Faithfulness and learning hypergraphs from discrete distributions

Anna Klimova

Institute of Science and Technology, Austria

`aklimova@ist.ac.at`

Caroline Uhler

Institute of Science and Technology, Austria

`caroline.uhler@ist.ac.at`

Tamás Rudas

Eötvös Loránd University, Budapest, Hungary

`rudas@tarki.hu`

Abstract

In this paper, we study the concepts of faithfulness and strong-faithfulness for discrete distributions. In the discrete setting, graphs are not sufficient for describing the association structure. So we consider hypergraphs instead, and introduce the concept of parametric (strong-) faithfulness with respect to a hypergraph. Assuming strong-faithfulness, we build uniformly consistent parameter estimators and corresponding procedures for a hypergraph search. The strength of association in a discrete distribution can be quantified with various measures, leading to different concepts of strong-faithfulness. We explore these by computing lower and upper bounds for the proportions of distributions that do not satisfy strong-faithfulness.

Keywords: contingency tables, directed acyclic graphs, hierarchical log-linear models, hypergraphs, (strong-) faithfulness

1 Introduction

A graphical model is a set of probability distributions whose association structure can be identified with a graph. Given a graph, the Markov property entails a set of conditional independence relations that are fulfilled by distributions in the model. Distributions in the model that obey no further conditional independence relations are called *faithful to the graph*. For each undirected graphical model, as well as for each directed acyclic graph (DAG) model, there is a distribution that is faithful to the graph [cf. Spirtes, Glymour, and Scheines, 2001]. Moreover, the Lebesgue measure of the set of parameters corresponding to distributions that are unfaithful to a graphical model is zero; this result was proven by Spirtes, Glymour, and Scheines [2001] for the case of multivariate normal distributions, by Meek [1995] for discrete distributions on multi-way contingency tables, and by Peña, Nilsson, Björkegren, and Tegnér [2009] for arbitrary sample spaces and dominating measures. It is also well-known, that a

arXiv:1404.6617v3 [stat.ME] 23 Jan 2015

DAG model may include distributions that are unfaithful to it but are not Markov to any nested DAG. This kind of unfaithfulness may occur due to path cancellation and can arise both in the discrete and in the multivariate normal settings [cf. Zhang and Spirtes, 2008, Uhler and Raskutti, 2013].

In the discrete case, the non-existence of a graph to which a distribution is faithful is related to the presence of higher than first order interactions in this distribution. Graph learning algorithms [cf. Spirtes et al., 2001], which do not recognize the presence of higher order interactions, may produce a graph which does not reveal the true association structure [cf. Studený, 2005]. In order to avoid such errors, graph learning algorithms usually assume the existence of a DAG to which the distribution is faithful. Since the Lebesgue measure of the set of parameters corresponding to the distributions that are unfaithful to the underlying graph is zero, the faithfulness assumption is not considered to be restrictive in the context of graphical search. While graph search procedures assuming faithfulness are pointwise consistent, they are not uniformly consistent and thus cannot simultaneously control Type I and Type II errors with a finite sample size [Robins, Scheines, Spirtes, and Wasserman, 2003]. To ensure existence of a uniformly consistent learning procedure, strong-faithfulness of a distribution to the underlying DAG is needed [Zhang and Spirtes, 2003]. Uhler, Raskutti, Bühlmann, and Yu [2013] analyzed the Gaussian setting and showed that the strong-faithfulness assumption may, in fact, be very restrictive and the corresponding proportions of distributions which do not satisfy strong-faithfulness may become very large as the number of nodes grows.

The concepts of faithfulness and strong-faithfulness were originally introduced in the causal search framework, where they are linked to identifiability of causal effects. However, as we show in this paper using the discrete setting, these concepts are also important for identifiability of more general parameters of association. In Section 2, we define the concept of a model class being closed under a faithfulness relation: for each positive distribution, there exists a model in such a class to which it is faithful. By giving examples of distributions that are not faithful to any directed or undirected graphical model, we show that these model classes are not closed under the faithfulness relation which is based on the corresponding Markov property. Further, we introduce the concept of parametric faithfulness of a distribution to a hypergraph (instead of a graph). This concept seems more adequate for categorical data, where hypergraphs can be used to represent hierarchical log-linear models. Indeed, we show that the class of models associated with hypergraphs is closed under a parametric faithfulness relation.

In Section 3, we describe two major difficulties with the concept of strong-faithfulness in the discrete case. First, in contrast to role of correlations in the multivariate normal case, there is no single standard measure of the strength of association in a joint distribution. Therefore, depending on the measure of association, different variants of strong-faithfulness may be considered. Second, the proportion of strong-faithful distributions depends on the parameterization used and can only be computed if the parameter space has finite volume. We explore the consequences of different parameterizations and measures of association for the case of the 2×2 contingency table. We define parametric strong-faithfulness with respect to a hypergraph under a parameterization based on the log-linear interaction parameters. Assuming strong-faithfulness, we show that the maximum likelihood estimators of the interaction parameters associated with the hyperedges are uniformly consistent. As a result,

we give a set of conditions, under which Type I and Type II errors can be controlled with a finite sample size. We also discuss the uniform consistency of model selection procedures for a hypergraph search, for example, using the approaches described by Edwards [2000, 2012].

In Section 4, we estimate the proportion of distributions that do not satisfy the parametric strong-faithfulness assumption with respect to a given hypergraph. We give an exact formulation of these proportions, under a parameterization based on conditional probabilities, for hypergraphs whose hyperedges form a decomposable set. The association structure of such distributions may be discovered incorrectly during a hypergraph learning procedure. Finally, we define the concept of projected strong-faithfulness, which applies to distributions which do not belong to the hypergraph, and estimate the proportions of projected strong-faithful distributions for several hypergraphs for the $2 \times 2 \times 2$ contingency table.

In Section 5, we conclude the paper with a brief discussion of our results and their implications.

2   Graphical and parametric faithfulness

In this section, we first review the concept of faithfulness with respect to a graph. We then introduce parametric faithfulness with respect to a hypergraph and show that this is a more relevant concept for categorical data.

2.1   Faithfulness with respect to a graph

Let $\mathcal{V}_1, \ldots, \mathcal{V}_K$ be random variables taking values in $\mathcal{I} = \mathcal{I}_1 \times \cdots \times \mathcal{I}_K$, a Cartesian product of finite sets. $\mathcal{I}$ describes a K-way contingency table and a vector $\boldsymbol{i} = (i_1, \ldots, i_K) \in \mathcal{I}$ forms a cell. A subset $M \subseteq \{1, \ldots, K\}$ specifies a marginal of the joint distribution of $\mathcal{V}_1, \ldots, \mathcal{V}_K$, and $M = \emptyset$ is the empty marginal. For $M = (k_1, \ldots k_t)$, the set $\mathcal{I}_M = \mathcal{I}_{k_1} \times \cdots \times \mathcal{I}_{k_t}$ is a *marginal table*, and the canonical projection $\boldsymbol{i}_M$ of the cell $\boldsymbol{i}$ onto the set $\mathcal{I}_M$ is a *marginal cell*. We parameterize the population distribution by cell probabilities $\boldsymbol{p} = (p_{\boldsymbol{i}})_{\boldsymbol{i} \in \mathcal{I}}$, where $p_{\boldsymbol{i}} \in (0, 1)$ and $\sum_{\boldsymbol{i} \in \mathcal{I}} p_{\boldsymbol{i}} = 1$, and denote by $\mathcal{P}$ the set of all distributions on $\mathcal{I}$. A subset of $\mathcal{P}$ is called a *model*. For simplicity of exposition, we assume that $\mathcal{V}_1, \ldots, \mathcal{V}_K$ are binary, $\mathcal{I}$ is treated as a sequence of cells ordered lexicographically, and a distribution $P \in \mathcal{P}$ is addressed by its parameter, $\boldsymbol{p}$.

A *graphical model* is a set of probability distributions, whose association structure can be identified with a graph with vertices $V = \{1, \ldots, K\}$, where each vertex i is associated with a random variable $\mathcal{V}_i$. In the following, we will identify each vertex with its associated random variable. The absence of an edge between two vertices means that the corresponding random variables satisfy some (conditional) independence relation. A detailed description of graphical models for discrete as well as for multivariate normal distributions can be found in Edwards [2000], among others. In the sequel, we only consider undirected graphical models and DAG models.

A graphical model identified with an undirected graph (also called a *graphical log-linear model* in the discrete setting) is a set of probability distributions on V that satisfy the *local undirected Markov property*: Every node is conditionally independent of its non-neighbors given its neighbors. In the discrete case, such models are a sublcass of hierarchical log-linear

models. A graphical model identified with a directed acyclic graph, a DAG model, is a set of probability distributions on V that satisfy the *directed Markov property*: Every node is conditionally independent of its non-descendants given its parents. A distribution that satisfies the Markov property with respect to a graph is called *Markov* to it.

A distribution which is Markov to a graph, is said to be *faithful* to it if all conditional independencies in this distribution can be derived from the graph. The faithfulness relation can be thought of as a decision rule that classifies a distribution $\boldsymbol{p}$ in a model $\mathcal{M}$ as faithful or unfaithful to it:

$$\mathbb{F}(\boldsymbol{p}, \mathcal{M}) = \begin{cases} 1 & \text{if } \boldsymbol{p} \text{ is faithful to } \mathcal{M}, \\ 0 & \text{otherwise.} \end{cases}$$

Definition 2.1. A class $\mathfrak{C}$ of models on $\mathcal{P}$, where $\mathfrak{C}$ is partially ordered with respect to inclusion, is said to be *closed* under the faithfulness relation indicated by $\mathbb{F}$, if for every non-empty $\mathcal{M} \in \mathfrak{C}$ and for every $\boldsymbol{p} \in \mathcal{M}$ such that $\mathbb{F}(\boldsymbol{p}, \mathcal{M}) = 0$, there exists an $\mathcal{M}' \in \mathfrak{C}$ with $\mathcal{M}' \subset \mathcal{M}$ and $\mathbb{F}(\boldsymbol{p}, \mathcal{M}') = 1$.

This definition implies that a class $\mathfrak{C}$ is closed under the faithfulness relation indicated by $\mathbb{F}$ if and only if for every $\boldsymbol{p} \in \mathcal{P}$ there exists an $\mathcal{M} \in \mathfrak{C}$, such that $\mathbb{F}(\boldsymbol{p}, \mathcal{M}) = 1$. Graphical log-linear models and DAG models are specified by a list of conditional independence relations which, in turn, comprise other conditional independencies. Thus, these model classes have a natural partial order implied by the conditional independence relation. We now show that these classes are not closed under the corresponding faithfulness relations.

Proposition 2.1. *The class of graphical log-linear models is not closed under the faithfulness relation defined by the local undirected Markov property.*

The following example is given as a proof.

Example 2.1. Let $V = \{A, B, C, D\}$ and consider the log-linear model $[ABC][ABD]$ [cf. Agresti, 2002]. This is the model of conditional independence of C and D given A and B. All distributions in this model are Markov to the graph in Figure 1. Consider the distribution parameterized by

$$\begin{aligned} \boldsymbol{p} = \ & (0.022, 0.062, 0.063, 0.103, 0.103, 0.063, 0.062, 0.022, \\ & 0.103, 0.063, 0.062, 0.022, 0.022, 0.062, 0.063, 0.103)', \end{aligned}$$

where the cell probabilities are ordered lexicographically. In this distribution, the conditional odds ratios ($\mathcal{COR}$) of C and D given the levels of A and B are equal to 1:

$$\mathcal{COR}(CD \mid A = i, B = j) = \frac{p_{ij00}p_{ij11}}{p_{ij01}p_{ij10}} = 1, \text{ for all } i, j \in \{0, 1\}.$$

Hence, the distribution is in the model. The (A, B)-marginal of this distribution is uniform:

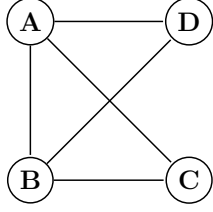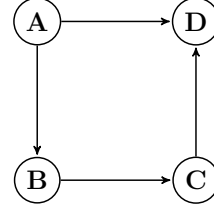|	$B = 0$	$B = 1$
$A = 0$	1/4	1/4
$A = 1$	1/4	1/4

,

Figure 1: $C \perp\!\!\!\perp D \mid A, B$.

Figure 2: $A \perp\!\!\!\perp C \mid B, \quad B \perp\!\!\!\perp D \mid A, C$.

and thus $A \perp\!\!\!\perp B$. So the distribution is unfaithful to the graph in Figure 1. In addition, since the conditional odds ratios of A, B and C given D and of A, B, and D given C are not equal to 1:

$$
\begin{aligned}
\mathcal{COR}(ABC \mid D = 0) &= \frac{p_{0000}p_{1100}p_{0110}p_{1010}}{p_{0100}p_{1000}p_{0010}p_{1110}} \approx 0.04418483, \\
\mathcal{COR}(ABC \mid D = 1) &= \frac{p_{0001}p_{1101}p_{0111}p_{1011}}{p_{0101}p_{1001}p_{0011}p_{1111}} \approx 0.04418483, \\
\mathcal{COR}(ABD \mid C = 0) &= \frac{p_{0000}p_{1100}p_{0101}p_{1001}}{p_{0100}p_{1000}p_{0001}p_{1101}} \approx 0.04710518, \\
\mathcal{COR}(ABD \mid C = 1) &= \frac{p_{0010}p_{1110}p_{0111}p_{1011}}{p_{0110}p_{1010}p_{0011}p_{1111}} \approx 0.04710518,
\end{aligned}
$$

the distribution cannot be Markov to any nested undirected graph. $\square$

The situation described in Example 2.1 is distinctive to discrete distributions. In the Gaussian setting, marginal independence of more than two variables implies their joint independence. Thus, a multivariate normal distribution whose components are pairwise independent is Markov and faithful to a graph with no edges. But in the discrete case, a joint distribution of pairwise independent variables may have a non-trivial structure of higher than first order interactions. Next, we prove that also the class of DAG models is not closed with respect to the faithfulness relation.

Proposition 2.2. *The class of DAG models is not closed under the faithfulness relation defined by the directed Markov property.*

As a proof, two examples are given. The second example only pertains to the discrete case.

Example 2.2. Let $V = \{A, B, C, D\}$ and consider the model specified by two conditional independence relations: $A \perp\!\!\!\perp C \mid B$ and $B \perp\!\!\!\perp D \mid A, C$. Any distribution in this model is Markov to the DAG in Figure 2. For example, the distribution parameterized by

$$
\begin{aligned}
\boldsymbol{p} = \ &(0.006, 0.006, 0.0288, 0.0192, 0.06, 0.06, 0.072, 0.048, 0.0056, \\
&0.0504, 0.187148, 0.0368516, 0.021, 0.189, 0.175452, 0.0345484)',
\end{aligned}
$$

is in the model. However, this distribution also satisfies the additional independence relation $A \perp\!\!\!\perp D$. This independence relation is not reflected in the graph. Thus the distribution is unfaithful to the graph in Figure 2.

Next, we show that there is no DAG that fulfills all three (conditional) independence relations $A \perp\!\!\!\perp D$, $A \perp\!\!\!\perp C \mid B$, $B \perp\!\!\!\perp D \mid A, C$. If such a DAG existed, then its skeleton would have three edges: AB, BC, CD. In order to satisfy faithfulness, $A \perp\!\!\!\perp D$ requires that $A \to B \leftarrow C$ or $B \to C \leftarrow D$. However, $A \perp\!\!\!\perp C \mid B$ is unfaithful to $A \to B \leftarrow C$ and $B \perp\!\!\!\perp D \mid A, C$ is unfaithful to $B \to C \leftarrow D$. $\qquad\square$

Remark 2.1. One can also construct an instance of Example 2.2 using multivariate normal distributions by choosing the partial correlations in such a way that the causal effect associated with the edge $A \to D$ cancels with the causal effect associated with the path $A \to B \to C \to D$ (see Figure 2). This shows that also Gaussian DAG models are not closed under the faithfulness relation defined by the directed Markov property.

The next example illustrates a situation that occurs only in the discrete case. To construct this example, we will use the fact that, in contrary to the Gaussian case, a discrete distribution with pairwise independent random variables can have non-vanishing interactions of higher than the first order.

Example 2.3. Let $V = \{A, B, C\}$. Consider the distribution parameterized by

$$\boldsymbol{p} = (1/8 - \delta, 1/8 + \delta, 1/8 + \delta, 1/8 - \delta, 1/8 + \delta, 1/8 - \delta, 1/8 - \delta, 1/8 + \delta)', \qquad (1)$$

where $\delta \in (-1/8, 1/8)$. Its marginals are uniform resulting in pairwise independence: $A \perp\!\!\!\perp B$, $A \perp\!\!\!\perp C$, and $B \perp\!\!\!\perp C$. The second order odds ratio of this distribution,

$$\frac{p_{000}p_{011}p_{101}p_{110}}{p_{001}p_{010}p_{100}p_{111}} = \left(\frac{1/8 - \delta}{1/8 + \delta}\right)^4,$$

does not vanish, implying that A, B, and C are not jointly independent. The distribution belongs to the graphical log-linear model that can be identified with the graph shown in row 1 of Table 5. Further, since each pairwise independence holds, the distribution is Markov to the DAGs shown in rows 2, 3, and 4 of Table 5. However, the distribution is not faithful to these DAGs and it is not Markov to any of the nested DAGs (rows 5, 6, 7 and 8). $\qquad\square$

The association structure of a distribution that is unfaithful to every model in a given class can be considered within a larger model class. We have described examples of discrete distributions for which there is no undirected graphical model or DAG model to which they are faithful. Graphical models (directed and undirected) for discrete distributions are a subclass of hierarchical marginal log-linear models [Bergsma and Rudas, 2002, Rudas, Bergsma, and Németh, 2006] and can be considered within this larger class. We revisit Example 2.2 to motivate the introduction of *parametric faithfulness*, a generalization of the concept of faithfulness that can be applied to the class of hierarchical marginal log-linear models. In the following, we show that under this natural generalization of faithfulness, we can find a model in the class of hierarchical marginal log-linear models to which the distribution described in Example 2.2 is faithful.

Example 2.2 (revisited): A marginal log-linear parameterization [Bergsma and Rudas, 2002] for the DAG in Figure 2 can be derived from the set of marginals

$$\mathcal{M} = \{(A, D), (A, B, C), (A, B, C, D)\}.$$

6

The corresponding parameters are:

$$\lambda_{\emptyset}^{AD}, \ \lambda_{A*}^{AD}, \ \lambda_{*D}^{AD}, \ \lambda_{AD}^{AD}, \ \lambda_{*B*}^{ABC}, \ \lambda_{**C}^{ABC}, \ \lambda_{*BC}^{ABC}, \ \lambda_{AB*}^{ABC}, \ \lambda_{A*C}^{ABC}, \ \lambda_{ABC}^{ABC},$$

$$\lambda_{*B*D}^{ABCD}, \ \lambda_{**CD}^{ABCD}, \ \lambda_{AB*D}^{ABCD}, \ \lambda_{A*CD}^{ABCD}, \ \lambda_{*BCD}^{ABCD}, \ \lambda_{ABCD}^{ABCD}. \tag{2}$$

The conditional independencies $A \perp\!\!\!\perp C \mid B$ and $B \perp\!\!\!\perp D \mid A, C$ are obtained by taking

$$\lambda_{A*C}^{ABC} = 0, \ \lambda_{ABC}^{ABC} = 0, \ \lambda_{*B*D}^{ABCD} = 0, \ \lambda_{AB*D}^{ABCD} = 0, \ \lambda_{*BCD}^{ABCD} = 0, \ \lambda_{ABCD}^{ABCD} = 0. \tag{3}$$

Any distribution that is Markov to the DAG in Figure 2 can be parameterized by the remaining marginal log-linear parameters. The faithfulness relation in the class of marginal log-linear models can be defined as a relationship between the parameters of a distribution and the parameters of a model that contains the distribution. A distribution which also satisfies the marginal independence $A \perp\!\!\!\perp D$, has $\lambda_{AD}^{AD} = 0$ and thus belongs to a nested marginal log-linear model, to which it is faithful in the parametric sense. $\qquad\square$

This example motivates taking a parametric approach (instead of a graphical approach) to faithfulness. In the next section we introduce the concept of parametric faithfulness for discrete distributions more formally.

2.2   Parametric faithfulness

Let $\mathcal{P}$ denote the full exponential family of distributions. We choose a mixed parameterization $(\boldsymbol{\mu}, \boldsymbol{\nu})$ of this family, where $\boldsymbol{\mu}$ denotes the vector of mean value parameters and $\boldsymbol{\nu}$ the vector of canonical parameters [cf. Barndorff-Nielsen, 1978]. Let $\mathfrak{C}$ be a class of partially ordered exponential families that are obtained by setting some of the components of $\boldsymbol{\mu}$ and/or some of the components of $\boldsymbol{\nu}$ to zero, and let $\mathcal{M} \in \mathfrak{C}$. Assume that $\mathcal{M}$ is parameterized by $(\boldsymbol{\mu}_{\mathcal{M}}, \boldsymbol{\nu}_{\mathcal{M}})$, where $\boldsymbol{\mu}_{\mathcal{M}} \subseteq \boldsymbol{\mu}$, $\boldsymbol{\nu}_{\mathcal{M}} \subseteq \boldsymbol{\nu}$, $\boldsymbol{\mu} \setminus \boldsymbol{\mu}_{\mathcal{M}} = \mathbf{0}$, and $\boldsymbol{\nu} \setminus \boldsymbol{\nu}_{\mathcal{M}} = \mathbf{0}$. We define faithfulness as a relationship between the parameters of a distribution and the parameters of a model containing the distribution under consideration.

Definition 2.2. A distribution $\boldsymbol{p} \in \mathcal{M}$ parameterized by $(\boldsymbol{\mu}_{\mathcal{M}}(\boldsymbol{p}), \boldsymbol{\nu}_{\mathcal{M}}(\boldsymbol{p}))$, satisfies the *parametric faithfulness relation with respect to* $\mathcal{M}$ if none of the components of $\boldsymbol{\mu}_{\mathcal{M}}(\boldsymbol{p})$ or $\boldsymbol{\nu}_{\mathcal{M}}(\boldsymbol{p})$ vanish.

The class of discrete exponential families, where the canonical parameters are the interactions of the variables in V of order up to $K - 1$, corresponds to the class of hierarchical log-linear models on V. More precisely, let $\mathcal{M} = \{M_1, \ldots, M_T\}$ be a set of incomparable subsets of V. Then the hierarchical log-linear model generated by $\mathcal{M}$ is the set of distributions in $\mathcal{P}$ that satisfy

$$\log p_{\boldsymbol{i}} = \sum_{M \subseteq V : M \subseteq M_j \in \mathcal{M}} \gamma_M(\boldsymbol{i}_M), \tag{4}$$

where $\gamma_{M'}(\boldsymbol{i}_{M'}) = 0$ implies $\gamma_{M''}(\boldsymbol{i}_{M''}) = 0$ for any $M'' \supseteq M'$, and γ_M are called the *interaction parameters* (interactions for short). Their identifiability is assumed in the sequel.

7

The set $\mathcal{M}$ partitions the power set of V into a descending class, consisting of subsets of $M_1, \ldots, M_T$, and a complementary ascending class. The partition induces a mixed parameterization of $\mathcal{P}$ with the canonical parameters equal to the conditional odds ratios (or their logarithms) of the subsets in the ascending class, given the remaining variables, and the mean value parameters equal to the marginal distributions of the subsets in the descending class. Under this parameterization, the canonical parameters of the distributions in the model generated by $\mathcal{M}$ are equal to 1 (or zero) and the distributions are parameterized by the mean value parameters [Rudas, 1998]. The structure of the highest order interactions of the distributions in the hierarchical log-linear model generated by $\mathcal{M} = \{M_1, \ldots, M_T\}$ is described next. In the sequel, $\bar{M}_t = V \setminus M_t$.

Lemma 2.3. *There exists a parameterization of $\mathcal{P}$ under which, for every $t = 1, \ldots, T$, the interaction parameter γ_t is equal to the logarithm of the conditional odds ratio of M_t given $\bar{M}_t = \boldsymbol{i}_{\bar{M}_t}$, and is invariant of the choice of $\boldsymbol{i}_{\bar{M}_t}$.*

Proof. There exists a marginal log-linear parameterization of $\mathcal{P}$ under which for every $t = 1, \ldots, T$, the interaction parameter γ_t, corresponding to the generating marginal M_t, is the average log conditional odds ratio of M_t conditioned on and averaged over $\bar{M}_t$ [cf. Rudas et al., 2006]:

$$\gamma_{M_t} = \frac{1}{|\mathcal{I}_{\bar{M}_t}|} \sum_{\boldsymbol{i}_{\bar{M}_t}} \log \mathcal{COR}(M_t \mid \bar{M}_t = \boldsymbol{i}_{\bar{M}_t}).$$

Since M_t is a maximal interaction, $\mathcal{COR}(M' \mid \bar{M}' = \boldsymbol{i}_{\bar{M}'}) = 1$, for any $M' \supsetneq M_t$.

Further, it can be shown by induction on the elements of the ascending class of $M_1, \ldots, M_T$, that

$$\mathcal{COR}(M' \mid \bar{M}' = \boldsymbol{i}_{\bar{M}'}) = \frac{\mathcal{COR}(M_t \mid (M' \setminus M_t) \cup \bar{M}' = (\boldsymbol{i}_{M' \setminus M_t}, \boldsymbol{i}_{\bar{M}'}))}{\mathcal{COR}(M_t \mid (M' \setminus M_t) \cup \bar{M}' = (\boldsymbol{j}_{M' \setminus M_t}, \boldsymbol{i}_{\bar{M}'}))} = \frac{\mathcal{COR}(M_t \mid \bar{M}_t = \boldsymbol{i}_{\bar{M}_t})}{\mathcal{COR}(M_t \mid \bar{M}_t = \boldsymbol{j}_{\bar{M}_t})},$$

and thus,

$$\log \mathcal{COR}(M_t \mid \bar{M}_t = \boldsymbol{i}_{\bar{M}_t}) = \log \mathcal{COR}(M_t \mid \bar{M}_t = \boldsymbol{j}_{\bar{M}_t}),$$

for any $\boldsymbol{i}_{\bar{M}_t}$ and $\boldsymbol{j}_{\bar{M}_t}$. Hence,

$$\gamma_{M_t} = \log \mathcal{COR}(M_t \mid \bar{M}_t = \boldsymbol{i}_{\bar{M}_t}),$$

for any $\boldsymbol{i}_{\bar{M}_t}$. $\qquad\qquad\square$

The association structure of a discrete distribution in a hierarchical log-linear model generated by $\mathcal{M}$ can be described with a hypergraph, $\mathcal{H} = \mathcal{H}(\mathcal{M})$ with vertices $V = \{1, \ldots, K\}$ and hyperedges equal to the generating marginals, or, equivalently, to the maximum non-vanishing interactions in $\mathcal{M}$. Faithfulness to a hypergraph is naturally defined as follows:

Definition 2.3. A distribution is *faithful to a hypergraph* $\mathcal{H}$ if the non-vanishing maximal interactions of this distribution coincide with hyperedges of $\mathcal{H}$.

This definition implies that a distribution in the log-linear model generated by $\mathcal{M} = \{M_1, \ldots M_T\}$ is faithful to the hypergraph with hyperedges $M_1, \ldots M_T$ if, for all $t \in \{1, \ldots, T\}$, none of the conditional odds ratios of M_t given the variables in $\bar{M}_t = V \setminus M_t$ is equal to 1. In the following result, we show that the class of hypergraphs is closed under the parametric faithfulness relation.

8

Theorem 2.4. *The class of hypergraphs in $\mathcal{P}$ is closed under the faithfulness relation specified by Definition 2.3.*

Proof. Let $P \in \mathcal{P}$. In the following, we show that there exists a hypergraph to which P is faithful. First, derive the ascending class, $\mathcal{A}$, of subsets of V such that the log conditional odds ratios of the elements of $\mathcal{A}$ given the remaining variables vanish on P. Next, find the maximal (with respect to inclusion) elements, $M_1, \ldots, M_T$, of the complement of $\mathcal{A}$. Then, by construction, P is faithful to the hypergraph with hyperedges $M_1, \ldots, M_T$. $\qquad\square$

Remark 2.2. This paper is solely concerned with discrete distributions. However, it is worth pointing out that Definition 2.2 makes sense for exponential families in general. In particular, multivariate normal distributions can be described using an exponential family whose canonical parameters correspond to pairwise interactions between the random variables in V. We mentioned in Remark 2.1 that there are examples of distributions in Gaussian DAG models that are not Markov to any nested DAG, and hence the class of Gaussian DAG models is not closed under the faithfulness relation. However, the class of multivariate normal exponential families is closed under the parametric faithfulness relation. This is the case since setting an additional canonical parameter to zero leads to a nested exponential family.

3   Parametric Strong-Faithfulness

In order to test statistical hypotheses when working with data, a stronger version of faithfulness is needed. In this section, we generalize the notion of parametric faithfulness to parametric strong-faithfulness and discuss difficulties arising with this concept in the discrete setting.

3.1   Strong-faithfulness in the discrete setting

A distribution in a model is faithful to it if the model fully describes the conditional independence structure in this distribution. It is further called *strong-faithful* if the conditional dependencies present in the distribution are strong enough. The concept of strong-faithfulness, originally defined by Zhang and Spirtes [2003], is usually applied to multivariate normal distributions: For a given $\lambda > 0$, a multivariate normal distribution in a DAG model is λ-*strong-faithful* with respect to this DAG if all non-zero partial correlations are bounded away from zero by λ. A formal definition of strong-faithfulness in the discrete case has not been proposed, although some analogies were used. For example, Zuk, Margel, and Domany [2006] made use of the assumption that the conditional probabilities in a Bayesian network are bounded between λ and $1 - \lambda$. This can be seen as a form of strong-faithfulness.

In the discrete setting, one problem is that many variants of strong-faithfulness relations can be considered. Whether a distribution is λ-strong-faithful to a model, depends on the choice of parameterization and the measure of association. This is illustrated in the following example for two binary random variables.

Example 3.1. Let $V = \{A, B\}$ and consider the saturated model $[AB]$, which allows for interaction between A and B. A distribution in which this interaction vanishes is unfaithful to $[AB]$ and belongs to the model of independence, $A \perp\!\!\!\perp B$. A distribution in which the

association between A and B is strong enough is called strong-faithful to $[AB]$. While in the multivariate normal setting the partial correlations are a standard measure of association, in the discrete setting there are many viable choices of association measures, see Goodman and Kruskal [1954]. Table 1 illustrates different possible definitions of strong-faithfulness based on three different measures of association, the log odds ratio, ϕ_1, Yule's coefficient of association, ϕ_2, and the absolute difference between the conditional probabilities, ϕ_3. In all three cases, the parameter space has finite volume. So it is possible to estimate the proportion (relative volume) of distributions that do not satisfy the λ-strong-faithfulness relation with respect to $[AB]$. Figure 3 shows that this proportion varies considerably depending on the chosen parameterization and association measure. □

The proportion of distributions in a model that do not satisfy the strong-faithfulness relation with respect to this model is of importance for model selection procedures, which are often based on the strong-faithfulness assumption. Lemma 2.3 justifies the use of ϕ_1 to define strong-faithfulness in the discrete case. In the following, we propose the concept of strong-faithfulness to a hypergraph and, assuming strong-faithfulness, prove existence of uniformly consistent estimators of the hypergraph parameters.

3.2   Strong-faithfulness with respect to a hypergraph

Let $\mathcal{H}$ be the hypergraph generated by a set of marginals $\mathcal{M} = \{M_1, \ldots, M_T\}$. For $\boldsymbol{p} \in \mathcal{H}$ let $\boldsymbol{\gamma}(\boldsymbol{p}) = (\gamma_1(\boldsymbol{p}), \ldots, \gamma_T(\boldsymbol{p}))$ denote the set of interaction parameters of $\boldsymbol{p}$ corresponding to the hyperedges of $\mathcal{H}$.

Table 1: Selected parameterizations, measures of association, and strong-faithfulness conditions for the 2×2 contingency table.

Parametrization	Parameter space	Variation independence	Association function; λ-strong-faithfulness condition
Cell probabilities: $p_{00}, p_{01},$ p_{10}, p_{11}	Simplex Δ_3	No	$\phi_1 = \left\lvert \log\left(\frac{p_{00}p_{11}}{p_{01}p_{10}}\right) \right\rvert > \lambda$ $\phi_2 = \left\lvert \frac{p_{00}p_{11} - p_{01}p_{10}}{p_{00}p_{11} + p_{01}p_{10}} \right\rvert > \lambda$
Conditional probabilities: $\theta_1 = \mathbb{P}(A = 0),$ $\theta_2 = \mathbb{P}(B = 0 \mid A = 0),$ $\theta_3 = \mathbb{P}(B = 0 \mid A = 1)$	$(0,1)^3$	Yes	$\phi_3 = \lvert \theta_2 - \theta_3 \rvert > \lambda$

Definition 3.1. For $\lambda > 0$, a distribution $\boldsymbol{p} \in \mathcal{H}$ is λ-*strong-faithful* to $\mathcal{H}$ if

$$\min\{|\gamma_1(\boldsymbol{p})|, \ldots, |\gamma_T(\boldsymbol{p})|\} > \lambda. \tag{5}$$

As described in Section 3.1, one can, in principle, use different measures of association to define strong-faithfulness. The advantage of the definition given here is that it generalizes the original definition of strong-faithfulness given by Zhang and Spirtes [2003]. For a hypergraph generated by two-way marginals the interactions $\boldsymbol{\gamma}(\boldsymbol{p})$ are analogous to partial correlations of a multivariate normal distribution [cf. Wermuth, 1976]. Therefore, the definition of strong-faithfulness to a hypergraph proposed here is consistent with the original definition of strong-faithfulness of a multivariate normal distribution with respect to a DAG given by Zhang and Spirtes [2003]. In addition, as we will show in Section 3.3, strong-faithfulness with respect to a hypergraph allows to build uniformly consistent algorithms for learning hypergraphs.

In the following example, we illustrate the concept of strong-faithfulness with respect to a hypergraph for distributions on the $2 \times 2 \times 2$ contingency table.

Example 3.2. Let $V = \{A, B, C\}$. A distribution of V can be parameterized by

$$\log \boldsymbol{p} = \mathbf{M}\boldsymbol{\gamma},$$

where

$$\mathbf{M} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix},$$
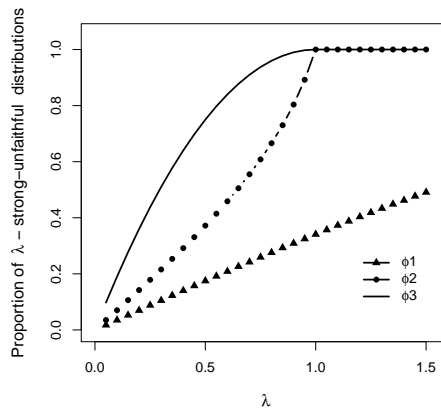

Figure 3: The proportions of distributions that are not λ-strong-faithful to the model $[AB]$ with respect to different association measures, see Example 3.1.

11

and
$$\boldsymbol{\gamma} = (\gamma^{\emptyset}, \gamma_1^A, \gamma_1^B, \gamma_1^C, \gamma_{11}^{AB}, \gamma_{11}^{AC}, \gamma_{11}^{BC}, \gamma_{111}^{ABC})$$

are the interaction parameters corresponding to the marginal distributions indicated in the superscript. The matrix $\mathbf{M}$ is of full rank, and it can easily be shown that

$$\gamma^{\emptyset} = \log p_{000}, \qquad \gamma_1^A = \log \frac{p_{100}}{p_{000}},$$

$$\gamma_1^B = \log \frac{p_{010}}{p_{000}}, \qquad \gamma_1^C = \log \frac{p_{001}}{p_{000}},$$

$$\gamma_{11}^{AB} = \log \frac{p_{000}p_{110}}{p_{010}p_{100}}, \quad \gamma_{11}^{AC} = \log \frac{p_{000}p_{101}}{p_{001}p_{100}},$$

$$\gamma_{11}^{BC} = \log \frac{p_{000}p_{011}}{p_{001}p_{010}}, \quad \gamma_{111}^{ABC} = \log \frac{p_{001}p_{010}p_{100}p_{111}}{p_{000}p_{011}p_{101}p_{110}}.$$

In the following table we give the λ-strong-faithfulness conditions for several hypergraph models:

Hyperedges	Strong-faithfulness constraints
$\{ABC\}$	$\|\gamma_{111}^{ABC}\| > \lambda$
$\{AB\}, \{AC\}, \{BC\}$	$\min\{\|\gamma_{11}^{AB}\|, \|\gamma_{11}^{AC}\|, \|\gamma_{11}^{BC}\|\} > \lambda$
$\{AC\}, \{BC\}$	$\min\{\|\gamma_{11}^{AC}\|, \|\gamma_{11}^{BC}\|\} > \lambda$
$\{A\}, \{BC\}$	$\min\{\|\gamma_1^A\|, \|\gamma_{11}^{BC}\|\} > \lambda$
$\{A\}, \{B\}, \{C\}$	$\min\{\|\gamma_1^A\|, \|\gamma_1^B\|, \|\gamma_1^C\|\} > \lambda$

3.3   Hypergraph search

In this section, we discuss how to construct hypothesis tests, when the association measure is based on the interaction parameters $\boldsymbol{\gamma}(\boldsymbol{p})$, and how to perform a hypergraph search based on these hypothesis tests.

Let $\mathcal{H}$ be a hypergraph generated by the marginals $M_1, \ldots, M_T$ and let $\gamma_1, \ldots, \gamma_T$ be the corresponding interaction parameters. We denote by $\mathcal{H}_\lambda$ the set of distributions that are λ-strong-faithful to the hypergraph $\mathcal{H}$, i.e.,

$$\mathcal{H}_\lambda = \{\boldsymbol{p} \in \mathcal{H} : \min\{|\gamma_1(\boldsymbol{p})|, \ldots, |\gamma_T(\boldsymbol{p})|\} > \lambda\},$$

and define

$$\mathcal{H}_{\lambda,\delta} = \mathcal{H}_\lambda \cap \{\boldsymbol{p} \in \mathcal{P} : p_{\boldsymbol{i}} \in [\delta, 1), \sum_{\boldsymbol{i} \in \mathcal{I}} p_{\boldsymbol{i}} = 1\},$$

where $\delta > 0$ is small enough so $\mathcal{H}_{\lambda,\delta}$ is not empty.

If M_t, for $t \in \{1, \ldots, T\}$, is an interaction of order h_t, then the conditional odds ratio of M_t given the variables in $\bar{M}_t$ is the ratio of the product of some 2^{h_t} cell probabilities and the product of a disjoint set of 2^{h_t} cell probabilities. Since $p_{\boldsymbol{i}} \in [\delta, 1)$ for all $\boldsymbol{i} \in \mathcal{I}$, the interaction parameter $|\gamma_t(\boldsymbol{p})| \leq 2^{h_t} \log((1-\delta)/\delta)$, and, therefore,

$$|\gamma_t(\boldsymbol{p})| \leq C(\delta), \quad \text{for } t = 1, \ldots, T,$$

where $C(\delta) = 2^{\max\{h_1,\ldots,h_T\}}\log((1-\delta)/\delta)$. Here, $C(\delta)$ is an upper bound on the interaction parameters (it plays the same role as the constant M in Assumption (A4) of Kalisch and Bühlmann [2007] for the Gaussian setting).

Theorem 3.1. *Let* $\mathbf{Y}$ *have a multinomial distribution with parameters N and $\boldsymbol{p}$. Assume that, under the log-linear model corresponding to $\mathcal{H}$, the maximum likelihood estimates of the interaction parameters*

$$\hat{\boldsymbol{\gamma}}^{(N)}(\boldsymbol{p}) = (\hat{\gamma}_1^{(N)}(\boldsymbol{p}),\ldots,\hat{\gamma}_T^{(N)}(\boldsymbol{p})) = (\gamma_1^{(N)}(\hat{\boldsymbol{p}}),\ldots,\gamma_T^{(N)}(\hat{\boldsymbol{p}}))$$

exist and are unique. Then, $\hat{\boldsymbol{\gamma}}^{(N)}(\boldsymbol{p})$ is a uniformly, over $\mathcal{H}_{\lambda,\delta}$, consistent estimator of $\boldsymbol{\gamma}(\boldsymbol{p})$.

Proof. For $t \in \{1,\ldots,T\}$, $\gamma_t(\boldsymbol{p}) = \boldsymbol{c}_t'\log\boldsymbol{p}$, where $\boldsymbol{c}_t$ is a vector in $\mathbb{Z}^{|\mathcal{I}|}$ whose components are comprised of equal number, 2^{h_t}, of 1's and -1's, and some 0's. By Theorem 14.6-4 in Bishop et al. [1975], as $N \to \infty$, $\gamma_t(\boldsymbol{p})$ is asymptotically normal with mean zero and variance

$$\text{var}(\hat{\gamma}_t) = \frac{1}{N}\boldsymbol{c}_t'\text{diag}^{-1}(\boldsymbol{p})\boldsymbol{c}_t. \tag{6}$$

For every $\boldsymbol{p} \in \mathcal{H}_{\lambda,\delta}$,

$$\text{var}(\hat{\gamma}_t) \le \frac{\boldsymbol{c}_t'\boldsymbol{c}_t}{N\delta}, \tag{7}$$

and thus, using the Chebyshev inequality,

$$\mathbb{P}\left(|\hat{\gamma}_t^{(N)}(\boldsymbol{p}) - \gamma_t(\boldsymbol{p})| < \epsilon\right) = \mathbb{P}\left(|Z| < \frac{\sqrt{N}\epsilon}{\sqrt{\boldsymbol{c}_t'\text{diag}^{-1}(\boldsymbol{p})\boldsymbol{c}_t}}\right) \ge \mathbb{P}\left(|Z| < \epsilon\sqrt{\frac{N\delta}{\boldsymbol{c}_t'\boldsymbol{c}_t}}\right)$$

$$\ge 1 - \frac{\boldsymbol{c}_t'\boldsymbol{c}_t}{N\delta\epsilon^2} \ge 1 - \frac{\max\limits_{t=1,\ldots,T}(\boldsymbol{c}_t'\boldsymbol{c}_t)}{N\delta\epsilon^2}, \quad \forall\epsilon > 0, \tag{8}$$

where Z is a random variable with a standard normal distribution. Therefore, $\mathbb{P}(|\hat{\gamma}_t^{(N)}(\boldsymbol{p}) - \gamma_t(\boldsymbol{p})| < \epsilon) \to 1$ for every $t = 1,\ldots,T$, uniformly over $\boldsymbol{p} \in \mathcal{H}_{\lambda,\delta}$. Since the lower bound in (8) does not depend on t, the proof is complete. $\square$

We now address the question of how to select the threshold λ in a hypergraph learning procedure. We fix a $t \in \{1,\ldots,T\}$ and consider testing the "one-hyperedge" hypothesis $H_{0t} : \gamma_t = 0$ versus $H_{1t} : \gamma_t \ne 0$ under a significance level α. Let $\hat{\boldsymbol{p}}$ be the observed distribution and let $\hat{\gamma}_t = \gamma_t(\hat{\boldsymbol{p}}) = \boldsymbol{c}_t\log\hat{\boldsymbol{p}}$ denote the corresponding interaction parameter. By Slutsky's Theorem,

$$\sqrt{N}\frac{\hat{\gamma}_t}{\sqrt{\boldsymbol{c}_t'\text{diag}^{-1}(\hat{\boldsymbol{p}})\boldsymbol{c}_t}} \to N(0,1), \quad \text{as } N \to \infty,$$

and thus we reject the null hypothesis if

$$\frac{|\hat{\gamma}_t|}{\sqrt{\frac{1}{N}\boldsymbol{c}_t'\text{diag}^{-1}(\hat{\boldsymbol{p}})\boldsymbol{c}_t}} > z_{1-\alpha/2}, \tag{9}$$

Table 2: Possible threshold values for the parameter λ.

Hyperedge	The order of the odds ratio	λ_N^*
$[AB]$	$h = 1$	$\lambda_N^* = \frac{z_{1-\alpha/2}}{N^{1/4}} \cdot 2$
$[ABC]$	$h = 2$	$\lambda_N^* = \frac{z_{1-\alpha/2}}{N^{1/4}} \cdot 2^{3/2}$
$[ABCD]$	$h = 3$	$\lambda_N^* = \frac{z_{1-\alpha/2}}{N^{1/4}} \cdot 4$
$[ABCDE]$	$h = 4$	$\lambda_N^* = \frac{z_{1-\alpha/2}}{N^{1/4}} \cdot 2^{5/2}$

where $z_{1-\alpha/2} = \Phi^{-1}(1 - \alpha/2)$ is the corresponding quantile of the standard normal distribution. With such a procedure, the probability of wrongly rejecting the null hypothesis does not exceed α.

Theorem 3.2. *Let $\epsilon \in (0, 1/2)$ and set*

$$\lambda_N^* = \frac{z_{1-\alpha/2}}{N^{1/2-\epsilon}} \min_{t=1,\dots,T} \sqrt{\boldsymbol{c}_t' \boldsymbol{c}_t}. \tag{10}$$

For the distributions that are λ_N^*-strong-faithful to $\mathcal{H}$, the power of the one-hyperedge test approaches 1 as $N \to \infty$.

Proof. The asymptotic variance of $\hat{\gamma}_t$ is bounded below by

$$\mathrm{var}(\hat{\gamma}_t) = \frac{1}{N} \boldsymbol{c}_t' \mathrm{diag}^{-1}(\hat{\boldsymbol{p}}) \boldsymbol{c}_t \geq \frac{1}{N} \boldsymbol{c}_t' \boldsymbol{c}_t.$$

Since for an h_t-order interaction, the vector $\boldsymbol{c}_t$ has 2^{h_t} components equal to 1, 2^{h_t} components equal to -1, and the remaining components equal to zero, we have $\sqrt{\boldsymbol{c}_t' \boldsymbol{c}_t} = 2^{(h_t+1)/2}$. The distributions that are λ_N^*-strong-faithful to $\mathcal{H}$ satisfy (9) for all $t = 1, \dots, T$. For these distributions the power of the one-hyperedge test is bounded below:

$$\Phi\left(\frac{|\hat{\gamma}_t|}{\sqrt{\frac{1}{N} \boldsymbol{c}_t' \mathrm{diag}^{-1}(\hat{\boldsymbol{p}}) \boldsymbol{c}_t}} - z_{1-\alpha/2}\right) \geq \Phi\left(z_{1-\alpha/2}\left(\frac{\sqrt{\boldsymbol{c}_t' \boldsymbol{c}_t} N^{-1/2+\epsilon}}{\sqrt{\frac{1}{N} \boldsymbol{c}_t' \mathrm{diag}^{-1}(\hat{\boldsymbol{p}}) \boldsymbol{c}_t}} - 1\right)\right) \quad \text{for all } t \in \{1, \dots T\}$$

and approaches 1 as $N \to \infty$. $\qquad \square$

Examples of λ_N^* computed for hyperedges of different sizes are shown in Table 2. In this paper, we do not investigate any multiple comparison issues arising with testing several one-hyperedge hypotheses at the same time.

For learning a hypergraph, any model selection procedure for hierarchical log-linear models can be applied. A review of such procedures can be found in Edwards [2000]. Backward selection which starts from the saturated model and, using the edge removal mechanism

described by Edwards [2012], goes through a sequence of nested hypergraphs, is a polynomial time algorithm that is appropriate for high dimensions. Uniform consistency of the maximum likelihood estimates for the maximal interactions of the distributions in $\mathcal{H}_{\lambda,\delta}$ entails that backward selection is a uniformly consistent procedure and the hypergraph will be determined correctly.

4 Proportions of strong-unfaithful distributions

As shown in the previous section, strong-faithfulness ensures the existence of uniformly consistent tests for developing methods for learning the underlying hypergraph. If the parameter space has finite volume it is possible to estimate the proportion (relative volume) of the distributions that are not λ-strong-faithful to a model of interest and thus whose association structure may be discovered incorrectly. Uhler et al. [2013] analyzed the proportion of distributions that are not λ-strong-faithful to a DAG in the Gaussian setting. Partial correlations define varieties and strong-unfaithful distributions correspond to the parameters that lie in a tube around these varieties. So the relative volume of these tubes corresponds to the proportion of distributions that don't satisfy the strong-faithfulness assumption, and lower bounds on these volumes were given for different classes of DAGs. The following example illustrates how one can estimate such volumes in the discrete case.

Example 3.1 (revisited): Consider a hierarchical log-linear parameterization of the distributions on the 2×2 contingency table:

$$
\begin{aligned}
\log p_{00} &= \gamma^{\emptyset}, \\
\log p_{01} &= \gamma^{\emptyset} + \gamma_1^B, \\
\log p_{10} &= \gamma^{\emptyset} + \gamma_1^A, \\
\log p_{11} &= \gamma^{\emptyset} + \gamma_1^A + \gamma_1^B + \gamma_{11}^{AB}.
\end{aligned}
\tag{11}
$$

The interaction parameter, γ_{11}^{AB}, which was denoted by ϕ_1 in Example 3.1 and in the corresponding Table 1 and Figure 3, can be expressed in terms of conditional probabilities $\theta_1 = \mathbb{P}(A = 0)$, $\theta_2 = \mathbb{P}(B = 0 \mid A = 0)$, and $\theta_3 = \mathbb{P}(B = 0 \mid A = 1)$:

$$
\gamma_{11}^{AB} = \log\left(\frac{p_{00}p_{11}}{p_{01}p_{10}}\right) = \log\frac{\theta_2(1-\theta_3)}{(1-\theta_2)\theta_3} = \log\frac{\theta_2}{1-\theta_2} - \log\frac{\theta_3}{1-\theta_3}.
$$

Let

$$
\mathcal{H}_\lambda = \left\{(\theta_1, \theta_2, \theta_3) \in (0,1)^3 : \left|\log\frac{\theta_2}{1-\theta_2} - \log\frac{\theta_3}{1-\theta_3}\right| > \lambda\right\}.
$$

The volume of its complement, $\bar{\mathcal{H}}_\lambda$, is equal to:

$$
\begin{aligned}
\mathrm{vol}(\bar{\mathcal{H}}_\lambda) &= \mathrm{vol}\left\{(\theta_1, \theta_2, \theta_3) \in (0,1)^3 : e^{-\lambda} < \frac{\theta_2}{1-\theta_2} \cdot \frac{1-\theta_3}{\theta_3} < e^\lambda\right\} \\
&= \int_0^1 d\theta_2 \left(\frac{\theta_2}{\theta_2(1-e^{-\lambda})+e^{-\lambda}} - \frac{\theta_2}{\theta_2(1-e^\lambda)+e^\lambda}\right) \\
&= \frac{e^{2\lambda} - 2\lambda e^\lambda - 1}{(1-e^\lambda)^2},
\end{aligned}
\tag{12}
$$

where the integral was computed by substitution. The parameter space $(0,1)^3$ has a unit volume. Hence, the relative proportion of distributions that are not λ-strong-faithful to $[AB]$ is equal to $\mathrm{vol}(\bar{\mathcal{H}}_\lambda)$. For small λ this proportion is approximately $\frac{\lambda}{3}$, which is consistent with the simulation results for ϕ_1 in Figure 3. $\qquad\square$

Theorem 4.1. *Let $\mathcal{H}$ be a hypergraph whose hyperedges $M_1, \ldots, M_T$ are interactions of order $h_1, \ldots, h_T$ respectively, and let $\bar{\mathcal{H}}_\lambda$ be the set of distributions that are not λ-strong-faithful to $\mathcal{H}$. Then,*

$$\mathrm{vol}(\bar{\mathcal{H}}_\lambda) \geq \max_{t \in \{1,\ldots,T\}} \mathrm{vol}\{\boldsymbol{p} \in \mathcal{H} : \; |\gamma_t(\boldsymbol{p})| < \lambda\} \geq \max_{t \in \{1,\ldots,T\}} \left(\frac{e^{2\mu} - 2\mu e^\mu - 1}{(1 - e^\mu)^2} \right)^{2^{h_t - 1}},$$

where $\mu = \lambda/2^{h_t - 1}$

Proof. By Lemma 2.3, the parameter γ_t for $t \in \{1,\ldots,T\}$ is the log conditional odds ratio of M_t, given the variables in $\bar{M}_t$. We can express γ_t using a corresponding set of variation independent conditional probabilities:

$$\gamma_t = \log \left(\frac{\theta_1}{1 - \theta_1} \cdots \frac{\theta_{2^{h_t - 1}}}{1 - \theta_{2^{h_t - 1}}} \cdot \frac{1 - \theta_{2^{h_t - 1} + 1}}{\theta_{2^{h_t - 1} + 1}} \cdots \frac{1 - \theta_{2^{h_t}}}{\theta_{2^{h_t}}} \right).$$

Therefore,

$$\mathrm{vol}\{\boldsymbol{p} \in \mathcal{H} : \; |\gamma_t(\boldsymbol{p})| < \lambda\} = \mathrm{vol}\{\boldsymbol{\theta} \in (0,1)^m : \; |\gamma_t| < \lambda\}$$

$$= \mathrm{vol}\left\{ (\theta_1, \ldots, \theta_{2^{h_t}}) \in (0,1)^{2^{h_t}} : \; \left| \log \left(\frac{\theta_1}{1 - \theta_1} \cdots \frac{\theta_{2^{h_t - 1}}}{1 - \theta_{2^{h_t - 1}}} \cdot \frac{1 - \theta_{2^{h_t - 1} + 1}}{\theta_{2^{h_t - 1} + 1}} \cdots \frac{1 - \theta_{2^{h_t}}}{\theta_{2^{h_t}}} \right) \right| < \lambda \right\}$$

$$\geq \left(\mathrm{vol}\left\{ (\zeta_1, \zeta_2) \in (0,1)^2 : \; \left| \log \frac{\zeta_1}{1 - \zeta_1} - \log \frac{\zeta_2}{1 - \zeta_2} \right| < \frac{\lambda}{2^{h_t - 1}} \right\} \right)^{2^{h_t - 1}}$$

$$= \left(\frac{e^{2\mu} - 2\mu e^\mu - 1}{(1 - e^\mu)^2} \right)^{2^{h_t - 1}},$$

where $\mu = \lambda/2^{h_t - 1}$, and for the last equation we used (12). Since $\bar{\mathcal{H}}_\lambda = \{\boldsymbol{p} \in \mathcal{H} : |\gamma_t(\boldsymbol{p})| < \lambda \text{ for at least one } t \in \{1,\ldots,T\}\}$,

$$\mathrm{vol}(\bar{\mathcal{H}}_\lambda) \geq \mathrm{vol}\{\boldsymbol{p} \in \mathcal{H} : \; |\gamma_t(\boldsymbol{p})| < \lambda \text{ for all } t \in \{1,\ldots,T\}\},$$

and thus

$$\mathrm{vol}(\bar{\mathcal{H}}_\lambda) \geq \max_{t \in \{1,\ldots,T\}} \left(\frac{e^{2\mu} - 2\mu e^\mu - 1}{(1 - e^\mu)^2} \right)^{2^{h_t - 1}}, \quad \text{for } \mu = \lambda/2^{h_t - 1}.$$

$\qquad\square$

As we will see in the following result, for hypergraphs whose hyperedges are variation independent we can in fact give an exact formulation of the proportion of distributions that don't satisfy strong-faithfulness.

Theorem 4.2. *Let $\mathcal{H}$ be the hypergraph generated by marginals $M_1, \ldots, M_T$. Assume that there exists a parameterization of $\mathcal{H}$ under which the interaction parameters corresponding to $M_1, \ldots, M_T$ are variation independent and, further, that the parameter space has finite volume. Then, the proportion of distributions that are not λ-strong-faithful to $\mathcal{H}$ is*

$$\mathrm{vol}(\bar{\mathcal{H}}_\lambda) = 1 - (1 - \boldsymbol{\nu}_{M_1}) \cdots (1 - \boldsymbol{\nu}_{M_T}), \tag{13}$$

where $\boldsymbol{\nu}_{M_t}$, for $t = 1, \ldots, T$, is the proportion of distributions that are not λ-strong-faithful to M_t.

Proof. Consider a parameterization of $\mathcal{H}$ under which the maximal interaction parameters corresponding to $M_1, \ldots, M_T$ are variation independent. For $t \in \{1, \ldots, T\}$, let $\boldsymbol{\nu}_{M_t}$ denote the proportion of distributions that do not satisfy λ-strong-faithfulness with respect to M_t. Since the joint range of the variation independent parameters is equal to the Cartesian product of the individual ranges, the proportion of distributions that are not λ-strong-faithful to $\mathcal{H}$ is equal to $1 - (1 - \boldsymbol{\nu}_{M_1}) \cdots (1 - \boldsymbol{\nu}_{M_T})$. $\square$

The maximal interaction parameters in decomposable log-linear models [Haberman, 1974] and in ordered decomposable marginal log-linear models [Bergsma and Rudas, 2002] are variation independent, and Theorem 4.2 applies. Let $\mathcal{H}$ be the hypergraph generated by a decomposable sequence of marginals $M_1, \ldots, M_T$. The interaction parameter associated with a hyperedge M_t can be expressed using a variation independent set of conditional probabilities. Under this parameterization, the proportion of distributions that do not satisfy λ-strong-faithfulness with respect to the hyperedge M_t is equal to

$$\boldsymbol{\nu}_{h_t} = \mathrm{vol}\left\{ (\theta_1, \ldots, \theta_{2^{h_t}}) : \left| \log \left(\frac{\theta_1}{1 - \theta_1} \cdots \frac{\theta_{2^{h_t-1}}}{1 - \theta_{2^{h_t-1}}} \cdot \frac{1 - \theta_{2^{h_t-1}+1}}{\theta_{2^{h_t-1}+1}} \cdots \frac{1 - \theta_{2^{h_t}}}{\theta_{2^{h_t}}} \right) \right| < \lambda \right\}, \tag{14}$$

where h_t denotes the order of interaction M_t. The proportion of distributions that are not λ-strong-faithful to $\mathcal{H}$ is calculated using (13).

Figure 4 shows values of $\boldsymbol{\nu}_h$ as functions of h and λ. The concrete computations involved in Theorem 4.2 are illustrated in Example 4.1. We next analyze hypergraphs with a special "chain" structure and show that in this case Equation (13) simplifies.

Definition 4.1. A hypergraph $\mathcal{H}$ is called a chain of order h if the generating sequence of marginals $\{M_1, \ldots M_T\}$, where $\cup_{t=1}^T M_t = V$, is decomposable and all of the hyperedges correspond to h-th order interactions of the joint distribution.

For example, the hypergraph generated by $\{A, B\}, \{B, C\}, \{C, D\}, \{D, E\}$ is a chain of order 1 of length 4, and the hypergraph generated by $\{A, B, C\}$ and $\{A, B, D\}$ is a chain of order 2 of length 2.

Corollary 4.3. *Let a hypergraph $\mathcal{H}$ be a chain of order h of length L. Then the proportion of distributions that are not λ-strong-faithful to $\mathcal{H}$ is equal to*

$$1 - (1 - \boldsymbol{\nu}_h)^L,$$

where

$$\boldsymbol{\nu}_h = \mathrm{vol}\left\{ (\theta_1, \ldots, \theta_{2^h}) : \left| \log \left(\frac{\theta_1}{1 - \theta_1} \cdots \frac{\theta_{2^{h-1}}}{1 - \theta_{2^{h-1}}} \cdot \frac{1 - \theta_{2^{h-1}+1}}{\theta_{2^{h-1}+1}} \cdots \frac{1 - \theta_{2^h}}{\theta_{2^h}} \right) \right| < \lambda \right\}. \tag{15}$$
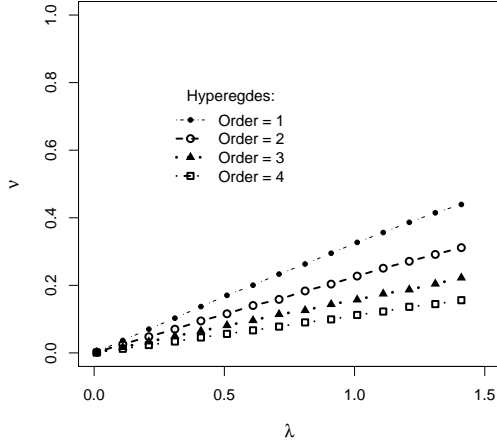
Figure 4: Proportions of distributions that do not satisfy strong-faithfulness with respect to a single hyperedge. See Equation (14).
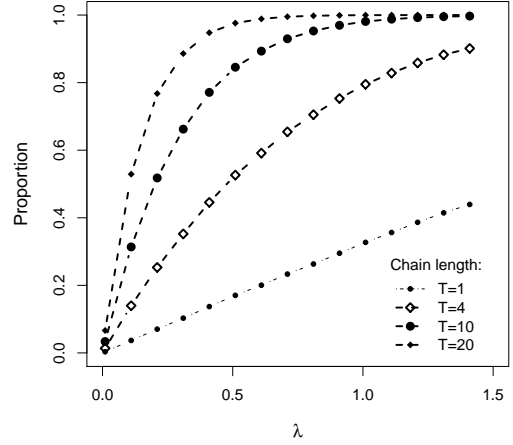
Figure 5: Proportions of distributions that are not λ-strong-faithful to a first order chain. See Corollary 4.3.

For a chain of order 1, the proportion of distributions that are not λ-strong-faithful to $\mathcal{H}$ is especially simple:

$$\text{vol}(\bar{\mathcal{H}}_\lambda) = 1 - (1 - \boldsymbol{\nu}_1)^T,$$

where

$$\boldsymbol{\nu}_1 = \frac{e^{2\lambda} - 2\lambda e^\lambda - 1}{(1 - e^\lambda)^2}. \tag{16}$$

The proportions for chains of several orders were estimated using Monte-Carlo method and are displayed in Figure 5.

Example 4.1. We demonstrate the volume computation using the chain $[AB][BC][CD]$ of order 1. The maximal interaction parameters corresponding to the hyperedges are:

$$
\begin{aligned}
\gamma_1 &= \log \mathcal{COR}(AB \mid CD) = \log \frac{p_{00kl}p_{11kl}}{p_{01kl}p_{10kl}}, \\
\gamma_2 &= \log \mathcal{COR}(BC \mid AD) = \log \frac{p_{i00l}p_{i11l}}{p_{i01l}p_{i10l}}, \\
\gamma_3 &= \log \mathcal{COR}(CD \mid AB) = \log \frac{p_{ij00}p_{ij11}}{p_{ij01}p_{ij10}},
\end{aligned}
$$

where $i, j, k, l \in \{0, 1\}$ are fixed categories of A, B, C, D respectively. The chain can be described by two conditional independence relations: $A \perp\!\!\!\perp C \mid B$ and $AB \perp\!\!\!\perp D \mid C$. Thus

the distributions in a chain model can be parameterized using the conditional probabilities:

$$\theta_0 = \mathbb{P}(B = 0),$$
$$\theta_{10} = \mathbb{P}(A = 0 \mid B = 0), \quad \theta_{11} = \mathbb{P}(A = 0 \mid B = 1),$$
$$\theta_{20} = \mathbb{P}(C = 0 \mid B = 0), \quad \theta_{21} = \mathbb{P}(C = 0 \mid B = 1),$$
$$\theta_{30} = \mathbb{P}(D = 0 \mid C = 0), \quad \theta_{31} = \mathbb{P}(D = 0 \mid C = 1).$$

The parameters $\boldsymbol{\theta}$ are variation independent, and, for $t = 1, 2, 3$,

$$\gamma_t = \log \left(\frac{\theta_{t0}}{1 - \theta_{t0}} \cdot \frac{1 - \theta_{t1}}{\theta_{t1}} \right).$$

Let

$$\boldsymbol{\nu}_1 = \mathrm{vol}\left\{ (\theta_1, \theta_2) \in (0, 1)^2 : \left| \log \frac{\theta_1}{1 - \theta_1} - \log \frac{\theta_2}{1 - \theta_2} \right| < \lambda \right\} = \frac{e^{2\lambda} - 2\lambda e^{\lambda} - 1}{(1 - e^{\lambda})^2}.$$

Using the binomial formula, we obtain that $\mathrm{vol}(\bar{\mathcal{H}}_\lambda) = 1 - (1 - \boldsymbol{\nu}_1)^3$. $\qquad\square$

Example 2.1 (revisited): The maximal non-vanishing interactions of a distribution that is faithful to a hypergraph $\mathcal{H}$ with hyperedges $\{A, B, C\}$ and $\{A, B, D\}$ can be described using the interaction parameters equal to the logarithm of the second order conditional odds ratios of ABC given D and of ABD, given C:

$$\gamma_0^{ABC} = \log \mathcal{COR}(ABC \mid D = 0),$$
$$\gamma_1^{ABC} = \log \mathcal{COR}(ABC \mid D = 1),$$
$$\gamma_0^{ABD} = \log \mathcal{COR}(ABD \mid C = 0),$$
$$\gamma_1^{ABD} = \log \mathcal{COR}(ABD \mid C = 1).$$

Using conditional probabilities,

$$\theta_1 = \mathbb{P}(C = 0 \mid A = 0, B = 0), \quad \theta_2 = \mathbb{P}(C = 0 \mid A = 0, B = 1),$$
$$\theta_3 = \mathbb{P}(C = 0 \mid A = 1, B = 0), \quad \theta_4 = \mathbb{P}(C = 0 \mid A = 1, B = 1),$$
$$\theta_5 = \mathbb{P}(D = 0 \mid A = 0, B = 0), \quad \theta_6 = \mathbb{P}(D = 0 \mid A = 0, B = 1),$$
$$\theta_7 = \mathbb{P}(D = 0 \mid A = 1, B = 0), \quad \theta_8 = \mathbb{P}(D = 0 \mid A = 1, B = 1),$$

one obtains

$$\gamma_0^{ABC} = \gamma_1^{ABC} = \log \frac{\theta_1}{1 - \theta_1} + \log \frac{\theta_4}{1 - \theta_4} - \log \frac{\theta_2}{1 - \theta_2} - \log \frac{\theta_3}{1 - \theta_3},$$
$$\gamma_0^{ABD} = \gamma_1^{ABD} = \log \frac{\theta_5}{1 - \theta_5} + \log \frac{\theta_8}{1 - \theta_8} - \log \frac{\theta_6}{1 - \theta_6} - \log \frac{\theta_7}{1 - \theta_7}.$$

A distribution is not λ-strong-faithful to the hypergraph $\mathcal{H}$ if at least one of the following inequalities holds:

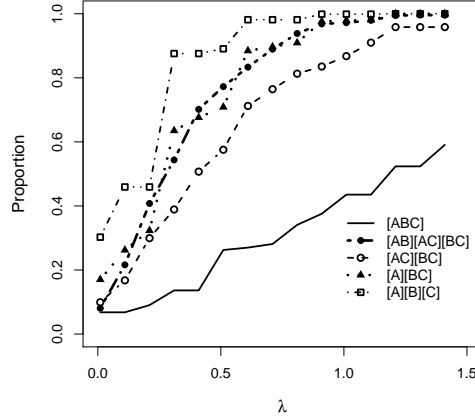$$|\gamma_0^{ABC}| < \lambda, \text{ or } |\gamma_0^{ABD}| < \lambda.$$

Figure 6: Proportions of distributions that are not projected-λ-strong-faithful computed for several hypergraphs on the $2 \times 2 \times 2$ contingency table.

Hence, the proportion of distributions that are not λ-strong-faithful to the hypergraph $\mathcal{H}$ is equal to $1 - (1 - \boldsymbol{\nu}_2)^2$, where

$$\boldsymbol{\nu}_2 = \mathrm{vol}\left\{(\theta_1, \ldots, \theta_4) \in (0,1)^4 : \left|\log \frac{\theta_1}{1 - \theta_1} + \log \frac{\theta_4}{1 - \theta_4} - \log \frac{\theta_2}{1 - \theta_2} - \log \frac{\theta_3}{1 - \theta_3}\right| < \lambda\right\}.$$

We were not able to find a closed-form expression for $\boldsymbol{\nu}_2$. It can be shown that $\boldsymbol{\nu}_2$ is bounded above by $\boldsymbol{\nu}_1$ and thus the volume of distributions that are not λ-strong-faithful to the hypergraph $\mathcal{H}$ is bounded above by the volume computed for the chain of the same length of order 1, that is, $\mathrm{vol}(\bar{\mathcal{H}}_\lambda) \leq 1 - (1 - \boldsymbol{\nu}_1)^2$. $\qquad \square$

Remark 4.1. The concept of strong-faithfulness can be extended to distributions that do not belong to a given hypergraph model. Let $\boldsymbol{p}, \boldsymbol{q} \in \mathcal{P}$, and let ρ be a divergence function. The distance from a distribution $\boldsymbol{p}$ to a hypergraph $\mathcal{H}$ can be defined as

$$\rho(\boldsymbol{p}, \mathcal{H}) = \min_{\boldsymbol{q} \in \mathcal{H}} \rho(\boldsymbol{p}, \boldsymbol{q}). \tag{17}$$

In particular, $\rho(\boldsymbol{p}, \mathcal{H}) = 0$ if and only if $\boldsymbol{p} \in \mathcal{H}$. We denote by $\boldsymbol{p}_{\mathcal{H},\rho}$ the projection of $\boldsymbol{p}$ onto the hypergraph $\mathcal{H}$, i.e.:

$$\boldsymbol{p}_{\mathcal{H},\rho} = \operatorname*{argmin}_{\boldsymbol{q} \in \mathcal{H}} \rho(\boldsymbol{p}, \boldsymbol{q}),$$

and call a distribution $\boldsymbol{p}$ *projected-λ-strong-faithful to* $\mathcal{H}$ with respect to ρ (for $\lambda > 0$) if $\boldsymbol{p}_{\mathcal{H},\rho}$ is λ-strong-faithful to $\mathcal{H}$. The concept of projected-strong-faithfulness is relevant in various estimation procedures.

We end by illustrating the concept of projected-strong-faithfulness by estimating the proportions of projected-λ-strong-faithful distributions for several hypergraph models on the $2 \times 2 \times 2$ contingency table.

Example 3.2 (revisited): To determine the distance from a distribution $\boldsymbol{p}$ to a hypergraph $\mathcal{H}$ we use the likelihood function under the corresponding log-linear model. Relative frequencies of distributions that do not satisfy the projected-λ-strong-faithfulness relation for different hypergraphs and different values of λ are displayed in Figure 6. $\qquad \square$

20

5   Conclusion

We demonstrated that the association structure of discrete data can be very complex, and some distributions are not faithful to any undirected graphical model or any DAG. Thus, the attractive simplicity of graphical models may be misleading. In Section 2, we proposed the concept of parametric faithfulness, which can be applied to any exponential family, including those which cannot be specified using Markov properties.

We considered the class of hypergraphs which can be identified with hierarchical log-linear models. We showed that for any distribution there exists a hypergraph to which it is parametrically faithful and suggested to conduct the search in this class. As the class also contains graphical models, if a model structure which can be described by a graph is appropriate, it will be discovered (see the consistency result in Section 3).

Our work is relevant for the popular causal search algorithms, referred to in Section 1, which assume (strong-) faithfulness. The findings described in Sections 3 and 4 imply that, depending on the quantitative expression for association and on the choice of the cut-off parameter λ, to define strong-faithfulness, the resulting model selection procedures may yield different results for the same data.

Acknowledgements
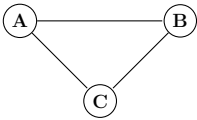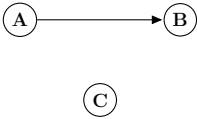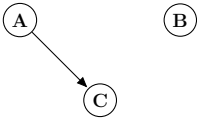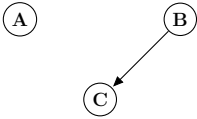
References
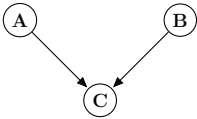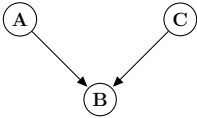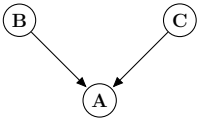
A. Agresti. *Categorical data analysis*. Wiley, New York, 2002.

O. E. Barndorff-Nielsen. *Information and exponential families*. Wiley, New York, 1978.

W. P. Bergsma and T. Rudas. Marginal models for categorical data. *Ann. Statist.*, 30: 140–159, 2002.

Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete multivariate analysis: theory and practice*. MIT, 1975.

D. Edwards. *Introduction to graphical modeling*. Springer, London, 2000.

D. Edwards. A note on adding and deleting edges in hierarchical log-linear models. *Comput. Stat.*, 27:799–803, 2012.

L. A. Goodman and W. H. Kruskal. Measures of association for cross classifications. *J. Amer. Statist. Assoc.*, 49:732–764, 1954.

S. J. Haberman. *The analysis of frequency data*. The University of Chicago Press, 1974.

M. Kalisch and P. Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8:613–636, 2007.

C. Meek. Strong completeness and faithfulness in Bayesian networks. In P. Besnard and S. Hanks, editors, *Proceedings of the Eleventh conference on Uncertainty in Artificial Intelligence (UAI 1995)*, pages 411–418. Morgan Kauffmann, 1995.

J. M. Peña, R. Nilsson, J. Björkegren, and J. Tegnér. An algorithm for reading dependencies from the minimal undirected independence map of a graphoid that satisfies weak transitivity. *Journal of Machine Learning Research*, 10:1071–1094, 2009.

J. M. Robins, R. Scheines, P. Spirtes, and L. Wasserman. Uniform consistency in causal inference. *Biometrika*, 90:491–515, 2003.

T. Rudas. *Odds ratios in the analysis of contingency tables*, volume 119 of *Quantative Applications in the Social Sciences*. Sage Publications, Inc., 1998.

T. Rudas, W. Bergsma, and R. Németh. Parameterization and estimation of path models for categorical data. In A. Rizzi and M. Vichi, editors, *COMPSTAT 2006 - Proceedings in Computational Statistics*, pages 383–394. Physica Verlag Heidelberg, 2006.

P. Spirtes, C. Glymour, and R. Scheines. *Causation, prediction and search*. MIT Press, 2001.

M. Studený. *Probabilistic conditional independence structures*. Springer, London, 2005.

C. Uhler and G. Raskutti. Learning directed acyclic graphs based on sparsest permutations. arXiv: 1307.0366, 2013.

C. Uhler, G. Raskutti, P. Bühlmann, and B. Yu. Geometry of faithfulness assumption in causal inference. *Ann. Statist.*, 41:436–463, 2013.

N. Wermuth. Analogies between multiplicative models in contingency tables and covariance selection. *Biometrics*, 32:95–108, 1976.

J. Zhang and P. Spirtes. Strong faithfulness and uniform consistency in causal inference. In U. Kjærulff and C. Meek, editors, *Proceedings of the Nineteenth conference on Uncertainty in Artificial Intelligence (UAI 2003)*, pages 632–639. Morgan Kaufmann, San Francisco, CA, USA, 2003.

J. Zhang and P. Spirtes. Detection of unfaithfulness and robust causal inference. *Minds and Machines*, 18:239–271, 2008.

O. Zuk, S. Margel, and E. Domany. On the number of samples needed to learn the correct structure of a Bayesian network. In R. Dechter and T. Richardson, editors, *Proceedings of the Twenty-Second international conference on Uncertainty in Artificial Intelligence (UAI 2006)*, pages 560–567. AUAI Press, Arlington, Virginia, 2006.

Table 3: Some graphical models on three nodes.

	Log-Linear Model	Conditional Independence	Graph
1	[ABC]	None	
2	[A][B]	$A \perp\!\!\!\perp B$	
3	[A][C]	$A \perp\!\!\!\perp C$	
4	[B][C]	$B \perp\!\!\!\perp C$	
5	[AB][C]	$AB \perp\!\!\!\perp C$	
6	[AC][B]	$AC \perp\!\!\!\perp B$	
7	[A][BC]	$A \perp\!\!\!\perp BC$	
8	[A][B][C]	$A \perp\!\!\!\perp B, \quad A \perp\!\!\!\perp C, \quad B \perp\!\!\!\perp C,$ $A \perp\!\!\!\perp B\|C, \quad A \perp\!\!\!\perp C\|B, \quad B \perp\!\!\!\perp C\|A$	