rI^{eB} achieves better efficiency. With the proposed methods, the number of independence tests decreases due to the reduced samples for all datasets, which is consistent with the conclusion we have from the local causal discovery. In addition, both RAI-BF and PC-stable show a preference of independence under insufficient data, leading to the decreased number of independence tests with reduced number of samples.

Since BF_{chi2} essentially is only an approximate of original BF, BF with the optimal threshold should outperform BF_{chi2} in principle. However, selecting the optimal threshold for BF can be challenging and incorrect thresholds can lead to

---

[3]https://github.com/honghaoli42/consistent_pcalg.

Dataset	Size	SHD				#Independence Test			
		rI^{eB}	rBF_{chi2}	RAI-BF	PC-Stable	rI^{eB}	rBF_{chi2}	RAI-BF	PC-Stable
CHILD	100	21.6±2.1	24.2±2.3	30.4±3.7	23.8±1.7	283	314	893	559
	300	19.9±2.7	17.7±1.8	23.5±4.4	22.6±1.9	342	546	997	986
	500	17.6±1.7	16.0±2.9	22.6±2.4	24.4±2.2	424	754	975	1317
	MEAN	19.7	**19.3**	25.5	23.6	**350**	538	955	954
INSURANCE	100	48.9±1.3	50.1±2.9	54.9±3.6	52.0 ±1.5	486	604	905	1217
	300	47.3±0.8	44.5±2.0	46.6±3.2	50.2±3.1	576	986	1011	1250
	500	49.5±1.8	39.4±3.0	47.1±2.2	50.7±2.5	662	1200	1120	2326
	MEAN	48.6	**44.7**	49.5	51.0	**575**	930	1012	1598
ALARM	100	44.5±2.2	42.7±2.3	48.4±5.8	45.8 ±4.9	891	958	1591	2215
	300	40.7±3.0	36.1±4.5	35.3±5.4	34.6±2.7	1158	1752	1881	3398
	500	40.0±3.1	29.8±5.1	29.8±5.2	36.5 ±5.7	1433	2018	2098	3992
	MEAN	41.7	**36.2**	37.8	39.0	**1161**	1576	1857	3202
HAILFINDER	500	88.0±2.0	98.3 ±1.5	118.0±1.0	91.6±1.0	2024	2587	6171	3267
	800	85.0±1.7	106.3 ±2.1	124.7 ±6.7	99.7±1.2	1983	3726	7847	3423
	1000	92.3±4.5	108.3 ±2.3	131.3 ±3.2	101.8±2.2	2638	3073	16618	3603
	MEAN	**88.4**	104.3	124.7	97.7	**2215**	3129	10212	3431
CHILD3	500	67.6±3.2	54.3±2.6	79.6±4.9	81.2±2.8	2693	3796	5422	4963
	800	65.8±2.5	52.9±2.8	74.0±3.7	79.9±2.4	3941	4587	5106	6026
	1000	61.5±3.8	52.3±3.9	71.0±6.5	81.4±2.7	4723	5170	5980	6846
	MEAN	65.0	**53.2**	74.9	80.8	**3786**	4518	5503	5945
CHILD5	500	122.0±2.6	109.3±5.1	134.0 ±2.6	113.9±2.4	6966	8646	10038	10253
	800	121.7±3.8	105.3±4.0	132.3±6.7	120.1±2.9	10249	10431	9337	10708
	1000	116.3±2.9	105.7±2.5	126.3±7.0	123.4±1.7	10375	10494	11174	11070
	MEAN	120.0	**106.8**	126.3	119.1	**9197**	9857	11174	10677

Table 2: Global causal discovery performance under insufficient data

inferior causal discovery performance. Instead of fixing the threshold of RAI-BF with its default value, we consider the optimal performance of RAI-BF with tuned thresholds for comparison. According to the results (details can be found in Appx. F), RAI-BF with the optimally tuned threshold at best achieves comparable performance compared to rBF_{chi2} in terms of both accuracy and efficiency, which is expected. While rBF_{chi2} only requires a fixed significance level (5% by default) without additional tuning process.

To further show that our proposed methods can consistently improve the causal discovery performance, we consider another two widely used DAG learning algorithms: PC [Spirtes *et al.*, 2000] and MMHC [Tsamardinos *et al.*, 2006b]. We incorporate the proposed methods into PC and MMHC for evaluation. From results (details are in Appx. F), our proposed methods can consistently improve the DAG learning performance, particularly with PC. For example, on ALARM, PC with BF_{chi2} achieves averaged SHD 40.5, while PC only achieves averaged SHD 58.2. Overall, BF_{chi2} achieves better accuracy and I^{eB} achieves better efficiency with both PC and MMHC on different datasets.

Global causal discovery without causal sufficiency assumption. To demonstrate that our robust independent tests can also be applied to causal discovery without causal sufficiency assumption, we employ the conservative FCI (cFCI) method [Ramsey *et al.*, 2012] as our baseline. cFCI is considered as the state-of-the-art causal discovery algorithm that identifies latent confounders. We denote cI^{eB} as the cFCI with empirical Bayesian MI estimation, and cBF_{chi2} as cFCI with BF_{chi2} independence test. We compare our approaches

Dataset	SHD			#Independence Test		
(**MEAN**)	cI^{eB}	cBF_{chi2}	cFCI	cI^{eB}	cBF_{chi2}	cFCI
CHILD	49.1	**35.1**	50.4	**109**	417	2289
INSURANCE	118.7	**94.9**	121.3	**147**	593	4836
ALARM	105.3	**78.2**	94.7	**397**	902	7361
HAILFINDER	**153.2**	220.2	339.0	**368**	2024	82683
CHILD3	204.3	135.2	**103.6**	**692**	2858	4009
CHILD5	250.7	**159.0**	178.8	**1145**	5161	7068

Table 3: Global causal discovery performance (with latent confounder) under insufficient data

to cFCI with default g^2 statistical based independence test[4]. As we can see from Table 3, cI^{eB} achieves best efficiency by performing the smallest number of independence tests. In terms of accuracy, cBF_{chi2} achieves overall better performance. The consistent performance improvement further demonstrates that the proposed independence test can improve the causal discovery performance under insufficient data, independent of the existence of latent confounders.

4.3 Bayesian Approaches for Independence Tests

To compare the proposed independence tests to state-of-the-art methods, we firstly perform a direct evaluation of proposed independence tests on synthetic data, and we then compare to state-of-the-art methods in terms of causal discovery performance on benchmark datasets. On synthetic data, we compare to three state-of-the-art independence tests: adaptive partition [Seok and Seon Kang, 2015], empirical Bayesian with fixed α [Hutter, 2002] and full Bayesian method [Archer *et al.*, 2013]. We evaluate the performance in terms of both

---

[4]https://github.com/striantafillou/causal-graphs.

accuracy and efficiency. Experimental results show that the proposed methods achieve better accuracy with significantly improved efficiency. Detailed experiment settings and results can be found in Appx.G. More importantly, we compare proposed independence tests to two state-of-the-art methods: adaptive partition and empirical Bayesian with fixed α methods in terms of causal discovery performance on benchmark datasets. Because the full Bayesian method is of high computational complexity, making it impractical to be applied to constraint-based causal discovery, we exclude the comparison to this method. We incorporate the adaptive partition method and the empirical Bayesian with fixed α method to RAI (denoted as rI^{AdP} and rI^{eBFix} respectively). As

Dataset	SHD				
(MEAN)	rI^{AdP}	rI^{eBFix}	rI^{eB}	rBF_{chi2}	RAI-BF
CHILD	26.5	23.9	19.7	**19.3**	25.5
INSURANCE	53.2	49.1	48.6	**44.7**	49.5
ALARM	46.9	40.9	41.7	**36.2**	37.8
HAILFINDER	**70.8**	91.2	88.4	104.3	124.7
CHILD3	81.6	66.3	65.0	**53.2**	74.9
CHILD5	129.9	121.6	120.0	**106.8**	126.3

Table 4: Accuracy comparison to SoTA independence tests

shown in Table 4, our methods achieve overall better accuracy than rI^{AdP} and rI^{eBFix} on different datasets. For example, on CHILD3, rBF_{chi2} achieves averaged SHD 53.2, significantly better than rI^{AdP} which achieves averaged SHD 81.6. In terms of efficiency evaluation, rI^{eB} also achieves competitive efficiency (details are in Appx. G). Overall, our proposed methods outperform other SoTA independence tests in terms of causal discovery performance. On HAILFINDER, because rI^{AdP} tends to declare independence, the learned DAG contains fewer false positive edges compared to other methods and thus its averaged SHD is the best.

5 Conclusion

In this paper, we introduce Bayesian methods for robust constraint-based causal discovery under insufficient data. Two Bayesian-augmented frequentist independence tests are proposed for reliable statistic estimation under a frequentist independence test framework. Specifically, we propose: 1) an effective empirical Bayesian method for accurate estimation of mutual information under limited data; 2) a Bayesian statistical testing method for independence test by formulating the Bayes Factor into the well-defined χ^2 statistical test. We apply the proposed methods to both local and global causal discovery algorithms and evaluate their performance against state-of-the-art methods on different benchmark datasets. The experiments show that, by introducing Bayesian approaches, the proposed methods not only outperform the competing methods in terms of accuracy, but also improve efficiency significantly.

Acknowledgments

Appendix
A. Full Bayesian MI method

We define the full Bayesian MI estimation (Eq. 1) as

$$
\begin{aligned}
\hat{MI}^{fB} &= \int \int MI(X;Y|\boldsymbol{\theta}, \alpha)p(\boldsymbol{\theta}, \alpha|D)d\boldsymbol{\theta}d\alpha \\
&= \int \int MI(X;Y|\boldsymbol{\theta})p(\boldsymbol{\theta}|\alpha)p(\alpha|D)d\boldsymbol{\theta}d\alpha
\end{aligned} \tag{10}
$$

To approximately solve for $\hat{MI}^{fB}$, we propose to approximate the integration over hyper-parameters by its mode α^* that maximizes the posterior of α given data D, and we obtain the proposed empirical Bayesian MI,

$$
\hat{MI}^{eB} = \int MI(X;Y|\boldsymbol{\theta})p(\boldsymbol{\theta}|D, \alpha^*)d\boldsymbol{\theta} \tag{11}
$$

where $\alpha^* = \arg\max p(\alpha|D)$. Apply Bayes's rule, we have

$$
p(\alpha|D) \propto p(D|\alpha)p(\alpha) \tag{12}
$$

By assuming uniform distribution for $p(\alpha)$, we have

$$
\alpha^* = \arg\max p(\alpha|D) = \arg\max p(D|\alpha) \tag{13}
$$

The likelihood distribution $p(D|\alpha)$ follows the Pólya distribution and is computed as

$$
p(D|\alpha) = \int p(D|\boldsymbol{\theta})p(\boldsymbol{\theta}|\alpha)d\boldsymbol{\theta} \tag{14}
$$

where $p(D|\boldsymbol{\theta})$ is a mutinomial distribution and $p(\boldsymbol{\theta}|\alpha)$ is a Dirichlet distribution. Eq. 14 can be solved analytically

$$
\begin{aligned}
p(D|\alpha) &= \int p(D|\theta)p(\boldsymbol{\theta}|\alpha)d\boldsymbol{\theta} \\
&= \int \frac{N!}{\prod_{i=1}^{K} n_i!} \prod_{i=1}^{K} \theta_i^{n_i} \frac{\Gamma(K\alpha)}{\prod_{i=1}^{K}\Gamma(\alpha)} \prod_{i=1}^{K} \theta_i^{\alpha-1}d\boldsymbol{\theta} \\
&= \frac{N!}{\prod_{i=1}^{K} n_i!} \int \frac{\Gamma(K\alpha)}{\prod_{i=1}^{K}\Gamma(\alpha)} \prod_{i=1}^{K} \theta_i^{n_i+\alpha-1}d\boldsymbol{\theta} \\
&= \frac{N!}{\prod_{i=1}^{K} n_i!} \frac{\Gamma(\alpha K)}{\Gamma(\alpha K+N)} \frac{\prod_{i=1}^{K}\Gamma(\alpha+n_i)}{\Gamma(\alpha)^K} \\
&= \frac{N!\Gamma(\alpha K)}{\Gamma(\alpha)^K\Gamma(\alpha K+N)} \prod_{i=1}^{K} \frac{\Gamma(\alpha+n_i)}{n_i!}
\end{aligned} \tag{15}
$$

We solve for α^* with a fixed-point update

$$
\alpha^{new} = \alpha^{old}\frac{\sum_{i,j}\psi(\alpha+n_{ij}) - K\psi(\alpha)}{K\psi(\alpha K+N) - \psi(K\alpha)} \tag{16}
$$

where ψ is the digamma function. Given α^*, $\hat{MI}^{eB}$ can be computed as (Eq. 4)

$$
\begin{aligned}
\hat{MI}^{eB} &= \int MI(X;Y|\boldsymbol{\theta})p(\boldsymbol{\theta}|D, \alpha^*)d\boldsymbol{\theta} \\
&= \psi(N+\alpha^* K+1) - \sum_{ij}\frac{n_{ij}+\alpha^*}{N+\alpha^* K}[\psi(n_i+\alpha^* K_y+1) \\
&\quad + \psi(n_j+\alpha^* K_x+1) - \psi(n_{ij}+\alpha^*+1)]
\end{aligned} \tag{17}
$$

where ψ is the digamma function.

B. Underlying distribution of a Likelihood Ratio Test

The likelihood ratio test is defined as

$$LR = -2\ln\frac{\sup_{\boldsymbol{\theta}\in\Theta_0}P(D|\boldsymbol{\theta})}{\sup_{\boldsymbol{\theta}\in\Theta}P(D|\boldsymbol{\theta})} \tag{18}$$

where a null hypothesis is corresponding to parameter $\boldsymbol{\theta} \in \Theta_0$ and Θ_0 is a subset of the whole parameter space Θ.

B.1. Underlying distribution in general

In this section, we will derive a general form of the underlying distribution for the likelihood ratio test statistic without explicitly providing a likelihood function. The likelihood ratio test statistic can be re-written as

$$LR = -2(L(\boldsymbol{\theta}^0) - L(\boldsymbol{\theta}^1)) \tag{19}$$

where $L(\boldsymbol{\theta}) = \ln P(D|\boldsymbol{\theta})$ is a log-likelihood function and $\boldsymbol{\theta}^0 \in \Theta_0$, $\boldsymbol{\theta}^1 \in \Theta$. We do a Taylor expansion for $L(\boldsymbol{\theta}^0)$ at point $\boldsymbol{\theta}^1$, and we have

$$L(\boldsymbol{\theta}^0) \approx L(\boldsymbol{\theta}^1) + (\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1)L'(\boldsymbol{\theta}^1) \\ + \frac{1}{2}(\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1)^T L''(\boldsymbol{\theta}^1)(\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1) \tag{20}$$

Note that $\boldsymbol{\theta}^1 = \arg\max P(D|\boldsymbol{\theta}, H_1)$ and thus $L'(\boldsymbol{\theta}^1) = \mathbf{0}$. Then we have

$$L(\boldsymbol{\theta}^0) = L(\boldsymbol{\theta}^1) + \frac{1}{2}(\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1)^T L''(\boldsymbol{\theta}^1)(\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1) \tag{21}$$

Plug equation(12) into equation(10) and we have

$$LR = (\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1)^T(-L''(\boldsymbol{\theta}^1))(\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1) \tag{22}$$

In order to derive a general form for the distribution of LR from equation(13), we need to compute the distribution of $(\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1)$. We do another Taylor expansion for $L'(\boldsymbol{\theta}^0)$ at point $\boldsymbol{\theta}^1$ and we have

$$L'(\boldsymbol{\theta}^0) = L'(\boldsymbol{\theta}^1) + (\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1)L''(\boldsymbol{\theta}^1) \tag{23}$$

where $L'(\boldsymbol{\theta}^1) = \mathbf{0}$. And we have

$$\boldsymbol{\theta}^0 - \boldsymbol{\theta}^1 = (L''(\boldsymbol{\theta}^1))^{-1}L'(\boldsymbol{\theta}^0) \tag{24}$$

We plug equation(15) into equation(13), and we have

$$LR = L'(\boldsymbol{\theta}^0)^T(L''(\boldsymbol{\theta}^1))^{-T}(-L''(\boldsymbol{\theta}^1))(L''(\boldsymbol{\theta}^1))^{-1}L'(\boldsymbol{\theta}^0) \\ = L'(\boldsymbol{\theta}^0)^T(-L''(\boldsymbol{\theta}^1))^{-1}L'(\boldsymbol{\theta}^0) \tag{25}$$

where $L''(\boldsymbol{\theta}^1))^{-T} = L''(\boldsymbol{\theta}^1))^{-1}$. Furthermore, we decompose the positive definite matrix $(-L''(\boldsymbol{\theta}^1))^{-1}$ as $(-L''(\boldsymbol{\theta}^1))^{-1} = Q^2$. In other words, we denote the decomposition as

$$(-L''(\boldsymbol{\theta}^1))^{-1} = (-L''(\boldsymbol{\theta}^1))^{-1/2}(-L''(\boldsymbol{\theta}^1))^{-1/2} \tag{26}$$

with $(-L''(\boldsymbol{\theta}^1))^{-1/2} = Q$. Then we have

$$LR \\ = L'(\boldsymbol{\theta}^0)^T(-L''(\boldsymbol{\theta}^1))^{-1/2}(-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0) \\ = L'(\boldsymbol{\theta}^0)^T(-L''(\boldsymbol{\theta}^1))^{-T/2}(-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0) \\ = [(-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0)]^T[(-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0)] \\ = Z^T Z \tag{27}$$

where $Z = (-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0)$. Assume Z follows a Gaussian distribution and we compute its expectation and variance in the following.

$$E[Z] = E[(-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0)] \\ = (-L''(\boldsymbol{\theta}^1))^{-1/2}E[L'(\boldsymbol{\theta}^0)] \tag{28}$$

where

$$E[L'(\boldsymbol{\theta}^0)] = E[\frac{d\log P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}}] \\ = E[\frac{dP(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}}\frac{1}{P(D|\boldsymbol{\theta}^0)}] \\ = \int \frac{dP(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}}\frac{1}{P(D|\boldsymbol{\theta}^0)}P(D|\boldsymbol{\theta}^0)dD \tag{29} \\ = \frac{d}{d\boldsymbol{\theta}}\int p(D|\boldsymbol{\theta}^0)dD \\ = \frac{d}{d\boldsymbol{\theta}}(1) = \mathbf{0}$$

Given the result in Eq. 29, we can show that $E[Z] = \mathbf{0}$. Now we consider the variance of Z. We have

$$Var[Z] = Var[(-L''(\boldsymbol{\theta}^1))^{-1/2}L'(\boldsymbol{\theta}^0)] \\ = (-L''(\boldsymbol{\theta}^1))^{-1}Var[L'(\boldsymbol{\theta}^0)] \tag{30}$$

where $Var[L'(\boldsymbol{\theta}^0)] = E[(L'(\boldsymbol{\theta}^0))^2] - E^2[L'(\boldsymbol{\theta}^0)] = E[(L'(\boldsymbol{\theta}^0))^2]$ because $E[L'(\boldsymbol{\theta}^0)] = \mathbf{0}$. We claim that $E[(L'(\boldsymbol{\theta}^0))^2] = E[-L''(\boldsymbol{\theta}^0)]$ and we prove this claim in the following:

$$E[-L''(\boldsymbol{\theta}^0)] \\ = E[-\frac{d^2\log P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}] \\ = E[-\frac{d}{d\boldsymbol{\theta}}(\frac{dP(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}}\frac{1}{P(D|\boldsymbol{\theta}^0)})] \\ = E[-\frac{d^2P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}\frac{1}{P(D|\boldsymbol{\theta}^0)} + (\frac{dP(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}}\frac{1}{P(D|\boldsymbol{\theta}^0)})^2] \\ = E[-\frac{d^2P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}\frac{1}{P(D|\boldsymbol{\theta}^0)}] + E[(\frac{d\log P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}})^2] \\ = E[(L'(\boldsymbol{\theta}^0)^2] \tag{31}$$

where $E[-\frac{d^2P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}\frac{1}{P(D|\boldsymbol{\theta}^0)}] = \mathbf{0}$ because

$$E[-\frac{d^2P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}\frac{1}{P(D|\boldsymbol{\theta}^0)}] \\ = -\int \frac{d^2P(D|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}\frac{1}{P(D|\boldsymbol{\theta}^0)}P(D|\boldsymbol{\theta}^0)dD \\ = -\frac{d^2}{d\boldsymbol{\theta}^2}\int p(D|\boldsymbol{\theta}^0)dD \tag{32} \\ = -\frac{d^2}{d\boldsymbol{\theta}^2}(1) = \mathbf{0}$$

After we prove the claim, we plug the claim into Eq. 30 and we have

$$Var[Z] = (-L''(\boldsymbol{\theta}^1))^{-1}E[-L''(\boldsymbol{\theta}^0)] \tag{33}$$

In the end, we derive a general form for the distribution of LR as

$$LR = Z^T Z \tag{34}$$

where $Z \sim N(\mathbf{0}, Var[Z])$ with $Var[Z] = (-L''(\boldsymbol{\theta}^1))^{-1}E[-L''(\boldsymbol{\theta}^0)]$.

B.2. G-test and χ^2 distribution

The value of the G test is derived from the LR where the underlying model is a multinomial model. Suppose we have $i.i.d$ samples $D = \{D_n\} = \{(X_n, Y_n)\}$ where $n = \{1, 2, ..., N\}$, given the multinomial model, we have

$$g = -2\ln \frac{\prod_{n=1}^{N} P(X_n|\boldsymbol{\theta}^0)P(Y_n|\boldsymbol{\theta}^0)}{\prod_{n=1}^{N} P((X_n, Y_n)|\boldsymbol{\theta}^1)} \tag{35}$$

where we apply the null hypothesis that $P(X, Y) = P(X)P(Y)$. Suppose we have $X = x_i$ where $i = \{1, 2, ..., K_x\}$ and $Y = y_j$ where $j = \{1, 2, ..., K_y\}$, then we have

$$g = -2\ln \frac{\prod_{i=1}^{K_x} \theta_i^{n_i} \prod_{i=1}^{K_x} \theta_j^{n_i}}{\prod_{i=1}^{K_x} \prod_{j=1}^{K_y} \theta_{ij}^{n_{ij}}} \tag{36}$$

Note that

$$\prod_{i=1}^{K_x} \theta_i^{n_i} = \prod_{i=1}^{K_x} \theta_i^{\sum_j n_{ij}} = \prod_{i=1}^{K_x}\prod_{j=1}^{K_y} \theta_i^{n_{ij}} \tag{37}$$

Similarly, we have $\prod_{j=1}^{K_y} \theta_j^{n_j} = \prod_{j=1}^{K_y}\prod_{i=1}^{K_x} \theta_j^{n_{ij}}$. Then we can write Eq. 36 as

$$
\begin{aligned}
g &= -2\ln \frac{\prod_{i=1}^{K_x}\prod_{j=1}^{K_y} \theta_i^{n_{ij}}\theta_i^{n_{ij}}}{\prod_{i=1}^{K_x}\prod_{j=1}^{K_y} \theta_{ij}^{n_{ij}}} \\
&= -2\sum_{i=1}^{K_x}\sum_{j=1}^{K_y} n_{ij} \ln \frac{\theta_i \theta_j}{\theta_{ij}}
\end{aligned}
\tag{38}
$$

which is the definition of the G statistic. Now we show that given a multinomial model which is applied in the G test, $Var[Z] = (-L''(\boldsymbol{\theta}^1))^{-1}E[-L''(\boldsymbol{\theta}^0)] = 1$ and thus g follows χ^2 distribution. Firstly, given a multinomial model, we have $L(\boldsymbol{\theta}) = \log P(D|\alpha) = \log \prod_{n=1}^{N} P(D_n|\alpha) = \sum_{n=1}^{N} \log P(D_n|\boldsymbol{\theta})$. Then we have

$$
\begin{aligned}
E[-L''(\boldsymbol{\theta}^0)] &= E[-\sum_{n=1}^{N} \frac{d^2 \log P(D_n|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}] \\
&= \sum_{n=1}^{N} E[-\frac{d^2 \log P(D_n|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}]
\end{aligned}
\tag{39}
$$

where $E[-\frac{d^2 \log P(D_n|\boldsymbol{\theta}^0)}{d\boldsymbol{\theta}^2}]$ is a function of the true underlying $\boldsymbol{\theta}$ and we denote it as $I_{\boldsymbol{\theta}}$ (which is known as the fisher information). We thus have $E[-L''(\boldsymbol{\theta}^0)] = NI_{\boldsymbol{\theta}}$. Then for

the $Var[Z]$, we have

$$
\begin{aligned}
Var[Z] &= (-L''(\boldsymbol{\theta}^1))^{-1}E[-L''(\boldsymbol{\theta}^0)] \\
&= (-L''(\boldsymbol{\theta}^1))^{-1}NI_{\boldsymbol{\theta}} \\
&= (\frac{-L''(\boldsymbol{\theta}^1)}{N})^{-1}I_{\boldsymbol{\theta}} \\
&= (\frac{1}{N}\sum_{n=1}^{N} -\frac{d^2 \log P(D_n|\boldsymbol{\theta}^1)}{d\boldsymbol{\theta}^2})^{-1}I_{\boldsymbol{\theta}} \\
&= I_{\boldsymbol{\theta}}^{-1}I_{\boldsymbol{\theta}} = \mathbb{1}
\end{aligned}
\tag{40}
$$

where we assume N is large enough so that the averaged $-\frac{d^2 \log P(D_n|\boldsymbol{\theta}^1)}{d\boldsymbol{\theta}^2}$ can well approximate the expectation $I_{\boldsymbol{\theta}}$. Then we can show that $g = Z^T Z \sim \chi^2$ with $Z \sim N(0, \mathbb{1})$.

C. Bayes Factor and χ^2 Distribution

C.1. Closed-form solution of Bayes Factor

We denote the Bayesian likelihood ratio value as BF with the definition (Eq. 5),

$$BF = \frac{P(D|H_0, \alpha^0)}{P(D|H_1, \alpha^1)} = \frac{\int P(D|\boldsymbol{\theta})P(\boldsymbol{\theta}|H_0, \alpha^0)d\boldsymbol{\theta}}{\int P(D|\boldsymbol{\theta})P(\boldsymbol{\theta}|H_1, \alpha^1)d\boldsymbol{\theta}} \tag{41}$$

where we assume the symmetric Dirichlet prior. Under the null hypothesis H_0, i.e., two random variables are independent, we set the posterior for each random variable X, Y separately, i.e.,

$$p(\boldsymbol{\theta}|H_0, \alpha^0) = p(\boldsymbol{\theta}_x|H_0, \alpha_x)p(\boldsymbol{\theta}_y|H_0, \alpha_y) \tag{42}$$

with $p(\boldsymbol{\theta}_x|H_0, \alpha_x) \sim Dir(\alpha_x)$ and $p(\boldsymbol{\theta}_y|H_0, \alpha_y) \sim Dir(\alpha_y)$. Given parameters under null hypothesis, we have the multinomial distribution for $p(D|\boldsymbol{\theta})$,

$$p(D|\boldsymbol{\theta}) = N! \prod_{i=1}^{K_x}\prod_{j=1}^{K_y} \frac{(\theta_{x_i}\theta_{y_j})^{n_{ij}}}{n_{ij}!} \tag{43}$$

Then we have $p(D|H_0, \alpha^0)$ as

$$
\begin{aligned}
&p(D|H_0, \alpha^0) = \int p(D|\boldsymbol{\theta})p(\boldsymbol{\theta}|H_0, \alpha^0)d\boldsymbol{\theta} \\
&= \int N! \prod_{i=1}^{K_x}\prod_{j=1}^{K_y} \frac{(\theta_{x_i}\theta_{y_j})^{n_{ij}}}{n_{ij}!}p(\boldsymbol{\theta}_x|H_0, \alpha_x)p(\boldsymbol{\theta}_y|H_0, \alpha_y)d\boldsymbol{\theta} \\
&= \frac{N!}{\prod_{i=1}^{K_x}\prod_{j=1}^{K_y} n_{ij}!} \\
&\quad \times \int \prod_{i=1}^{K_x}\prod_{j=1}^{K_y}(\theta_{x_i}\theta_{y_j})^{n_{ij}}p(\boldsymbol{\theta}_x|H_0, \alpha_x)p(\boldsymbol{\theta}_y|H_0, \alpha_y)d\boldsymbol{\theta}
\end{aligned}
\tag{44}
$$

Notice that

$$
\begin{aligned}
\prod_{i=1}^{K_x}\prod_{j=1}^{K_y}(\theta_{x_i}\theta_{y_j})^{n_{ij}} &= \prod_{i=1}^{K_x}\theta_{x_i}^{\sum_j n_{ij}} \prod_{j=1}^{K_y}\theta_{y_j}^{\sum_i n_{ij}} \\
&= \prod_{i=1}^{K_x}\theta_{x_i}^{n_i} \prod_{j=1}^{K_y}\theta_{y_j}^{n_j}
\end{aligned}
\tag{45}
$$

Plug Eq. 45 back into Eq. 44 and we have

$$p(D|H_0, \alpha^0)$$

$$= M \int \prod_{i=1}^{K_x} \theta_{x_i}^{n_i} \prod_{j=1}^{K_y} \theta_{y_j}^{n_j} p(\boldsymbol{\theta}_x|H_0, \alpha_x) p(\boldsymbol{\theta}_y|H_0, \alpha_y) d\boldsymbol{\theta}$$

$$= M \int \prod_{i=1}^{K_x} \theta_{x_i}^{n_i} p(\boldsymbol{\theta}_x|H_0, \alpha_x) d\boldsymbol{\theta}_x \int \prod_{j=1}^{K_y} \theta_{y_j}^{n_j} p(\boldsymbol{\theta}_y|H_0, \alpha_y) d\boldsymbol{\theta}_y$$

$$(46)$$

where $M = \frac{N!}{\prod_{i=1}^{K_x} \prod_{j=1}^{K_y} n_{ij}!}$. As the integration over $\boldsymbol{\theta}_x$ and $\boldsymbol{\theta}_y$ in Eq. 46 are identical, we only derive for $\boldsymbol{\theta}_x$,

$$\int \prod_{i=1}^{K_x} \theta_{x_i}^{n_i} p(\boldsymbol{\theta}_x|H_0, \alpha_x) d\boldsymbol{\theta}_x$$

$$= \int \prod_{i=1}^{K_x} \theta_{x_i}^{n_i} \frac{\Gamma(K_x \alpha_x)}{\prod_{i=1}^{K_x} \Gamma(\alpha_x)} \prod_{i=1}^{K_x} \theta_{x_i}^{\alpha_x - 1} d\boldsymbol{\theta}_x \quad (47)$$

$$= \frac{\Gamma(K_x \alpha_x)}{\Gamma(K_x \alpha_x + N)} \prod_{i=1}^{K_x} \frac{\Gamma(\alpha_x + n_i)}{\Gamma(\alpha_x)}$$

The integration over $\boldsymbol{\theta}_y$ can be done in the same way,

$$\int \prod_{j=1}^{K_y} \theta_{y_j}^{n_j} p(\boldsymbol{\theta}_y|\alpha_y) d\boldsymbol{\theta}_y = \frac{\Gamma(K_y \alpha_y)}{\Gamma(K_y \alpha_y + N)} \prod_{j=1}^{K_y} \frac{\Gamma(\alpha_y + n_j)}{\Gamma(\alpha_y)}$$

$$(48)$$

Given Eq. 47 and Eq. 48, we have

$$p(D|H_0, \alpha^0) = \int p(D|\boldsymbol{\theta}) p(\boldsymbol{\theta}|H_0, \alpha^0) d\boldsymbol{\theta}$$

$$= \frac{M \Gamma(K_x \alpha_x) \Gamma(K_y \alpha_y)}{\Gamma(K_x \alpha_x + N) \Gamma(K_y \alpha_y + N)} \quad (49)$$

$$\times \prod_{i=1}^{K_x} \frac{\Gamma(\alpha_x + n_i)}{\Gamma(\alpha_x)} \prod_{j=1}^{K_y} \frac{\Gamma(\alpha_y + n_j)}{\Gamma(\alpha_y)}$$

where $M = \frac{N!}{\prod_{i=1}^{K_x} \prod_{j=1}^{K_y} n_{ij}!}$. Under the alternative hypothesis, i.e., two random variables are dependent, we set posterior for the joint distribution of two random variables, i.e.,

$$p(\boldsymbol{\theta}|H_1, \alpha^1) \sim Dir(\alpha_{xy}) \quad (50)$$

with dimension $K = K_x K_y$. Following the similar procedure as we did for the null hypothesis, we can show that

$$p(D|H_1, \alpha^1) = \int p(D|\boldsymbol{\theta}) p(\boldsymbol{\theta}|H_1, \alpha_{xy}) d\boldsymbol{\theta}$$

$$= M \frac{\Gamma(\alpha_{xy} K)}{\Gamma(\alpha_{xy} K + N)} \prod_{i}^{K} \frac{\Gamma(\alpha_{xy} + n_i)}{\Gamma(\alpha_{xy})}$$

$$(51)$$

where $M = \frac{N!}{\prod_{i=1}^{K} n_i!}$. Given the marginal likelihood for the null hypothesis (Eq. 49) and and the marginal likelihood for

the alternative hypothesis (Eq. 51)), we calculate the likelihood ratio as

$$BF = \frac{\Gamma(\alpha_x K_x) \Gamma(\alpha_y K_y) \Gamma(\alpha_{xy} K + N)}{\Gamma(\alpha_{xy} K) \Gamma(\alpha_x K_x + N) \Gamma(\alpha_y K_y + N)}$$

$$\times \frac{\Gamma(\alpha_{xy})^K}{\Gamma(\alpha_x)^{K_x} \Gamma(\alpha_y)^{K_y}} \quad (52)$$

$$\times \frac{\prod_{i=1}^{K_x} \Gamma(\alpha_x + n_i) \prod_{j=1}^{K_y} \Gamma(\alpha_y + n_j)}{\prod_{i=1}^{K_x} \prod_{j=1}^{K_y} \Gamma(\alpha_{xy} + n_{ij})}$$

Given a pre-defined threshold η, if $BF > \eta$, the null hypothesis is more likely to be supported by the data and two variables are declared to be independent. Otherwise, we accept alternative hypothesis and declare the variables to be dependent.

C.2. BF and χ^2 distribution

The Bayes factor is a likelihood ratio of the marginal likelihood of two competing hypothesis, usually the null and alternative. The Bayes factor is defined as

$$BF = \frac{P(D|H_0, \alpha^0)}{P(D|H_1, \alpha^1)} = \frac{\int P(D|\boldsymbol{\theta}) P(\boldsymbol{\theta}|H_0, \alpha^0) d\boldsymbol{\theta}}{\int P(D|\boldsymbol{\theta}) P(\boldsymbol{\theta}|H_1, \alpha^1) d\boldsymbol{\theta}} \quad (53)$$

Different from the G-test which calculates statistics based on one set of parameters, the Bayes factor is considering all possible sets of parameters given the hypothesis. Note that $P(D|\alpha)$ is a likelihood function with $\alpha = \arg \max P(D|\alpha)$, and we can apply the general form of the underlying distribution of the likelihood ratio test to BF. We modify the value BF as

$$LR_{BF} = -2 \ln BF = -2 \ln \frac{P(D|H_0, \alpha^0)}{P(D|H_1, \alpha^0)} = Z^T Z \quad (54)$$

where $Z = (-L''(\alpha^1))^{-1/2} L'(\alpha^0)$. And $Z \sim N(0, Var[Z])$ with $Var[Z] = (-L''(\alpha^1))^{-1} E[-L''(\alpha^0)]$. We can analytically compute the likelihood function by integrating out parameter $\boldsymbol{\theta}$ as

$$P(D|\alpha) = \int P(D|\boldsymbol{\theta}) P(\boldsymbol{\theta}|\alpha) d\boldsymbol{\theta}$$

$$= \frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0 + N)} \prod_{k}^{K} \frac{\Gamma(\alpha_k + n_k)}{\Gamma(\alpha_k)} \quad (55)$$

where $\alpha_0 = \sum_{k=1}^{K} \alpha_k$ and $P(D|\alpha)$ is the probability mass function for the Polya distribution. And the log-likelihood function is

$$L(\alpha) = \ln P(D|\alpha)$$

$$= \ln \Gamma(\alpha_0) - \ln \Gamma(\alpha_0 + N)$$

$$+ \sum_{k}^{K} \ln \Gamma(\alpha_k + n_k) - \ln \Gamma(\alpha_k) \quad (56)$$

In this case, we should treat D as one sample data and thus we can't expect $-L''(\alpha^1)$ approach to $E[-L''(\alpha^0)]$. In other words, we can't naturally have $Var[Z] = \mathbb{1}$ and thus $\tilde{BF}$ doesn't follow χ^2 distribution.

D. Bayesian G statistic and its distribution

D.1. An approximated Polya Distribution

We denote the proposed likelihood ratio value as BF with the definition (Eq. 5),

$$BF = \frac{P(D|H_0, \alpha^0)}{P(D|H_1, \alpha^1)} = \frac{\int P(D|\boldsymbol{\theta})P(\boldsymbol{\theta}|H_0, \alpha^0)d\boldsymbol{\theta}}{\int P(D|\boldsymbol{\theta})P(\boldsymbol{\theta}|H_1, \alpha^1)d\boldsymbol{\theta}} \quad (57)$$

where we assume the symmetric Dirichlet prior. To better decide the threshold, we propose to combine the Bayes Factor BF with a well-defined statistical test χ^2 via an approximated Pólya distribution, i.e.,

$$P(D|\alpha) \approx P(D|\tilde{\boldsymbol{\theta}}) = \frac{N!}{\prod_{i=1}^{K} n_i!} \prod_{k=1}^{K} \tilde{\theta}_k^{\ n_k} \quad (58)$$

Because $\frac{N!}{\prod_{i=1}^{K} n_i!}$ is a common component that exists in both $P(D|\alpha)$ and $P(D|\tilde{\boldsymbol{\theta}})$, we ignore this term in the following derivations. We re-write the log-likelihood function $\ln P(D|\alpha)$ as

$$\ln P(D|\alpha) = \sum_{k}^{K} (\ln \Gamma(n_k + \alpha) - \ln \Gamma(\alpha)) \\ - (\ln \Gamma(N + K\alpha) - \ln \Gamma(K\alpha)) \quad (59)$$

On the other hand, we have the log-likelihood of $P(D|\tilde{\boldsymbol{\theta}})$ as

$$\ln P(D|\tilde{\boldsymbol{\theta}}) = \sum_{k}^{K} n_k \ln \tilde{\theta}_k \quad (60)$$

Furthermore, we assume that $\tilde{\theta}_k$ has the form $\frac{g(n_k, \alpha)}{g(N, K\alpha)}$ where $g(n_k, \alpha)$ is a function with unknown parameters that need to be estimated. We plug $\tilde{\theta}_k = \frac{g(n_k, \alpha)}{g(N, K\alpha)}$ into Eq. 60 and we have

$$\ln P(D|\tilde{\boldsymbol{\theta}}) = \sum_{k}^{K} n_k \ln g(n_k, \alpha) - N \ln g(N, K\alpha) \quad (61)$$

Comparing Eq. 59 and Eq. 60, we can see that to make $P(D|\alpha) \approx P(D|\tilde{\boldsymbol{\theta}})$, the desired property of the function $g(n_k, \alpha)$ is

$$n_k \ln g(n_k, \alpha) = \ln \Gamma(n_k + \alpha) - \ln \Gamma(\alpha) \quad (62)$$

Based on the fact that $\sum_{k}^{K} \tilde{\theta}_k = 1$, the function $g(n_k, \alpha)$ should subject to the constraint that

$$\sum_{k}^{K} g(n_k, \alpha) = g(N, K\alpha) = g(\sum_{k}^{K} n_k, \sum_{k}^{K} \alpha) \quad (63)$$

In other words, the function $g(n_k, \alpha)$ is a linear function with respect to both n_k and α. Thus, the form of the function $g(n_k, \alpha)$ should be

$$g(n_k, \alpha) = a n_k + b\alpha \quad (64)$$

where parameters $\Lambda = \begin{pmatrix} a \\ b \end{pmatrix}$ are unknowns that need to be computed. To estimate two parameters, we have

$$\begin{pmatrix} n_1, \alpha \\ n_2, \alpha \\ ... \\ n_K, \alpha \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} t(n_1, \alpha) \\ t(n_2, \alpha) \\ ... \\ t(n_K, \alpha) \end{pmatrix} \quad (65)$$

where $t(n_k, \alpha) = \exp(\frac{1}{n_k}(\ln \Gamma(n_k + \alpha) - \ln \Gamma(\alpha)))$. Denote

$$M = \begin{pmatrix} n_1, \alpha \\ n_2, \alpha \\ ... \\ n_K, \alpha \end{pmatrix}, T = \begin{pmatrix} t(n_1, \alpha) \\ t(n_2, \alpha) \\ ... \\ t(n_K, \alpha) \end{pmatrix}, \text{ and we re-write Eq. 65}$$

as $M\Lambda = T$. We always have $K \geq 2$ and we solve for Λ with the least square error, i.e.,

$$\Lambda^* = \arg \min ||M\Lambda - T||^2 \quad (66)$$

with the solution $\Lambda^* = (M^T M)^{-1} M^T T$. Given Λ^*, we have $\tilde{\theta}_k$ as

$$\tilde{\theta}_k = \frac{g(n_k, \alpha)}{g(N, K\alpha)} = \frac{a^* n_k + b^* \alpha}{a^* N + b^* K\alpha} \quad (67)$$

and $P(D|\tilde{\boldsymbol{\theta}})$ can well approximate $P(D|\alpha)$.

We demonstrate the effectiveness of approximating the Polya probability $p(D|\alpha)$ with modified parameter $\tilde{\boldsymbol{\theta}}$ by comparing it with Minka's approach which approximates the Polya distribution by modified counts. Synthetic data is generated following the procedure stated in the paper. In particular, we consider independency and dependency separately and synthetic data is generated without assuming symmetric Dirichlet prior. The relative absolute error between the estimated probability $\tilde{p}$ and the true probability p, i.e., $\frac{|p - \tilde{p}|}{p}$ is applied as the measurement given each sample set D. We report the average error over 1000 runs for each sample size. As we can see from Figure 1, our approach approximates the
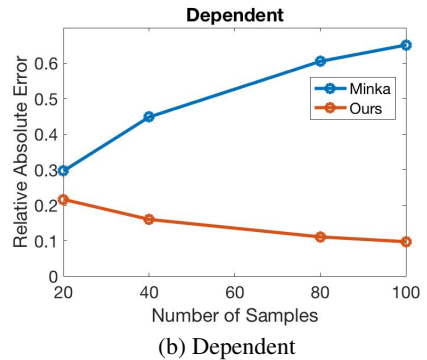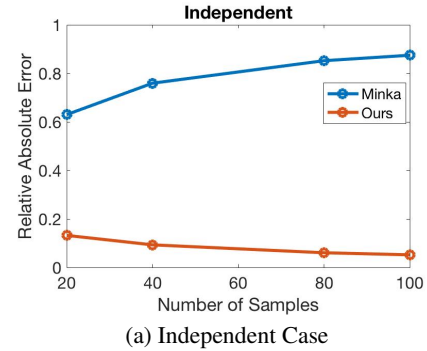

(a) Independent Case

(b) Dependent

Figure 1: Relative Absolute Error of Polya distribution

true polya probability much better than Minka's approach. The reason is that Minka's approach requires the estimation

of Dirichlet hyper-parameters and can't work well with symmetric Dirichlet prior. In addition, the visualization of the polya distribution is shown in Figure 2 and Figure 3.
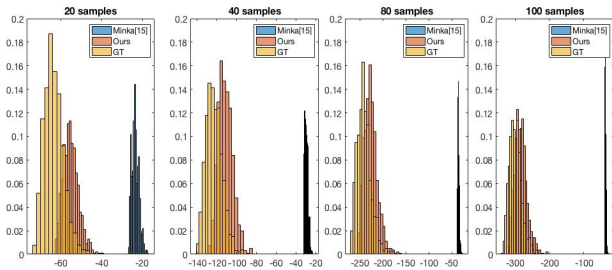

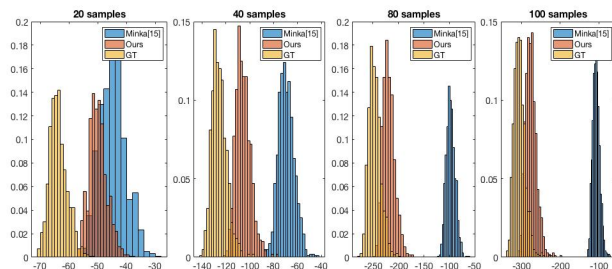Figure 2: Visualization of Polya distribution Estimation (Independent Case)

Figure 3: Visualization of Polya distribution Estimation (Dependent Case)

D.2. Distribution of Bayesian G statistic

We study the distribution of the proposed Bayes Factor statistic BF_{chi2} to show that it asymptomatically follows the χ^2 distribution given the null hypothesis being true, i.e., two variables are independent. We perform experiments on the synthetic data and follow the procedure stated in the paper to generate the synthetic data. We consider both the uniform prior $\alpha = 1$ and Jeffrey's prior $\alpha = 0.5$ to study the effect of the Dirichlet hyper-parameter. As we estimate the distribution under the true null hypothesis, we set two random variables X and Y to be independent. We estimate the statistic distribution based on frequencies. For comparison, we show the classical G statistic. We visualize the distribution in Figure 4, Figure 5 and Figure 6. As we can see from Figure 4, with sufficient data, i.e., 1000, both BF_{chi2} and G follow χ^2 distribution. Under insufficient data, the probability of BF_{chi2} statistic tends to be overestimated with χ^2 distribution bias towards independence declaration. Compared to uniform prior, BF_{chi2} statistic with Jeffrey's prior produces the distribution that is closer to the χ^2 distribution.

E. Theoretical Guarantee on Bayesian Estimation Improvement over MLE

To theoretically prove that the Bayesian estimation gives better estimation then MLE, we compare the uncertainty of two estimations via measuring their variances. We consider the
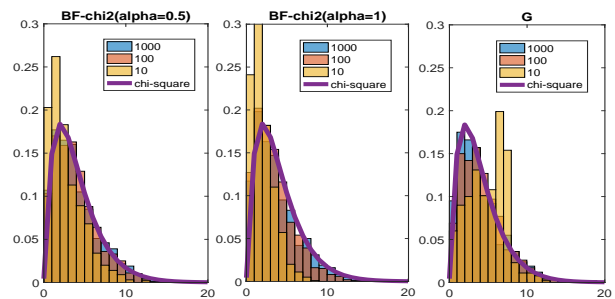

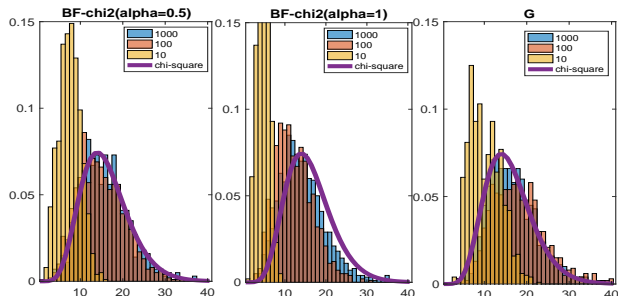Figure 4: Distribution of the Bayesian G statistic (3 by 3)

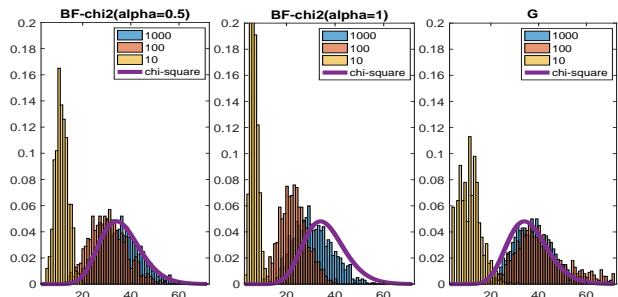Figure 5: Distribution of the Bayesian G statistic (5 by 5)

Figure 6: Distribution of the Bayesian G statistic (7 by 7)

parameter estimation. Given data D, the closed-form solution for MLE is $\theta_i^{MLE} = \frac{n_i}{N}$, where n_i is the number of samples for state i, with $\sum_{i=1}^{K} n_i = N$. The variance is then computed as

$$Var(\theta_i^{MLE}) = Var(\frac{n_i}{N}) = \frac{Var(n_i)}{N^2} \qquad (68)$$

The variance computed with Eq. 68 captures the variance of estimator caused by the randomness of data. To better show this point, we start from the definition of the variance of the MLE by considering the randomness of data:

$$Var(\theta^{MLE}) = \int_D (\theta^{MLE} - E[\theta^{MLE}])^2 p(D)dD \qquad (69)$$

Since probability parameters are independent, we can simply compute each of them θ_i^{MLE} separately, i.e.,

$$Var(\theta_i^{MLE})$$
$$= \int_D (\theta_i^{MLE} - E[\theta_i^{MLE}])^2 p(D)dD$$
$$= \int_D ((\theta_i^{MLE})^2 - 2\theta_i^{MLE}E[\theta_i^{MLE}] + E^2[\theta_i^{MLE}])p(D)dD \tag{70}$$

For MLE estimation, we have $\theta_i^{MLE} = \frac{n_i}{N}$, and $E[\theta_i^{MLE}] = \frac{E[n_i]}{N}$ which is independent of data D. The variance of θ_i^{MLE} then becomes

$$Var(\theta_i^{MLE}) = \int_D (\frac{n_i^2}{N^2} - 2\frac{n_i}{N}\frac{E[n_i]}{N} + \frac{E^2[n_i]}{N^2})p(D)dD$$
$$= \int_D \frac{(n_i^2 - 2n_iE[n_i] + E^2[n_i])}{N^2}p(D)dD$$
$$= \frac{\int_D (n_i - E[n_i])^2 p(D)dD}{N^2}$$
$$= \frac{Var(n_i)}{N} \tag{71}$$

Given the fact that n_i follows the multinomial distribution with GT parameter θ_i, the variance is $Var(n_i) = N\theta_i(1 - \theta_i)$. In the end, we have

$$Var(\theta_i^{MLE}) = \frac{N\theta_i(1 - \theta_i)}{N^2} \tag{72}$$

Similarly, given the closed-form solution for Bayesian estimation is $\theta_i^{Bayes} = \frac{n_i + \alpha_i}{N + \sum_i \alpha_i}$, we compute the variance as

$$Var(\theta_i^{Bayes}) = Var(\frac{n_i + \alpha_i}{N + \sum_i \alpha_i})$$
$$= \frac{Var(n_i + \alpha_i)}{(N + \sum_i \alpha_i)^2} = \frac{Var(n_i)}{(N + \sum_i \alpha_i)^2} \tag{73}$$

where the variance is $Var(n_i) = N\theta_i(1 - \theta_i)$. Furthermore, for hyper-parameters α_i, we always have $\alpha_i > 0$. In the end, we have

$$Var(\theta_i^{MLE})$$
$$= \frac{N\theta_i(1 - \theta_i)}{N^2} > \frac{N\theta_i(1 - \theta_i)}{(N + \sum_i \alpha_i)^2} \tag{74}$$
$$= Var(\theta_i^{Bayes})$$

Here we show that variance of Bayesian estimation θ_i^{Bayes} is always smaller than the variance of MLE θ_i^{MLE} for parameter estimation.

F. Evaluations on Constraint-based Causal Discovery

F.1. Experiment Settings on Algorithms and Dataset

Statistical information on datasets: We employ six benchmark datasets that are widely used for causal discovery evaluation: CHILD, INSURANCE, ALARM, HAILFINDER, CHILD3 and CHILD5. CHILD, INSURANCE and ALARM are medium networks with the number of variables being 20,27 and 37 respectively. HAILFINDER, CHILD3 and CHILD5 are larger and more challenging networks with the number of variables being 56, 60 and 100 respectively. Their information is shown in Table 5.

Dataset	#Variables	#Edges	Maximum #States
CHILD	20	25	6
INSURANCE	27	52	5
ALARM	37	46	4
HAILFINDER	56	66	11
CHILD3	60	75	6
CHILD5	100	125	6

Table 5: Benchmark DAG Information

Algorithm setting for local causal discovery: For the local causal discovery, we employ Causal Markov Blanket (CMB) [Gao and Ji, 2015], which is the state-of-the-art method. For hyper-parameters, CMB applies G-test and the significance level is set to be 5% as suggested. For cI^{eB}, threshold is set based on empirical analysis with synthetic data[5]. For cBF_{chi2}, we apply the Jeffreys prior ($\alpha^0 = \alpha^1 = 0.5$) and significance level is 5%. To verify the performance of original CMB algorithm, we perform CMB on ALARM dataset with 5000 samples. The averaged SHD is 1.63, which is comparable with the reported result, i.e., SHD=1.81 [Gao and Ji, 2015].

Algorithm setting for global causal discovery: We employ RAI as our baseline algorithm and incorporate the proposed independence tests. For hyper-parameters, RAI-BF applies BF with Jeffreys prior for independence test and the threshold is set to 1. PC-stable applies χ^2 independence test, and the significance level is 5% by default. For rI^{eB}, threshold is determined based on our empirical analysis with synthetic data. For rBF_{chi2} test, we apply the Jeffreys prior and the significance level is 5%. To compare to RAI-BF, we implement BF independence test and incorporate it into RAI[6]. To verify the implementation, we perform RAI-BF on ALARM dataset with 10,000 samples. The averaged SHD is 25.2, which is comparable to the 25.3 in their reported result [Natori et al., 2017]. For PC-stable, we directly apply the algorithm provided by bnlearn[7].

F.2. Running Time on CPU

All the experiments in the paper are performed on a laptop with a 2.3 GHz 8-Core Intel Core i9 processor using CPU only. Specifically, we compare the running time of rI^{eB}, rBF_{chi2} and RAI-BF since they are all based on the RAI algorithm. We report the average running time over different sample sizes for each dataset. From the results shown in Table 6, we can see that rBF_{chi2} requires less running time than rI^{eB} and RAI-BF, particularly on large graphs. Though the number of independence tests performed by rI^{eB} is the smallest, the running time per independence test is large, and

---

[5]Threshold is estimated as a function of the number of samples per each configuration using synthetic data.

[6]https://github.com/benzione/FixRAI.

[7]https://www.bnlearn.com/documentation/man/constraint.html.

hence the total running time for rI^{eB} is not competitive. We will incorporate detailed results into the paper.

Datasets	rI^{eB}	rBF_{chi2}	RAI-BF
CHILD	5.83s	1.09s	0.72s
INSURANCE	10.08s	1.53s	0.80s
ALARM	24.05s	2.73s	1.63s
HAILFINDER	52.96s	9.67s	57.07s
CHILD3	68.86s	12.67s	18.06s
CHILD5	181.59s	23.82s	37.01s

Table 6: Running Time on CPU

F.3. Performance of RAI-BF with tuned threshold

Since the proposed BF_{chi2} essentially is only an approximate of original BF and thus the BF with the optimal threshold, in principle, should outperform BF_{chi2}. But selecting the optimal threshold for the BF can be challenging and incorrect thresholds can lead to inferior performance. BF_{chi2}, in contrast, only needs a default significant level (5%) to perform the test. This may explain why BF_{chi2} outperforms BF in our experiments.

Instead of fixing the threshold for RAI-BF with its default value, we tune its threshold and report the best SHD for comparison. The threshold is selected from [0.2, 2.0]. The corresponding number of independence tests is also reported. From the results shown in Table 7, we can see that rI^{eB} achieves overall better efficiency by performing a smaller number of independence tests, which is consistent with our conclusion stated in the paper. Comparing the RAI-BF with the optimally tuned threshold and rBF_{chi2}, they achieve comparable performance in terms of both accuracy and efficiency. On Child3 and Child5, RAI-BF, with an optimally

Dataset		SHD			#Independence Test	
	rI^{eB}	rBF_{chi2}	RAI-BF	rI^{eB}	rBF_{chi2}	RAI-BF
CHILD	19.7	19.3	19.4	350	538	493
INSURANCE	48.6	44.7	45.1	575	930	757
ALARM	41.7	36.2	32.9	1161	1576	1708
HAILFINDER	88.4	104.3	103.9	2215	3129	2531
CHILD3	65.0	53.2	63.4	3786	4518	2260
CHILD5	120.0	106.8	116.7	9197	9857	8028

Table 7: Comparison to optimal RAI-BF with tuned threshold

selected threshold, performs a smaller number of independence tests as the optimal threshold is small ($\sim$0.2), leading to more independence declarations.

F.4. Performance Improvement Consistency with PC and MMHC

To further show that our proposed methods can consistently improve the causal discovery performance, we consider another two widely used DAG learning algorithms: PC [Spirtes et al., 2000] and MMHC [Tsamardinos et al., 2006b]. We incorporate the proposed methods into PC and MMHC for evaluation. PC with empirical Bayesian MI estimation and BF_{chi2} independence test are denoted as pcI^{eB} and $pcBF_{chi2}$ respectively. MMHC with empirical Bayesian MI estimation and BF_{chi2} independence test are denoted as mI^{eB} and mBF_{chi2} respectively.

Method	PC					
Dataset		SHD			#Independence Test	
(MEAN)	pcI^{eB}	$pcBF_{chi2}$	PC	pcI^{eB}	$pcBF_{chi2}$	PC
CHILD	**22.0**	**22.0**	27.4	**331**	382	610
INSURANCE	50.4	**50.0**	57.4	**548**	699	1067
ALARM	41.9	**40.5**	58.2	**1098**	1355	3585
HAILFINDER	**84.1**	101.1	119.0	**2511**	3959	33352
CHILD3	77.2	**75.2**	88.8	**2099**	2548	3880
CHILD5	108	**95.3**	107.8	13180	11009	**10348**

Table 8: Evaluation of proposed methods with PC

Method	MMHC					
Dataset		SHD			#Independence Test	
(MEAN)	mI^{eB}	mBF_{chi2}	MMHC	mI^{eB}	mBF_{chi2}	MMHC
CHILD	22.5	**21.9**	22.4	**14**	**14**	37
INSURANCE	52.1	**49.9**	52.9	**19**	23	45
ALARM	42.0	**38.7**	40.1	**24**	26	33
HAILFINDER	**78.3**	87.0	92.0	10	**5**	34
CHILD3	66.9	**63.7**	64.9	**8**	20	24
CHILD5	103	**102**	104	**9**	25	24

Table 9: Evaluation of proposed methods with MMHC

As shown in Table 8 and Table 9, our proposed methods can consistently improve the DAG learning performance, particularly with PC. For example, on ALARM, $pcBF_{chi2}$ achieves averaged SHD 40.5, while PC only achieves averaged SHD 58.2. MMHC is a hybrid approach, where a constraint-based algorithm is only used to obtain an initial graph for a score-based algorithm. Thus, the performance of MMHC doesn't completely reflect the performance of independence tests, and different independence tests don't introduce much difference to DAG learning performance. Overall, BF_{chi2} achieves better accuracy and I^{eB} achieves better efficiency with both PC and MMHC on different datasets.

G. Evaluation of Independence Tests

G.1. Evaluation on Synthetic Data

We perform experiments on synthetic data to study the performance of the proposed independence tests. We firstly evaluate the proposed empirical Bayesian MI estimation. We then analyze the proposed BF_{chi2} independence test. For the synthetic data, we consider two multi-state random variables X and Y. The underlying probabilistic dependency between X and Y is randomly generated. The probability parameters are randomly generated given the dependency with the symmetric Dirichlet prior. We generate synthetic data of different small sizes for evaluation.

Mutual Information Estimation We compare our proposed empirical Bayesian MI estimation with state-of-the-art MI estimation methods. Specifically, we consider the adaptive partition method [Seok and Seon Kang, 2015], the empirical Bayesian method [Hutter, 2002] and the full Bayesian method [Archer et al., 2013]. To measure the accuracy of the MI estimation, we apply the absolute error between the estimated MI $\hat{MI}$ and the ground truth MI MI, i.e., $|\hat{MI} - MI|$. We report the averaged absolute error over 1000 runs for each sample size. From Figure 7, we can see that our approach gives the best estimation compared to others. In particular, our proposed empirical Bayesian MI via a MAP estimation of the hyper-parameter α performs better than the empirical Bayesian method with fixed α [Hutter, 2002]. Addition-
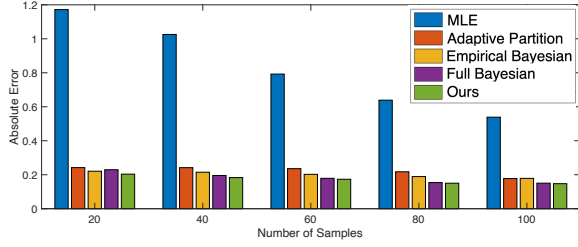
Figure 7: Mutual Information Estimation

Dataset	#Independence Test				
(**MEAN**)	rI^{AdP}	rI^{eBFix}	rI^{eB}	rBF_{chi2}	RAI-BF
CHILD	**267**	552	350	538	955
INSURANCE	**516**	589	575	930	1012
ALARM	**915**	1236	1161	1576	1857
HAILFINDER	**1806**	2486	2215	3129	10212
CHILD3	**2642**	4299	3786	4518	5503
CHILD5	**7278**	9889	9197	9857	11174

Table 10: Efficiency comparison to SoTA independence tests

ally, we achieve comparable accuracy compared to the full Bayesian method [Archer *et al.*, 2013]. Without requiring the integration over the hyper-parameter space, our methods only takes 0.002 seconds on average to finish one run, while the full Bayesian method [Archer *et al.*, 2013] needs on average 10.481 seconds to finish one run. Thus, our empirical Bayesian MI estimation is more computationally efficient to be applied in causal discovery.

Hypothesis testing based Independence Test We compare the proposed BF_{chi2} independence test against the standard G test and the Bayes Factor (BF) [Natori *et al.*, 2017] which represents the state-of-the-art independence test that matches with our approach. We follow the BF [Natori *et al.*, 2017] and apply Jeffreys prior. We consider Type-1 error and Power as measurements. Type-1 error is rejecting the null hypothesis H_0 when it is true. The power is the probability of correctly
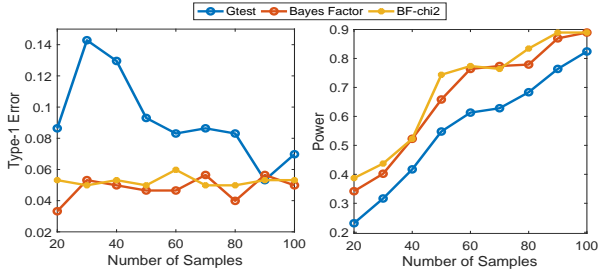

Figure 8: Performance of BF-chi2

rejecting H_0. We set the significance level of BF_{chi2} test, G test, and manually tune the threshold of BF to make their corresponding Type-1 error near 5%. And we compare the power of different methods. From the results shown in Figure 8, we can see that Bayesian approaches BF_{chi2} and BF achieve higher power than frequentist-based approach G test under insufficient data with lower Type-1 error. And our approach achieves slightly better power compared to the BF [Natori *et al.*, 2017], without any threshold tuning.

G.2. Evaluation on Benchmark Datasets

We compare proposed independence tests to two state-of-the-art methods: adaptive partition and empirical Bayesian with fixed α methods in terms of causal discovery performance on benchmark datasets. We incorporate the adaptive partition method and the empirical Bayesian with fixed α method to RAI (denoted as rI^{AdP} and rI^{eBFix} respectively). In terms

of efficiency, as shown in Table 10, rI^{AdP} achieves best efficiency by performing the smallest number of independence tests. Though rI^{AdP} is efficient, its accuracy is compromised and rI^{AdP} performs worse than other methods on almost all the datasets. Our proposed method rI^{eB} is the second most efficient method, and achieves competitive accuracy at the same time. Overall, our proposed methods outperform other SoTA independence tests in terms of causal discovery performance. On HAILFINDER, because rI^{AdP} tends to declare independence, the learned DAG contains fewer false positive edges compared to other methods and thus its averaged SHD is the best.

H. Classification Performance under Imbalanced Data

To further demonstrate the effectiveness of the proposed methods, we consider real world imbalanced datasets where samples for certain classes are insufficient. Particularly, we construct a structured classifier by using the learned DAG from global causal discovery methods. We perform the evaluation on four UCI datasets [Dua and Graff, 2017] that are benchmark imbalanced datasets [Jiang *et al.*, 2014] [Jiang *et al.*, 2013]. The statistic information of datasets is shown in Table 11. F1-score and AUC are applied as the evaluation

Table 11: Information of UCI datasets

Dataset	#Samples	#Attributes	#Majority/#Minority
Breast-w	699	9	458/241
Spect	267	22	212/55
Diabetes	768	8	500/268
Parkinsons	195	23	147/48

metrics to measure the classification accuracy. We apply 3-fold cross-validation. The settings of hyper-parameters for independence tests remain the same as we did in the previous section. Results are shown in Table 12.

Table 12: Predictive performance under imbalanced data

	F1-score / AUC			
Dataset	rI^{eB}	rBF_{chi2}	RAI-BF	PC-stable
Breast-w	0.42/**0.52**	**0.84**/0.50	0.40/0.49	0.47/0.24
Spect	**0.59**/0.74	0.58/**0.77**	0.57/**0.74**	0.41/**0.55**
Diabetes	0.63/0.44	**0.68**/0.43	0.50/**0.51**	0.67/0.44
Parkinsons	0.61/**0.74**	**0.70**/0.64	0.38/0.50	0.48/0.54

We can see that rBF_{chi_2} achieves better performance on most of the datasets, and rI^{eB} also achieves improved performance. For example, on Breast-w dataset, rBF_{chi_2} im-

proves the F1-score by 44% and 37% compared to RAI-BF method and PC-stable method respectively. Considering AUC, we can see that both rI^{eB} and rBF_{chi2} achieve better AUC than baseline methods RAI-BF and PC-stable on most of the datasets. These results further demonstrate that, with proposed independence tests, we can learn DAGs that better capture underlying structure among variables under imbalanced data, leading to improved structure classification performance.

References

[Archer *et al.*, 2013] Evan Archer, Il Park, and Jonathan Pillow. Bayesian and quasi-bayesian estimators for mutual information from discrete data. *Entropy*, 15(5):1738–1755, 2013.

[Claassen and Heskes, 2012] Tom Claassen and Tom Heskes. A bayesian approach to constraint based causal inference. *arXiv preprint arXiv:1210.4866*, 2012.

[Colombo and Maathuis, 2014] Diego Colombo and Marloes H Maathuis. Order-independent constraint-based causal structure learning. *The Journal of Machine Learning Research*, 15(1):3741–3782, 2014.

[Colombo *et al.*, 2012] Diego Colombo, Marloes H Maathuis, Markus Kalisch, and Thomas S Richardson. Learning high-dimensional directed acyclic graphs with latent and selection variables. *The Annals of Statistics*, pages 294–321, 2012.

[Dua and Graff, 2017] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.

[Gao and Ji, 2015] Tian Gao and Qiang Ji. Local causal discovery of direct causes and effects. *Advances in Neural Information Processing Systems*, 28:2512–2520, 2015.

[Geweke and Singleton, 1980] John F Geweke and Kenneth J Singleton. Interpreting the likelihood ratio statistic in factor models when sample size is small. *Journal of the American Statistical Association*, 75(369):133–137, 1980.

[Glymour *et al.*, 2019] Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.

[Hausser and Strimmer, 2009] Jean Hausser and Korbinian Strimmer. Entropy inference and the james-stein estimator, with application to nonlinear gene association networks. *Journal of Machine Learning Research*, 10(7), 2009.

[Hutter, 2002] Marcus Hutter. Distribution of mutual information. In *Advances in neural information processing systems*, pages 399–406, 2002.

[Jiang *et al.*, 2013] L. Jiang, C. Li, Z. Cai, and H. Zhang. Sampled bayesian network classifiers for class-imbalance and cost-sensitive learning. pages 512–517, 2013.

[Jiang *et al.*, 2014] Liangxiao Jiang, Chaoqun Li, and Shasha Wang. Cost-sensitive bayesian network classifiers. *Pattern Recognition Letters*, 45:211–216, 2014.

[Kass and Raftery, 1995] Robert E Kass and Adrian E Raftery. Bayes factors. *Journal of the american statistical association*, 90(430):773–795, 1995.

[Kruschke, 2013] John K Kruschke. Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142(2):573, 2013.

[Li *et al.*, 2019] Honghao Li, Vincent Cabeli, Nadir Sella, and Hervé Isambert. Constraint-based causal structure learning with consistent separating sets. In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, 2019.

[Marx and Vreeken, 2018] Alexander Marx and Jilles Vreeken. Stochastic complexity for testing conditional independence on discrete data. In *NeurIPS 2018 Workshop on Causal Learning*, 2018.

[McDonald, 2009] John H McDonald. *Handbook of biological statistics*, volume 2. sparky house publishing Baltimore, MD, 2009.

[Minka, 2000] Thomas Minka. Estimating a dirichlet distribution, 2000.

[Mukherjee and Speed, 2007] Sach Mukherjee and Terence P Speed. Markov chain monte carlo for structural inference with prior information. *University of California; Berkeley*, 2007.

[Natori *et al.*, 2017] Kazuki Natori, Masaki Uto, and Maomi Ueno. Consistent learning bayesian networks with thousands of variables. In *Advanced Methodologies for Bayesian Networks*, pages 57–68, 2017.

[Ramsey *et al.*, 2012] Joseph Ramsey, Jiji Zhang, and Peter L Spirtes. Adjacency-faithfulness and conservative causal inference. *arXiv preprint arXiv:1206.6843*, 2012.

[Rohekar *et al.*, 2018] Raanan Y Rohekar, Yaniv Gurwicz, Shami Nisimov, Guy Koren, and Gal Novik. Bayesian structure learning by recursive bootstrap. In *Advances in Neural Information Processing Systems*, pages 10525–10535, 2018.

[Rohekar *et al.*, 2020] Raanan Y Rohekar, Yaniv Gurwicz, Shami Nisimov, and Gal Novik. A single iterative step for anytime causal discovery. *arXiv preprint arXiv:2012.07513*, 2020.

[Sechidis *et al.*, 2019] Konstantinos Sechidis, Laura Azzimonti, Adam Pocock, Giorgio Corani, James Weatherall, and Gavin Brown. Efficient feature selection using shrinkage estimators. *Machine Learning*, 108(8):1261–1286, 2019.

[Sen *et al.*, 2017] Rajat Sen, Ananda Theertha Suresh, Karthikeyan Shanmugam, Alexandros G Dimakis, and Sanjay Shakkottai. Model-powered conditional independence test. *arXiv preprint arXiv:1709.06138*, 2017.

[Seok and Seon Kang, 2015] Junhee Seok and Yeong Seon Kang. Mutual information between discrete variables with many categories using recursive adaptive partitioning. *Scientific Reports*, 5, 06 2015.

[Spirtes *et al.*, 2000] Peter Spirtes, Clark N Glymour, Richard Scheines, David Heckerman, Christopher Meek, Gregory Cooper, and Thomas Richardson. *Causation, prediction, and search*. MIT press, 2000.

[Spirtes, 2010] Peter Spirtes. Introduction to causal inference. *Journal of Machine Learning Research*, 11(5), 2010.

[Strobl *et al.*, 2019] Eric V Strobl, Kun Zhang, and Shyam Visweswaran. Approximate kernel-based conditional independence tests for fast non-parametric causal discovery. *Journal of Causal Inference*, 7(1), 2019.

[Tsamardinos *et al.*, 2006a] Ioannis Tsamardinos, Laura E. Brown, and Constantin F. Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine Learning*, 65(1):31–78, Oct 2006.

[Tsamardinos *et al.*, 2006b] Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006.

[Yang *et al.*, 2021] Shuai Yang, Hao Wang, Kui Yu, Fuyuan Cao, and Xindong Wu. Towards efficient local causal structure learning, 2021.

[Yehezkel and Lerner, 2009] Raanan Yehezkel and Boaz Lerner. Bayesian network structure learning by recursive autonomy identification. *Journal of Machine Learning Research*, 10(Jul):1527–1570, 2009.

[Yu *et al.*, 2019] Yue Yu, Jie Chen, Tian Gao, and Mo Yu. Dag-gnn: Dag structure learning with graph neural networks. *arXiv preprint arXiv:1904.10098*, 2019.