
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Hämäläinen, Wilhelmiina; Webb, Geoffrey I.
A tutorial on statistically sound pattern discovery

Published in:
Data Mining and Knowledge Discovery

DOI:
[10.1007/s10618-018-0590-x](https://doi.org/10.1007/s10618-018-0590-x)

Published: 01/01/2019

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Hämäläinen, W., & Webb, G. I. (2019). A tutorial on statistically sound pattern discovery. *Data Mining and Knowledge Discovery*, 33(2), 325–377. <https://doi.org/10.1007/s10618-018-0590-x>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



A tutorial on statistically sound pattern discovery

Wilhelmiina Hämmäläinen^{1,2}  · Geoffrey I. Webb³

Received: 1 July 2017 / Accepted: 17 September 2018 / Published online: 20 December 2018
© The Author(s) 2018

Abstract

Statistically sound pattern discovery harnesses the rigour of statistical hypothesis testing to overcome many of the issues that have hampered standard data mining approaches to pattern discovery. Most importantly, application of appropriate statistical tests allows precise control over the risk of false discoveries—patterns that are found in the sample data but do not hold in the wider population from which the sample was drawn. Statistical tests can also be applied to filter out patterns that are unlikely to be useful, removing uninformative variations of the key patterns in the data. This tutorial introduces the key statistical and data mining theory and techniques that underpin this fast developing field. We concentrate on two general classes of patterns: dependency rules that express statistical dependencies between condition and consequent parts and dependency sets that express mutual dependence between set elements. We clarify alternative interpretations of statistical dependence and introduce appropriate tests for evaluating statistical significance of patterns in different situations. We also introduce special techniques for controlling the likelihood of spurious discoveries when multitudes of patterns are evaluated. The paper is aimed at a wide variety of audiences. It provides the necessary statistical background and summary of the state-of-the-art for any data mining researcher or practitioner wishing to enter or understand statistically sound pattern discovery research or practice. It can serve as a general introduction to the field of statistically sound pattern discovery for any reader with a general background in data sciences.

Keywords Pattern discovery · Statistical significance · Hypothesis testing · Dependency rule · Dependency set · Association rule

Responsible editor: Johannes Fürnkranz.

This research was supported by the Academy of Finland (Decision Number 307026).

✉ Wilhelmiina Hämmäläinen
wilhelmiina.hamalainen@iki.fi

Extended author information available on the last page of the article

1 Introduction

Pattern discovery is a core technique of data mining that aims at finding all patterns of a specific type that satisfy certain constraints in the data (Agrawal et al. 1993; Cooley et al. 1997; Rigoutsos and Floratos 1998; Kim et al. 2010). Common pattern types include frequent or correlated sets of variables, association and correlation rules, frequent subgraphs, subsequencies, and temporal patterns. Traditional pattern discovery has emphasized efficient search algorithms and computationally well-behaving constraints and pattern types, like frequent pattern mining (Aggarwal and Han 2014), and less attention has been paid to the statistical validity of patterns. This has also restricted the use of pattern discovery in many applied fields, like bioinformatics, where one would like to find certain types of patterns without risking costly false or suboptimal discoveries. As a result, there has emerged a new trend towards *statistically sound pattern discovery* with strong emphasis on statistical validity. In statistically sound pattern discovery, the first priority is to find genuine patterns that are likely to reflect properties of the underlying population and hold also in future data. Often the pattern types are also different, because they have been dictated by the needs of application fields rather than computational properties.

The problem of statistically sound pattern discovery is illustrated in Fig. 1. Usually, the analyst has a sample of data drawn from some population of interest. This sample is typically only a very small proportion of the total population of interest and may contain noise. The pattern discovery tool is applied to this sample, finding some set of patterns. It is unrealistic to expect this set of discovered patterns to directly match the ideal patterns that would be found by direct analysis of the real

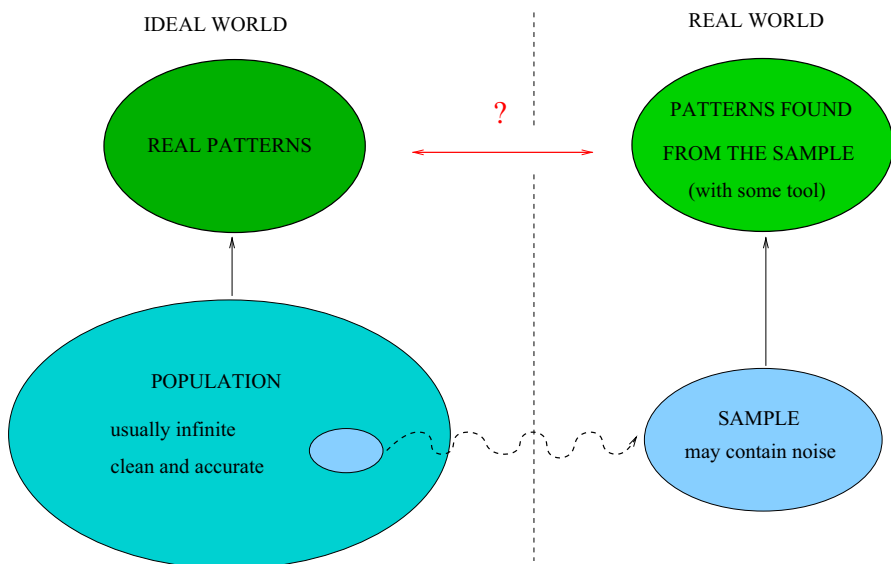


Fig. 1 An illustration of the problem of finding true patterns from sample data

population rather than a sample thereof. Indeed, it is clear that in at least some cases, the application of naive techniques results in the majority of patterns found being only spurious artifacts. An extreme example of this problem arises with the popular minimum support and minimum confidence technique (Agrawal et al. 1993) when applied to the well-known Covtype benchmark dataset from the UCI repository (Lichman 2013). The minimum support and minimum confidence technique seeks to find the frequent positive dependencies in data using thresholds for minimum frequency ('support') and precision ('confidence'). For the Covtype dataset, the top 197,183,686 rules found by minimum support and minimum confidence are in fact negative dependencies (Webb 2006). This gives rise to the suggestion that the oft cited problem of pattern discovery finding unmanageably large numbers of patterns is largely due to standard techniques returning results that are dominated by spurious patterns (Webb 2011).

There is a rapidly growing body of pattern discovery techniques being developed in the data science community that utilize statistics to control the risk of such spurious discoveries. This tutorial paper grew out of tutorials presented at ECML PKDD 2013 (Hämäläinen and Webb 2013) and KDD-14 (Hämäläinen and Webb 2014). It introduces the relevant statistical theory and key techniques for statistically sound pattern discovery. We concentrate on pattern types that express statistical dependencies between categorical attributes, such as *dependency rules* (dependencies between condition and consequent parts) and *dependency sets* (mutual dependencies between set elements). The same techniques of testing statistical significance of dependence also apply to situations where one would like to test dependencies in other types of patterns, like dependencies between subgraphs and classes or between frequent episodes.

To keep the scope manageable, we do not describe actual search algorithms but merely the statistical techniques that are employed during the search. We aim at a generic presentation that is not bound to any specific search method, pattern type or school of statistics. Instead, we try to clarify alternative interpretations of statistical dependence and the underlying assumptions on the origins of data, because they often lead to different statistical methods and also different patterns to be selected. We describe the preconditions, limitations, strengths, and shortcomings of different approaches to help the reader to select a suitable method for the problem at hand. However, we do not make any absolute recommendations, as there is no one correct way to test statistical significance or reliability of patterns. Rather, the appropriate choice is always problem-dependent.

The paper is aimed at a wide variety of audiences. The main goal is to offer a general introduction to the field of statistically sound pattern discovery for any reader with a general background in data sciences. Knowledge on the main principles, important concerns, and alternative techniques is especially useful for practical data miners (how to improve the quality or test the reliability of discovered patterns) and algorithm designers (how to target the search into the most reliable patterns). Another goal is to introduce possibilities of pattern discovery to researchers in other fields, like bio-scientists, for whom statistical significance of findings is the main concern and who would like to find new useful information from large data masses. As a prerequisite, we assume knowledge of the basic concepts of probability theory. The paper pro-

vides the necessary statistical background and summary of the state of the art, but a knowledgeable reader may well skip preliminaries (Sects. 2.2–2.3) and the overview of multiple hypothesis testing (Sect. 6.1).

In the paper we have tried to use terminology that is consistent with statistics for two reasons. First, knowing statistical terms makes it easier to consult external sources, like textbooks in statistics, for further knowledge. Second, common terminology should make the paper more readable to wider audience, like researchers from applied science who would like to extend their repertoire of statistical analysis with pattern discovery techniques. To achieve this goal, we have avoided some special terms originated in pattern discovery that have another meaning in statistics or may become easily confused in this context (see Appendix).

The rest of the paper is organized as follows. In Sect. 2 we give definitions of various types of statistical dependence and introduce the main principles and approaches of statistical significance testing. In Sect. 3 we investigate how statistical significance of dependency rules is evaluated under different assumptions. Especially, we contrast two alternative interpretations of dependency rules that are called *variable-based* and *value-based interpretations* and introduce appropriate tests for different situations. In Sect. 4 we discuss how to evaluate statistical significance of the improvement of one rule over another one. In Sect. 5 we survey the key techniques that have been developed for finding different types of statistically significant dependency sets. In Sect. 6 we discuss the problem of multiple hypothesis testing. We describe the main principles and popular correction methods and then introduce some special techniques for increasing power in the pattern discovery context. Finally, in Sect. 7 we summarize the main points and present conclusions.

2 Preliminaries

In this paper, we consider patterns that express statistical dependence. Dependence is usually defined in the negative, as absence of independence. Therefore, we begin by defining different types of statistical independence that are needed in defining dependence patterns and their relationships, like improvement of one pattern over another one. After that, we give an overview of the main principles and approaches of statistical significance testing. These approaches are applicable to virtually any pattern type, but we focus on how they are used in independence testing. In the subsequent sections, we will describe in detail how to evaluate statistical significance of dependency patterns or their improvement under different sets of assumptions.

2.1 Notations

The mathematical notations used in this paper are given in Table 1. We note that the sample space spanned by variables A_1, \dots, A_k is $\mathcal{S} = \text{Dom}(A_1) \times \dots \times \text{Dom}(A_k)$. When A_i s are binary variables $\mathcal{S} = \{0, 1\}^k$. Sample points $\mathbf{r} \in \mathcal{S}$ correspond to atomic events and all other events can be presented as their disjunctions. Often these can be presented in a reduced form, for example, event $(A_1=1, A_2=1) \vee (A_1=1, A_2=0)$ reduces to

($A_1=1$). In this paper, we focus on events that can be presented as conjunctions $\mathbf{X} = \mathbf{x}$, where $\mathbf{X} \subseteq \{A_1, \dots, A_k\}$. When it is clear from the context, we notate elements of \mathbf{X} , $|\mathbf{X}| = m$, by A_1, \dots, A_m instead of more complicated A_{i_1}, \dots, A_{i_m} , $\{i_1, \dots, i_m\} \subseteq \{1, \dots, k\}$. We also note that data set \mathcal{D} is defined as a vector of data points so that duplicate rows (i.e., rows $\mathbf{r}_i, \mathbf{r}_j$ where $\mathbf{r}_i = \mathbf{r}_j$ but $i \neq j$) can be distinguished by their index numbers.

2.2 Statistical dependence

The notion of statistical dependence is equivocal and even the simplest case, dependence between two events, is subject to alternative interpretations. Interpretations of statistical dependence between more than two events or variables are even more various. In the following, we introduce the main types of statistical independence that are needed for defining dependency patterns and evaluating their statistical significance and mutual relationships.

2.2.1 Dependence between two events

Definitions of statistical dependence are usually based on the classical notion of statistical independence between two events. We begin from a simple case where the events are variable-value combinations, $A = a$ and $B = b$.

Definition 1 (*Statistical independence between two events*) Let $A = a$ and $B = b$ be two events, $P(A = a)$ and $P(B = b)$ their marginal probabilities, and $P(A = a, B = b)$ their joint probability. Events $(A = a)$ and $(B = b)$ are *statistically independent*, if

$$P(A = a, B = b) = P(A = a)P(B = b). \quad (1)$$

Statistical dependence is seldom defined formally, but in practice, there are two approaches. If dependence is considered as a Boolean property, then any departure from complete independence (Eq. 1) is defined as dependence. Another approach, prevalent in statistical data analysis, is to consider dependence as a continuous property ranging from complete independence to complete dependence. Complete dependence itself is an ambiguous term, but usually it refers to equivalence of events: $P(A = a, B = b) = P(A = a) = P(B = b)$ (perfect positive dependence) or mutual exclusion of events: $P(A = a, B \neq b) = P(A = a) = P(B \neq b)$ (perfect negative dependence).

The strength of dependence between two events can be evaluated with several alternative measures. In pattern discovery, two of the most popular measures are *leverage* and *lift*.

Leverage is equivalent to Yule's δ (Yule 1912), Piatetsky-Shapiro's unnamed measure (Piatetsky-Shapiro 1991), and Meo's 'dependence value' (Meo 2000). It measures the absolute deviation of the joint probability from its expectation under independence:

Table 1 Notations

$\mathbb{N} = \mathbb{Z}^+ \cup \{0\}$	Natural numbers (including 0)
N_X, N_A, N_{XA}, N	Random variables in \mathbb{N}
A, B, C, \dots	Variables
$Dom(A)$	Domain of variable A
$a_1, a_2, a_3, \dots \in Dom(A)$	Values of variable A
$A_1=a_1, A_2=a_2$	Value assignment, where $A_1=a_1$ and $A_2=a_2$
$\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{Q}$	Variable sets (vector-valued variables)
$ \mathbf{X} = m$	Number of variables in set \mathbf{X} (its cardinality)
$\mathbf{x} = (a_1, \dots, a_m)$	Vector of variable values
$(\mathbf{X} = \mathbf{x}) = ((A_1=a_1), \dots, (A_I=a_m))$	Value assignment of set $\mathbf{X} = \{A_1, \dots, A_m\}$, also interpreted as a conjunction of assignments $A_i=a_i$
$I_{\mathbf{X}=\mathbf{x}}$	An indicator variable for $\mathbf{X} = \mathbf{x}$; $I_{\mathbf{X}=\mathbf{x}} = 1$, if $\mathbf{X} = \mathbf{x}$, and $I_{\mathbf{X}=\mathbf{x}} = 0$ otherwise
$A, \neg A$	Short hand notations for $A=1$ and $A=0$, when A is binary
$\mathbf{X} = ((A=1), \dots, (A_m=1))$	Short hand notation for a conjunction of positive-valued assignments
$\neg \mathbf{X} = ((A_1=0) \vee \dots \vee (A_m=0))$	And its complement when \mathbf{X} is a set of binary variables
$\mathcal{S} = Dom(A_1) \times \dots \times Dom(A_k)$	Sample space spun by variables A_1, \dots, A_k
$\mathbf{r} = ((A_1=a_1), \dots, (A_k=a_k)) \in \mathcal{S}$	A data point; corresponds to an atomic event in \mathcal{S}
$\mathcal{D} = (\mathbf{r}_1, \dots, \mathbf{r}_n), \mathbf{r}_i \in \mathcal{S}$	Data set, a vector of data points, whose elements are also called rows or tuples of data
$ \mathcal{D} = n$	Number of data points in \mathcal{D}
$P(\mathbf{X} = \mathbf{x})$	Probability of event $\mathbf{X} = \mathbf{x}$
$fr(\mathbf{X} = \mathbf{x})$	Absolute frequency of event $\mathbf{X} = \mathbf{x}$; number of data points where $\mathbf{X} = \mathbf{x}$
$\phi(\mathbf{X} = \mathbf{x} \rightarrow A = a) = P(A = a \mathbf{X} = \mathbf{x})$	Precision of rule $\mathbf{X} = \mathbf{x} \rightarrow A = a$
$\gamma(\mathbf{X} = \mathbf{x}, A = a) = \frac{P(\mathbf{X} = \mathbf{x}, A = a)}{P(\mathbf{X} = \mathbf{x})P(A = a)}$	Lift of rule $\mathbf{X} = \mathbf{x} \rightarrow A = a$
$\delta(\mathbf{X} = \mathbf{x}, A = a) = P(\mathbf{X} = \mathbf{x}, A = a) - P(\mathbf{X} = \mathbf{x})P(A = a)$	Leverage of rule $\mathbf{X} = \mathbf{x} \rightarrow A = a$
Δ, Γ	Random variables for the leverage and lift
H_0, H_A	Null hypothesis and an alternative hypothesis
$\{H_1, \dots, H_m\}$	A set of null hypotheses
p	Probability value or a p -value
p_i	Observed p -value of H_i
P_i	Random variable for the p -value of hypothesis H_i
p_F	p -value defined by Fisher's exact test
$p_A, p_X, p_{XA}, p_{A X}$	Parameter values of discrete probability distributions

Table 1 continued

\mathcal{M}	Statistical model, a ‘sampling scheme’
T	Test statistic
t_i	Observed value of the test statistic of H_i
T_i	Random variable for the value of the test statistic of H_i
$L(\cdot)$	Likelihood function
$MI(\cdot)$	Mutual information
z	z -score

$$\delta(A = a, B = b) = P(A = a, B = b) - P(A = a)P(B = b). \quad (2)$$

We note that this is the same as covariance between binary variables A and B .

Lift has also been called ‘interest’ (Brin et al. 1997), ‘dependence’ (Wu et al. 2004), and ‘degree of independence’ (Yao and Zhong 1999). It measures the ratio of the joint probability and its expectation under independence:

$$\gamma(A = a, B = b) = \frac{P(A = a, B = b)}{P(A = a)P(B = b)}. \quad (3)$$

For perfectly independent events, leverage is $\delta = 0$ and lift is $\gamma = 1$, for positive dependencies $\delta > 0$ and $\gamma > 1$, and for negative dependencies, $\delta < 0$ and $\gamma < 1$.

If the real probabilities of events were known, the strength of dependence could be determined accurately. However, in practice, the probabilities are estimated from the data. The most common method is to approximate the real probabilities with relative frequencies (maximum likelihood estimates) but other estimation methods are also possible. The accuracy of these estimates depends on how representative and error-free the data is. The size of the data affects also precision, because continuous probabilities are approximated with discrete frequencies. Therefore, it is quite possible that two independent events express some degree of dependence in the data (i.e., $\hat{P}(A = a, B = b) \neq \hat{P}(A = a)\hat{P}(B = b)$, where \hat{P} is the estimated probability, even if $P(A = a, B = b) = P(A = a)P(B = b)$ in the population). In the worst case, two events always co-occur in the data, indicating maximal dependence, even if they are actually independent. To some extent the probability of such false discoveries can be controlled by statistical significance testing, which is discussed in Sect. 2.3. In the other extreme, two dependent events may appear independent in the data (i.e., $\hat{P}(A = a, B = b) = \hat{P}(A = a)\hat{P}(B = b)$). However, this is not possible if the actual dependence is sufficiently strong (i.e., $P(A = a, B = b) = P(A = a)$ or $P(A = a, B = b) = P(B = b)$), assuming that the data is error-free. Such missed discoveries are harder to detect, but to some extent the problem can be alleviated by using *powerful* methods in significance testing (Sect. 2.3).

2.2.2 Dependence between two variables

For each variable, we can define several events which describe its values. If the variable is categorical, it is natural to consider each variable-value combination as a possible

	B	$\neg B$	Σ
A	$fr(AB) = fr(A)P(B) + n\delta$	$fr(A\neg B) = fr(A)P(\neg B) - n\delta$	$fr(A)$
$\neg A$	$fr(\neg AB) = fr(\neg A)P(B) - n\delta$	$fr(\neg A\neg B) = fr(\neg A)P(\neg B) + n\delta$	$fr(\neg A)$
Σ	$fr(B)$	$fr(\neg B)$	n

Fig. 2 A contingency table for two binary variables A and B expressing absolute frequencies of events AB , $A\neg B$, $\neg AB$ and $\neg A\neg B$ using leverage, $\delta = \delta(A, B)$

event. Then, the independence between two categorical variables can be defined as follows:

Definition 2 (*Statistical independence between two variables*) Let A and B be two categorical variables, whose domains are $Dom(A)$ and $Dom(B)$. A and B are *statistically independent*, if for all $a \in Dom(A)$ and $b \in Dom(B)$ $P(A=a, B=b) = P(A=a)P(B=b)$.

Once again, dependence can be defined either as a Boolean property (lack of independence) or a continuous property. However, there is no standard way to measure the strength of dependence between variables. In practice, the measure is selected according to data and modelling purposes. Two commonly used measures are the χ^2 -measure

$$\chi^2(A, B) = \sum_{a \in Dom(A)} \sum_{b \in Dom(B)} \frac{n(P(A=a, B=b) - P(A=a)P(B=b))^2}{P(A=a)P(B=b)} \quad (4)$$

and *mutual information*

$$MI(A, B) = \sum_{a \in Dom(A)} \sum_{b \in Dom(B)} P(A=a, B=b) \log \frac{P(A=a, B=b)}{P(A=a)P(B=b)}. \quad (5)$$

If the variables are binary, the notions of independence between variables and the corresponding events coincide. Now independence between any of the four value combinations AB , $A\neg B$, $\neg AB$, $\neg A\neg B$ means independence between variables A and B and vice versa. In addition, the absolute value of leverage is the same for all value combinations and can be used to measure the strength of dependence between binary variables. This is shown in the corresponding contingency table (Fig. 2). Unfortunately, this handy property does not hold for multivalued variables. Figure 3 shows an example contingency table for two three-valued variables where some value combinations are independent and others dependent.

2.2.3 Dependence between many events or variables

The notion of statistical independence can be generalized to three or more events or variables in several ways. The most common types of independence are mutual inde-

	$B=b_1$	$B=b_2$	$B=b_3$	Σ
$A=a_1$	$fr(A=a_1B=b_1) = fr(A=a_1)P(B=b_1)$	$fr(A=a_1B=b_2) = fr(A=a_1)P(B=b_2) + n\delta$	$fr(A=a_1B=b_3) = fr(A=a_1)P(B=b_3) - n\delta$	$fr(A=a_1)$
$A=a_2$	$fr(A=a_2B=b_1) = fr(A=a_2)P(B=b_1)$	$fr(A=a_2B=b_2) = fr(A=a_2)P(B=b_2) - n\delta$	$fr(A=a_2B=b_3) = fr(A=a_2)P(B=b_3) + n\delta$	$fr(A=a_2)$
$A=a_3$	$fr(A=a_3B=b_1) = fr(A=a_3)P(B=b_1)$	$fr(A=a_3B=b_2) = fr(A=a_3)P(B=b_2)$	$fr(A=a_3B=b_3) = fr(A=a_3)P(B=b_3)$	$fr(A=a_3)$
Σ	$fr(B=b_1)$	$fr(B=b_2)$	$fr(B=b_3)$	n

Fig. 3 An example contingency table where some value combinations of A and B express independence and others dependence. The frequencies are expressed using leverage, $\delta = \delta(A = a_1, B = b_2)$

pendence, bipartition independence, and conditional independence (see e.g., Agresti 2002, p. 318). In the following, we give general definitions for these three types of independence.

In statistics and probability theory, mutual independence of a set of events is classically defined as follows (see e.g., Feller 1968, p. 128):

Definition 3 (*Mutual independence*) Let $\mathbf{X} = \{A_1, \dots, A_m\}$ be a set of variables, whose domains are $Dom(A_i)$, $i = 1, \dots, m$. Let $a_i \in Dom(A_i)$ notate a value of A_i . A set of events $(A_1 = a_1, \dots, A_m = a_m)$ is called mutually independent if for all $\{i_1, \dots, i_{m'}\} \subseteq \{1, \dots, m\}$ holds

$$P(A_{i_1} = a_{i_1}, \dots, A_{i_{m'}} = a_{i_{m'}}) = \prod_{j=1, \dots, m'} P(A_{i_j} = a_{i_j}). \quad (6)$$

If variables $A_i \in \mathbf{X}$ are binary, the conjunction of true-valued variables $(A_1=1, \dots, A_m=1)$ can be expressed as A_1, \dots, A_m and the condition for mutual independence reduces to $\forall \mathbf{Y} \subseteq \mathbf{X} \ P(\mathbf{Y}) = \prod_{A_i \in \mathbf{Y}} P(A_i)$. An equivalent condition is to require that for all truth value combinations $(a_1, \dots, a_m) \in \{0, 1\}^m$ holds $P(A_1=a_1, \dots, A_m=a_m) = \prod_{i=1}^m P(A_i=a_i)$ (Feller 1968, p. 128). We note that in data mining this property has sometimes been called independence of binary variables and independence of events has referred to a weaker condition (e.g., Silverstein et al. 1998)

$$P(\mathbf{X}) = \prod_{A_i \in \mathbf{X}} P(A_i). \quad (7)$$

The difference is that in the latter it is not required that all $\mathbf{Y} \subseteq \mathbf{X}$ should express independence. Both definitions have been used as a starting point to define interesting set-formed dependency patterns (e.g., Webb 2010; Silverstein et al. 1998). In this paper we will call this type of patterns *dependency sets* (Sect. 5).

In addition to mutual independence, a set of events or variables can express independence between different partitions of the set. The only difference to the basic definition of statistical independence is that now single events or variables have been replaced by sets of events or variables. In this paper we call this type of independence bipartition independence.

Definition 4 (*Bipartition independence*) Let \mathbf{X} be a set of variables. For any partition $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z}$, where $\mathbf{Y} \cap \mathbf{Z} = \emptyset$, possible value combinations are notated by $\mathbf{y} \in \text{Dom}(\mathbf{Y})$ and $\mathbf{z} \in \text{Dom}(\mathbf{Z})$.

- (i) Event $\mathbf{Y} = \mathbf{y}$ is independent of event $\mathbf{Z} = \mathbf{z}$, if $P(\mathbf{Y} = \mathbf{y}, \mathbf{Z} = \mathbf{z}) = P(\mathbf{Y} = \mathbf{y})P(\mathbf{Z} = \mathbf{z})$.
- (ii) Set of variables \mathbf{Y} is independent of \mathbf{Z} , if $P(\mathbf{Y} = \mathbf{y}, \mathbf{Z} = \mathbf{z}) = P(\mathbf{Y} = \mathbf{y})P(\mathbf{Z} = \mathbf{z})$ for all $\mathbf{y} \in \text{Dom}(\mathbf{Y})$ and $\mathbf{z} \in \text{Dom}(\mathbf{Z})$.

Now one can derive a large number of different dependence patterns from a single set \mathbf{X} or event $\mathbf{X} = \mathbf{x}$. There are $2^{m-1} - 1$ ways to partition set \mathbf{X} , $|\mathbf{X}| = m$, into two subsets \mathbf{Y} and $\mathbf{Z} = \mathbf{X} \setminus \mathbf{Y}$ ($|\mathbf{Y}| = 1, \dots, \lceil \frac{m-1}{2} \rceil$). In data mining, patterns expressing bipartition dependence between sets of events are often expressed as *dependency rules* $\mathbf{Y} = \mathbf{y} \rightarrow \mathbf{Z} = \mathbf{z}$. Because both the rule antecedent and consequent are binary conditions, the rule can be interpreted as dependence between two new binary (indicator) variables $I_{\mathbf{Y}=\mathbf{y}}$ and $I_{\mathbf{Z}=\mathbf{z}}$ ($I_{\mathbf{Y}=\mathbf{y}} = 1$ if $\mathbf{Y} = \mathbf{y}$ and $I_{\mathbf{Y}=\mathbf{y}} = 0$ otherwise). In statistical terms, this is the same as collapsing a multidimensional contingency table into a simple 2×2 table. In addition to statistical dependence, dependency rules are often required to fulfil other criteria like sufficient frequency, strength of dependency or statistical significance. Corresponding patterns between sets of variables are less often studied, because the search is computationally much more demanding. In addition, collapsed contingency tables can reveal interesting and statistically significant dependencies between composed events, when no significant dependencies could be found between variables.

The third main type of independence is conditional independence between events or variables:

Definition 5 (*Conditional independence*) Let \mathbf{X} be a set of variables. For any partition $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z} \cup \mathbf{Q}$, where $\mathbf{Y} \cap \mathbf{Z} = \emptyset$, $\mathbf{Y} \cap \mathbf{Q} = \emptyset$, $\mathbf{Z} \cap \mathbf{Q} = \emptyset$, possible value combinations are notated by $\mathbf{y} \in \text{Dom}(\mathbf{Y})$, $\mathbf{z} \in \text{Dom}(\mathbf{Z})$, and $\mathbf{q} \in \text{Dom}(\mathbf{Q})$.

- (i) Events $\mathbf{Y} = \mathbf{y}$ and $\mathbf{Z} = \mathbf{z}$ are conditionally independent given $\mathbf{Q} = \mathbf{q}$, if $P(\mathbf{Y} = \mathbf{y}, \mathbf{Z} = \mathbf{z} \mid \mathbf{Q} = \mathbf{q}) = P(\mathbf{Y} = \mathbf{y} \mid \mathbf{Q} = \mathbf{q})P(\mathbf{Z} = \mathbf{z} \mid \mathbf{Q} = \mathbf{q})$.
- (ii) Sets of variables \mathbf{Y} and \mathbf{Z} are conditionally independent given \mathbf{Q} , if $P(\mathbf{Y} = \mathbf{y}, \mathbf{Z} = \mathbf{z} \mid \mathbf{Q} = \mathbf{q}) = P(\mathbf{Y} = \mathbf{y} \mid \mathbf{Q} = \mathbf{q})P(\mathbf{Z} = \mathbf{z} \mid \mathbf{Q} = \mathbf{q})$ for all $\mathbf{y} \in \text{Dom}(\mathbf{Y})$, $\mathbf{z} \in \text{Dom}(\mathbf{Z})$, and $\mathbf{q} \in \text{Dom}(\mathbf{Q})$.

Conditional independence can be defined also for more than two sets of events or variables, given a third one. For example, in set $\{A, B, C, D\}$ we can find four conditional independencies given D : $A \perp BC$, $B \perp AC$, $C \perp AB$, and $A \perp B \perp C$. However, these types of independence are seldom needed in practice. In pattern discovery notions of conditional independence and dependence between events are used for inspecting improvement of a dependency rule $\mathbf{YQ} \rightarrow C$ over its generalization $\mathbf{Y} \rightarrow C$ (Sect. 4). In machine learning conditional independence between variables or sets of variables is an important property for constructing full probability models, like Bayesian networks or log-linear models.




Fig. 4 Different approaches to statistical significance testing

2.3 Statistical significance testing

Often when searching dependency rules and sets the aim is to find dependencies that hold in the population from which the sample is drawn (cf. Fig. 1). Statistical significance tests are the tools that have been created to control the risk that such inferences drawn from sample data do not hold in the population. This section introduces the key concepts that underlie significance testing and gives an overview of the main approaches that can be applied in testing dependency rules and sets. The same principles can be applied in testing other types of patterns, but a reader would be well advised to consult a statistician with regard to which tests to apply and how to apply them.

The main idea of statistical significance testing is to estimate the probability that the observed discovery would have occurred by chance. If the probability is very small, we can assume that the discovery is genuine. Otherwise, it is considered spurious and discarded. The probability can be estimated either analytically or empirically. The analytical approach is used in the *traditional significance testing*, while *randomization tests* estimate the probability empirically. Traditional significance testing can be further divided into two main classes: the *frequentist* and *Bayesian approaches*. These main approaches to statistical significance testing are shown in Fig. 4.

2.3.1 Frequentist approach

The frequentist approach of significance testing is the most commonly used and best studied (see e.g. Freedman et al. 2007, Ch. 26, or Lindgren 1993, Ch. 10.1). The approach is actually divided into two opposing schools, Fisherian and Neyman–Pearsonian, but most textbooks present a kind of synthesis (see e.g., Hubbard and Bayarri 2003). The main idea is to estimate the probability of the observed or a more




Fig. 5 An example distribution of test statistic T under the null hypothesis. If the observed value of T is t , the p -value is probability $P(T \geq t)$ (directional hypothesis) or $P(T \leq -t \vee T \geq t)$ (non-directional hypothesis)

extreme phenomenon O under some *null hypothesis*, H_0 . In general, the null hypothesis is a statement on the value of some statistic or statistics S in the population. For example, when the objective is to test the significance of dependency rule $\mathbf{X} \rightarrow A$, the null hypothesis H_0 is the independence assumption: $N_{XA} = nP(\mathbf{X})P(A)$, where N_{XA} is a random variable for the absolute frequency of $\mathbf{X}A$. (Equivalently, H_0 could be $\Delta = 0$ or $\Gamma = 1$, where Δ and Γ are random variables for the leverage and lift.) In independence testing the null hypothesis is usually an equivalence statement, $S = s_0$ (nondirectional hypothesis), but in other contexts it can also be of the form $S \leq s_0$ or $S \geq s_0$ (directional hypothesis). Often, one also defines an explicit alternative hypothesis, H_A , which can be either directional or nondirectional. For example, in pattern discovery dependency rules $\mathbf{X} \rightarrow A$ are assumed to express positive dependence, and therefore it is natural to form a directional hypothesis $H_A: N_{XA} > nP(\mathbf{X})P(A)$ (or $\Delta > 0$ or $\Gamma > 1$).

When the null hypothesis has been defined, one should select a test statistic T (possibly S itself) and define its distribution (*null distribution*) under H_0 . The p -value is defined from this distribution as the probability of the observed or a more extreme T -value, $P(T \geq t \mid H_0)$, $P(T \leq t \mid H_0)$, or $P(T \leq -t \text{ or } T \geq t \mid H_0)$ (Fig. 5). In the case of independence testing, possible test statistics are, for example, leverage, lift, and the χ^2 -measure (Eq. 4). The distribution under independence is defined according to the selected *sampling model*, which we will introduce in Sect. 3. The probability of observing positive dependence whose strength is at least $\delta(\mathbf{X}, A)$ is $P_{\mathcal{M}}(\Delta \geq \delta(\mathbf{X}, A) \mid H_0)$, where $P_{\mathcal{M}}$ is the complementary cumulative distribution function for the assumed sampling model \mathcal{M} .

Up to this point, all frequentist approaches are more or less in agreement. The differences appear only when the p -values are interpreted. In the classical (Neyman–Pearsonian) hypothesis testing, the p -value is compared to some predefined threshold α . If $p \leq \alpha$, the null hypothesis is rejected and the discovery is called significant at level α . Parameter α (also known as the *test size*) defines the probability of committing a *type I error*, i.e., accepting a spurious pattern (and rejecting a correct null hypothesis). Another parameter, β , is used to define the probability of committing a *type II error*, i.e., rejecting a genuine pattern as non-significant (and keeping a false null hypothesis). The complement $1 - \beta$ defines the *power* of the test, i.e., the probability that a genuine pattern passes the test. Ideally, one would like to minimize the test size and maximize

its power. Unfortunately, this is not possible, because β increases when α decreases and vice versa. As a solution it has been recommended (e.g., Lehmann and Romano 2005, p. 57) to select appropriate α and then to check that the power is acceptable given the sample size. However, the power analysis can be difficult and all too often it is skipped altogether.

The most controversial problem in hypothesis testing is how to select an appropriate significance level. A convention is to use always the same standard levels, like $\alpha = 0.05$ or $\alpha = 0.01$. However, these values are quite arbitrary and widely criticized (see e.g., Lehmann and Romano 2005, p. 57; Lecoutre et al. 2001; Johnson 1999). Especially in large data sets, the p -values tend to be very small and hypotheses get too easily rejected with conventional thresholds. A simple alternative is to report only p -values, as advocated by Fisher and also many recent statisticians (e.g., Lehmann and Romano 2005, pp. 63–65; Hubbard and Bayarri 2003). Sometimes, this is called ‘significance testing’ in distinction from ‘hypothesis testing’ (with fixed α s), but the terms are not used systematically. Reporting only p -values may often be sufficient, but there are still situations where one should make concrete decisions and a binary judgement is needed.

Deciding threshold α is even harder in data mining where numerous patterns are tested. For example, if we use threshold $\alpha = 0.05$, then there is up to 5% chance that a spurious pattern passes the significance test. If we test 10 000 spurious patterns, we can expect up to 500 of them to pass the test erroneously. This so called *multiple testing problem* is inherent in knowledge discovery, where one often performs an exhaustive search over all possible patterns. We will return to this problem in Sect. 6.

2.3.2 Bayesian approach

Bayesian approaches are becoming increasingly popular in both statistics and data mining (see e.g., Corani et al. 2016). However, to date there has been little uptake of them in statistically sound pattern discovery. We include here a brief summary of the Bayesian approach for completeness and in the hope that it will stimulate further investigation of this promising approach.

The idea of Bayesian significance testing (see e.g., Lee 2012, Ch. 4; Albert 1997; Jamil et al. 2017) is quite similar to the frequentist approach, but now we assign some prior probabilities $P(H_0)$ and $P(H_A)$ to the null hypothesis H_0 and the alternative research hypothesis H_A . Next, the conditional probabilities, $P(O|H_0)$ and $P(O|H_A)$, of the observed or a more extreme phenomenon O under H_0 and H_A are estimated from the data. Finally, the probabilities of both hypotheses are updated by the Bayes’ rule and the acceptance or rejection of H_0 is decided by comparing posterior probabilities $P(H_0|O)$ and $P(H_A|O)$. The resulting conditional probabilities $P(H_0|O)$ are asymptotically similar (under some assumptions even identical) to the traditional p -values, but Bayesian testing is sensitive to the selected prior probabilities (Agresti and Min 2005). One attractive feature of the Bayesian approach is that it allows to quantify the evidence for and against the null hypothesis. However, the procedure tends to be more complicated than the frequentist one; specifying prior distributions may require a plethora of parameters and the posterior probabilities cannot always be evaluated analytically (Agresti and Hitchcock 2005; Jamil et al. 2017).

2.3.3 Randomization testing

Randomization testing (see e.g., Edgington 1995) offers a relatively assumption-free approach for testing statistical dependencies. Unlike traditional significance testing, there is no need to assume that the data would be a random sample from the population or to define what type of distribution the test statistic has under the null hypothesis. Instead, the significance is estimated empirically, by generating random data sets under the null hypothesis and checking how often the observed or a more extreme phenomenon occurs in them.

When independence between A and B is tested, the null hypothesis is *exchangeability* of the A -values on rows when B -values are kept fixed, or vice versa. This is the same as stating that all permutations of A -values in the data are equally likely. A similar null hypothesis can be formed for mutual independence in a set of variables. If only a single dependency set \mathbf{X} is tested, it is enough to generate random data sets $\mathcal{D}_1, \dots, \mathcal{D}_b$ by permuting values of each A_i , $A_i \in \mathbf{X}$. Usually, it is required that all marginal probabilities $P(A_i)$ remain the same as in the original data, but there may be additional constraints, defined by the *permutation scheme*. Test statistic T that evaluates goodness of the pattern is calculated in each random data set. For simplicity, we assume that the test statistic T is increasing by goodness (a higher value indicates a better pattern). If the original data set produced T -value t_0 and b random data sets produced T -values t_1, \dots, t_b , the empirical p -value of the observed pattern is

$$p_{em} = \frac{|\{\mathcal{D}_j \mid t_j \geq t_0, j = 1, \dots, b\}| + 1}{b + 1}. \quad (8)$$

If the data set is relatively small and \mathbf{X} is simple, it is possible to enumerate all possible permutations where the marginal probabilities hold. This leads to an *exact permutation test*, which gives an exact p -value. On the other hand, if the data set is large and/or \mathbf{X} is more complex, all possibilities cannot be checked, and the empirical p -value is less accurate. In this case, the test is called a *random permutation test* or an *approximate permutation test*. There are also some special cases, like testing a single dependency rule, where it is possible to express the permutation test in a closed form that is easy to evaluate exactly (see Fisher's exact test in Sect. 3.2).

An advantage of randomization testing is that the test statistic can have any kind of distribution, which is especially handy when the statistic is new or poorly known. With randomization one can test also such null hypotheses for which no closed form test exists. Randomization tests are technically valid even if the data are not a random sample because strictly speaking the population to which the null hypotheses relate is the set of all permutations of the sample defined by the permutation scheme. However, the results can be generalized to the reference population only to the extent of how representative the sample was for that population (Legendre and Legendre 1998, p. 24). One critical problem with randomization testing is that it is not always clear how the data should be permuted, and different permutation schemes can produce quite different results in their assessment of patterns (see e.g., Hanhijärvi 2011). The number of random permutations plays also an important role in testing. The more random permutations are performed, the more accurate the empirical p -values are,

but in practice, extensive permuting can be too time consuming. Computational costs restrict also the use of randomization testing in search algorithms especially in large data sets.

The idea of randomization tests can be extended for estimating the overall significance of all mining results or even for tackling the multiple testing problem. For example, one may test the significance of the number of all frequent sets (given a minimum frequency threshold) or the number of all sufficiently strong pair-wise correlations (given a minimum correlation threshold) using randomization tests (Gionis et al. 2007). In this case, it is necessary to generate complete data sets randomly for testing. The difficulty is to decide what properties of the original data set should be maintained. One common solution in pattern mining is to keep both the column margins ($fr(A_i)$ s) and the row margins (numbers of 1s on each row) fixed and generate new data sets by *swap randomization* (Cobb and Chen 2003). A prerequisite for this method is that the attributes are semantically similar (e.g. occurrence or absence of species) and it is sensible to swap their values. In addition, there are some pathological cases, where no or only a few permutations exist with the given row and column margins, resulting in a large p -value, even if the original data set contains a significant pattern (Gionis et al. 2007).

3 Statistical significance of dependency rules

Dependency rules are a famous pattern type that expresses bipartition dependence between the rule antecedent and the consequent. In this section, we discuss how statistical significance of dependency rules is evaluated under different assumptions. Especially, we contrast two alternative interpretations of dependency rules that are called variable-based and value-based interpretations and introduce appropriate tests in different sampling models.

3.1 Dependency rules

Dependency rules are maybe the simplest type of statistical dependency patterns. As a result, it has been possible to develop efficient exhaustive search algorithms. With these, dependency rules can reveal arbitrarily complex bipartite dependencies from categorical or discretized numerical data without any additional assumptions. This makes dependency rule analysis an attractive starting point for any data mining task. In medical science, for example, an important task is to search for statistical dependencies between gene alleles, environmental factors, and diseases. We recall that statistical dependencies are not necessarily causal relationships, but still they can help to form causal hypotheses and reveal which factors predispose or prevent diseases (see e.g., Jin et al. 2012; Li et al. 2016). Interesting dependencies do not necessarily have to be strong or frequent, but instead, they should be statistically valid, i.e., genuine dependencies that are likely to hold also in future data. In addition, it is often required that the patterns should not contain any superfluous variables which would only obscure the real dependencies. Based on these considerations, we will first

give a general definition of dependency rules and then discuss important aspects of genuine dependencies.

Definition 6 (*Dependency rule*) Let \mathbf{R} be a set of categorical variables, $\mathbf{X} \subseteq \mathbf{R}$, and $\mathbf{Y} \subseteq \mathbf{R} \setminus \mathbf{X}$. Let us denote value vectors of \mathbf{X} and \mathbf{Y} by $\mathbf{x} \in \text{Dom}(\mathbf{X})$ and $\mathbf{y} \in \text{Dom}(\mathbf{Y})$. Rule $\mathbf{X} = \mathbf{x} \rightarrow \mathbf{Y} = \mathbf{y}$ is a dependency rule, if $P(\mathbf{X} = \mathbf{x}, \mathbf{Y} = \mathbf{y}) \neq P(\mathbf{X} = \mathbf{x})P(\mathbf{Y} = \mathbf{y})$.

The dependency is (i) positive, if $P(\mathbf{X} = \mathbf{x}, \mathbf{Y} = \mathbf{y}) > P(\mathbf{X} = \mathbf{x})P(\mathbf{Y} = \mathbf{y})$, and (ii) negative, if $P(\mathbf{X} = \mathbf{x}, \mathbf{Y} = \mathbf{y}) < P(\mathbf{X} = \mathbf{x})P(\mathbf{Y} = \mathbf{y})$. Otherwise, the rule expresses independence.

It is important to recognize that while the convention is to specify the antecedent and consequent and use a directed arrow to distinguish them, statistical dependence is a symmetric relation and strictly speaking the direction is arbitrary. Often, the rule is expressed with the antecedent and consequent selected so that the precision ('confidence') of the rule ($\phi(\mathbf{X} = \mathbf{x} \rightarrow \mathbf{Y} = \mathbf{y}) = P(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x})$ or $\phi(\mathbf{Y} = \mathbf{y} \rightarrow \mathbf{X} = \mathbf{x}) = P(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y})$) is maximal. An exception is supervised descriptive rule discovery (including class association rules (Li et al. 2001), subgroup discovery (Herrera et al. 2011), emerging pattern mining (Dong and Li 1999) and contrast set mining (Bay and Pazzani 2001)), where the consequent is fixed (Novak et al. 2009).

For simplicity, we will concentrate on a common special case of dependency rules where 1) all variables are binary, 2) the consequent $\mathbf{Y} = \mathbf{y}$ consists of a single variable-value combination, $A = i$, $i \in \{0, 1\}$, and 3) the antecedent $\mathbf{X} = \mathbf{x}$ is a conjunction of true-valued attributes, i.e., $\mathbf{X} = \mathbf{x} \equiv (A_1 = 1, \dots, A_l = 1)$, where $\mathbf{X} = \{A_1, \dots, A_l\}$. With these restrictions the resulting rules can be expressed in a simpler form $\mathbf{X} \rightarrow A = i$, where $i \in \{0, 1\}$, or $\mathbf{X} \rightarrow A$ and $\mathbf{X} \rightarrow \neg A$. Allowing negated consequents means that it is sufficient to represent only positive dependencies (a positive dependency between \mathbf{X} and $\neg A$ is the same as a negative dependency between \mathbf{X} and A). We note that this restriction is purely representational and the following theory is easily extended to general dependency rules as well. Furthermore, we recall that this simpler form of rules can still represent all dependency rules after suitable data transformations (i.e., creating new binary variables for all values of the original variables).

Finally, we note that dependency rules deviate from traditional *association rules* (Agrawal et al. 1993) in their requirement of statistical dependence. Traditional association rules do not necessarily express any statistical dependence but relations between frequently occurring attribute sets. However, there has been research on association rules where the requirement of minimum frequency ('minimum support') has been replaced by requirements of statistical dependence (see e.g., Webb and Zhang 2005; Webb 2008, 2007; Hämmäläinen 2012, 2010b; Li 2006; Morishita and Sese 2000; Nijssen and Kok 2006; Nijssen et al. 2009). For clarity, we will use here the term 'dependency rule' for all rule type patterns expressing statistical dependencies, even if they had been called association rules, classification rules, or other similar patterns in the original publications.

Statistical dependence is a necessary requirement of a dependency rule, but in addition, it is frequently useful to impose further constraints like that of statistical significance and absence of superfluous variables. The following example illustrates some of these properties of dependency rules.

Table 2 An imaginary example of dependency rules related to heart disease

	Rule	fr	ϕ	δ	γ
1	smoking \rightarrow heart disease	120	0.400	0.030	1.333
2	sports $\rightarrow \neg$ heart disease	400	0.800	0.050	1.143
3	coffee $\rightarrow \neg$ heart disease	240	0.700	0.000	1.000
4	stress \rightarrow heart disease	150	0.300	0.000	1.000
5	pine bark extract $\rightarrow \neg$ heart disease	1	1.000	< 0.001	1.429
6	female $\rightarrow \neg$ heart disease	352	0.704	0.002	1.006
7	female, stress \rightarrow heart disease	100	0.385	0.022	1.282
8	stress, smoking \rightarrow heart disease	100	0.500	0.040	1.667
9	smoking, coffee \rightarrow heart disease	96	0.400	0.024	1.333
10	smoking, sports \rightarrow heart disease	20	0.333	0.020	1.111
11	female, sports $\rightarrow \neg$ heart disease	203	0.808	0.027	1.154

fr = frequency, ϕ = precision, δ = leverage, γ = lift

Example 1 Let us consider an imaginary database consisting of 1000 patients (50% female, 50% male), 30% of them with heart disease. The database contains information on patients and their life style like smoking status, drinking coffee, having stress, going for sports, and using natural products. Table 2 lists some candidate dependency rules related to heart disease together with their frequency, precision, leverage, and lift.

The first two rules are examples of simple positive and negative dependencies (predisposing and protecting factors for heart disease). Rules 3 and 4 are included as examples of so called independence rules that express statistical independence between the antecedent and consequent. Normally, such rules would be pruned out by dependency rule mining algorithms.

Rule 5 is an example of a spurious rule, which is statistically insignificant and likely due to chance. The database contains only one person who uses pine bark extract regularly and who does not have heart disease. Note that the lift is still quite large, the maximal possible for that consequent. Rule 6 is also statistically insignificant, but for a different reason. The rule is very common, but the difference in the prevalence of heart disease among female and male patients is so small (148 vs. 152) that it can be explained by chance.

Rule 7 demonstrates non-monotonicity of statistical dependence. The combination of stress and female gender correlates positively with heart disease, even though stress alone was independent of heart disease and the female gender was negatively correlated with it.

The last four rules illustrate the problem of superfluous variables. In rule 8, neither of the condition attributes is superfluous, because the dependency is stronger and more significant than simpler dependencies involving only stress or only smoking. However, rules 9–11 demonstrate three types of superfluous rules where extra factors (i) have no effect on the dependency, (ii) weaken it, or (iii) apparently improve it but not significantly. Rule 9 is superfluous, because coffee has no effect on the dependency between smoking and heart disease (coffee consumption and heart disease are condi-

tionally independent given smoking). Rule 10 is superfluous, because going for sports weakens the dependency between smoking and heart disease. This kind of modifying effect might be interesting in some contexts, if it were statistically significant. However, dependency rule mining algorithms do not usually perform such analysis. Rule 11 is the most difficult to judge, because the dependence is itself significant and the rule has larger precision and lift than either of simpler dependencies involving only the female gender or only sports. However, the improvement with respect to rule 2 is so small ($\phi = 0.808$ vs. $\phi = 0.800$) that it is likely due to chance.

In the previous example we did not state which measure should be preferred for measuring the strength of dependence or how the statistical significance should be evaluated. The reason is that *the selection of these measures as well as evaluation of statistical significance and superfluousness depend on the interpretation of dependency rules*. In principle there are two alternative interpretations for rule $\mathbf{X} \rightarrow A = i$, $i \in \{0, 1\}$: either it can represent a dependency between events \mathbf{X} (or $I_{\mathbf{X}}=1$) and $A = i$ or between variables $I_{\mathbf{X}}$ and A , where $I_{\mathbf{X}}$ is an indicator variable for event \mathbf{X} . These two interpretations have sometimes been called *value-based* and *variable-based* semantics (Blanchard et al. 2005) of the rule. Unfortunately, researchers have often forgotten to mention explicitly which interpretation they follow. This has caused much confusion and, in the worst case, led to missed or inappropriate discoveries. The following example demonstrates how variable- and value-based interpretations can lead to different results.

Example 2 Let us consider a database of 100 apples describing their colour (green or red), size (big or small), and taste (sweet or bitter). Let us notate $A=\text{sweet}$, $\neg A=\text{bitter}$ (not sweet), $\mathbf{Y}=\{\text{red}\}$, $\neg \mathbf{Y}=\{\text{green}\}$ (not red), $\mathbf{X}=\{\text{red, big}\}$ and $\neg \mathbf{Y}=\neg\{\text{red, big}\}$ (i.e., green or small).

We would like to find strong dependencies related to either variable ‘taste’ (variable-based interpretation) or value ‘sweet’ (value-based interpretation). Figure 6 represents two such rules: $\mathbf{Y} \rightarrow A$ ($\text{red} \rightarrow \text{sweet}$) and $\mathbf{X} \rightarrow A$ ($\text{red and big} \rightarrow \text{sweet}$).

The first rule expresses a strong dependency between binary variables $I_{\mathbf{Y}}$ and A (i.e., colour and taste) with $P(A|\mathbf{Y}) = 0.92$, $P(\neg A|\neg \mathbf{Y}) = 1.0$, $\delta(\mathbf{Y}, A) = 0.22$, and $\gamma(\mathbf{Y}, A) = 1.67$. So, with this rule we can divide the apples into two baskets according to colour. The first basket contains 60 red apples, 55 of which are sweet, and the second basket contains 40 green apples, which are all bitter. This is quite a good rule if the goal is to classify well both sweet apples (for eating) and bitter apples (for juice and cider).

The second rule expresses a strong dependency between the value combination \mathbf{X} (red and big) and value $A = 1$ (sweet) with $P(A|\mathbf{X}) = 1.0$, $P(\neg A|\neg \mathbf{X}) = 0.75$, $\delta(\mathbf{X}, A) = 0.18$, $\gamma(\mathbf{X}, A) = 1.82$. This rule produces a basket of 40 big, red apples, all of them sweet, and another basket of 60 green or small apples, 45 of them bitter. This is an excellent rule if we would like to predict sweetness better (e.g., get a basket of sweet apples for our guests) without caring how well bitterness is predicted.

So, the choice between variable-based and value-based interpretation results in a preference for a different rule. Either one can be desirable for different modelling purposes. This decision affects also which goodness measure should be used. Leverage suits the variable-based interpretation, because its absolute value is the same for all




Fig. 6 Apple baskets corresponding to rules $red \rightarrow sweet$ (top) and $red \text{ and } big \rightarrow sweet$ (bottom) (Color figure online)

truth value combinations ($\mathbf{X}A$, $\mathbf{X}\neg A$, $\neg\mathbf{X}A$, $\neg\mathbf{X}\neg A$), but it may miss interesting dependencies related to particular values. Lift, on the other hand, suits the value-based interpretation, because it favours rules where the given values are strongly dependent. However, it is not a reliable measure alone, because it ranks well also coincidental ‘noise rules’ (e.g., *apple maggot* \rightarrow *bitter*). Therefore, it has to be accompanied with statistical significance tests.

In general, the variable-based interpretation tends to produce more reliable patterns, in the sense that the discovered dependencies hold well in future data (see e.g., Hämäläinen 2010a, Ch.5). However, there are applications where the value-based interpretation may better identify interesting dependency rules. One example could be analysis of predisposing factors (like gene alleles) for a serious disease. Some factors \mathbf{X} may be rare, but still their occurrence could strongly predict the onset of some disease D . Medical scientists would certainly want to find such dependencies $\mathbf{X} \rightarrow D$, even if the overall dependency between variables $I_{\mathbf{X}}$ and D would be weak or insignificant.

In the following sections, we will examine how statistical significance is tested in the variable-based and value-based interpretations.

3.2 Sampling models for the variable-based interpretation

In the variable-based interpretation, the significance of dependency rule $\mathbf{X} \rightarrow A$ is determined by classical independence tests. The task is to estimate the probability of the observed or a more ‘extreme’ contingency table, assuming that variables $I_{\mathbf{X}}$ and A were actually independent. There is no consensus how the extremeness relation should be defined, but intuitively, contingency table τ_i is more extreme than table τ_j , if the dependence between \mathbf{X} and A is stronger in τ_i than in τ_j . So, any measure for the strength of dependence between variables can be used as a *discrepancy measure*, to order contingency tables. The simplest such measure is leverage, but also *odds ratio*

$$\text{odds}(N_{XA}, N_{X\rightarrow A}, N_{\neg XA}, N_{\neg X\rightarrow A}) = \frac{N_{XA}N_{\neg X\rightarrow A}}{N_{X\rightarrow A}N_{\neg XA}} \quad (9)$$

is commonly used. We note that odds ratio is not defined when $N_{X\rightarrow A}N_{\neg XA} = 0$ and some special policy is needed for these cases. In the following, we will notate the relation “table τ_i is equally or more extreme to table τ_j ” by $\tau_i \succeq \tau_j$.

The probability of each contingency table τ_i depends on the assumed statistical model \mathcal{M} . Model \mathcal{M} defines the space of all possible contingency tables $\mathcal{T}_{\mathcal{M}}$ (under the model assumptions) and the probability $P(\tau_i \mid \mathcal{M})$ of each table $\tau_i \in \mathcal{T}_{\mathcal{M}}$. Because the task is to test independence, the assumed model should satisfy the independence assumption $P(\mathbf{X}A) = P(\mathbf{X})P(A)$ in some form. For the probabilities $P(\tau_i \mid \mathcal{M})$ holds

$$\sum_{\tau_i \in \mathcal{T}_{\mathcal{M}}} P(\tau_i \mid \mathcal{M}) = 1.$$

Now the probability of the observed contingency table τ_o or any τ_i , $\tau_i \succeq \tau_o$, is the desired p -value

$$p = \sum_{\tau_i \succeq \tau_o} P(\tau_i \mid \mathcal{M}). \quad (10)$$

Classically, statistical models for independence testing have been divided into three main categories (sampling schemes) (Barnard 1947; Pearson 1947), which we call *multinomial*, *double binomial*, and *hypergeometric* models. In the statistics literature (e.g. Barnard 1947; Upton 1982), the corresponding sampling schemes are called *double dichotomy*, *2 × 2 comparative trial*, and *2 × 2 independence trial*.

In the following we describe the three models using the classical urn metaphor. However, because there are two binary variables of interest, $I_{\mathbf{X}}$ and A , we cannot use the basic urn model with white and black balls. Instead, we will use an apple basket model, with red and green, sweet and bitter apples, like in Example 2.

3.2.1 Multinomial model

In the multinomial model, it is assumed that the real probabilities of sweet red apples, bitter red apples, sweet green apples, and bitter green apples are defined by parameters

p_{XA} , $p_{X\bar{A}}$, $p_{\bar{X}A}$, and $p_{\bar{X}\bar{A}}$. The probability of red apples is p_X and of green apples $1 - p_X$. Similarly, the probability of sweet apples is p_A and of bitter apples $1 - p_A$. According to the independence assumption, $p_{XA} = p_X p_A$, $p_{X\bar{A}} = p_X(1 - p_A)$, $p_{\bar{X}A} = (1 - p_X)p_A$, and $p_{\bar{X}\bar{A}} = (1 - p_X)(1 - p_A)$. A sample of n apples is taken *randomly* from an infinite basket (or from a finite basket with replacement). Now the probability of obtaining N_{XA} sweet red apples, $N_{X\bar{A}}$ bitter red apples, $N_{\bar{X}A}$ sweet green apples, and $N_{\bar{X}\bar{A}}$ bitter green apples is defined by multinomial probability

$$\begin{aligned} P(N_{XA}, N_{X\bar{A}}, N_{\bar{X}A}, N_{\bar{X}\bar{A}} \mid n, p_X, p_A) \\ = \binom{n}{N_{XA}, N_{X\bar{A}}, N_{\bar{X}A}, N_{\bar{X}\bar{A}}} \cdot p_X^{N_X} (1 - p_X)^{n - N_X} p_A^{N_A} (1 - p_A)^{n - N_A}. \end{aligned} \quad (11)$$

Since data size n is given, the contingency tables can be defined by triplets $\langle N_{XA}, N_{X\bar{A}}, N_{\bar{X}A} \rangle$ or, equivalently, triplets $\langle N_X, N_A, N_{XA} \rangle$. Therefore, the space of all possible contingency tables is

$$\begin{aligned} \mathcal{T}_M = \{ \langle N_X, N_A, N_{XA} \rangle \mid N_X = 0, \dots, n; N_A = 0, \dots, n; \\ N_{XA} = 0, \dots, \min\{N_X, N_A\} \}. \end{aligned}$$

For estimating the p -value with Eq. (10), we should still solve two problems. First, the parameters p_X and p_A are unknown. The most common solution is to estimate them by the observed relative frequencies (maximum likelihood estimates). Second, we should decide when a contingency table τ_i is equally or more extreme than the observed contingency table τ_o . For this purpose, we have to select the discrepancy measure, which evaluates the overall dependence in a contingency table, when only the data size n is fixed. Examples of such measures are leverage and the odds ratio.

In practice, the multinomial test is seldom used, but the multinomial model is an important theoretical model, from which other models can be derived as special cases.

3.2.2 Double binomial model

In the double binomial model, it is assumed that we have two infinite baskets, one for red and one for green apples. Let us call these the red and the green basket. In the red basket the probability of sweet apples is $p_{A|X}$ and of bitter apples $1 - p_{A|X}$, and in the green basket the probabilities are $p_{A|\bar{X}}$ and $1 - p_{A|\bar{X}}$. According to the independence assumption, the probability of sweet apples is the same in both baskets: $p_A = p_{A|X} = p_{A|\bar{X}}$. A sample of $fr(\mathbf{X})$ apples is taken randomly from the red basket and another random sample of $fr(\bar{\mathbf{X}})$ apples is taken from the green basket. The probability of obtaining N_{XA} sweet apples among the selected $fr(\mathbf{X})$ red apples is defined by the binomial probability

$$P(N_{XA} \mid fr(\mathbf{X}), p_A) = \binom{fr(\mathbf{X})}{N_{XA}} p_A^{N_{XA}} (1 - p_A)^{fr(\mathbf{X}) - N_{XA}}.$$

Similarly, the probability of obtaining $N_{\neg XA}$ sweet apples among the selected green apples is

$$P(N_{\neg XA} \mid fr(\neg \mathbf{X}), p_A) = \binom{fr(\neg \mathbf{X})}{N_{\neg XA}} p_A^{N_{\neg XA}} (1 - p_A)^{fr(\neg \mathbf{X}) - N_{\neg XA}}.$$

Because the two samples are independent from each other, the probability of obtaining N_{XA} sweet apples from $fr(\mathbf{X})$ red apples and $N_{\neg XA}$ sweet apples from $fr(\neg \mathbf{X})$ green apples is the product of the two binomials

$$P(N_{XA}, N_{\neg XA} \mid n, fr(\mathbf{X}), p_A) = \binom{fr(\mathbf{X})}{N_{XA}} \binom{fr(\neg \mathbf{X})}{N_{\neg XA}} p_A^{N_A} (1 - p_A)^{n - N_A}, \quad (12)$$

where $N_A = N_{XA} + N_{\neg XA}$ is the total number of the obtained sweet apples. (Here $fr(\neg \mathbf{X})$ was dropped from the condition, because n is given.) We note that the double binomial probability is not *exchangeable* with respect to the roles of \mathbf{X} and A , i.e., generally $P(N_{XA}, N_{\neg XA} \mid n, fr(\mathbf{X}), p_A) \neq P(N_{XA}, N_{X \neg A} \mid n, fr(A), p_X)$. In practice, this means that the probability of obtaining $fr(\mathbf{X}A)$ sweet red apples, $fr(\mathbf{X} \neg A)$ bitter red apples, $fr(\neg \mathbf{X}A)$ sweet green apples, and $fr(\neg \mathbf{X} \neg A)$ bitter green apples is (nearly always) different in the model of the red and green baskets from the model of the sweet and bitter baskets.

Since $fr(\mathbf{X})$ and $fr(\neg \mathbf{X})$ are given, each contingency table is defined as a pair $\langle N_{XA}, N_{\neg XA} \rangle$ or, equivalently, $\langle N_A, N_{XA} \rangle$. The space of all possible contingency tables is

$$\mathcal{T}_{\mathcal{M}} = \{ \langle N_{XA}, N_{\neg XA} \rangle \mid N_{XA} = 0, \dots, fr(\mathbf{X}); N_{\neg XA} = 0, \dots, fr(\neg \mathbf{X}) \}.$$

We note that N_A is not fixed, and therefore N_A is generally not equal to the observed $fr(A)$.

For estimating the significance with Eq. (10), we should estimate the unknown parameter p_A and select a discrepancy measure, like leverage or odds ratio. Then the exact p -value is obtained by summing over all possible values of N_{XA} and $N_{\neg XA}$ where the dependence is sufficiently strong. However, often this is considered impractical and the p -value is approximated with asymptotic tests, which are discussed later.

3.2.3 Hypergeometric model

In the hypergeometric model, there is no sampling from an infinite basket. Instead, we can assume that we are given a finite basket of n apples, containing exactly $fr(\mathbf{X})$ red apples and $fr(\neg \mathbf{X})$ green apples. We test all n apples and find that $fr(\mathbf{X}A)$ of red apples and $fr(\neg \mathbf{X}A)$ of green apples are sweet. The question is how probable is our basket, or the set of all at least equally extreme baskets, among all possible apple baskets with $fr(\mathbf{X})$ red apples, $fr(\neg \mathbf{X})$ green apples, $fr(A)$ sweet apples, and $fr(\neg A)$ bitter apples.

Now the baskets correspond to contingency tables. The number of all possible baskets with the fixed totals $fr(\mathbf{X})$, $fr(\neg\mathbf{X})$, $fr(A)$, and $fr(\neg A)$ is

$$\sum_{i=0}^{fr(A)} \binom{fr(\mathbf{X})}{i} \binom{fr(\neg\mathbf{X})}{fr(A)-i} = \binom{n}{fr(A)}.$$

(We recall that customarily $\binom{m}{l} = 0$, when $l > m$.) Assuming that all baskets with these fixed totals are equally likely, the probability of a basket with N_{XA} sweet red apples is

$$P(N_{XA} \mid n, fr(\mathbf{X}), fr(A)) = \binom{n}{fr(A)}^{-1}.$$

Because all totals are fixed, the extremeness relation is also easy to define. Positive dependence is stronger than observed, when $N_{XA} > fr(\mathbf{X}A)$. For the p -value it is enough to sum the probabilities of baskets containing at least $fr(\mathbf{X}A)$ sweet red apples. The resulting p -value is

$$p_F = \sum_{i=0}^{J_1} \frac{\binom{fr(\mathbf{X})}{fr(\mathbf{X}A)+i} \binom{fr(\neg\mathbf{X})}{fr(\neg\mathbf{X}\neg A)+i}}{\binom{n}{fr(A)}}, \quad (13)$$

where $J_1 = \min\{fr(\mathbf{X}\neg A), fr(\neg\mathbf{X}A)\}$. (Instead of J_1 we could give an upper range $fr(A)$, because the zero terms disappear.) This p -value is known as *Fisher's p* , because it is used in *Fisher's exact test*, an exact permutation test. We give it a special symbol p_F , because it will be used later. For negative dependence between red and sweet apples (or positive dependence between green and sweet apples) the p -value is

$$p_F = \sum_{i=0}^{J_2} \frac{\binom{fr(\mathbf{X})}{fr(\mathbf{X}A)-i} \binom{fr(\neg\mathbf{X})}{fr(\neg\mathbf{X}\neg A)-i}}{\binom{n}{fr(A)}}, \quad (14)$$

where $J_2 = \min\{fr(\mathbf{X}A), fr(\neg\mathbf{X}\neg A)\}$.

3.2.4 Asymptotic measures

We have seen that the p -values in the multinomial and double binomial models are quite difficult to calculate. However, the p -value can often be approximated easily using asymptotic measures. With certain assumptions, the resulting p -values converge to the correct p -values, when the data size n (or $fr(\mathbf{X})$ and $fr(\neg\mathbf{X})$) tend to infinity. In the following, we introduce two commonly used asymptotic measures for independence testing: the χ^2 -measure and mutual information. In statistics, the latter corresponds to the *log likelihood ratio* (Neyman and Pearson 1928).

The main idea of asymptotic tests is that instead of estimating the probability of the contingency table as such, we calculate some better behaving test statistic T . If T

gets value t , we estimate the probability of $P(T \geq t)$ (assuming that large T -values indicate a strong dependency).

In the case of the χ^2 -test, the test statistic is the χ^2 -measure. Now the variables are binary and Eq. (4) reduces into a simpler form:

$$\begin{aligned}\chi^2 &= \sum_{i=0}^1 \sum_{j=0}^1 \frac{n(P(I_{\mathbf{X}}=i, A=j) - P(I_{\mathbf{X}}=i)P(A=j))^2}{P(I_{\mathbf{X}}=i)P(A=j)} \\ &= \frac{n(P(\mathbf{X}, A) - P(\mathbf{X})P(A))^2}{P(\mathbf{X})P(\neg\mathbf{X})P(A)P(\neg A)} = \frac{n\delta^2(\mathbf{X}, A)}{P(\mathbf{X})P(\neg\mathbf{X})P(A)P(\neg A)}. \quad (15)\end{aligned}$$

So, in principle, each term measures how much the observed frequency $fr(I_{\mathbf{X}}=i, A=j)$ deviates from its expectation $nP(I_{\mathbf{X}}=i)P(A=j)$ under the independence assumption. If the data size n is sufficiently large and none of the expected frequencies is too small, the χ^2 -measure follows approximately the χ^2 -distribution with one degree of freedom. As a classical rule of thumb (Fisher 1925), the χ^2 -measure can be used only, if all expected frequencies $nP(I_{\mathbf{X}}=i)P(A=j)$, $i, j \in \{0, 1\}$, are at least 5. However, the approximations can still be poor in some situations, when the underlying binomial distributions are skewed, e.g., if $P(A)$ is near 0 or 1, or if $fr(\mathbf{X})$ and $fr(\neg\mathbf{X})$ are far from each other (Yates 1984; Agresti 1992). According to Carriere (2001), this is quite typical for data in medical science.

One reason for the inaccuracy of the χ^2 -measure is that the original binomial distributions are discrete while the χ^2 -distribution is continuous. A common solution is to make a *continuity correction* and subtract 0.5 from the expected frequency $nP(\mathbf{X})P(A)$. According to Yates (1984) the resulting continuity corrected χ^2 -measure can give a good approximation to Fisher's p_F , if the underlying hypergeometric distribution is not markedly skewed. However, according to Haber (1980) the resulting χ^2 -value can underestimate the significance, while the uncorrected χ^2 -value overestimates it.

Mutual information is another popular asymptotic measure, which has been used to test independence. For binary variables Eq. (5) becomes

$$MI = \log \frac{P(\mathbf{X}A)^{P(\mathbf{X}A)} P(\mathbf{X}\neg A)^{P(\mathbf{X}\neg A)} P(\neg\mathbf{X}A)^{P(\neg\mathbf{X}A)} P(\neg\mathbf{X}\neg A)^{P(\neg\mathbf{X}\neg A)}}{P(\mathbf{X})^{P(\mathbf{X})} P(\neg\mathbf{X})^{P(\neg\mathbf{X})} P(A)^{P(A)} P(\neg A)^{P(\neg A)}}. \quad (16)$$

Mutual information is actually an information theoretic measure, but in statistics $2n \cdot MI$ is known as log likelihood ratio or the G -test of independence. It follows asymptotically the χ^2 -distribution (Wilks 1935) and often it gives similar results to the χ^2 -measure (Vilalta and Oblinger 2000). However, sometimes the two tests can give totally different results (Agresti 1992).

3.2.5 Selecting the right model

Selecting the right sampling model and defining the extremeness relation is a controversial problem, which statisticians have argued for the last century (see e.g., Yates

1984; Agresti 1992; Lehmann 1993; Upton 1982; Howard 1998). Therefore, we cannot give any definite recommendations which model to select but each situation should be judged in its own context.

The main decision is whether the analysis should be done conditionally or unconditionally and which variables N , N_X , or N_A should be considered fixed. In the multinomial model all variables except $N=n$ are randomized. However, if the model is conditioned with $N_X = fr(\mathbf{X})$, it leads to the double binomial model. If the double binomial model is conditioned with $N_A = fr(A)$, it leads to the hypergeometric model. For completeness, we could also consider the *Poisson model* where all variables, including N , are unfixed Poisson variables. If the Poisson model is conditioned with the given data size, $N=n$, it leads to the multinomial model (Lehmann and Romano 2005, ch. 4.6–4.7).

In principle, the sampling scheme should be decided before the data is gathered. However, in pattern discovery the data may not be sampled according to a particular scheme. In this situation the main choices are to perform an unconditional analysis where none of the margins are considered fixed or a conditional analysis where all margins are considered fixed. The main argument of the unconditional approach is that the results are better generalizable outside the data set, if some variables are kept unfixed. However, both multinomial and double binomial models are computationally demanding, and in practice the corresponding asymptotic tests have been used instead. The opponents have argued that the unconditional approach is also conditional on the data, since the unknown parameters (p_X and/or p_A) are anyway estimated from the observed counts ($fr(\mathbf{X})$ and/or $fr(A)$). Therefore, Fisher and his followers have suggested that we should always assume both N_X and N_A fixed and use Fisher's exact test or—when it is heavy to compute—a suitable asymptotic test.

In pattern discovery the most popular choices for evaluating dependency rules and other similar bipartition dependence patterns in the variable-based interpretation have been Fisher's exact test (e.g., Hämmäläinen 2012; Terada et al. 2013b, 2015; Llinares López et al. 2015; Jabbar et al. 2016) and the χ^2 -test (e.g., Morishita and Sese 2000; Morishita and Nakaya 2000; Nijssen and Kok 2006; Hämmäläinen 2011; Jin et al. 2012; Terada et al. 2015). Both of these tests have also been used for evaluating significance of improvement (see Sect. 4). According to our cross-validation experiments (Hämmäläinen 2012), the χ^2 -measure can be quite unreliable, in the sense that the discovered dependency rules may not hold in the test data at all or their lift and leverage values differ significantly between the training and test sets. The problem is alleviated to some extent when the continuity correction is used, but still the errors can be considerable. On the contrary, Fisher's p has turned out to be a very robust and reliable measure in the dependency rule search and we recommend it as a first choice whenever applicable. There is also an accurate approximation of Fisher's p when faster evaluation is needed (Hämmäläinen 2016). Mutual information is also a good alternative and it often produces the same rules as p_F .

3.3 Sampling models for the value-based interpretation

In the value-based interpretation the idea is that we would like to find events $\mathbf{X}A$ or $\mathbf{X}\neg A$, which express a strong positive dependency, even if the dependency between

variables I_X and A were relatively weak. In this case the strength of the dependency is usually measured by lift, because leverage has the same absolute value for all events $\mathbf{X}A$, $\mathbf{X}\neg A$, $\neg\mathbf{X}A$, $\neg\mathbf{X}\neg A$. However, lift alone is not a reliable measure, because it obtains its maximum value also when $fr(\mathbf{X}A=i) = fr(\mathbf{X}) = fr(A=i) = 1$ ($i \in \{0, 1\}$)—i.e., when the rule occurs on just one row (Hahsler et al. 2006). Such a rule is quite likely due to chance and hardly interesting (see Example 1). Therefore, we should evaluate the probability of observing such a large lift value, if \mathbf{X} and A were actually independent (independence testing, $H_0: \Gamma = 1$ (Benjamini and Leshno 2005)) or, alternatively, that the lift is at most some threshold $\gamma_0 > 1$ ($H_0: \Gamma \leq \gamma_0$ (Lallich et al. 2007)).

The p -value is defined like in the variable-based testing by Eq. (10). The only difference is how to define the extremeness relation $\tau_i \geq \tau_j$. A necessary condition for the extremeness of table τ_i over τ_j is that in τ_i the lift is larger than in τ_j . However, since the lift is largest, when N_X and/or N_A are smallest (and $N_{XA} = N_X$ or $N_{XA} = N_A$), it is sensible to require that also N_{XA} is larger in τ_i than in τ_j . If both N_X and N_A are fixed, then the lift is larger than observed if and only if the leverage is larger than observed, and it is enough to consider tables where $N_{XA} \geq fr(\mathbf{X}A)$. However, if either N_X , N_A , or both are unfixed, then we should always check the lift $\Gamma = \frac{nN_{XA}}{N_X N_A}$ and compare it to the observed lift $\gamma(\mathbf{X}, A)$.

In the following, we will describe different approaches for evaluating statistical significance of dependency rules in the value-based interpretation. The approaches fall into two categories depending on whether the dependence is tested only in the part of data where the rule antecedent holds or in the whole data. We will call these main strategies *partial* and *complete evaluation of significance* according to corresponding measures that are called *partial* and *complete evaluators* (Vilalta and Oblinger 2000). We introduce three approaches: partial evaluation with a single binomial test, complete evaluation under the classical sampling models, and complete evaluation with a single binomial test. Finally, we discuss the problem of selecting the right model.

3.3.1 Partial evaluation with a single binomial test

In the previous research on association rules, some authors (Dehaspe and Toivonen 2001; Lallich et al. 2007, 2005; Bruzzese and Davino 2003; Megiddo and Srikant 1998) have speculated how to test the null hypothesis $\Gamma = 1$. For some reason, it has often been taken for granted that one should perform partial evaluation and evaluate significance of rule $\mathbf{X} \rightarrow A$ in the part of the data where \mathbf{X} is true. As a solution, it has been suggested to use only a single binomial from the double binomial model. This is equivalent to assuming two infinite baskets of apples, the red and green one, but taking only a sample of $fr(\mathbf{X})$ apples from the red basket and trying to decide whether there is a dependency between the red colour and sweetness. It is assumed that $N_{XA} \sim \text{Bin}(fr(\mathbf{X}), p_A)$ and the unknown parameter p_A is estimated from the data, as usual. For positive dependence the p -value is defined as (Dehaspe and Toivonen 2001)

$$p = \sum_{i=fr(\mathbf{X}A)}^{fr(\mathbf{X})} \binom{fr(\mathbf{X})}{i} P(A)^i P(\neg A)^{fr(\mathbf{X})-i} \quad (17)$$

and for negative dependence as (Dehaspe and Toivonen 2001)

$$p = \sum_{i=0}^{fr(\mathbf{X}A)} \binom{fr(\mathbf{X})}{i} P(A)^i P(\neg A)^{fr(\mathbf{X})-i}. \quad (18)$$

We see that $N_X = fr(\mathbf{X})$ is the only variable that has to be fixed—even N can be unfixed. We note that since N_A is unfixed, i goes from $fr(\mathbf{X}A)$ to $fr(\mathbf{X})$ (and not to $\min\{fr(\mathbf{X}), fr(A)\}$) in the case of positive dependence (Lallich et al. 2005). The idea is that when $N_{XA} \geq fr(\mathbf{X}A)$, then $\frac{N_{XA}}{fr(\mathbf{X})} \geq P(A|\mathbf{X})$, and since $p_A = P(A)$ was fixed, then also $\Gamma \geq \gamma(\mathbf{X}, A)$. Similarly, in the negative case $\Gamma \leq \gamma(\mathbf{X}, A)$. So, the test checks correctly all cases where the lift is at least as large (or as small) as observed.

Since the cumulative binomial probability is quite difficult to calculate, it is common to estimate it asymptotically by the z -score. The z -score measures how many standard deviations the observed frequency deviates from its expectation. In the case of positive dependence, the binomial variable N_{XA} has expected value $\hat{\mu} = fr(\mathbf{X})P(A)$ and standard deviation $\hat{\sigma} = \sqrt{fr(\mathbf{X})P(A)P(\neg A)}$. The corresponding z -score is (Lallich et al. 2005; Bruzese and Davino 2003)

$$z = \frac{fr(\mathbf{X}A) - \hat{\mu}}{\hat{\sigma}} = \frac{fr(\mathbf{X}A) - fr(\mathbf{X})P(A)}{\sqrt{fr(\mathbf{X})P(A)P(\neg A)}} = \frac{\sqrt{n}\delta(\mathbf{X}, A)}{\sqrt{P(\mathbf{X})P(A)P(\neg A)}}. \quad (19)$$

If $fr(\mathbf{X})$ is sufficiently large and $P(A)$ is not too near to 1 or 0, the z -score follows the standard normal distribution. However, when the expected frequency $fr(\mathbf{X})P(A)$ is low (as a rule of thumb < 5), the binomial distribution is positively skewed. This means that the z -score overestimates the significance.

It is also possible to construct a partial evaluator from the mutual information (Eq. 16) by ignoring terms related to $\neg \mathbf{X}$. The result is known as *J-measure* (Smyth and Goodman 1992):

$$J = P(\mathbf{X}A) \log \frac{P(\mathbf{X}A)}{P(\mathbf{X})P(A)} + P(\mathbf{X}\neg A) \log \frac{P(\mathbf{X}\neg A)}{P(\mathbf{X})P(\neg A)}. \quad (20)$$

However, it is an open problem how the corresponding p -value could be evaluated and whether the *J-measure* could be used for estimating statistical significance.

The problem of all partial evaluators is that two rules with different antecedents \mathbf{X} are not comparable. So, all rules (with different \mathbf{X}) are thought to be from different populations and are tested in different parts of the data. We also note that the single binomial probability (like the double binomial probability) is not an exchangeable measure in the sense that generally $p(\mathbf{X} \rightarrow A) \neq p(A \rightarrow \mathbf{X})$. The same holds for the corresponding z -score and *J-measure*. This can be counter-intuitive when the task is to search for statistical dependencies, and these measures should be used with care. In addition, with this binomial model the significance of the positive dependence between \mathbf{X} and A is generally not the same as the significance of the negative dependence between \mathbf{X} and $\neg A$. With the corresponding z -score the significance values are related,

and

$$z_{pos}(\mathbf{X} \rightarrow A) = -z_{neg}(\mathbf{X} \rightarrow \neg A),$$

where z_{pos} denotes the z -score of positive dependence and z_{neg} the z -score of negative dependence. With the J -measure the significance of positive dependence between \mathbf{X} and A and the significance of negative dependence between \mathbf{X} and $\neg A$ are equal.

3.3.2 Complete evaluation under the classical sampling models

Let us now analyze the value-based significance of dependency rules using the classical statistical models. For simplicity we consider only positive dependence. We assume that the extremeness relation is defined by lift Γ and frequency N_{XA} , i.e., a contingency table is more extreme than the observed contingency table, if it has $\Gamma \geq \gamma(\mathbf{X}, A)$ and $N_{XA} \geq fr(\mathbf{X}A)$.

In the multinomial model only the data size $N = n$ is fixed. Each contingency table, described by triplet $\langle N_X, N_A, N_{XA} \rangle$, has probability $P(N_{XA}, N_X - N_{XA}, N_A - N_{XA}, n - N_X - N_A + N_{XA} \mid n, p_X, p_A)$, defined by Eq. (11). The p -value is obtained when we sum over all possible triplets where $\Gamma \geq \gamma(\mathbf{X}, A)$:

$$p = \sum_{N_X=0}^n \sum_{N_{XA}=fr(\mathbf{X}A)}^{N_X} \sum_{N_A=N_{XA}}^{Q_1} P(N_{XA}, N_X - N_{XA}, N_A - N_{XA}, n - N_X - N_A + N_{XA} \mid n, p_X, p_A),$$

where $Q_1 = \frac{nN_{XA}}{\gamma(\mathbf{X}, A)N_X}$. (We note that the terms are zero, if $N_X < N_{XA}$.)

In the double binomial model $N_X = fr(\mathbf{X})$ is also fixed. Each contingency table, described by pair $\langle N_A, N_{XA} \rangle$, has probability $P(N_{XA}, N_A - N_{XA} \mid n, fr(\mathbf{X}), p_A)$ by Eq. (12). Now we should sum over all possible pairs, where $\Gamma \geq \gamma(\mathbf{X}, A)$:

$$p = \sum_{N_{XA}=fr(\mathbf{X}A)}^n \sum_{N_A=N_{XA}}^{Q_2} P(N_{XA}, N_A - N_{XA} \mid n, fr(\mathbf{X}), p_A),$$

where $Q_2 = \frac{nN_{XA}}{\gamma(\mathbf{X}, A)fr(\mathbf{X})}$.

In the hypergeometric model also $N_A = fr(A)$ is fixed. As noted before, the extremeness relation is now the same as in the variable-based case and the p -value is defined by Eq. (13). This is an important observation, because it means that *Fisher's exact test tests significance also in the value-based interpretation*. The same is not true for the first two models, where rule $\mathbf{X} \rightarrow A$ can get a different p -value in variable-based and value-based interpretations.

3.3.3 Complete evaluation with a single binomial test

When N_X and/or N_A are unfixed the p -values are quite heavy to compute. Therefore, we will now introduce a simple binomial model (suggested by Hämmäläinen (2010b))

and as model 2 by Lallich et al. (2005)), where it is enough to sum over just one variable. The binomial probability can be further estimated by an equivalent z -score or the z -score can be used as an asymptotic test measure as such. Contrary to the previously described binomial test, this test performs a complete evaluation in the whole data set, which means that the p -values of different rules are comparable.

Let us suppose that we have an infinite basket of apples where the probability of red and sweet apples is p_{XA} . According to the independence assumption $p_{XA} = p_X p_A$. A sample of n apples is taken randomly from the basket. The probability of obtaining N_{XA} sweet red apples among all n apples is defined by binomial probability

$$P(N_{XA} | n, p_X, p_A) = \binom{n}{N_{XA}} (p_X p_A)^{N_{XA}} (1 - p_X p_A)^{n - N_{XA}}.$$

Since N_{XA} is the only variable which occurs in the probability, the extremeness relation is defined simply by $\tau_i \geq \tau_o \Leftrightarrow N_{XA} \geq fr(\mathbf{X}A)$. When the unknown parameters p_X and p_A are estimated from the data, the p -value of rule $\mathbf{X} \rightarrow A$ becomes

$$p_{bin} = \sum_{i=fr(\mathbf{X}A)}^n \binom{n}{i} (P(\mathbf{X})P(A))^i (1 - P(\mathbf{X})P(A))^{n-i}. \quad (21)$$

Since N_{XA} is a binomial variable with expected value $\hat{\mu} = nP(\mathbf{X})P(A)$ and standard deviation $\hat{\sigma} = \sqrt{nP(\mathbf{X})P(A)(1 - P(\mathbf{X})P(A))}$, the corresponding z -score is

$$z(\mathbf{X} \rightarrow A) = \frac{fr(\mathbf{X}A) - nP(\mathbf{X})P(A)}{\sqrt{nP(\mathbf{X})P(A)(1 - P(\mathbf{X})P(A))}} = \frac{\sqrt{n}\delta(\mathbf{X}, A)}{\sqrt{P(\mathbf{X})P(A)(1 - P(\mathbf{X})P(A))}}. \quad (22)$$

Because the discrete binomial distribution is approximated by the continuous normal distribution, the continuity correction can be useful, like with the χ^2 -measure.

We note that this binomial probability and the corresponding z -score are exchangeable, which is intuitively a desired property. However, the statistical significance of positive dependence between \mathbf{X} and A is generally not the same as the significance of negative dependence between \mathbf{X} and $\neg A$. For example, the z -score for negative (or, equally, positive) dependence between \mathbf{X} and $\neg A$ is

$$\begin{aligned} z(\mathbf{X} \rightarrow \neg A) &= \frac{fr(\mathbf{X}\neg A) - nP(\mathbf{X})P(\neg A)}{\sqrt{nP(\mathbf{X})P(\neg A)(1 - P(\mathbf{X})P(\neg A))}} \\ &= \frac{-\sqrt{n}\delta(\mathbf{X}, A)}{\sqrt{P(\mathbf{X})P(\neg A)(1 - P(\mathbf{X})P(\neg A))}}. \end{aligned}$$

3.3.4 Selecting the right model

The main decision in the value-based interpretation is whether the significance of dependency rule $\mathbf{X} \rightarrow A$ is evaluated in the whole data or only in the part of data where \mathbf{X} holds. This decision is critical, because partial and complete evaluators can

Table 3 Comparison of p -values and asymptotic measures for example rules $\mathbf{X} \rightarrow A$ and $\mathbf{Y} \rightarrow A$

	$\mathbf{X} \rightarrow A$	$\mathbf{Y} \rightarrow A$
p_{bin1}	1.06e-4	2.21e-5
p_{mul}	8.86e-13	1.01e-19
p_F	1.60e-12	7.47e-19
p_{double}	2.05e-13	7.35e-20
z_1	4.20	4.36
p_{bin2}	9.31e-10	8.08e-8
z_2	5.48	5.16
J	0.36	0.15

disagree significantly in their ranking and selection of rules. This is demonstrated in the following example.

Example 3 Let us compare two rules, $\mathbf{X} \rightarrow A$ and $\mathbf{Y} \rightarrow A$, in the value-based interpretation. The frequencies are $n = 100$, $fr(A) = 50$, $fr(\mathbf{X}) = fr(\mathbf{X}A) = 30$, $fr(\mathbf{Y}) = 60$, and $fr(\mathbf{Y}A) = 50$, i.e., $P(A|\mathbf{X}) = 1$ and $P(\mathbf{Y}|A) = 1$. The p -values, z -scores, and J -values are given in Table 3.

All of the traditional association rule measures (binomial p_{bin2} , Eq. (17), its z -score z_2 , and J -measure) favour rule $\mathbf{X} \rightarrow A$, while all the other measures (binomial p_{bin1} , Eq. (21), its z -score z_1 , multinomial p_{mul} , double binomial p_{double} , and Fisher's p_F) rank rule $\mathbf{Y} \rightarrow A$ better. In the three classical models, the difference between the rules is quite remarkable.

In general, we do not recommend partial evaluation for dependency rule mining. The main problem is that the p -values of discovered rules are not comparable, because each of them has been tested in a different part of data. In addition, the measures are not exchangeable, which means that $\mathbf{X} \rightarrow A$ can get a totally different ranking than $A \rightarrow \mathbf{X}$, even if they express the same dependency between events.

When the classical statistical models are used, the only difference to the variable-based interpretation is that now the discrepancy measure is lift. Computationally, the only practical choices are the hypergeometric model and asymptotic measures. The hypergeometric model produces reliable results, but it tends to favour large leverage instead of lift, which might be more interesting in the value-based interpretation. In addition, one should check for each rule $\mathbf{X} \rightarrow A$ that the dependency is due to strong $\gamma(\mathbf{X}, A)$ and not due to $\gamma(\neg\mathbf{X}, \neg A)$. With this checking the χ^2 -measure can also be used. According to our experiment (Hämläinen 2010a, Ch.5), the χ^2 -measure and the z -score (Eq. 22) tend to find rules with the strongest lift (among all compared measures), but at the same time the results are also the most unreliable. Robustness of the χ^2 -measure can be improved with the continuity correction, but with the z -score it has only a marginal effect. One solution is to use the z -score only for preliminary pruning and select the rules with the corresponding binomial probability (Hämläinen 2010b). Based on these considerations, we cannot give a universal recommendation, but Fisher's exact test is always a safe choice, if there is no specific need to maximize

lift. If large lift values are desired, one could consider either the χ^2 -measure or the z -score accompanied by an exact binomial test.

4 Redundancy and significance of improvement

An important task in dependency rule discovery is to identify redundant rules, which add little or no additional information on statistical dependencies to other rules. In this section we consider an important type of redundancy called *superfluosness*, where a more specific dependency rule does not improve its generalization significantly. We present statistical significance tests for evaluating superfluosness in the value-based and variable-based interpretations. Finally, we will briefly discuss the relationship to a more general approach of speciousness testing.

4.1 Redundant and superfluous rules

According to a classical definition (Bastide et al. 2000) “An association rule is redundant if it conveys the same information or less general information than the information conveyed by another rule of the same usefulness and the same relevance.” However, what is considered useful or relevant depends on the modelling purpose, and numerous definitions of redundant or uninformative rules have been proposed.

In traditional association rule research, the goal has been to find all sufficiently frequent and ‘confident’ (high precision) rules. Thus, if the sufficient frequency or precision of a rule can be derived from other rules, the rule can be considered redundant (e.g., Aggarwal and Yu 2001; Goethals et al. 2005; Cheng et al. 2008; Li and Hamilton 2004; see also a good overview by Balcazar (2010)). On the other hand, when the goal is to find statistical dependency rules, then rules that are merely side-products of other dependencies can be considered uninformative. An important type of such dependencies are superfluous specializations ($\mathbf{X} \rightarrow A$) of more general dependency rules ($\mathbf{Y} \rightarrow A$, $\mathbf{Y} \subsetneq \mathbf{X}$). This concept of superfluous rules covers earlier notions of non-optimal or superfluous classification rules (Li 2006), (statistically) redundant rules (Hu and Rao 2007; Hämmäläinen 2012) and unproductive rules (Webb 2007).

Superfluous rules are a common problem, because rules ‘inherit’ dependencies from their ancestor rules unless their extra factors reverse the dependency. This is regrettable, because undetected superfluous rules may lead to quite serious misconceptions. For example, if disease D is caused by gene group \mathbf{Y} (i.e., $\mathbf{Y} \rightarrow D$), we are likely to find a large number of other dependency rules $\mathbf{YQ} \rightarrow D$ where \mathbf{Q} contains coincidental genes. Now one could make a conclusion that the combination \mathbf{YQ}_1 (with some arbitrary \mathbf{Q}_1) predisposes to disease D and begin preventive care only with these patients.

Intuitively, the idea of superfluosness is clear. A superfluous rule $\mathbf{X} \rightarrow A$ contains extraneous variables $\mathbf{Q} \subsetneq \mathbf{X}$ which have no effect or only weaken the original dependency $\mathbf{X} \setminus \mathbf{Q} \rightarrow A$. It is also possible that \mathbf{Q} apparently improves the dependency but the improvement is spurious (due to chance). In this case the apparent improvement occurs only in the sample, and it may be detected with appropriate statistical signifi-

cance tests. We recall that significance tests do not necessarily detect all superfluous rules but we can always adjust the significance level to prune more or less potentially superfluous rules. Formalizing the idea of superfluosness is more difficult, because it depends on the used measure, assumed statistical model, required significance level, and—most of all—whether we are using the value-based or variable-based interpretation. Therefore, we give here only a tentative, generic definition of superfluosness.

Definition 7 (*Superfluous dependency rules*) Let T be a goodness measure which is used to evaluate dependency rules. Let us assume that T is increasing by goodness and rule $\mathbf{X} \rightarrow A=i$ improves $\mathbf{Y} \rightarrow A=i$, $i \in \{0, 1\}$, when $T(\mathbf{X} \rightarrow A=i) > T(\mathbf{Y} \rightarrow A=i)$ (for decreasing measures, $T(\mathbf{X} \rightarrow A=i) < T(\mathbf{Y} \rightarrow A=i)$). Let \mathcal{M} be a statistical model which is used for determining the statistical significance and α the selected significance level.

Rule $\mathbf{X} \rightarrow A=i$ is superfluous (given T , \mathcal{M} , and α) if there exists rule $\mathbf{Y} \rightarrow A=i$, $\mathbf{Y} \subsetneq \mathbf{X}$, such that

- (i) $T(\mathbf{X} \rightarrow A=i) \leq T(\mathbf{Y} \rightarrow A=i)$ (vs. $T(\mathbf{X} \rightarrow A=i) \geq T(\mathbf{Y} \rightarrow A=i)$) or
- (ii) Improvement of rule $\mathbf{X} \rightarrow A=i$ over rule $\mathbf{Y} \rightarrow A=i$ is not significant at level α (value-based interpretation) or
- (iii) Improvement of rule $\mathbf{X} \rightarrow A=i$ over rule $\mathbf{Y} \rightarrow A=i$ is less significant than the improvement of rule $\neg\mathbf{X} \rightarrow A \neq i$ over rule $\neg\mathbf{Y} \rightarrow A \neq i$ (variable-based interpretation).

We note that in a special case where $P(\mathbf{X}) = P(\mathbf{Y})$, rules $\mathbf{X} \rightarrow A=i$ and $\mathbf{Y} \rightarrow A=i$, $\mathbf{Y} \subsetneq \mathbf{X}$, have equivalent contingency tables and they obtain the same measure value with all commonly used goodness measures (that are functions of N_X , N_A , N_{XA} and n). Otherwise, if $P(\mathbf{X}) < P(\mathbf{Y})$, the contingency tables are different and rule $\mathbf{X} \rightarrow A=i$ may or may not improve $\mathbf{Y} \rightarrow A=i$ depending on the observed counts and the selected goodness measure. The special case $\mathbf{Y} \subsetneq \mathbf{X}$, $P(\mathbf{X}) = P(\mathbf{Y})$, is closely connected to the notions of *closed itemsets* (\mathbf{X} such that $\forall \mathbf{Z} \supsetneq \mathbf{X}: P(\mathbf{X}) > P(\mathbf{Z})$) and their *minimal generators* ($\mathbf{Y} \subseteq \mathbf{X}$ such that $P(\mathbf{Y}) = P(\mathbf{X})$ and $\nexists \mathbf{Y}' \subsetneq \mathbf{Y}: P(\mathbf{Y}') = P(\mathbf{Y})$) (Pasquier et al. 1999; Bastide et al. 2000). If the rule antecedents \mathbf{X} are selected only among closed sets, some of them may have distinct minimal generators $\mathbf{Y} \subsetneq \mathbf{X}$ and are necessarily superfluous. This is avoided, if the rule antecedents are selected only among minimal generators (also called *free sets* (Boulicaut et al. 2000)), but the rules may still be superfluous when tested against more general rules.

4.2 Testing superfluosness in the value-based interpretation

Let us first consider the problem of superfluosness in the value-based interpretation, where the significance tests are somewhat simpler. To simplify notations, we will consider only rule $\mathbf{X} \rightarrow A$ with a positive-valued consequent. For $\mathbf{X} \rightarrow \neg A$ the tests are analogous, except A is replaced by $\neg A$.

In traditional association rule research, the goodness measure T is precision (or, equivalently, lift, because the antecedent is fixed). Rule $\mathbf{X} \rightarrow A$ is called *productive*, if $P(A|\mathbf{X}) > P(A|\mathbf{Y})$ for all $\mathbf{Y} \subsetneq \mathbf{X}$ (e.g., Bayardo et al. 2000; Webb 2007). The

significance of productivity is tested separately for all $Y \rightarrow A$, $Y \subsetneq X$, and all p -values should be below some fixed threshold α .

Let us now notate $X=YQ$ (i.e., $Q = X \setminus Y$, $Q \neq \emptyset$) so that we can compare rule $YQ \rightarrow A$ to a simpler rule $Y \rightarrow A$. In each test, the null hypothesis is that there is no improvement in the precision: $P(A|YQ) = P(A|Y)$. The condition means that Q and A are conditionally independent given Y . The significance is estimated by calculating $p(Q \rightarrow A|Y)$, i.e., the p -value of rule $Q \rightarrow A$ in the set where Y holds. Now it is quite natural to assume $fr(Y)$, $fr(YQ)$, and $fr(YA)$ fixed, which leads to the hypergeometric model. The corresponding test is Fisher's exact test for conditional independence, and the significance of productivity of $YQ \rightarrow A$ over $Y \rightarrow A$ is (Webb 2007)

$$p(Q \rightarrow A | Y) = \sum_{j=0}^{J_1} \frac{\binom{fr(YQ)}{fr(YQA)+j} \binom{fr(Y-Q)}{fr(Y-QA)-j}}{\binom{fr(Y)}{fr(YA)}}, \quad (23)$$

where $J_1 = \min\{fr(YQ \neg A), fr(Y \neg QA)\}$. When the χ^2 -measure is used to estimate the significance of productivity, the equation is (Liu et al. 1999)

$$\chi^2(Q \rightarrow A | Y) = \frac{fr(Y)(P(Y)P(YQA) - P(YQ)P(YA))^2}{P(YQ)P(Y \neg Q)P(YA)P(Y \neg A)}. \quad (24)$$

In principle, measure T can be any goodness measure for statistical dependence between values, including the previously introduced binomial probabilities and corresponding z -scores. However, different measures can disagree on their ranking of rules and which rules are considered superfluous. For example, leverage has a strong bias in favour of general rules, when compared to lift or precision. This is clearly seen from expression $\delta(Y, A) = P(Y)(P(A|Y) - P(A)) = P(Y)(\gamma(Y, A) - 1)$. On the other hand, asymptotic measures like the z -score and the χ^2 -measure tend to overestimate the significance, when the frequencies are small. Therefore, it is possible that a rule is not superfluous, when evaluated with an asymptotic measure, but superfluous, when the exact p -values are calculated.

4.3 Testing superfluosness in the variable-based interpretation

In the variable-based interpretation, superfluosness of dependency rules is more difficult to judge, because there may be two kinds of improvement into opposite directions in the same time. Improvement of rule $YQ \rightarrow A$ over rule $Y \rightarrow A$ is tested as in the value-based interpretation. However, in the same time rule $\neg Y \rightarrow \neg A$ may improve a more general rule $\neg(YQ) \rightarrow \neg A$, and one should weigh which improvement is more significant.

The significance of improvement of rule $\neg Y \rightarrow \neg A$ over $\neg(YQ) \rightarrow \neg A$ is tested in the same way as productivity of $YQ \rightarrow A$ over $Y \rightarrow A$. However, now the null hypothesis is conditional independence between $\neg Y$ and $\neg A$ given $\neg(YQ) = \neg Y \vee Y \neg Q$. It is natural to assume $fr(\neg(YQ))$, $fr(\neg Y)$, and $fr(\neg(YQ) \neg A)$ fixed, which leads to an exact test

$$p(\neg Y \rightarrow \neg A \mid \neg(YQ)) = \sum_{j=0}^{J_2} \frac{\binom{fr(\neg Y)}{fr(\neg Y \neg A) + j} \binom{fr(Y \neg Q)}{fr(Y \neg Q \neg A) - j}}{\binom{fr(\neg(YQ))}{fr(\neg(YQ) \neg A)}}, \quad (25)$$

where $J_2 = \min\{fr(\neg YA), fr(Y \neg Q \neg A)\}$. The corresponding χ^2 -test is

$$\begin{aligned} \chi^2(\neg Y \rightarrow \neg A \mid \neg(YQ)) \\ = \frac{fr(\neg(YQ))(P(\neg(YQ))P(\neg Y \neg A) - P(\neg Y)P(\neg(YQ) \neg A))^2}{P(\neg Y)P(Y \neg Q)P(\neg(YQ)A)P(\neg(YQ) \neg A)}. \end{aligned} \quad (26)$$

An important property of variable-based superfluosity testing is that sometimes significance tests can be avoided altogether. This is possible with such goodness measures, for which any improvement is significant improvement. One such measure is p_0 , the first and largest term of p_F . It can be shown (Hämäläinen and Webb 2017) that for dependency rules $Y \rightarrow A$ and $YQ \rightarrow A$ holds

$$\frac{p_0(YQ \rightarrow A)}{p_0(Y \rightarrow A)} = \frac{p_0(Q \rightarrow A \mid Y)}{p_0(\neg Y \rightarrow \neg A \mid \neg(YQ))}.$$

It is an open problem whether the equality holds exactly for the cumulative probability, p_F , but at least it holds approximately. This is also the justification for the simpler superfluosity testing in Kingfisher (Hämäläinen 2012), where a dependency rule is considered superfluous if it has a larger (poorer) p_F -value than some of its ancestor rules.

Previously, we have already shown that goodness measures for the variable-based and value-based interpretations can diverge quite much in their ranking of rules. The same holds for superfluosity testing. The following example demonstrates that the same rule may or may not be superfluous depending on the interpretation.

Example 4 Let us reconsider the rules $X \rightarrow A$ ($= YQ \rightarrow A$) and $Y \rightarrow A$ in Example 3. Rule $X \rightarrow A$ is clearly productive with respect to $Y \rightarrow A$ ($P(A|Y) = 1.00$ vs. $P(A|X) = 0.83$). Similarly, rule $\neg Y \rightarrow \neg A$ is productive with respect to $\neg X \rightarrow \neg A$ ($P(\neg A|\neg Y) = 1.00$ vs. $P(\neg A|\neg X) = 0.71$).

Let us now calculate the significance of productivity using Fisher's exact test. In the value-based interpretation, we evaluate only the first improvement:

$$p(Q \rightarrow A \mid Y) = \frac{\binom{fr(Y \neg Q)}{fr(Y \neg Q \neg A)}}{\binom{fr(Y)}{fr(YA)}} = \frac{\binom{30}{10}}{\binom{60}{50}} = 3.99 \cdot 10^{-4}.$$

The value is so small that we can assume that the productivity is significant and $X \rightarrow A$ is not superfluous.

In the variable-based interpretation, we evaluate also the second improvement:

$$p(Y \rightarrow \neg A \mid \neg(YQ)) = \frac{\binom{fr(Y \neg Q)}{fr(Y \neg Q \neg A)}}{\binom{fr(\neg(YQ))}{fr(\neg(YQ) \neg A)}} = \frac{\binom{30}{10}}{\binom{70}{50}} = 1.86 \cdot 10^{-10}.$$

This value is much smaller than the previous one, which means that the improvement of $\neg \mathbf{Y} \rightarrow \neg A$ over $\neg \mathbf{X} \rightarrow \neg A$ is more significant than the improvement of $\mathbf{X} \rightarrow A$ over $\mathbf{Y} \rightarrow A$. Thus, we would consider rule $\mathbf{X} \rightarrow A$ superfluous. We would have ended up into the same conclusion, if we had simply compared the p_F -values of both rules: $p_F(\mathbf{Y} \rightarrow A) = 7.47 \cdot 10^{-19} < 1.60 \cdot 10^{-12} = p_F(\mathbf{X} \rightarrow A)$.

4.4 Relationship to speciousness

The concept of superfluosness is closely related to *speciousness* (Yule 1903; Hämmäläinen and Webb 2017), where an observed unconditional dependency vanishes or changes its sign when conditioned on other variables, called *confounding factors*. The latter phenomenon, reversal of the direction of the dependency, is also known as *Yule-Simpson's paradox*. In the context of dependency rules, rule $\mathbf{X} \rightarrow A$ is considered specious if there is another rule $\mathbf{Y} \rightarrow A$ or $\mathbf{Y} \rightarrow \neg A$ such that \mathbf{X} and A are either independent or negatively dependent *in the population* when conditioned on \mathbf{Y} and $\neg \mathbf{Y}$. In the sample either of the conditional dependencies may also appear as weakly positive, and one has to test their significance with a suitable test, like Birch's exact test (Birch 1964), conditional mutual information (Hämmäläinen and Webb 2017) or various χ^2 -based tests.

It is noteworthy that the confounding factor \mathbf{Y} does not necessarily share any attributes with \mathbf{X} . However, in a special case when $\mathbf{Y} \subsetneq \mathbf{X}$, Birch's exact test for speciousness of $\mathbf{X} \rightarrow A$ with respect to $\mathbf{Y} \rightarrow A$ reduces to Eq. (23) (significance of productivity). On the other hand, Birch's exact test for speciousness of $\mathbf{Y} \rightarrow A$ with respect to $\mathbf{X} \rightarrow A$ is equivalent to Eq. (25). So, testing superfluosness of $\mathbf{X} \rightarrow A$ with respect to $\mathbf{Y} \rightarrow A$ in a variable-based interpretation can be considered as a special case of testing if $\mathbf{X} \rightarrow A$ is specious by $\mathbf{Y} \rightarrow A$ or vice versa.

5 Dependency sets

Dependency rules capture the most common conception of dependence as a relationship between two elements. Often, however, multiple elements will all interact with each other, and the mutual dependency structure is better represented by set-type of patterns. Dependency sets are a general name for set-type patterns that express interdependence between the elements of the set. In this section we will first give a short overview of set dependency patterns and then describe key approaches for evaluating their statistical significance.

5.1 Overview

Approaches to finding dependency sets differ in terms of the forms of interdependence that they seek to capture. A common starting point is to assume mutual dependence among the elements of the set, i.e., absence of mutual independence (Definition 3). However, this notion is very inclusive because it suffices that \mathbf{X} contains at least one subset $\mathbf{Y} \subseteq \mathbf{X}$ where $P(\mathbf{Y}) \neq \prod_{A_i \in \mathbf{Y}} P(A_i)$. This means that the property of

mutual dependence is monotonic, i.e., all supersets of a mutually dependent set are also mutually dependent. To avoid an excessive number of patterns, dependency sets usually represent only some of all mutually dependent sets, like minimal mutually dependent sets (Brin et al. 1997), sets that present new dependencies in comparison to their subsets (for some $A \in \mathbf{X}$, $\delta(\mathbf{X} \setminus \{A\}, A) \neq 0$) (Meo 2000), or sets for which all bipartitions express statistical dependence (for all $\mathbf{Y} \subsetneq \mathbf{X}$ $\delta(\mathbf{Y}, \mathbf{X} \setminus \mathbf{Y}) \neq 0$) (Webb 2010). We note that the latter two approaches assume bipartition dependence (absence of bipartition independence, Definition 4), which is a stronger condition than mutual dependence.

Compared to dependency rules, dependency sets offer a more compact presentation of dependencies, and in some contexts the reduction in the number of patterns can be quite drastic. This is evident when we recall that any set \mathbf{X} can give rise up to $|\mathbf{X}|$ rules of the form $\mathbf{X} \setminus \{A_i\} \rightarrow A_i$ and up to $2^{|\mathbf{X}|} - 2$ rules of the form $\mathbf{X} \setminus \mathbf{Y} \rightarrow \mathbf{Y}$. In many cases these permutation rules reflect the same statistical dependency. This is always true when $|\mathbf{X}| = 2$ ($A \rightarrow B$ and $B \rightarrow A$ present the same dependency), but the same phenomenon can occur also with more complex sets as the following observation demonstrates.

Observation 1 *Let \mathbf{X} be a set of binary attributes such that for all $\mathbf{Y} \subsetneq \mathbf{X}$ $P(\mathbf{Y}) = \prod_{A_i \in \mathbf{Y}} P(A_i)$ (i.e., attributes are mutually independent). Then for all $\mathbf{Z} \subsetneq \mathbf{X}$ $\delta(\mathbf{X} \setminus \mathbf{Z}, \mathbf{Z}) = P(\mathbf{X}) - \prod_{A_i \in \mathbf{X} \setminus \mathbf{Z}} P(A_i) \prod_{A_i \in \mathbf{Z}} P(A_i) = P(\mathbf{X}) - \prod_{A_i \in \mathbf{X}} P(A_i)$.*

This means that when all proper subsets of \mathbf{X} express only mutual independence, then all permutation rules of $\mathbf{X} \setminus \mathbf{Z} \rightarrow \mathbf{Z}$ have the same leverage, frequency and expected frequency, and many goodness measures would rank them equally good. In real world data, the condition holds seldom precisely, but the same phenomenon tends to occur to some extent also when all subsets express at most weak dependence. In this case, it is intuitive to report only set \mathbf{X} instead of listing all of its permutation rules.

In principle, all dependency rules could be represented by dependency sets without losing any other information than the division to an antecedent and a consequent. The reason is that for any dependency rule $\mathbf{X} \setminus \mathbf{Y} \rightarrow \mathbf{Y}$, set \mathbf{X} is mutually dependent. This follows immediately from the fact that mutual independence of \mathbf{X} (Definition 3) implies bipartition independence between $\mathbf{X} \setminus \mathbf{Y}$ and \mathbf{Y} (Definition 4) for any $\mathbf{Y} \subsetneq \mathbf{X}$. However, as explained below, some set dependency approaches have more stringent constraints which may exclude interesting dependency rules selected under other schemes. Further, if mutual independence is violated only by a single bipartition, or if the objective is to find dependencies with a specific element of interest, a dependency rule between the relevant partitions will more concisely convey the relevant information. Which dependency rule or set scheme is most appropriate depends entirely on the analytic objective.

The approaches for finding statistically significant dependency sets can be roughly divided into two categories: (i) selecting dependency sets among frequent item sets and testing their statistical significance afterwards and (ii) searching directly all sufficiently strong and significant dependency sets using appropriate goodness measures and significance tests. In the following subsections we describe the main methods for evaluating statistical significance in these approaches.

5.2 Statistically significant dependency sets derived from candidate frequent itemsets

Frequent itemsets (Agrawal et al. 1996) are undoubtedly the most popular type of set patterns in knowledge discovery. A frequent itemset is a set of true-valued binary attributes (called items, according to the original market-basket setting) whose frequency exceeds some user-specified minimum frequency threshold ('minimum support'). However, being frequent does not ensure that the elements in an itemset express statistical dependence. For example, consider two elements A and B that each occur in all but one example such that the examples in which A and B do not occur differ. In this case itemset $\{A, B\}$ will occur in all but two examples and thus be frequent, but it will represent negative dependence rather than positive dependence such as association discovery typically seeks.

Frequent itemsets have been employed as an initial step in dependency set discovery in order to constrain the number of patterns that must be considered. The idea is to search first all frequent itemsets and then select statistically significant dependency set patterns among them. A limitation is that this approach will fail to discover statistically significant but infrequent dependency sets, which can be the most significant.

The most common null hypothesis used for significance testing of dependency sets is mutual independence between all attributes of the data (Definition 3). Statistical significance of a set is defined as the probability that its frequency is at least as large as observed, given the mutual independence assumption. In principle, any significance testing approach could be used, but often this is done with randomization testing. In swap randomization (e.g., Gionis et al. 2007; Cobb and Chen 2003), both column margins (attribute frequencies) and row margins (numbers of items on each row) are kept fixed. The latter requirement allows suppression of dependencies that are due to co-occurrence of items only due to their appearing solely in rows that contain many items. A variant is iterative randomization (Hanhijärvi et al. 2009). This approach begins with fixed row and column margins, but on each iteration it adds the most significant frequent itemset as a new constraint. The randomization problem is computationally very hard, and thus it is sufficient that the frequencies of itemsets hold only approximately. The process is repeated until no more significant itemsets can be found.

Vreeken and Tatti (2014) have proposed identifying statistically significant dependency sets \mathbf{X} using the binomial test for the null hypothesis of independence of events $A_i \in \mathbf{X}$ (Eq. 7). Under the independence assumption, the probability of \mathbf{X} in the population is $p_X = \prod_{A_i \in \mathbf{X}} p_{A_i}$, where p_{A_i} s can be estimated by observed $P(A_i)$ s as usual. Then the probability of observing $N_X \geq fr(\mathbf{X})$ in a sample of n rows is

$$p_{bin} = \sum_{j=fr(\mathbf{X})}^n \binom{n}{j} \left[\prod_{A_i \in \mathbf{X}} P(A_i) \right]^j \left[1 - \prod_{A_i \in \mathbf{X}} P(A_i) \right]^{n-j}. \quad (27)$$

We note that this test assumes a weaker notion of independence than mutual independence (Definition 3). In consequence, it may find fewer dependency patterns than the previously described randomization test for mutual independence.

5.3 Direct search for significant dependency sets

An alternative approach is to search directly for sets that satisfy specific criteria of statistical dependence and significance, using those criteria to prune the search space and support efficient search. Sometimes, these set patterns are still called ‘rules’ or are represented by the best rule that can be derived from the set. Examples are correlation rules (Brin et al. 1997), strictly non-redundant association rules (Hämmäläinen 2010b, 2011), and—the most rigorous of all—self-sufficient itemsets (Webb 2010; Webb and Vreeken 2014). All these pattern types have three common requirements: dependency set \mathbf{X} expresses mutual dependence, it adds new dependencies to its subsets $\mathbf{Y} \subsetneq \mathbf{X}$, and the dependency is significant with the selected measure.

Correlation rules (Brin et al. 1997) are defined as minimal sets \mathbf{X} , where \mathbf{X} expresses mutual dependence (at least for some $\mathbf{x} \in \text{Dom}(\mathbf{X})$ $P(\mathbf{X}=\mathbf{x})$ is greater than expected under independence) but all $\mathbf{Y} \subsetneq \mathbf{X}$ express mutual independence. The significance is evaluated with the χ^2 -measure

$$\chi^2 = \sum_{(a_1, \dots, a_m) \in \{0,1\}^m} \frac{n \left(P(\mathbf{X}=(a_1, \dots, a_m)) - \prod_{A_i \in \mathbf{X}} P(A_i=a_i) \right)^2}{\prod_{A_i \in \mathbf{X}} P(A_i=a_i)} \quad (28)$$

with one degree of freedom. Since all supersets of \mathbf{X} can only increase the χ^2 -value, only minimal sets whose χ^2 -value exceeds a specified threshold are presented.

Strictly non-redundant association rules are an intermediate form between set type and rule type patterns, where each set is presented by its best rule, whose significance is evaluated in the desired sampling model. The discovered patterns are mutually dependent sets \mathbf{X} that express bipartition dependence between some $A \in \mathbf{X}$ and $\mathbf{X} \setminus \{A\}$ and the bipartition dependence is more significant than any bipartition dependence in simpler sets $\mathbf{Y} \subsetneq \mathbf{X}$ (between any $B \in \mathbf{Y}$ and $\mathbf{Y} \setminus \{B\}$). In the significance testing, one can assume either value- or variable-based interpretation and use any of the sampling models presented in Sect. 3. For search purposes, feasible choices are the binomial model and the corresponding z -score (Hämmäläinen 2010b) for the value-based interpretation, the double binomial test and the corresponding χ^2 -measure (Hämmäläinen 2011) for the variable-based interpretation, and Fisher’s exact test that can be used in both interpretations.

Self-sufficient itemsets are a pattern type that imposes much stronger requirements. The core idea is that an itemset should only be considered interesting if its frequency cannot be explained by assuming independence between any partition of the items. That is, there should be no partition $\mathbf{Q} \subsetneq \mathbf{X}$, $\mathbf{X} \setminus \mathbf{Q}$ such that $P(\mathbf{X}) \approx P(\mathbf{Q})P(\mathbf{X} \setminus \mathbf{Q})$. For example, being male (M) and having prostate cancer (P) are associated and hence should form a dependency set $\{M, P\}$. Suppose that having a name containing a ‘G’ (G) is independent of both factors. Then $\{M, P, G\}$ should not be a dependency set. However, it is more frequent than would be expected by assuming independence between $\{M\}$ and $\{P, G\}$ or between $\{M, G\}$ and $\{P\}$, and hence most interestingness measures would assess both $\{P, G\} \rightarrow \{M\}$ and $\{M, G\} \rightarrow \{P\}$ as interesting. Nonetheless, under the self-sufficient itemset approach $\{M, P, G\}$ can be discarded

because it is not more frequent than would be expected by assuming independence between $\{G\}$ and $\{M, P\}$.

In self-sufficient itemsets this requirement is formalized as a test for *productivity*. It is required that there is a significant positive dependency between every partition of the itemset, when evaluated with Fisher's exact test. In addition, self-sufficient itemsets have two additional criteria: they have to be non-redundant and independently productive.

In the context of self-sufficient itemsets, set \mathbf{X} is considered redundant, if

$$\exists \mathbf{Y} \subsetneq \mathbf{X}, \mathbf{Z} \subsetneq \mathbf{Y} \text{ fr}(\mathbf{Y}) = \text{fr}(\mathbf{Z}). \quad (29)$$

The motivation is that if A is a necessary consequent of another set of items \mathbf{Z} , then $\mathbf{Y} = \{A\} \cup \mathbf{Z}$ should be associated with everything with which \mathbf{Z} is associated. For example $\mathbf{Z} = \{\text{pregnant}\}$ entails $A = \text{female}$ ($\mathbf{Y} = \{\text{female}, \text{pregnant}\}$) and *pregnant* is associated with *oedema*. In consequence, $\mathbf{X} = \{\text{female}, \text{pregnant}, \text{oedema}\}$ is not likely to be interesting if $\{\text{pregnant}, \text{oedema}\}$ is known.

A final form of test that can be employed is whether the frequency of an itemset \mathbf{X} can be explained by the frequency of its productive and nonredundant supersets $\mathbf{Y} \supsetneq \mathbf{X}$. For example, if A, B and C are jointly necessary and sufficient for D then all subsets of $\{A, B, C, D\}$ that include D should be productive and nonredundant. However, they may be misleading, as they fail to capture the full conditions necessary for D . Webb (2008) proposes that if $\mathbf{Y} \supsetneq \mathbf{X}$ is productive and nonredundant, \mathbf{X} should only be considered potentially interesting if it is *independently productive*, meaning that it passes tests for productivity when data covered by $\mathbf{Y} \setminus \mathbf{X}$ are not considered.

6 Multiple testing problem

The goal of pattern discovery is to find all sufficiently good patterns among exponentially many possible candidates. This leads inexorably to the problem of multiple hypothesis testing. The core of this problem is that as the number of tested patterns increases, it becomes ever more likely that spurious patterns pass their tests, causing type I error.

In this section, we will first describe the main principles and popular correction methods that the statistical community has developed to remedy the problem. Such understanding is critical to addressing this issue in the pattern discovery context. We then introduce some special techniques for increasing the power to detect true patterns while controlling the number of false discoveries in the pattern discovery context.

6.1 Overview

The problem of multiple testing is easiest to demonstrate in the classical Neyman–Pearsonian hypothesis testing. Let us suppose we are testing m true null hypotheses (spurious patterns) and in each test the probability of type I error is exactly the selected significance level α . (In general, the probability is at most α , but with increasing power it approaches α .) In this case the expectation is that in every $m \cdot \alpha$ tests a type I error is

	spurious pattern H_0 true	genuine pattern H_1 true	All
declared significant	V false positives	S true positives	R
declared insignificant	U true negatives	T false negatives	$m - R$
All	m_0	$m - m_0$	m

Fig. 7 A contingency table for m significance tests. Here m_0 is an unknown parameter for the number of true null hypotheses and R is an observable random variable for the number of rejected hypotheses. S , T , U , and V are unobservable random variables

committed and a spurious pattern passes the test. With normal significance levels this can be quite a considerable number. For example if $\alpha = 0.05$ and 100,000 spurious patterns are tested, we can expect 5000 of them to pass the test.

Solutions to the multiple testing problem try to control type I errors among all tests. In practice, there are two main approaches: The traditional approach is to control the *familywise error rate* which is the probability of accepting at least one false discovery (rejecting a true null hypothesis). Using the notations of Fig. 7 the familywise error rate is $FWER = P(V \geq 1)$. Another, less stringent approach is to control the *false discovery rate*, which is the expected proportion of false discoveries, $FDR = E\left(\frac{V}{\max\{R, 1\}}\right) = E\left(\frac{V}{R} \mid R > 0\right) P(R > 0)$.

Since $FDR \leq FWER$, control of $FWER$ subsumes control of FDR . In a special case where all null hypotheses are true ($m=m_0$), $FWER = FDR$. The latter means that a FDR controlling method controls $FWER$ in a *weak sense*, when the probability of type I errors is evaluated under the *global null hypothesis* $H_0^C = \cap_{i=1}^m H_i$ (all m hypotheses are true). However, usually it is required that $FWER$ should be controlled in a *strong sense*, under any set of true null hypotheses (for details, see e.g., Ge et al. 2003).

In general, $FWER$ control is preferred when false discoveries are intolerable (e.g., accepting a new medical treatment) or when it is expected that most null hypotheses would be true, while FDR control is often preferred in exploratory research, where the number of potential patterns is large and false discoveries are less serious (e.g., Goeman and Solari 2011).

6.2 Methods for multiple hypothesis testing

The general idea of multiple hypothesis testing methods is to make rejection of individual hypotheses more difficult by adjusting the significance level α (or the corresponding critical value of some test statistic) or, equivalently, adjusting individual p -values, p_1, \dots, p_m , corresponding to hypotheses H_1, \dots, H_m . When the goal is to control $FWER$ at level α , the procedure determines for each hypothesis H_i an adjusted significance level $\hat{\alpha}_i$ (possibly $\hat{\alpha}_1 = \dots = \hat{\alpha}_m$) such that H_i is rejected if and only if $p_i \leq \hat{\alpha}_i$. Alternatively, the p -value of H_i can be adjusted and the adjusted p -value \hat{p}_i

Table 4 Multiple hypothesis testing procedures for controlling *FWER* or *FDR* at level α and assumptions on the dependency structures between true H_i s

Method	Control	Criterion	Assumptions on H_i s
Bonferroni	<i>FWER</i>	$k = \min(i) : p_{i+1} > \frac{\alpha}{m}$	None
Šidák	<i>FWER</i>	$k = \min(i) : p_{i+1} > 1 - (1 - \alpha)^{\frac{1}{m}}$	Independence (or positive dependence)
Holm–Bonferroni	<i>FWER</i>	$k = \min(i) : p_{i+1} > \frac{\alpha}{m-i}$	None
Hochberg	<i>FWER</i>	$k = \max(i) : p_i \leq \frac{\alpha}{m-i+1}$	Independence or certain types of positive dependence
Single-step <i>minP</i>	<i>FWER</i>	$k = \min(i) : P(\min_{1 \leq j \leq m} P_j \leq p_i) > \alpha$	None
Step-down <i>minP</i>	<i>FWER</i>	$k = \min(i) : P(\min_{i \leq j \leq m} P_j \leq p_i) > \alpha$	None
Benjamini–Hochberg	<i>FDR</i>	$k = \max(i) : p_k \leq \frac{k \cdot \alpha}{m}$	Independence or certain types of positive dependence
Benjamini–Hochberg–Yekutieli	<i>FDR</i>	$k = \max(i) : p_k \leq \frac{k \cdot \alpha}{m \cdot c(m)}$, where $c(m) = \sum_{i=1}^m \frac{1}{i} \approx \ln(m) + 0.58$ if negative dependence and $c(m) = 1$ otherwise	None

The methods are presented in a uniform manner leading to the rejection of H_1, \dots, H_k (keeping H_{k+1}, \dots, H_m). It is assumed that the hypotheses are ordered by their p -values and $p_1 \leq \dots \leq p_k \leq \dots \leq p_m$. P_j is a random variable for the p -value of hypothesis H_j

is compared to the original significance level α . Now *FWER* can be expressed as

$$FWER = P(\cup_{i \in K} \{P_i \leq \hat{\alpha}_i\}) = P(\cup_{i \in K} \{\hat{P}_i \leq \alpha\}), \quad (30)$$

where K is the set of indices of true null hypotheses and P_i and \hat{P}_i denote random variables for the original and adjusted p -values.

The correction procedures are designed such that $FWER \leq \alpha$ holds at least asymptotically, when the underlying assumptions are met. In addition, it is usually required that the adjusted p -values have the same order as the original p -values, i.e., $p_i \leq p_j \Leftrightarrow \hat{p}_i \leq \hat{p}_j$. This ‘monotonicity of p -values’ is by no means necessary for *FWER* control, but it is in line with the statistical intuition according to which a pattern should not be declared significant if a more significant pattern (with smaller p) is declared spurious (Westfall and Young 1993, p. 65). Reporting adjusted p -values (together with original unadjusted p -values) is often recommended, since they are more informative than binary rejection decisions. However, individual p -values cannot be interpreted separately because \hat{p}_i is the smallest α -level of the entire test procedure that rejects H_i given p -values or tests statistics of all hypotheses (Ge et al. 2003).

Famous multiple testing procedures and their assumptions are listed in Table 4. Bonferroni and Šidák corrections as well as the single-step *minP* method are exam-

ples of *single-step methods*, where the same adjusted significance level is applied to all hypotheses. All the other methods in the table are *step-wise methods* that determine individual significance levels for each hypothesis, depending on the order of p -values and rejection of other hypotheses. Step-wise methods can be further divided into *step-down methods* (Holm–Bonferroni method and the step-down $\min P$ method by Westfall and Young (1993)) that process hypotheses in the ascending order by their p -values and *step-up methods* (Hochberg, Benjamini–Hochberg and Benjamini–Hochberg–Yekutieli methods) that proceed in the opposite order. In general, single-step methods are least powerful and step-up methods most powerful, with the exception of the powerful $\min P$ methods.

Powerful methods are always preferable since they can detect most true patterns, but the selection of the method depends on also other factors like availability of all p -values during evaluation, the expected proportion of true patterns (false null hypotheses), seriousness of false discoveries, and assumptions on the dependency structure between hypotheses. In the following we will briefly discuss these issues. For more details, we refer the interested reader to e.g., (Goeman and Solari 2014; Ge et al. 2003).

The least powerful method for controlling $FWER$ is the popular Bonferroni correction. The lack of power is due to two pessimistic assumptions: m_0 is estimated by its upperbound m and the probability of type I error by upperbound $P((P_1 \leq \frac{\alpha}{m}) \vee \dots \vee (P_m \leq \frac{\alpha}{m})) \leq \sum_{i=1}^m P(P_i \leq \frac{\alpha}{m})$ (Boole’s or Bonferroni’s inequality). Therefore, the Bonferroni correction is least powerful when many null hypotheses are false or the hypotheses are positively associated. The Šídák correction (Šídák 1967) is slightly more powerful, because it assumes independence of true null hypotheses and can thus use a lower upperbound for the probability of type I error. However, the method gives exact control of $FWER$ only under the independence assumption. The control is not guaranteed if the true hypotheses are negatively dependent and the method may be overly conservative if they are positively dependent. The Bonferroni and Šídák corrections are attractive for pattern discovery where the size of the space of alternatives can be predetermined, because they impose minimal computational burden requiring simply that the value of α be appropriately decreased.

The Holm–Bonferroni method (Holm 1979) is a sequential variant of the Bonferroni method. It proceeds in a step-wise manner, by comparing the smallest p -value to $\frac{\alpha}{m}$ like the Bonferroni method, but the largest to α . Therefore, it rejects always at least as many null hypotheses as the Bonferroni method and the gain in power is greatest when most null hypotheses are false. The Hochberg method (Hochberg 1988) can be considered as a step-up variant of the Holm–Bonferroni method. It is a more powerful method especially if there are many false null hypotheses or the p -values of false null hypotheses are positively associated. However, it has extra requirements for the dependency structure among hypotheses. Sufficient conditions for the Hochberg method (and the underlying Simes inequality) are independence and certain types of positive dependence (e.g., *positive regression dependence on a subset* (Benjamini and Yekutieli 2001)) between true hypotheses. Šídák’s method can also be implemented in a similar step-wise manner by using criterion $p_i > 1 - (1 - \alpha)^{1/(m-i+1)}$ in the Holm–Bonferroni method. However, the resulting Holm–Šídák method assumes also independence of hypotheses.

The Benjamini–Hochberg method (Benjamini and Hochberg 1995) and the Benjamini–Hochberg–Yekutieli method (Benjamini and Yekutieli 2001) are also step-up methods, but they control FDR instead of $FWER$. The Benjamini–Hochberg method is always at least as powerful as the Hochberg method, and the difference is most pronounced when there are many false null hypotheses. The Benjamini–Hochberg method is also based on the Simes inequality and has the same requirements for the dependency structure between true hypotheses (independence or certain types of positive dependence). The Benjamini–Hochberg–Yekutieli method allows also negative dependencies, but it is less powerful and may sometimes be even more conservative than the Holm–Bonferroni method (Goeman and Solari 2014).

These step-wise approaches are problematic in pattern discovery unless statistical testing is applied as a postprocessing step. This is because they require all null hypotheses to be sorted on p -value which implies that all p -values must be known before the corrections are applied. However, step-wise methods are easily applied in multi-stage procedures (e.g., Webb 2007; Komiyama et al. 2017) that first select constrained sets of candidate patterns which are then subsequently subjected to statistical testing.

The $minP$ methods (Westfall and Young 1993) present a different approach to multiple hypothesis testing. These methods are usually implemented with permutation testing or other resampling methods, and thus they adapt to any dependency structure between null hypotheses. This makes them powerful methods and they have been shown to be asymptotically optimal for a broad class of testing problems (Meinshausen et al. 2011).

The $minP$ methods are based on an alternative expression of $FWER$ (Eq. 30): $FWER = P(\cup_{i \in K} \{P_i \leq \hat{\alpha}\} \mid H_K) = P(\min_{i \in K} \{P_i \leq \hat{\alpha}\} \mid H_K)$, where H_K is an intersection of all true hypotheses and $\hat{\alpha}$ is an adjusted significance level. Therefore, an optimal $\hat{\alpha}$ can be determined as an α -quantile from the distribution of the minimum p -value among the true null hypotheses, i.e.,

$$\hat{\alpha} = \max\{a \mid P(\min P_i \leq a \mid H_K) \leq \alpha\}.$$

In principle, any technique for estimating the α -quantile can be used, but analytical methods are seldom available. However, the evaluation can be done also empirically, with resampling methods.

For strong control of $FWER$ the probability should be evaluated under H_K , which is unknown. Therefore, the estimation is done under the complete null hypotheses H_0^C . Strong control (at least *partial strong control* (Rempala and Yang 2013)) can still be obtained under certain extra conditions. One such condition is *subset pivotality* (Westfall and Young 1993, p. 42) that requires the raw p -values (or other test statistics) of true null hypotheses to have the same joint distribution under H_0^C and any other set of hypotheses. Since the true null hypotheses are unknown, the minimum p -value is determined among all null hypotheses (the single-step method) or among all unrejected null hypotheses (the step-down method). The resulting single-step adjustment is $\hat{p}_i = P(\min_{1 \leq j \leq m} P_j \leq p_i \mid H_0^C)$. A similar adjustment can be done with other test statistics T , if subset pivotality or other required conditions hold. Assuming that high T -values are more significant, the adjusted p -value is $\hat{p}_i = P(\max_{1 \leq j \leq m} T_j \leq t_i \mid H_0^C)$, where T_i is a

random variable for the test statistic of H_i and t_i is its observed value. The probability under H_0^C can be estimated with permutation testing, by permuting the data under H_0^C and calculating the proportion of permuted data sets where $\min P$ or $\max T$ -value is at least as extreme as the observed p_i or t_i . In pattern discovery complete permutation testing is often infeasible, but there are more efficient approaches combining the $\min P$ correction with approximate permutation testing (e.g., Hanhijärvi 2011; Minato et al. 2014; Llinares López et al. 2015). However, the time and space requirements may still be too large for many practical pattern mining purposes.

6.3 Increasing power in pattern discovery

In pattern discovery the main problem of multiple hypothesis testing is the huge number of possible hypotheses. This number is the same as the number of all possible patterns or the size of the search space that is usually exponential. If correction is done with respect to all possible patterns, the adjusted critical values may become so small that few patterns can be declared as significant. This means that one should always use as powerful correction methods as possible or control FDR instead of $FWER$ when applicable, but this may still be insufficient. A complementary strategy is to reduce the number of hypotheses or otherwise target more power to those hypotheses that are likely to be interesting or significant. In the following we describe three general techniques designed for this purpose: hold-out evaluation, filtering hypotheses and weighted hypothesis testing.

The idea of *hold-out evaluation* (Webb 2007) (also known as *two-stage testing*, e.g., (Miller et al. 2001)) is to use only a part of the data for pattern discovery and save the rest for testing significance of patterns. The method consists of three steps:

- (i) Divide the data into an exploratory set \mathcal{D}_E and a hold-out set \mathcal{D}_H .
- (ii) Search for patterns in \mathcal{D}_E and select K patterns for testing. Note that the selection process at this step can use any principle suited to the application and need not involve hypothesis testing.
- (iii) Test the significance of the K patterns in \mathcal{D}_H using any multiple hypothesis testing procedure.

Now the number of hypotheses is only K which is typically much less than the size of the search space. This makes the method powerful, even if the p -values in the hold-out set are likely larger than they would have been in the whole data set. The power can be further enhanced by selecting powerful multiple testing methods, including methods that control FDR . A potential problem of hold-out evaluation is that the results may depend on how the data is partitioned. In a pathological case a pattern may occur only in the exploratory set or only in the hold-out set and thus remain undiscovered (Liu et al. 2011).

Another approach is to use *filtering methods* (see e.g., Bourgon et al. 2010) to select only promising hypotheses for significance testing. Ideally, the filter should prune out only true null hypotheses without compromising control of false discoveries. In practice, the true nulls are unknown and the filter uses some data characteristics to detect low power hypotheses that are unlikely to become rejected. Unfortunately, some

filtering methods affect also the distribution of the test statistic and can violate strong control of $FWER$. As a solution it has been suggested that the filtering statistic and the actual test statistic should be independent given a true null hypothesis (Bourgon et al. 2010).

In pattern discovery one useful filtering method is to prune out so called *untestable hypotheses* (Terada et al. 2013a; Mantel 1980) that cannot cause type I errors. This approach can be used when hypothesis testing is done conditionally on some data characteristics, like marginal frequencies. The idea is to determine a priori, using only the given conditions, if a hypothesis H_i can ever achieve sufficiently small p -value to become rejected at the adjusted level $\hat{\alpha}_i$. If this is not possible (i.e., if the smallest possible p -value, p_i^* , would be too large, $p_i^* > \hat{\alpha}_i$), then the hypothesis is called ‘untestable’. Untestable hypotheses cannot contribute to $FWER$ and they can be safely ignored when determining corrections for multiple hypothesis testing.

For example, the lowest possible p -value with Fisher’s exact test for rule $A \rightarrow B$ with any contingency table having marginal frequencies $fr(A) = 10$, $fr(B) = 4$ and $n = 20$ is $p^* = 0.043$. If we test just one hypothesis with $\alpha = 0.05$, then this pattern can pass the test and the hypothesis is considered testable. However, if we test two hypotheses and use Bonferroni correction, then the corrected α -level is $\hat{\alpha} = 0.025$ and the hypothesis is considered untestable ($p^* = 0.043 > 0.025 = \hat{\alpha}$).

In practice, selecting testable hypotheses can improve the power of the method considerably. In pattern discovery the idea of testability has been utilized successfully in the search algorithms, including the LAMP procedure (Terada et al. 2013a; Minato et al. 2014) that controls $FWER$ with the Bonferroni correction and Westfall–Young light (Llinares López et al. 2015) that implements a *min P* method.

A third approach is to use a *weighted multiple testing procedure* (e.g., Finos and Salmaso 2007; Holm 1979) that gives more power to those hypotheses that are likely to be most interesting. Usually, the weights are given a priori according to assumed importance of hypotheses, but it is also possible to determine optimal weights from the data to maximize power of the test (see e.g., Roeder and Wasserman 2009). The simplest approach is an *allocated Bonferroni procedure* (Rosenthal and Rubin 1983) that allocates total α among all m hypotheses according to their importance. Each hypothesis H_i is determined an individual significance level $\hat{\alpha}_i$ such that $\sum_{i=1}^m \hat{\alpha}_i \leq \alpha$. This is equivalent to a weighted Bonferroni procedure, where one determines weights w_i such that $\sum_{i=1}^m w_i = m$ and reject H_i if $p_i \leq \frac{w_i \alpha}{m}$. There are also weighted variants of other multiple correction procedures like the weighted Holm–Bonferroni procedure (Holm 1979) and the weighted Benjamini–Hochberg procedure (Benjamini and Hochberg 1997). Usually, these procedures do not respect the monotonicity of p -values, which means that the most significant patterns may be missed if they were deemed uninteresting.

In the pattern discovery context one natural principle is to base the weighting on the complexity of patterns and favour simple patterns. This approach is used in the *method of layered critical values* (Webb 2008) where simpler patterns are tested with looser thresholds and strictest thresholds are reserved to most complex patterns. The motivation is that simpler patterns tend to have higher proportions of significant patterns and can be expected to be more interesting. In addition, this weighting strategy supports efficient search, because it helps to prune deeper levels of the search space.

When dependency patterns are searched, the complexity can be characterized by the number of attributes in the pattern which is the same as the level of the search space. Webb (2008) has suggested an allocation strategy where all patterns at level L are tested with threshold $\hat{\alpha}_L$ such that $\sum_{L=1}^{L_{\max}} \hat{\alpha}_L \cdot S_L \leq \alpha$, where L_{\max} is the maximum level and S_L is the number of all possible patterns at level L . One such allocation is to set

$$\hat{\alpha}_L = \frac{\alpha}{L_{\max} S_L}.$$

The method was originally proposed for the breadth-first search of classification rules, but it can be applied to other pattern types and depth-first search as well. The only critical requirement is that the bias towards simple patterns fits the research problem. In a pathological case, the method may miss the most significant patterns if they are too complex. However, the same patterns might remain undetected also with a weaker but more balanced testing procedure.

7 Conclusions

Pattern discovery is a fundamental form of exploratory data analysis. In this tutorial we have covered the key theory and techniques for finding statistically significant dependency patterns that are likely to represent true dependencies in the underlying population. We have concentrated on two general classes of patterns: dependency rules that express statistical dependencies between condition and consequent parts and dependency sets that express mutual dependence between set elements.

Techniques for finding true statistical dependencies are based on statistical tests of different types of independence. The general idea is to evaluate how likely it is that the observed or a stronger dependency pattern would have occurred in the given sample data, if the independence assumption had been true. If this probability is too large, the pattern can be discarded as having a high risk of being spurious. In this tutorial we have presented the core relevant statistical theory and specific statistical tests for different notions of dependence under various assumptions on the underlying sampling model.

However, in many applications it is often desirable to use stronger filters to the discovered patterns than a simple test for independence. Statistically significant dependency rules and sets can be generated by adding unrelated or even negatively associated elements to existing patterns. Unless further tests are also satisfied, such as tests for productivity and significant improvement, the discovered rules and sets are likely to be dominated by many superfluous or redundant patterns. Fortunately, statistical significance testing can also be employed to control the risk of ‘discovering’ these and other forms of superfluous patterns. We have also surveyed the key such techniques.

The final major issue that we have covered is that of multiple testing. Each statistical hypothesis test controls the risk that its null hypothesis would be rejected if that hypothesis were false. However, typical pattern discovery tasks explore exceptionally large numbers of potential hypotheses, and even if the risks for each of the individual hypotheses are extremely small, they can accumulate until the cumulative risk of false

discoveries approaches almost certainty. We also survey multiple testing methods that can control this cumulative risk.

The field of statistically sound pattern discovery is in its infancy and there are numerous open problems. Most work in the field has been restricted to attribute-value or transactional data. Patterns over more complex data types like sequences and graphs would benefit also from statistically sound techniques but may require new statistical tests to be feasible. The possibilities of allowing for untestable hypotheses are also opening many possibilities for substantially increasing the power of multiple testing procedures. The field has been dominated by frequentist approaches of significance testing, but there is much scope for application of Bayesian techniques. But perhaps the two biggest challenges are determining the right statistical tests to identify patterns of interest for specific applications and then developing efficient search algorithms that find the most significant patterns under those tests.

It is important to remember that statistical significance testing controls only the risk of false discoveries—type I error. It does not control the risk of type II error—of failing to discover a pattern. When sample sizes are reasonably large, it is reasonable to expect that statistically sound pattern discovery techniques will find all real strong patterns in the data and will not find spurious weak patterns. However, it is important to recognize that in some circumstances it will be more appropriate to explore alternative techniques that trade off the risks of type I and type II error.

Statistically sound pattern discovery has brought the field of pattern mining to a new level of maturity, providing powerful and robust methods for finding useful sets of key patterns from sample data. We hope that this tutorial will help bring the power of these techniques to a wider group of users.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix: Terminology

The following lists some similar looking terms that have a different meaning in the traditional pattern discovery and statistics.

Association (Statistics)

Generally, statistical dependence between two (or more) random variables. In a narrower sense, it refers to statistical dependence between categorical variables, while the word ‘correlation’ is used for numerical variables.

Association rule (Pattern discovery)

Traditionally, an association rule $X \rightarrow Y$ merely means sufficiently frequent co-occurrence of two attribute sets, X and Y . The sufficient frequency is defined by a user-specified threshold min_{fr} . In the traditional definition (Agrawal et al. 1993), it has also been required that the ‘association’ should be sufficiently strong, measured by precision (‘confidence’). As such, an association rule does not necessarily express

any statistical dependence or the dependence may be negative, instead of the assumed positive dependence. Therefore, it has become more common to require that the rule expresses positive statistical dependence, i.e. $P(\mathbf{XY}) > P(\mathbf{X})P(\mathbf{Y})$, and use statistical dependence measures like lift or leverage instead of precision. In this paper, we call rule-formed statistical dependencies without any necessary minimum frequency requirements as ‘dependency rules’.

Confidence (Pattern discovery)

A traditional measure for the strength of an association rule $\mathbf{X} \rightarrow \mathbf{Y}$ defined as $\phi(\mathbf{X} \rightarrow \mathbf{Y}) = P(\mathbf{Y}|\mathbf{X})$. In pattern recognition, information retrieval, and binary classification, this measure is called ‘precision’ or ‘positive predictive value’.

Confidence interval and confidence level (Statistics)

Confidence interval is an interval estimate of some unknown population parameter. Confidence level (e.g., 95%) determines the proportion of confidence intervals that contain the true value of the parameter. The concepts are closely related to statistical hypothesis testing: a confidence interval with confidence level $1 - \alpha$ contains all values s_0 for which the corresponding null hypothesis $S = s_0$ is not rejected at significance level α .

Support (Pattern discovery)

A term used for frequency in frequent itemset and association rule mining. The support of an attribute set (itemset) \mathbf{X} can mean either absolute frequency, $fr(\mathbf{X})$, or relative frequency, $P(\mathbf{X})$. The support of association rule $\mathbf{X} \rightarrow \mathbf{Y}$ usually means $fr(\mathbf{XY})$ or $P(\mathbf{XY})$ but sometimes it can refer to $fr(\mathbf{X})$ or $P(\mathbf{X})$. The latter are also called ‘coverage’ of the rule, although coverage can sometimes refer to $fr(\mathbf{XY})$ or $P(\mathbf{XY})$.

Support (Mathematics, Statistics)

In general, the support of a function is the set of points where the function is not zero-valued or the closure of that set. In the probability theory and statistics, the support of a distribution whose density function is f , is the smallest closed set S such that $f(x) = 0$ for all $x \notin S$.

References

- Aggarwal C, Han J (2014) Frequent pattern mining. Springer, Cham
- Aggarwal C, Yu P (2001) A new approach to online generation of association rules. *IEEE Trans Knowl Data Eng* 13(4):527–540
- Agrawal R, Imielinski T, Swami A (1993) Mining association rules between sets of items in large databases. In: Buneman P, Jajodia S (eds) *Proceedings of the 1993 ACM SIGMOD international conference on management of data*, ACM Press, New York, pp 207–216
- Agrawal R, Mannila H, Srikant R, Toivonen H, Verkamo AI (1996) Fast discovery of association rules. In: Fayyad UM, Piatetsky-Shapiro G, Smyth P, Uthurusamy R (eds) *Advances in knowledge discovery and data mining*. AAAI Press, Menlo Park, pp 307–328
- Agresti A (1992) A survey of exact inference for contingency tables. *Stat Sci* 7(1):131–153
- Agresti A (2002) *Categorical data analysis*, 2nd edn. Wiley series in probability and statistics. Wiley, Hoboken
- Agresti A, Hitchcock D (2005) Bayesian inference for categorical data analysis. *Stat Methods Appl* 14:297–330
- Agresti A, Min Y (2005) Frequentist performance of Bayesian confidence intervals for comparing proportions in 2×2 contingency tables. *Biometrics* 61:515–523

- Albert J (1997) Bayesian testing and estimation of association in a two-way contingency table. *J Am Stat Assoc* 92:685–693
- Balcazar J (2010) Redundancy, deduction schemes, and minimum-size bases for association rules. *Log Methods Comput Sci* 6(2). <http://arxiv.org/abs/1002.4286>
- Barnard G (1947) Significance tests for 2×2 tables. *Biometrika* 34(1/2):123–138
- Bastide Y, Pasquier N, Taoüil R, Stumme G, Lakhal L (2000) Mining minimal non-redundant association rules using frequent closed itemsets. In: Lloyd J, , Dahl V, Furbach U, Kerber M, Lau KK, Palamidessi C, Pereira L, Sagiv Y, Stuckey P (eds) Proceedings of the first international conference on computational logic (CL'00). Lecturer notes in computer science, vol 1861. Springer, Berlin, pp 972–986
- Bay SD, Pazzani MJ (2001) Detecting group differences: mining contrast sets. *Data Min Knowl Discov* 5(3):213–246
- Bayardo R, Agrawal R, Gunopulos D (2000) Constraint-based rule mining in large, dense databases. *Data Min Knowl Discov* 4(2/3):217–240
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc, Ser B* 57(1):289–300
- Benjamini Y, Hochberg Y (1997) Multiple hypotheses testing with weights. *Scand J Stat* 24(3):407–418
- Benjamini Y, Leshno M (2005) Statistical methods for data mining. In: Maimon O, Rokach L (eds) The data mining and knowledge discovery handbook. Springer, New York, pp 565–87
- Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. *Ann Stat* 29(4):1165–1188
- Birch M (1964) The detection of partial association, I: the 2×2 case. *J R Stat Soc Ser B (Methodol)* 26(2):313–324
- Blanchard J, Guillet F, Gras R, Briand H (2005) Using information-theoretic measures to assess association rule interestingness. In: Han J, Wah B, Raghavan V, Wu X, Rastogi R (eds) Proceedings of the fifth IEEE international conference on data mining (ICDM'05). IEEE Computer Society, Washington, USA, pp 66–73
- Boulicaut JF, Bykowski A, Rigotti C (2000) Approximation of frequency queries by means of free-sets. In: Proceedings of the 4th European conference principles of data mining and knowledge discovery (PKDD'00). Lecture notes in computer science, vol 1910. Springer, Berlin, pp 75–85
- Bourgon R, Gentleman R, Huber W (2010) Independent filtering increases detection power for high-throughput experiments. *Proc Natl Acad Sci* 107(21):9546–9551
- Brin S, Motwani R, Silverstein C (1997) Beyond market baskets: generalizing association rules to correlations. In: Peckham J (ed) Proceedings ACM SIGMOD international conference on management of data. ACM Press, New York, pp 265–276
- Bruzzese D, Davino C (2003) Visual post-analysis of association rules. *J Vis Lang Comput* 14:621–635
- Carriere K (2001) How good is a normal approximation for rates and proportions of low incidence events? *Commun Stat Simul Comput* 30:327–337
- Cheng J, Ke Y, Ng W (2008) Effective elimination of redundant association rules. *Data Min Knowl Discov* 16(2):221–249
- Cobb G, Chen YP (2003) An application of Markov chain Monte Carlo to community ecology. *Am Math Mon* 110:265–288
- Cooley R, Mobasher B, Srivastava J (1997) Web mining: Information and pattern discovery on the world wide web. In: Proceedings of the ninth IEEE international conference on tools with artificial intelligence, IEEE, Los Alamitos, pp 558–567
- Corani G, Benavoli A, Demsar J (2016) Comparing competing algorithms: Bayesian versus frequentist hypothesis testing. In: Tutorial in the European conference on machine learning and principles and practice of knowledge discovery in databases (ECML-PKDD 2016). <http://ipg.idsia.ch/tutorials/2016/bayesian-tests-ml/>
- Dehaspe L, Toivonen H (2001) Discovery of relational association rules. In: Džeroski S, Lavrač N (eds) Relational data mining. Springer, Berlin, pp 189–212
- Dong G, Li J (1999) Efficient mining of emerging patterns: discovering trends and differences. In: Proceedings of the fifth ACM SIGKDD international conference on knowledge discovery and data mining, ACM, New York pp 43–52
- Edgington E (1995) Randomization tests, 3rd edn. Marcel Dekker Inc, New York
- Feller W (1968) An introduction to probability theory and its applications, vol I, 3rd edn. Wiley, New York
- Finos K, Salmaso L (2007) FDR- and FWE-controlling methods using data-driven weights. *J Stat Plan Inference* 137:3859–3870

- Fisher R (1925) Statistical methods for research workers. Oliver and Boyd, Edinburgh
- Freedman D, Pisani R, Purves R (2007) Statistics, 4th edn. Norton & Company, London
- Ge Y, Dudoit S, Speed TP (2003) Resampling-based multiple testing for microarray data analysis. *TEST: Off J Span Soc Stat Oper Res* 12(1):1–44
- Gionis A, Mannila H, Mielikäinen T, Tsaparas P (2007) Assessing data mining results via swap randomization. *ACM Trans Knowl Discov Data* 1(3):14:1–14:32
- Goeman JJ, Solari A (2011) Multiple testing for exploratory research. *Stat Sci* 26(4):584–597
- Goeman JJ, Solari A (2014) Multiple hypothesis testing in genomics—tutorial in biostatistics. *Stat Med* 33(11):1946–1978
- Goethals B, Muhonen J, Toivonen H (2005) Mining non-derivable association rules. In: Kargupta H, Sri-vastava J, Kamath C, Goodman A (eds) Proceedings of the 2005 SIAM international conference on data mining, SIAM, pp 239–249
- Haber M (1980) A comparison of some continuity corrections for the chi-squared test on 2×2 tables. *J Am Stat Assoc* 75(371):510–515
- Hahsler M, Hornik K, Reutterer T (2006) Implications of probabilistic data modeling for mining association rules. In: Spiliopoulou M, Kruse R, Borgelt C, Nürnberger A, Gaul W (eds) From data and information analysis to knowledge engineering. Proceedings of the 29th annual conference of the Gesellschaft für Klassifikation, Studies in classification, data analysis, and knowledge organization, Springer, Berlin, pp 598–605
- Hämäläinen W (2010a) Efficient search for statistically significant dependency rules in binary data. Ph.D. thesis, Department of Computer Science, University of Helsinki, Finland, series of Publications A, Report A-2010-2
- Hämäläinen W (2010b) Statapriori: an efficient algorithm for searching statistically significant association rules. *Knowl Inf Syst: Int J (KAIS)* 23(3):373–399
- Hämäläinen W (2011) Efficient search methods for statistical dependency rules. *Fund Inf* 113(2):117–150 (A Special issue on Statistical and Relational Learning in Bioinformatics)
- Hämäläinen W (2012) Kingfisher: an efficient algorithm for searching for both positive and negative dependency rules with statistical significance measures. *Knowl Inf Syst: Int J (KAIS)* 32(2):383–414
- Hämäläinen W (2016) New upper bounds for tight and fast approximation of Fisher’s exact test in dependency rule mining. *Comput Stat Data Anal* 93:469–482
- Hämäläinen W, Webb G (2013) Statistically sound pattern discovery. In: Tutorial in the European conference on machine learning and principles and practice of knowledge discovery in databases (ECML-PKDD 2013). <http://www.cs.joensuu.fi/~whamalai/ecmlpkdd13/sspdtutorial.html>
- Hämäläinen W, Webb G (2014) Statistically sound pattern discovery. In: Proceedings of the 20th ACM SIGKDD international conference on knowledge discovery and data mining (KDD14), ACM, New York, p 1976
- Hämäläinen W, Webb GI (2017) Specious rules: an efficient and effective unifying method for removing misleading and uninformative patterns in association rule mining. In: Chawla N, Wang W (eds) Proceedings of the 2017 SIAM international conference on data mining, SIAM, pp 309–317
- Hanhijärvi S (2011) Multiple hypothesis testing in pattern discovery. In: Elomaa T, Hollmén J, Mannila H (eds) Proceedings on the 14th international conference on discovery science. Lecture notes in artificial intelligence, vol 6926. Springer, Berlin, pp 122–134
- Hanhijärvi S, Ojala M, Vuokko N, Puolamäki K, Tatti N, Mannila H (2009) Tell me something I don’t know: randomization strategies for iterative data mining. In: Proceedings of the 15th ACM SIGKDD international conference on knowledge discovery and data mining, ACM, New York, NY, pp 379–388
- Herrera F, Carmona CJ, González P, Del Jesus MJ (2011) An overview on subgroup discovery: foundations and applications. *Knowl Inf Syst* 29(3):495–525
- Hochberg Y (1988) A sharper Bonferroni procedure for multiple tests of significance. *Biometrika* 75:800–802
- Holm S (1979) A simple sequentially rejective multiple test procedure. *Scand J Stat* 6:65–70
- Howard JV (1998) The 2×2 table: a discussion from a Bayesian viewpoint. *Stat Sci* 13(4):351–367
- Hu S, Rao J (2007) Statistical redundancy testing for improved gene selection in cancer classification using microarray data. *Cancer Inf* 3:29–41
- Hubbard R, Bayarri M (2003) Confusion over measures of evidence (p ’s) versus errors (α ’s) in classical statistical testing. *Am Stat* 57(3):171–178


- Jabbar M, Shazan M, Zaïane O (2016) Learning statistically significant contrast sets. In: Khoury R, Drummond C (eds) *Advances in artificial intelligence: Proceedings of the 29th Canadian conference on artificial intelligence*. Lecture notes in artificial intelligence, vol 9673. Springer, Cham, 237–242
- Jamil T, Ly A, Morey R, Love J, Marsman M, Wagenmakers EJ (2017) Default Gunel and Dickey Bayes factors for contingency tables. *Behav Res Methods* 49:638–652
- Jin Z, Li J, Liu L, Le TD, Sun B, Wang R (2012) Discovery of causal rules using partial association. In: Zaki M, Siebes A, Yu J, Goethals B, Webb G, Wu X (eds) *Proceedings of the 12th IEEE international conference on data mining (ICDM 2012)*, IEEE Computer Society, Los Alamitos, pp 309–318
- Johnson D (1999) The insignificance of statistical significance testing. *J Wildl Manag* 63:763–772
- Kim E, Helal S, Cook D (2010) Human activity recognition and pattern discovery. *IEEE Pervasive Comput* 9(1):48–53
- Komiyama J, Ishihata M, Arimura H, Nishibayashi T, Minato S (2017) Statistical emerging pattern mining with multiple testing correction. In: *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining (KDD'17)*, ACM, New York, pp 897–906
- Lallich S, Vaillant B, Lenca P (2005) Parametrised measures for the evaluation of association rule interestingness. In: Janssen J, Lenca P (eds) *Proceedings of the 11th symposium on applied stochastic models and data analysis (ASMDA'05)*, ASMDA International Society, pp 220–229
- Lallich S, Teytaud O, Prudhomme E (2007) Association rule interestingness: measure and statistical validation. In: Guillet F, Hamilton H (eds) *Quality measures in data mining, studies in computational intelligence*, vol 43. Springer, Berlin, pp 251–275
- Lecoutre B, Lecoutre MP, Poitevineau J (2001) Uses, abuses and misuses of significance tests in the scientific community: won't the Bayesian choice be unavoidable? *Int Stat Rev/Revue Internationale de Statistique* 69(3):399–417
- Lee P (2012) *Bayesian statistics: an introduction*, 4th edn. Wiley, Chichester
- Legendre P, Legendre L (1998) *Numerical ecology*. Elsevier Science, Amsterdam
- Lehmann E (1993) The Fisher, Neyman–Pearson theories of testing hypotheses: one theory or two? *J Am Stat Assoc* 88:1242–1249
- Lehmann E, Romano J (2005) *Testing statistical hypotheses*, 3rd edn. Texts in statistics. Springer, New York
- Li G, Hamilton H (2004) Basic association rules. In: Berry M, Dayal U, Kamath C, Skillicorn D (eds) *Proceedings of the fourth SIAM international conference on data mining*, SIAM, Philadelphia, pp 166–177
- Li J (2006) On optimal rule discovery. *IEEE Trans Knowl Data Eng* 18(4):460–471
- Li J, Le T, Liu L, Liu J, Jin Z, Sun B, Ma S (2016) From observational studies to causal rule mining. *ACM Trans Intell Syst Technol* 7(2):14:1–14:27
- Li W, Han J, Pei J (2001) CMAR: Accurate and efficient classification based on multiple class-association rules. In: Cercone N, Lin T, Wu X (eds) *Proceedings of the 2001 IEEE international conference on data mining*, IEEE, Los Alamitos, pp 369–376
- Lichman M (2013) UCI machine learning repository. <http://archive.ics.uci.edu/ml>
- Lindgren B (1993) *Statistical theory*, 4th edn. Chapman & Hall, Boca Raton
- Liu B, Hsu W, Ma Y (1999) Pruning and summarizing the discovered associations. In: *Proceedings of the fifth ACM SIGKDD international conference on knowledge discovery and data mining (KDD'99)*, ACM Press, New York, pp 125–134
- Liu G, Zhang H, Wong L (2011) Controlling false positives in association rule mining. *Proc VLDB Endow* 5(2):145–156
- Llinares López F, Sugiyama M, Papaxanthos L, Borgwardt K (2015) Fast and memory-efficient significant pattern mining via permutation testing. In: *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, ACM, New York, pp 725–734
- Mantel N (1980) Assessing laboratory evidence for neoplastic activity. *Biometrics* 36:381–399
- Megiddo N, Srikant R (1998) Discovering predictive association rules. In: Agrawal R, Stolorz P (eds) *Proceedings of the 4th international conference on knowledge discovery in databases and data mining*, AAAI Press, Cambridge, pp 274–278
- Meinshausen N, Maathuis MH, Bhlmann P (2011) Asymptotic optimality of the Westfall–Young permutation procedure for multiple testing under dependence. *Ann Stat* 39(6):3369–3391
- Meo R (2000) Theory of dependence values. *ACM Trans Database Syst* 25(3):380–406
- Miller R, Galecki A, Shmookler-Reis R (2001) Interpretation, design, and analysis of gene array expression experiments. *J Gerontol Ser A, Biol Sci Med Sci* 56:B52–B57

- Minato S, Uno T, Tsuda K, Terada A, Sese J (2014) A fast method of statistical assessment for combinatorial hypotheses based on frequent itemset enumeration. In: Calders T, Esposito F, Hüllermeier E, Meo R (eds) Proceedings of the European conference on machine learning and knowledge discovery in databases, Part II. Lecture notes in artificial intelligence, vol 8725. Springer, Berlin, pp 422–436
- Morishita S, Nakaya A (2000) Parallel branch-and-bound graph search for correlated association rules. In: Zaki M, Ho CT (eds) Large-scale parallel data mining, revised papers from the workshop on large-scale parallel KDD systems, in the 5th ACM SIGKDD international conference on knowledge discovery and data mining (KDD'00). Lecture notes in computer science, vol 1759. Springer, London, pp 127–144
- Morishita S, Sese J (2000) Transversing itemset lattices with statistical metric pruning. In: Proceedings of the nineteenth ACM SIGMOD-SIGACT-SIGART symposium on principles of database systems (PODS'00), ACM Press, New York, pp 226–236
- Neyman J, Pearson E (1928) On the use and interpretation of certain test criteria for purposes of statistical inference: Part II. *Biometrika* 20A(3/4):263–294
- Nijssen S, Kok J (2006) Multi-class correlated pattern mining. In: Bonchi F, Boulicaut JF (eds) Proceedings of the 4th international workshop on knowledge discovery in inductive databases. Lecture notes in computer science, vol 3933. Springer, Berlin, pp 165–187
- Nijssen S, Guns T, Raedt LD (2009) Correlated itemset mining in ROC space: a constraint programming approach. Proceedings the 15th ACM SIGKDD conference on knowledge discovery and data mining (KDD'09). ACM Press, New York, pp 647–656
- Novak P, Lavrac N, Webb G (2009) Supervised descriptive rule discovery: a unifying survey of contrast set, emerging pattern and subgroup mining. *J Machine Learn Res* 10:377–403
- Pasquier N, Bastide Y, Taouil R, Lakhal L (1999) Discovering frequent closed itemsets for association rules. In: Proceedings of the 7th international conference on database theory (ICDT'99). Lecture notes in computer science, vol 1540. Springer, Berlin, pp 398–416
- Pearson E (1947) The choice of statistical tests illustrated on the interpretation of data classed in a 2×2 table. *Biometrika* 34(1/2):139–167
- Piatetsky-Shapiro G (1991) Discovery, analysis, and presentation of strong rules. In: Frawley W (ed) Knowledge discovery in databases. MIT Press, Cambridge, pp 229–248
- Rempala GA, Yang Y (2013) On permutation procedures for strong control in multiple testing with gene expression data. *Stat Its Interface* 6(1):79–89
- Rigoutsos I, Floratos A (1998) Combinatorial pattern discovery in biological sequences: the TEIRESIAS algorithm. *Bioinformatics* 14(1):55–67
- Roeder K, Wasserman L (2009) Genome-wide significance levels and weighted hypothesis testing. *Stat Sci* 24(4):398–413
- Rosenthal R, Rubin D (1983) Ensemble-adjusted p values. *Psychol Bull* 94(3):540–541
- Šídlák ZK (1967) Rectangular confidence regions for the means of multivariate normal distributions. *J Am Stat Assoc* 62:626–633
- Silverstein C, Brin S, Motwani R (1998) Beyond market baskets: generalizing association rules to dependence rules. *Data Min Knowl Discov* 2(1):39–68
- Smyth P, Goodman R (1992) An information theoretic approach to rule induction from databases. *IEEE Trans Knowl Data Eng* 4(4):301–316
- Terada A, Okada-Hatakeyama M, Tsuda K, Sese J (2013a) Statistical significance of combinatorial regulations. *Proc Natl Acad Sci* 110(32):12,996–13,001
- Terada A, Tsuda K, Sese J (2013b) Fast Westfall–Young permutation procedure for combinatorial regulation discovery. In: Li GZ, Kim S, Hughes M, McLachlan G, Sun H, Hu X, Ransom H, Liu B, Lieberman M (eds) Proceedings of the 2013 IEEE international conference on bioinformatics and biomedicine, IEEE computer society, pp 153–158
- Terada A, Kim H, Sese J (2015) High-speed Westfall–Young permutation procedure for genome-wide association studies. In: Proceedings of the 6th ACM conference on bioinformatics, computational biology and health informatics (BCB'15), ACM, New York, pp 17–26
- Upton G (1982) A comparison of alternative tests for the 2×2 comparative trial. *J R Stat Soc Ser A (General)* 145(1):86–105
- Vilalta R, Oblinger D (2000) A quantification of distance bias between evaluation metrics in classification. In: Langley P (ed) Proceedings of the seventeenth international conference on machine learning (ICML'00), Morgan Kaufmann Publishers Inc., San Francisco, pp 1087–1094

- Vreeken J, Tatti N (2014) Interesting patterns. In: Aggarwal C, Han J (eds) Frequent pattern mining. Springer International Publishing, Cham, Switzerland, pp 105–134
- Webb G (2006) Discovering significant rules. In: Proceedings of the 12th ACM SIGKDD international conference on knowledge discovery and data mining (KDD'06), ACM Press, New York, pp 434–443
- Webb G (2007) Discovering significant patterns. *Mach Learn* 68(1):1–33
- Webb G (2008) Layered critical values: a powerful direct-adjustment approach to discovering significant patterns. *Mach Learn* 71(2–3):307–323
- Webb G (2010) Self-sufficient itemsets: an approach to screening potentially interesting associations between items. *Trans Knowl Discov Data* 4:3:1–3:20
- Webb G (2011) Filtered-top-k association discovery. *WIREs Data Min Knowl Discov* 1(3):183–192
- Webb G, Vreeken J (2014) Efficient discovery of the most interesting associations. *Trans Knowl Discov Data* 8(3):15:1–15:31
- Webb G, Zhang S (2005) K-optimal rule discovery. *Data Min Knowl Discov* 10(1):39–79
- Westfall PH, Young SS (1993) Resampling-based multiple testing: examples and methods for p-value adjustment. Wiley, New York
- Wilks S (1935) The likelihood test of independence in contingency tables. *Ann Math Stat* 6(4):190–196
- Wu X, Zhang C, Zhang S (2004) Efficient mining of both positive and negative association rules. *ACM Trans Inf Syst* 22(3):381–405
- Yao Y, Zhong N (1999) An analysis of quantitative measures associated with rules. In: Zhong N, Zhou L (eds) Proceedings of the third Pacific-Asia conference on methodologies for knowledge discovery and data mining (PAKDD'99). Lecture notes in computer science, vol 1574. Springer, London, pp 479–488
- Yates F (1984) Test of significance for 2×2 contingency tables. *J R Stat Soc Ser A (General)* 147(3):426–463
- Yule G (1903) Notes on the theory of association of attributes in statistics. *Biometrika* 2:121–134
- Yule G (1912) On the methods of measuring the association between two attributes. *J R Stat Soc* 75:579–652

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Wilhelmiina Hämäläinen^{1,2}  · Geoffrey I. Webb³

Geoffrey I. Webb
geoff.webb@monash.edu

- ¹ Department of Computer Science, Aalto University, P.O.Box 15400, 00076 Aalto, Finland
- ² Present Address: Institute of Biomedicine, University of Eastern Finland, P.O. Box 1627, 70211 Kuopio, Finland
- ³ Faculty of Information Technology, Monash University, P.O. Box 63, Clayton, VIC 3800, Australia