

A General Measure of Rule Interestingness

Szymon Jaroszewicz and Dan A. Simovici

University of Massachusetts at Boston
Department of Mathematics and Computer Science
Boston, Massachusetts 02446, USA
{sj,dsim}@cs.umb.edu

Abstract. The paper presents a new general measure of rule interestingness. Many known measures such as chi-square, gini gain or entropy gain can be obtained from this measure by setting some numerical parameters, including the amount of trust we have in the estimation of the probability distribution of the data. Moreover, we show that there is a continuum of measures having chi-square, Gini gain and entropy gain as boundary cases. Therefore our measure generalizes both conditional and unconditional classical measures of interestingness. Properties and experimental evaluation of the new measure are also presented.

Keywords: interestingness measure, distribution, Kullback-Leibler divergence, Csiszar divergence, rule.

1 Introduction

Determining the interestingness of rules is an important data mining problem since many data mining algorithms produce enormous amounts of rules, making it difficult for the user to analyze them manually. Thus, it is important to establish some numerical interestingness measure for rules, which can help users to sort the discovered rules. A survey of such measures can be found in [1]. Here we concentrate on measures that assess how much knowledge we gain about the joint distribution of a set of attributes Q by knowing the joint distribution of some set of attributes P . Examples of such measures are *entropy gain*, *mutual information*, *Gini gain*, χ^2 [9,10,4,1,12,11]. The rules considered here are different from classical association rules studied in data mining, since we consider full joint distributions of both antecedent and consequent, while association rules consider only the probability of all attributes having some specified value. This approach has the advantage of applicability to multivalued attributes.

We show that all the above mentioned measures are special cases of a more general parametric measure of interestingness, and by varying two parameters, a family of measures can be obtained containing several well-known classical measures as special cases.

There is work done in machine learning and information theory literature [5, 6,15,3] on generalizing information-theoretical measures. However, all previous work is concerned with either unconditional or conditional measures, while this

paper presents a generalization which includes family of intermediate measures, between conditional and unconditional ones, and shows a relation between these measures and the amount of trust we have in the estimate of probabilities from data. For example, we present a continuum of measures between χ^2 (unconditional measure), and the Gini gain (conditional measure). We show that the intermediate measures have many interesting properties which make them useful for rule evaluation.

Next, we give some essential definitions.

Definition 1. A probability distribution is a matrix of the form

$$\Delta = \begin{pmatrix} x_1 & \cdots & x_m \\ p_1 & \cdots & p_m \end{pmatrix},$$

where $p_i \geq 0$ for $1 \leq i \leq m$ and $\sum_{i=1}^m p_i = 1$.

Δ is an uniform distribution if $p_1 = \cdots = p_m = \frac{1}{m}$. An m -valued uniform distribution will be denoted by \mathcal{U}_m .

Let $\tau = (T, H, \rho)$ be a database table, where T is the name of the table, H is its heading, and ρ is its content. If $A \in H$ is an attribute of τ , the domain of A in τ is denoted by $\text{dom}(A)$. The projection of a tuple $t \in \rho$ on a set of attributes $L \subseteq H$ is denoted by $t[L]$. For more on relational notation and terminology see [14].

Definition 2. The distribution of a set of attributes $L = \{A_1, \dots, A_n\}$ is the matrix

$$\Delta_{L,\tau} = \begin{pmatrix} \ell_1 & \cdots & \ell_r \\ p_1 & \cdots & p_r \end{pmatrix}, \tag{1}$$

where $r = \prod_{j=1}^n |\text{dom}(A_j)|$, $\ell_i \in \text{dom}(A_1) \times \cdots \times \text{dom}(A_n)$, and $p_i = \frac{|t \in \rho | t[L] = \ell_i|}{|\rho|}$ for $1 \leq i \leq r$.

The subscript τ will be omitted when the table τ is clear from context.

The *Havrda-Charvát α -entropy* of the attribute set L (see [7]) is defined as:

$$\mathcal{H}_\alpha(L) = \frac{1}{1 - \alpha} \left(\sum_{j=1}^r p_j^\alpha - 1 \right).$$

The limit case, when α tends towards 1 yields the Shannon entropy $\mathcal{H}(L) = -\sum_{j=1}^r p_j \log p_j$. Another important case, the Gini index, is obtained when $\alpha = 2$ (see [1]) and is given by $\text{gini}(L) = 1 - \sum_{j=1}^r p_j^2$.

If L, K are two sets of attributes of a table τ that have the distributions

$$\Delta_L = \begin{pmatrix} l_1 & \cdots & l_m \\ p_1 & \cdots & p_m \end{pmatrix}, \text{ and } \Delta_K = \begin{pmatrix} k_1 & \cdots & k_n \\ q_1 & \cdots & q_n \end{pmatrix},$$

then the conditional Shannon entropy of L conditioned upon K is given by

$$\mathcal{H}(L|K) = - \sum_{i=1}^m \sum_{j=1}^n p_{ij} \log \frac{p_{ij}}{q_j},$$

where $p_{ij} = \frac{|\{t \in \rho | t[L]=\ell_i \text{ and } t[K]=k_j\}|}{|\rho|}$ for $1 \leq i \leq m$ and $1 \leq j \leq n$. Similarly, the Gini conditional index of these distributions is:

$$\text{gini}(L|K) = 1 - \sum_{i=1}^m \sum_{j=1}^n \frac{p_{ij}^2}{q_j}.$$

These definitions allow us to introduce the Shannon gain (called entropy gain in literature [9]) and the Gini gain defined as:

$$\begin{aligned} \text{gain}_{\text{gini}}(L, K) &= \text{gini}(L) - \text{gini}(L|K), \\ \text{gain}_{\text{shannon}}(L, K) &= \mathcal{H}(L) - \mathcal{H}(L|K) \\ &= \mathcal{H}(L) + \mathcal{H}(K) - \mathcal{H}(LK), \end{aligned} \tag{2}$$

respectively, where LK is an abbreviation for $L \cup K$. Note that the Shannon gain is identical to the *mutual information* between attribute sets P and Q [9]. For the Gini gain we can write:

$$\text{gain}_{\text{gini}}(L, K) = \sum_{i=1}^m \sum_{j=1}^n \frac{p_{ij}^2}{q_j} - \sum_{i=1}^m p_i^2 \tag{3}$$

The product of the distributions Δ_P, Δ_Q , where

$$\Delta_P = \begin{pmatrix} x_1 & \cdots & x_m \\ p_1 & \cdots & p_m \end{pmatrix}, \text{ and } \Delta_Q = \begin{pmatrix} y_1 & \cdots & y_n \\ q_1 & \cdots & q_n \end{pmatrix},$$

is the distribution

$$\Delta_P \times \Delta_Q = \begin{pmatrix} (x_1, y_1) & \cdots & (x_m, y_n) \\ p_1 q_1 & \cdots & p_m q_n \end{pmatrix}.$$

The attribute sets P, Q are *independent* if $\Delta_{PQ} = \Delta_P \times \Delta_Q$.

Definition 3. A rule is a pair of attribute sets (P, Q) . If $P, Q \subseteq H$, where $\tau = (T, H, \rho)$ is a table, then we refer to (P, Q) as a rule of τ .

If (P, Q) is a rule, then we refer to P as the antecedent and to Q as the consequent of the rule. A rule (P, Q) will be denoted, following the prevalent convention in the literature, by $P \rightarrow Q$.

This broader definition of rules originates in [4], where rules were replaced by dependencies in order to capture statistical dependence in both the presence and absence of items in itemsets. The significance of this dependence was measured by the χ^2 test, and our approach is a further extension of that point of view.

The notion of *distribution divergence* is central to the rest of the paper.

Definition 4. Let \mathcal{D} be the class of distributions. A distribution divergence is a function $D : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ such that:

1. $D(\Delta, \Delta') \geq 0$ and $D(\Delta, \Delta') = 0$ if and only if $\Delta = \Delta'$ for every $\Delta, \Delta' \in \mathcal{D}$.
2. When Δ' is fixed, $D(\Delta, \Delta')$ is a convex function of Δ ; in other words, if $\Delta = a_1\Delta_1 + \dots + a_k\Delta_k$, where $a_1 + \dots + a_k = 1$, then

$$D(\Delta, \Delta') \geq \sum_{i=1}^k a_i D(\Delta_i, \Delta').$$

An important class of distribution divergences was obtained by Csiszar in [5] as: $D_\phi(\Delta, \Delta') = \sum_{i=1}^n q_i \phi\left(\frac{p_i}{q_i}\right)$, where

$$\Delta = \begin{pmatrix} k_1 & \dots & k_n \\ p_1 & \dots & p_n \end{pmatrix}, \text{ and } \Delta' = \begin{pmatrix} l_1 & \dots & l_n \\ q_1 & \dots & q_n \end{pmatrix},$$

are two distributions and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is a twice differentiable convex function such that $\phi(1) = 0$. We will also make an additional assumption that $0 \cdot \phi\left(\frac{0}{0}\right) = 0$ to handle the case when for some i both p_i and q_i are zero. If for some i , $p_i > 0$, and $q_i = 0$ the value of $D_\phi(\Delta, \Delta')$ is undefined.

The Csiszar divergence satisfies properties (1) and (2) given above (see [7]).

The following result shows the invariance of Csiszar divergence with respect to distribution product:

Theorem 1. For any distributions Γ, Δ, Δ' and any Csiszar divergence measure D_ϕ we have $D_\phi(\Gamma \times \Delta, \Gamma \times \Delta') = D_\phi(\Delta, \Delta')$.

Depending on the choice of the function ϕ we obtain the divergences shown in the table below:

$\phi(x)$	$D(\Delta, \Delta')$	Divergence
$x \log x$	$p_i \log \frac{p_i}{q_i}$	Kullback-Leibler
$x^2 - x$	$\sum_{i=1}^n \frac{p_i^2}{q_i} - 1$	D_{χ^2}

Both the Kullback-Leibler divergence (also known as *crossentropy*), which we will denote by D_{KL} and the χ^2 -divergence denoted by D_{χ^2} are special cases of the Havrda-Charvát divergence $D_{\mathcal{H}_\alpha}$ generated by $\phi(x) = \frac{x^\alpha - x}{\alpha - 1}$ [7]; specifically, D_{χ^2} is obtained by taking $\alpha = 2$, while D_{KL} is obtained as a limit case, when α tends towards 1.

It is easy to verify that

$$D_{\chi^2}(\Delta, \Delta') = \sum_{i=1}^n \frac{(p_i - q_i)^2}{q_i}.$$

Note that $|\rho|D_{\chi^2}$ equals the χ^2 dependency measure, well known from statistics [1].

2 Interestingness of Rules

The main goal of this paper is to present a unified approach to the notion of interestingness of rules. Let $r = P \rightarrow Q$ be a rule in a table $\tau = (T, H, \rho)$. To construct an interestingness measure we will use a Bayesian approach, in that we will consider an a posteriori distribution Θ of the consequent set of attributes Q .

The definition of an interestingness measure of r will be guided by two main considerations:

- The more the observed joint distribution of PQ diverges from the product distribution of P and the a posteriori distribution Θ of Q the more interesting the rule is. Note that $\Delta_{PQ} = \Delta_P \times \Theta$ corresponds to the situation when P and Q are independent and the observed distribution of Q is identical to the a posteriori distribution.
- The rule is not interesting if P, Q are independent. Therefore, we need to consider a correcting term in the definition of an interestingness measure that will decrease its value when Δ_Q is different from the a posteriori distribution.

The choice of the distribution Θ of the consequent Q of rules of the form $P \rightarrow Q$ can be made starting either from the content of the table, that is, adopting Δ_Q for Θ , or from some exterior information. For example, if Q is the sex attribute for a table that contains data concerning some experiment subjects, we can adopt as the a posteriori distribution either

$$\Delta_{sex} = \begin{pmatrix} \text{'F'} & \text{'M'} \\ 0.45 & 0.55 \end{pmatrix},$$

assuming that 45% of the individuals involved are female, or the distribution

$$\Delta_{gen-pop} = \begin{pmatrix} \text{'F'} & \text{'M'} \\ 0.51 & 0.49 \end{pmatrix},$$

consistent with the general distribution of the sexes in the general population.

Moreover, we can use the Laplace estimator [10,16] (also known in literature as the m -estimate of probability) to obtain the a posteriori distribution

$$\Theta_m = \frac{|\rho|\Delta_Q + m\Theta_0}{|\rho| + m},$$

where Δ_Q is the distribution of Q that is extracted from a table τ , Θ_0 is the apriori distribution, $|\rho|$ is the size of the database, and the integer m represents the amount of trust we have in the prior distribution Θ_0 . If $m = 0$, this means we completely ignore the a priori distribution, and $m \rightarrow \infty$ means that we have no trust in the data, and totally rely on the prior distribution. To avoid using limits, we denote $a = \frac{|\rho|}{|\rho|+m}$, and write the Laplacian as a convex combination of the two distributions:

$$\Theta_a = a\Delta_Q + (1 - a)\Theta_0.$$

Now $a = 1$, and $a = 0$ correspond to cases $m = 0$, and $m \rightarrow \infty$ respectively.

Definition 5. Let $r : P \rightarrow Q$ be a rule, D be some measure of divergence between distributions, and let Θ be a distribution.

The measure of interestingness generated by D and Θ is defined by

$$\Upsilon_{D,\Theta}(r) = D(\Delta_{PQ}, \Delta_P \times \Theta) - D(\Delta_Q, \Theta).$$

In the above definition Θ represents the a posteriori distribution of Q , while Δ_Q is the distribution of Q observed from the data. The term $D(\Delta_Q, \Theta)$ measures the degree to which Δ_Q diverges from the prior distribution Θ , and $D(\Delta_{PQ}, \Delta_P \times \Theta)$ measures how far Δ_{PQ} diverges from the joint distribution of P and Q in case they were independent, and Q was distributed according to Θ .

The justification for the correcting term $D(\Delta_Q, \Theta)$ is given in the following theorem:

Theorem 2. If P and Q are independent, and D is a Csiszar measure of divergence then $\Upsilon_{D,\Theta}(P \rightarrow Q) = 0$.

Observe that if D is a Csiszar divergence $D = D_\phi$, then the invariance of these divergences implies:

$$\Upsilon_{D_\phi,\Theta}(P \rightarrow Q) = D_\phi(\Delta_{PQ}, \Delta_P \times \Theta) - D_\phi(\Delta_P \times \Delta_Q, \Delta_P \times \Theta).$$

3 Properties of the General Measure of Interestingness

Initially, we discuss several basic properties of the proposed measure.

Theorem 3. If D is a Csiszar divergence, then

$$\Upsilon_{D,\Delta_Q}(P \rightarrow Q) = \Upsilon_{D,\Delta_P}(Q \rightarrow P)$$

The above property means that when the a posteriori distribution of the consequent is always assumed equal to the distribution observed from data, then the measure is symmetric with respect to the direction of the rule, i.e. exchanging the antecedent and consequent does not change the value of the interestingness.

Theorem 4. Let D be a Csiszar divergence. If R is a set of attributes independent of P , and jointly of PQ , then, for any Θ we have $\Upsilon_{D,\Theta}(RP \rightarrow Q) = \Upsilon_{D,\Theta}(P \rightarrow Q)$.

If R is a set of attributes independent of Q , and jointly of PQ , then $\Upsilon_{D,\Delta_{RQ}}(P \rightarrow RQ) = \Upsilon_{D,\Delta_Q}(P \rightarrow Q)$.

The previous result gives a desirable property of $\Upsilon_{D,\Theta}$ since adding independent attributes should not affect rule's interestingness. In particular, when Θ equals the observed distribution of the consequent, then Υ is not affected by adding independent attributes to either the antecedent or the consequent.

Next, we consider several important special cases of the interestingness measure.

If the divergence D and the a priori distribution used in the definition of the interestingness measure are chosen appropriately, then the interestingness

$\Upsilon_{D,\Theta}(P \rightarrow Q)$ is proportional to a gain of the set of attributes of the consequent Q of the rule relative to the antecedent P . Both the Gini gain, $\text{gain}_{\text{gini}}(Q, P)$, and the entropy gain, $\text{gain}_{\text{shannon}}(Q, P)$, can be obtained by appropriate choice of D . Moreover a measure proportional to the χ^2 statistic can be obtained in that way.

Suppose that the attribute sets P, Q have the distributions

$$\Delta_P = \begin{pmatrix} x_1 & \cdots & x_m \\ p_1 & \cdots & p_m \end{pmatrix}, \text{ and } \Delta_Q = \begin{pmatrix} y_1 & \cdots & y_n \\ q_1 & \cdots & q_n \end{pmatrix}.$$

Let $\rho_{ij} = \{t \in \rho \mid t[P] = x_i \text{ and } t[Q] = y_j\}$ and let $p_{ij} = \frac{|\rho_{ij}|}{|\rho|}$ for $1 \leq i \leq m$ and $1 \leq j \leq n$.

Theorem 5. *Let $P \rightarrow Q$ be a rule in the table $\tau = (T, H, \rho)$. If $D = D_{\text{KL}}$ then*

$$\Upsilon_{D,\Theta}(P \rightarrow Q) = \text{gain}_{\text{shannon}}(Q, P),$$

regardless of the choice of Θ .

The above theorem means that for the case D_{KL} the family of measures generated by Θ reduces to a single measure: the Shannon gain (mutual information). This is not the case for other divergences.

Theorem 6. *Let $P \rightarrow Q$ be a rule in the table $\tau = (T, H, \rho)$. If $D = D_{\chi^2}$ and $\Theta = \mathcal{U}_n$, where $n = |\text{dom}(Q)|$, then*

$$\Upsilon_{D,\Theta}(P \rightarrow Q) = n \cdot \text{gain}_{\text{gini}}(Q, P).$$

Theorem 7. *We have $\Upsilon_{D_{\chi^2}, \Delta_Q}(P \rightarrow Q)$ is proportional to $\chi^2(P, Q)$, the chi-squared statistics [1] for attribute sets P, Q .*

Note that above we treat attribute sets $P = \{A_1, \dots, A_r\}$ and $Q = \{B_1, \dots, B_s\}$ as single attributes with the domains given by (1). This is appropriate, since we are interested in how one set of attributes P influences another set of attributes Q . Another way, used in [4], is to compute $\chi^2(A_1, \dots, A_r, B_1, \dots, B_s)$, however this is not what we want.

The case when $D = D_{\chi^2}$ is of practical interest since it includes two widely used measures (χ^2 , and $\text{gain}_{\text{gini}}$) as special cases, and allows for obtaining a continuum of measures “in between” the two.

Theorem 8 stated below shows that the generalized measure interestingness $\Upsilon_{D,\Theta}(P \rightarrow Q)$ is minimal when P and Q are independent and thus, it justifies our definition of this measure through variational considerations.

Theorem 8. *Let $\Upsilon_{D,\Theta}$ be the measure of interestingness generated by the a posteriori distribution Θ and the Kullback-Leibler divergence, or the χ^2 -divergence and let $P \rightarrow Q$ be a rule. For any fixed attribute distribution Δ_P, Δ_Q and a fixed distribution Θ , the value of $\Upsilon_{D,\Theta}(P \rightarrow Q)$ is minimal (and equal to 0) if only if $\Delta_{PQ} = \Delta_P \times \Delta_Q$, i.e., when P and Q are independent.*

We saw that $\text{gain}_{\text{shannon}}$ and $\text{gain}_{\text{gini}}$ are equivalent to $\mathcal{Y}_{D_{\text{KL}}, \mathcal{U}_n}$ and $\mathcal{Y}_{D_{\chi^2}, \mathcal{U}_n}$, respectively. It is thus natural to define a notion of *gain* for any divergence D as

$$\text{gain}_D(P \rightarrow Q) = \mathcal{Y}_{D, \mathcal{U}_n}(P \rightarrow Q).$$

Let $\Delta_Q|p_i$ denote the probability distribution of Q conditioned on $P = p_i$. For any Csiszar measure D_ϕ we have:

$$\begin{aligned} \text{gain}_{D_\phi}(P \rightarrow Q) &= D_\phi(\Delta_{PQ}, \Delta_P \times \mathcal{U}_n) - D_\phi(\Delta_Q, \mathcal{U}_n) \\ &= \sum_{i=1}^m p_i \sum_{j=1}^n \frac{1}{n} \phi\left(\frac{p_{ij}}{p_i \cdot \frac{1}{n}}\right) - D_\phi(\Delta_Q, \mathcal{U}_n) \\ &= - \left[D_\phi(\Delta_Q, \mathcal{U}_n) - \sum_{i=1}^m p_i D_\phi(\Delta_Q|p_i, \mathcal{U}_n) \right]. \end{aligned}$$

As special cases $\text{gain}_{\text{gini}} \equiv \text{gain}_{\chi^2}$, and $\text{gain}_{\text{shannon}} \equiv \text{gain}_{\text{KL}}$.

A parameterized version of $\hat{\mathcal{Y}}$ that takes into account the degree of confidence in the distribution of the consequent as it results from the data is introduced next.

Let us define the probability distribution $\Theta_a, a \in [0, 1]$ by

$$\Theta_a = a\Delta_Q + (1 - a)\mathcal{U}_n.$$

The value of a expresses the amount of confidence we have in Δ_Q estimated from the data. The value $a = 1$ means total confidence, we assume the probability estimated from data as the true probability distribution of Q . On the other hand, $a = 0$ means that we have no confidence in the estimate and use some prior distribution of Q instead. In our case, the prior is the uniform distribution \mathcal{U}_n . Note that $\Theta_1 = \Delta_Q$, and $\Theta_0 = \mathcal{U}_n$.

We can now define

$$\mathcal{Y}_{D,a} = \mathcal{Y}_{D, \Theta_a}.$$

Note that when $D = D_{\chi^2}$, we have (up to a constant factor) both $\chi^2(P \rightarrow Q)$ and $\text{gini}_{\text{gain}}(P \rightarrow Q)$ as special cases of $\mathcal{Y}_{D_{\chi^2}, a}$. Moreover by taking different values of parameter a we can obtain a continuum of measures in between the two.

As noted before, both D_{χ^2} and D_{KL} divergence measures are special cases of Havrda-Charvát divergence $D_{\mathcal{H}_\alpha}$ for $\alpha \rightarrow 1$, and $\alpha = 2$ respectively. We can thus introduce $\mathcal{Y}_{\alpha,a} = \mathcal{Y}_{D_{\mathcal{H}_\alpha}, \Theta_a}$, which allows us to obtain a family of interestingness measures, including (up to a constant factor) all three measures given in Sect. 3 as special cases, by simply changing two real valued parameters α and a .

Also note that for $a = 0$, we obtain a family of gains (as defined in Sect. 3) for all the Havrda-Charvát divergences.

4 Experimental Results

We evaluated the new measure on a simple synthetic dataset and on data from the UCI machine learning repository [2]. We concentrated on the case $D = D_{\chi^2}$, as potentially most useful in practice, and found interestingness of rules for different values of parameter a (see Sect. 3).

To ensure measures throughout the family handle obvious cases correctly, and to make it easy to observe properties of the measure for different values of parameter a we first evaluated the rules on a synthetic dataset with 3 attributes A, B, C and with known probabilistic dependencies between them.

Values of attributes A and B have been generated from known probability distributions:

$$\Delta_A = \begin{pmatrix} 0 & 1 & 2 \\ 0.1 & 0.5 & 0.4 \end{pmatrix}, \Delta_B = \begin{pmatrix} 0 & 1 \\ 0.2 & 0.8 \end{pmatrix}.$$

Attribute C depends on attribute A . Denote $\Delta_C|i$ the distribution of C conditioned upon $A = i$. We used

$$\Delta_C|0 = \begin{pmatrix} 0 & 1 \\ 0.2 & 0.8 \end{pmatrix}, \Delta_C|1 = \begin{pmatrix} 0 & 1 \\ 0.5 & 0.5 \end{pmatrix}, \Delta_C|2 = \begin{pmatrix} 0 & 1 \\ 0.7 & 0.3 \end{pmatrix},$$

One million data points have been generated according to this distribution, for a few values of a we sorted all possible rules based on their $\mathcal{Y}_{D_{\chi^2}, a}$ interestingness values. Results are given in Table 1.

1. Attribute B is totally independent of both A and C , so any rule containing only B as the antecedent or consequent should have interestingness 0. The experiments confirm this, for all values of parameter a such rules have interestingness close to zero, significantly lower than the interestingness of any other rules.
2. For $a = 0$ (the first quarter of the table) \mathcal{Y} becomes the Gini gain, a measure that is strongly asymmetric (and could thus suggest the direction of the dependence) and strongly affected by adding extra independent attributes to the consequent (which is undesirable).
3. For $a = 1$ (the last quarter of the table) \mathcal{Y} becomes (up to a constant factor) the χ^2 measure of dependence. This measure is totally symmetric and not affected by presence of independent attributes in either antecedent or consequent. Indeed, it can be seen that all rules involving A and C have the same interestingness regardless of the presence of B in the antecedent or consequent.
4. As a varies from 0 to 1 the intermediate measures can be seen to become more and more symmetric. Measures for a being close to but less than 1 could be of practical interest since they seem to ‘combine the best of the two worlds’, that is, are still asymmetric and pretty insensitive to presence of independent attributes in the consequent. E.g. for $a = 0.9$ all rules having A in the antecedent and C in the consequent have interestingness close to

Table 1. Rules on synthetic data ordered by $\Upsilon_{D_{\chi^2},a}$ for different values of a

rule	$\Upsilon_{D_{\chi^2},0}$	rule	$\Upsilon_{D_{\chi^2},0.5}$
$A \rightarrow BC$	0.122061	$A \rightarrow BC$	0.0989161
$C \rightarrow AB$	0.0896776	$AB \rightarrow C$	0.0898611
$AB \rightarrow C$	0.0896287	$A \rightarrow C$	0.089861
$A \rightarrow C$	0.0896287	$C \rightarrow AB$	0.0769886
$BC \rightarrow A$	0.065851	$BC \rightarrow A$	0.0683164
$C \rightarrow A$	0.0658484	$C \rightarrow A$	0.0683142
$B \rightarrow AC$	3.16585e-06	$B \rightarrow AC$	2.50502e-06
$B \rightarrow A$	2.7369e-06	$B \rightarrow A$	2.35091e-06
$AC \rightarrow B$	1.37659e-06	$AC \rightarrow B$	1.51849e-06
$A \rightarrow B$	1.32828e-06	$A \rightarrow B$	1.46355e-06
$B \rightarrow C$	1.70346e-07	$B \rightarrow C$	1.72781e-07
$C \rightarrow B$	1.10069e-07	$C \rightarrow B$	1.22814e-07

rule	$\Upsilon_{D_{\chi^2},0.9}$	rule	$\Upsilon_{D_{\chi^2},1}$
$A \rightarrow BC$	0.0908769	$BC \rightarrow A$	0.0905673
$AB \rightarrow C$	0.0903859	$A \rightarrow BC$	0.0905673
$A \rightarrow C$	0.0903859	$C \rightarrow AB$	0.0905654
$C \rightarrow AB$	0.0834734	$AB \rightarrow C$	0.0905654
$BC \rightarrow A$	0.082009	$A \rightarrow C$	0.0905653
$C \rightarrow A$	0.082007	$C \rightarrow A$	0.0905653
$B \rightarrow AC$	2.19739e-06	$AC \rightarrow B$	2.15872e-06
$B \rightarrow A$	2.12646e-06	$B \rightarrow AC$	2.15872e-06
$AC \rightarrow B$	1.95101e-06	$A \rightarrow B$	2.08117e-06
$A \rightarrow B$	1.87986e-06	$B \rightarrow A$	2.08017e-06
$B \rightarrow C$	1.73782e-07	$C \rightarrow B$	1.74126e-07
$C \rightarrow B$	1.57306e-07	$B \rightarrow C$	1.74126e-07

0.09, while rules having C in the antecedent and A in the consequent have all interestingness close to 0.082 regardless of the presence or absence of B in the consequents. So for $a = 0.9$ the intermediate measure correctly ranked the rules indicating the true direction of the relationship.

We then repeated the above experiment on data from the UCI machine learning repository [2]. Here we present results for the *agaricus-lepiota* database containing data on North American Mushrooms. To make the ruleset size manageable we restrict ourselves to rules involving the *class* attribute indicating whether the mushroom is edible or poisonous.

In the experiment we enumerated all rules involving up to 3 attributes and ranked them by interestingness for different values of parameter a . Top ten rules for each value of a are shown in Table 2. For $a = 1$ the symmetric rules were removed.

We noticed that for any value of a most of the rules involve the *odor* attribute. Indeed the inspection of data revealed that knowing the mushroom’s odor allows for identifying its class with 98.5% accuracy, far better than for any other attribute.

Table 2. Rules on mushroom dataset ordered by $\Upsilon_{D_{\chi^2},a}$ for different values of a

rule	$\Upsilon_{D_{\chi^2},0}$
class→odor ring-type	9.84024
class→odor spore-print-color	9.16709
class→odor veil-color	8.22064
class→odor gill-attachment	8.2026
class→gill-color spore-print-color	7.82161
class→ring-type spore-print-color	7.62564
class→odor stalk-root	7.60198
class→gill-color ring-type	7.28972
class→odor stalk-color-above-ring	7.19584
class→odor stalk-color-below-ring	7.14197
rule	$\Upsilon_{D_{\chi^2},0.9}$
odor→class stalk-root	3.61877
class stalk-root→odor	3.2782
odor→class cap-color	2.59777
odor→class ring-type	2.54896
odor→class spore-print-color	2.54864
stalk-color-above-ring→class stalk-color-below-ring	2.47669
class cap-color→odor	2.46105
odor→class gill-color	2.45027
stalk-color-below-ring→class stalk-color-above-ring	2.38593
class spore-print-color→odor	2.35384
rule	$\Upsilon_{D_{\chi^2},1}$
class stalk-root→odor	4.11701
class stalk-color-below-ring→stalk-color-above-ring	3.38287
stalk-color-below-ring→class stalk-color-above-ring	3.37968
class ring-type→odor	2.98764
class cap-color→odor	2.85308
odor→class gill-color	2.82423
odor→class spore-print-color	2.56331
odor→class stalk-color-below-ring	2.44004
class stalk-color-above-ring→odor	2.42725
class gill-color→spore-print-color	2.42224

We note also that similar rules are ranked close to the top for all values of a , which proves that measures throughout the family identify dependencies correctly. From data omitted in the tables it can be observed that, as in the case of synthetic data, when a approaches 1 the measures become more and more symmetric and unaffected by independent attributes in the consequent.

It has been shown experimentally that measures throughout the Υ family are useful for discovering interesting dependencies among data attributes. By modifying a numerical attribute we can obtain a whole spectrum of measure of varying degree of symmetry and dependence on the presence of extra attributes in the rule consequent. Especially interesting seem to be measures with a parameter close to, but less than 1, which combine the relative robustness against extra

independent attributes, while retaining the asymmetry suggesting the direction of the dependence.

5 Open Problems and Future Directions

Above we assumed complete confidence in the estimate of the distribution of P from the data. We may want to relax this restriction and assume that P has some a posteriori distribution Ψ (not necessarily equal to Δ_P), and Q the prior distribution Θ . We can then generalize Υ as $\Upsilon'_{D,\Theta,\Psi}(P \rightarrow Q) = D(\Delta_{PQ}, \Psi \times \Theta) - D(\Delta_Q, \Theta) - D(\Delta_P, \Psi)$. When $\Psi = \Delta_P$, Υ' reduces to Υ defined above. Some of the properties of Υ are preserved by this new definition. For example, if $D = D_{\text{KL}}$, and P, Q be independent, then $\Upsilon'_{D,\Theta,\Psi}(P \rightarrow Q) = 0$. Also, if $P \rightarrow Q$ is a rule in the table $\tau = (T, H, \rho)$ and $D = D_{\text{KL}}$ then $\Upsilon'_{D,\Theta,\Psi}(P \rightarrow Q) = \text{gain}_{\text{shannon}}(Q, P)$ regardless of the choice of Θ and Ψ .

Further theoretical and experimental evaluation of the new measure is necessary. It would be of practical interest to find a modified general definition of gain that, being asymmetric, is not affected by adding independent attributes to the consequent.

As a primary application, we envision using the measure in association rule mining systems for sorting the discovered association rules. For this purpose it would be necessary to generalize the measure to express the interestingness of a rule with respect to a system of beliefs (that could be represented for example by a set of rules). Then, the rule would be considered interesting if its probability distribution would be significantly different from the one expected based on the set of beliefs. See [13] for a discussion of a similar problem.

Further work is necessary to assess the impact that the generalized measure would have on other common datamining tasks like attribute selection in decision trees. It might, for example, be beneficial to use values of parameter a close to 1 in the upper parts of the tree when large amount of data is still available, and decreasing the value of a at lower levels, where the amount of data is small and thus we have less confidence in the estimates of probabilities.

References

1. Bayardo R.J. and R. Agrawal, *Mining the Most Interesting Rules*, Proc. of the 5th ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining, pp. 145-154, August 1999.
2. Blake C.L., and Merz C.J. *UCI Repository of machine learning databases* Irvine, CA: University of California, Department of Information and Computer Science, <http://www.ics.uci.edu/~mllearn/MLRepository.html>
3. Chou P.A., *Optimal Partitioning for classification and regression trees* IEEE Trans. on Pattern Analysis and Machine Intelligence, PAMI-13(14):340-354, 1991.
4. Silverstein C., Brin S. and Motwani R., *Beyond Market Baskets: Generalizing Association Rules to Dependence Rules* Data Mining and Knowledge Discovery, 2(1998), pp. 39-68

5. Cziszar I., *A Class of Measures of Informativity of Observation Channels*, Periodic Math. Hungarica, 2:191-213, 1972.
6. Havrda J.H., Charvát F., *Quantification Methods of Classification Processes: Concepts of Structural α Entropy*, Kybernetika, 3:30-35, 1967.
7. Kapur J.N. and Kesavan H.K., *Entropy Optimization Principles with Applications*, Academic Press, San Diego, 1992.
8. Kvålseth T.O., *Entropy and Correlation: Some comments*, IEEE Trans. on Systems, Man and Cybernetics, SMC-17(3):517-519, 1987.
9. McEliece R.J. *The Theory of Information and Coding. A mathematical Framework for Communication*, Encyclopedia of Mathematics and its Applications, Addison-Wesley, Reading Massachusetts, 1977.
10. Mitchell T.M.. *Machine Learning*, McGraw-Hill, ISBN: 0070428077.
11. Morimoto Y., Fukuda T., Matsuzawa H., Tokuyama T. and Yoda K. *Algorithms for Mining Association Rules for Binary Segmentations of Huge Categorical Databases*, Proc. of the 24th Conf. on Very Large Databases, pp. 380-391, 1998
12. Morishita S., *On Classification and Regression* Proc. of the First Int'l Conf. on Discovery Science — Lecture Notes in Artificial Intelligence 1532:40-57, 1998
13. Padmanabhan B. and Tuzhilin A. *Unexpectedness as a measure of interestingness in knowledge discovery* Decision and Support Systems 27(1999), pp. 303-318
14. Simovici, D. A. and Tenney R. L. *Relational Database Systems*, Academic Press, 1995, San Diego.
15. Wehenkel L., *On uncertainty Measures Used for Decision Tree Induction*, Info. Proc. and Manag. of Uncertainty in Knowledge-Based Systems (IPMU'96), July 1-5, 1996, Granada Spain, pp. 413-418.
16. Witten I.H., and Eibe F., *Data Mining, Practical Machine Learning and Techniques with JAVA Implementations*, Academic Press, San Diego, CA, 2000, ISBN: 1558605525.