



Title	Standardizing interestingness measures for association rules
Authors(s)	Shaikh, Mateen, McNicholas, Paul D., Antonie, M. Luiza, Murphy, Thomas Brendan
Publication date	2018-12
Publication information	Shaikh, Mateen, Paul D. McNicholas, M. Luiza Antonie, and Thomas Brendan Murphy. "Standardizing Interestingness Measures for Association Rules." Wiley, December 2018. https://doi.org/10.1002/sam.11394 .
Publisher	Wiley
Item record/more information	http://hdl.handle.net/10197/10544
Publisher's statement	This is the peer reviewed version of the following article:Shaikh MR, McNicholas PD, Antonie ML, Murphy TB. Standardizing interestingness measures for association rules. Stat Anal Data Min: The ASA Data Sci Journal 2018;11:282–295., which has been published in final form at http://onlinelibrary.wiley.com/doi/10.1002/sam.11394 . This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.
Publisher's version (DOI)	10.1002/sam.11394

Downloaded 2026-05-02 00:27:24

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Standardizing Interestingness Measures for Association Rules

Mateen Shaikh, Paul D. McNicholas*, M. Luiza Antonie and T. Brendan Murphy

Department of Mathematics & Statistics, University of Guelph

Abstract

Interestingness measures provide information that can be used to prune or select association rules. A given value of an interestingness measure is often interpreted relative to the overall range of the values that the interestingness measure can take. However, properties of individual association rules restrict the values an interestingness measure can achieve. An interesting measure can be standardized to take this into account, but this has only been done for one interestingness measure to date, i.e., the lift. Standardization provides greater insight than the raw value and may even alter researchers' perception of the data. We derive standardized analogues of three interestingness measures and use real and simulated data to compare them to their raw versions, each other, and the standardized lift.

1 Introduction

Association rules have been the subject of considerable research since the seminal paper by Agrawal et al (1993). They are used to identify trends and patterns in transactional data such as market-basket analyses, where researchers investigate common purchasing trends. Once mined, association rules are often compared to each other using some interestingness measure. Different interestingness measures quantify different characteristics of an association rule (see Piatetsky-Shapiro, 1991; Freitas, 1999; Tan et al, 2002; Geng and Hamilton, 2006, for discussions). Each interestingness measure takes values on some range; however, to only consider the value of an interestingness measure without reference to this range can be misleading, cf. Section 2.2. To overcome this issue, interestingness measures can be standardized to provide values that reflect the value achieved relative to the values that are obtainable. To date, only one interestingness measure has been standardized — the

*Department of Mathematics & Statistics, University of Guelph, Guelph, Ontario, N1G 2W1, Canada.
E-mail: pmcnicho@uoguelph.ca.

lift (McNicholas et al, 2008) — and so the effects of standardizing interestingness measures are not well understood in general. Herein, we standardize three additional interestingness measures and explore the corresponding effects through the analysis of three data sets.

Association rules are defined in the context of a non-empty set of items I , called the itemset. An association rule is a statement of the form $A \Rightarrow B$, where $A \subset I, B \subset I$ such that $A \neq \emptyset, B \neq \emptyset$, and $A \cap B = \emptyset$. The set A is called the antecedent and B is called the consequent. Association rules are mined over many transactions. A transaction is a single instance wherein each item is present or absent. The support of $X \subset I$, denoted $P(X)$, is the proportion of transactions that contain every item in X . Consider a rule $A \Rightarrow B$. $P(A)$ denotes the support of the antecedent, $P(B)$ denotes the support of the consequent, and $P(A, B)$ (or equivalently $P(A \Rightarrow B)$) denotes the support of the rule itself. The confidence of a rule is

$$P(B | A) = \frac{P(A, B)}{P(A)}. \quad (1)$$

There are many instances when the absence of items within transactions is also of interest. $P(\bar{X})$ denotes the proportion of transactions that do not contain every item in $X \subset I$. For example, one may be interested in the proportion of transactions that contain every item in A but do not contain the items in B , i.e., the rules $A \Rightarrow \bar{B}$ and $\bar{B} \Rightarrow A$. The set \bar{A} is called the negation of A . Although we do not explicitly consider rules containing negations herein, our calculations are more concise and understandable using this notation.

2 Interestingness Measures

2.1 General Remarks and Properties

There are many different interestingness measures for association rules and some are more suitable for some purposes than others. Piatetsky-Shapiro (1991) outlined three important properties every interestingness measure should have. The first property (Prop. 1) is that the interestingness measure should achieve the value 0 at independence, i.e., when $P(A, B) = P(A)P(B)$. A non-zero constant, if known in advance, satisfies the same purpose, e.g., a value of one indicates independence for lift (cf. Section 2.2). The second desirable property (Prop. 2) is that the value of the interestingness measure should be larger when $P(A, B)$ is larger but $P(A)$ and $P(B)$ remain the same. The third property (Prop. 3) is related to the second: the interestingness measure should be smaller when either $P(A)$ or $P(B)$ increases but $P(A, B)$ is unchanged, or when $P(B)$ increases but $P(A)$ and $P(A, B)$ are unchanged.

In addition to these three properties, five others have been identified by Tan et al (2002) and considered by Geng and Hamilton (2006). These properties are not necessarily desirable but can be important to help distinguish interestingness measures from one another. One property is symmetry (Prop. 4): the interestingness measure remains the same when the antecedent and consequent are interchanged. Another property is antisymmetry under one negation (Prop. 5), i.e., the measure’s value is multiplied by (-1) when the antecedent is

replaced by its negation or the consequent by its negation. A related property (Prop. 6) is symmetry under two negations, i.e., the measure’s value is multiplied by (-1) when the antecedent is replaced by its negation and the consequent by its negation. Prop. 5 trivially implies Prop. 6; however, some interestingness measures, such as the odds ratio and the Gini index, satisfy Prop. 6 but not Prop. 5. Another suggested property (Prop. 7) is that the interestingness measure should achieve the same value under row-wise or column-wise scaling (or both). The final property is null invariance (Prop. 8): additional transactions where neither the antecedent nor the consequent appear do not change the value of the interestingness measure.

A total of 21 interestingness measures are considered by Tan et al (2002). These are discussed in the specific context of association rules by Geng and Hamilton (2006), who outline methods of choosing suitable interestingness measures. Part of their analysis includes identifying interestingness measures that satisfied the three properties of Piattetsky-Shapiro (1991) as well as the five additional properties discussed by Tan et al (2002). This analysis includes a table (their Table 6) that shows whether or not each of the 21 interestingness measures satisfies these eight properties. In an attempt to group these interestingness measures, their analysis also contains a figure (their Figure 3) identifying levels of similarity, by correlation, between interestingness measures. It is from this figure, where all values of support are considered, that we identify at least three separate groupings of interestingness measures. Interestingness measures that appear in the same group often do not satisfy the same subset of the eight properties but are well correlated, in that they have correlation values of least 0.85 with any other interestingness measure in the same group. Because of this similarity, we choose one representative from each of the three groups to analyze: cosine similarity, Yule’s Q, and the Gini index. Table 1 shows a summary of the three interestingness measures in the same fashion as Table 6 from Tan et al (2002).

Table 1: Properties of four selected interestingness measures.

	Prop. 1	Prop. 2	Prop. 3	Prop. 4	Prop. 5	Prop. 6	Prop. 7	Prop. 8
Lift	Y	Y	Y	Y	N	N	N	N
Cosine Similarity	N	Y	Y	Y	N	N	N	Y
Yule’s Q	Y	Y	Y	Y	Y	Y	Y	N
Gini Index	Y	N	N	N	N	Y	N	N

2.2 Lift

The lift of a rule, which has also been called the ‘interest’ of a rule (Brin et al, 1997), is denoted here as $\mathcal{L}(A \Rightarrow B)$ and is defined as

$$\mathcal{L}(A \Rightarrow B) = \frac{P(A, B)}{P(A)P(B)}. \quad (2)$$

Lift takes on non-negative values and equals one when the antecedent and consequent are statistically independent. As an example of when a measure such as lift can be problematic, we consider an instance where the support of the antecedent of a rule, $A \Rightarrow B$, is 0.5 and the overall lift value for this rule is 1.95. Because $P(A) = 0.5$, we know that the value of lift for this rule cannot exceed 2. Suppose, also, that for another rule, $C \Rightarrow D$, the lift of the rule is 1.95 but the support of the antecedent, $P(C)$, is 0.1. In this second rule, the lift value was not bounded above by 2 but by 10. Even though these two rules have the same lift value of 1.95, the interpretation of the two lift values should not be the same because the maximum attainable value is different. Such discrepancies occur when using raw interestingness measures; standardizing lift (McNicholas et al, 2008) addresses the issue described here. In this example, standardized lift values for rules $A \Rightarrow B$ and $C \Rightarrow D$ are 0.975 and 0.195, respectively, all else being equal.

To show just how bounded lift is, consider the values of lift where $P(A) = P(B)$. We can see (Figure 1) that for larger supports, the value of lift is bounded near one. Without any other considerations, the small differences in the values of lift between rules may be overlooked; however, these differences may be more significant after considering the narrow range of possible values. Though we only highlight this issue in the case when $P(A) = P(B)$, such problems exist in many cases.




Figure 1: The bounds on lift when $P(A) = P(B)$.

The standardized lift of a rule, $\mathcal{L}^*(A \Rightarrow B)$, is obtained by finding the suitable upper bound, v , and lower bound, λ , in the following definition:

$$\mathcal{L}^*(A \Rightarrow B) = \frac{\mathcal{L}(A \Rightarrow B) - \lambda}{v - \lambda}.$$

McNicholas et al (2008) derive the following values of v and λ :

$$v = \frac{1}{\max\{P(A), P(B)\}},$$

$$\lambda = \max \left\{ \frac{P(A) + P(B) - 1}{P(A)P(B)}, \frac{4\sigma}{(1 + \sigma)^2}, \frac{\sigma}{P(A)P(B)}, \frac{\kappa}{P(B)} \right\}.$$

These values are functions of the minimum support and confidence thresholds, denoted in this paper by σ and κ , respectively. These thresholds are assumed to be at least $1/n$, where n is the number of transactions. Analogous approaches are used herein to standardize three other interestingness measures.

2.3 Cosine Similarity

The definition of the cosine similarity resembles the definition of the lift of a rule and is:

$$\mathcal{C}(A \Rightarrow B) = \frac{P(A, B)}{\sqrt{P(A)P(B)}}. \quad (3)$$

The cosine similarity lies on the closed unit interval and no constant value is achieved when the antecedent and consequent are statistically independent of one another. In such cases, the cosine similarity attains the value $\sqrt{P(A)P(B)}$. The standardized cosine similarity of a rule $A \Rightarrow B$ is given by

$$\mathcal{C}^*(A \Rightarrow B) = \frac{\mathcal{C}(A \Rightarrow B) - \lambda}{v - \lambda},$$

where

$$v = \min \left\{ \sqrt{\frac{P(A)}{P(B)}}, \sqrt{\frac{P(B)}{P(A)}} \right\}$$

and

$$\lambda = \max \left\{ \frac{2\sigma}{1 + \sigma}, \frac{\sigma}{\sqrt{P(A)P(B)}}, \frac{P(A) + P(B) - 1}{\sqrt{P(A)P(B)}}, \sqrt{\kappa \frac{\sigma}{P(B)}}, \kappa \sqrt{\frac{P(A)}{P(B)}} \right\}.$$

Derivations for v and λ are found in Appendix A.1.

2.4 Yule's Q

Some interestingness measures take supports of negations into account. One such interestingness measure is Yule's Q (Yule, 1903), which is based on the odds ratio. Unlike the odds ratio, however, Yule's Q remains defined in all cases where $P(A) > 0$ and $P(B) > 0$. Yule's Q takes values on $[-1, 1]$ and is defined as

$$\mathcal{Q}(A \Rightarrow B) = \frac{P(A, B)P(\bar{A}, \bar{B}) - P(\bar{A}, B)P(A, \bar{B})}{P(A, B)P(\bar{A}, \bar{B}) + P(\bar{A}, B)P(A, \bar{B})}. \quad (4)$$

This interestingness measure obtains a value of zero when the antecedent and consequent are statistically independent. Positive values indicate that the antecedent and consequent are more likely to occur together than if they were independent, and negative values indicate that the antecedent and consequent are less likely to occur together than if they were independent. Akin to lift and the cosine similarity, Yule's Q is symmetric.

The standardized Yule's Q of a rule $A \Rightarrow B$ is

$$\mathcal{Q}^*(A \Rightarrow B) = \frac{\mathcal{Q}(A \Rightarrow B) - \lambda}{v - \lambda},$$

where $v = 1$ and

$$\lambda = \max \left\{ -1, \frac{\sigma - P(A)P(B)}{\sigma + P(A)P(B) - 2\sigma(P(A) + P(B) - \sigma)}, \frac{\kappa - P(B)}{\kappa + P(B) - 2\kappa(P(A) + P(B) - \kappa P(A))} \right\}.$$

Derivations for v and λ are found in Appendix A.2.

2.5 Gini Index

The Gini index is the only measure discussed here that is not symmetric. That is, the Gini index for the rule $A \Rightarrow B$ may not achieve the same value as the Gini index for the rule $B \Rightarrow A$. Used in a variety of contexts, the Gini index quantifies impurity and is defined as

$$\begin{aligned} \mathcal{G}(A \Rightarrow B) = & P(A) [P(B | A)^2 + P(\bar{B} | A)^2] \\ & + P(\bar{A}) [P(B | \bar{A})^2 + P(\bar{B} | \bar{A})^2] - P(B)^2 - P(\bar{B})^2. \end{aligned} \quad (5)$$

This interestingness measure achieves a value of zero when the antecedent and consequent are statistically independent, and achieves a value between zero and 0.5 when the antecedent and consequent are not independent. The standardized Gini index of a rule $A \Rightarrow B$ is

$$\mathcal{G}^*(A \Rightarrow B) = \frac{\mathcal{G}(A \Rightarrow B) - \lambda}{v - \lambda},$$

where

$$v = \begin{cases} 2 \frac{(\min\{P(A), P(B)\} - P(A)P(B))^2}{P(A)(1-P(A))} & \text{if } P(A, B) \geq P(A)P(B), \\ 2 \frac{(\max\{\sigma, \kappa P(A), P(A)+P(B)-1\} - P(A)P(B))^2}{P(A)(1-P(A))} & \text{if } P(A, B) < P(A)P(B), \end{cases}$$

and

$$\lambda = \begin{cases} 2 \frac{(\max\{\sigma, \kappa P(A), P(A)+P(B)-1, P(A)P(B)\} - P(A)P(B))^2}{P(A)(1-P(A))} & \text{if } P(A, B) \geq P(A)P(B), \\ 0 & \text{if } P(A, B) < P(A)P(B). \end{cases}$$

Derivations for v and λ are given in Appendix A.3.

3 Experimental Results and Evaluation

We explore the effects of standardizing interestingness measures with three different data sets. In the first data set, values are obtained for lift, cosine similarity, Yule’s Q, and the Gini index. We see that values of each interestingness measure are clearly reordered when standardized, making previously unremarkable rules more noteworthy. The second data set that we analyze is comparably more sparse, and comparing the effect of standardizing interestingness measures with results from the first data set is interesting. We also consider a third data set that is comprised of completely simulated, random data, with independent items. These data draw a contrast between standardized interestingness measures and their non-standardized analogues. Specifically, they highlight that by taking into account the low probabilities of co-occurrences, standardized interestingness measures do not suggest that rules in random transactions are interesting.

For every data set, rules are mined using the Apriori algorithm (Agrawal and Srikant, 1994). This method uses the Apriori principle, also known as downward closure; i.e., if a set of items is deemed to be frequent relative to a threshold, then every subset of this set will also be (at least as) frequent relative to the same threshold. Conversely, if a set of items is judged to be infrequent relative to a threshold, then every superset of this set will also be infrequent relative to the threshold. There are many alternatives to this algorithm, such as SETM (Houtsma and Swami, 1995); EClat (Zaki, 2000), which divides the itemset into equivalence classes; frequency pattern trees (Han et al, 2000), which avoid many passes over the data; and variations to the Apriori algorithm itself such as AprioriTID, AprioriHybrid (Agrawal and Srikant, 1994), and direct hashing and pruning (Park et al, 1995). Because the focus of our work is to explore the effects of standardizing interestingness measures, we employ the established and well-understood Apriori algorithm. Results for each data set can be visualized as a pair of scatterplot matrices (Appendix A.4); the first scatterplot matrix shows the raw interestingness measures and the second the standardized interestingness measures.

We note that the thresholds chosen in our exploratory analyses may not be typical for specific applications, which often use higher support thresholds and impose higher confidence thresholds. Such thresholds typically result in fewer rules, most often those indicating a positive correlation between the antecedent and consequent. To appropriately explore the effect of standardization, we intentionally keep these rules in our analysis when they existed, without the burden of a plethora of rules.

One method of comparing rule orderings is through Goodman and Kruskal’s gamma (Goodman and Kruskal, 1954). In this context, it is calculated by determining the difference in the proportions of concordant pairs and discordant pairs while ignoring ties. Kendall’s tau-b is a similar method of measuring the rankings that considers ties to be discordant, making it a more conservative description of relative ordering (cf. Agresti, 2002; Blaikie, 2003). We use Kendall’s tau-b herein. If the order of rules by raw interestingness measure is the same as the order of rules by the corresponding standardized interestingness measure,

a value of 1 results. If the order of one is the opposite of the other, a value of -1 results. A value of zero results when the two rule orderings have no apparent relationship, such as the case when one of the rule orders have been (uniformly) randomly ordered. In addition to calculating the tau-b value for all rules in each data set, we calculated tau-b by decile in each data set. Rather than numerically reporting these values, plots have been made available in Figure 9 of Appendix A.5.

4 Traffic Accidents

Information from over 340,000 traffic accidents from Belgium’s National Institute of Statistics (Geurts et al, 2003) is analyzed by considering up to 572 factors that may have been involved. These factors include the type of road, intersection, features of the road, conditions, time, type of vehicle, direction, severity of accident, age of road user, obstacles, and other relevant criteria recorded by a police officer. As Table 2 indicates, information identifying the actual factors has not been made available; rather, factors are simply enumerated. We set minimum support and confidence thresholds to 0.3 and 0.4, respectively. These thresholds were chosen because approximately 90,000 rules were produced when we considered entire rules that were no more than five items in length; we deemed this adequate for analysis.

To illustrate how just one interestingness measure is standardized, Figure 2 shows ten randomly selected rules sorted by \mathcal{C}^* . Each vertical bar represents the range of values of \mathcal{C} that were attainable by the rule based on the bounds discussed in Section 2.3. The horizontal tick within the bar shows where the actual value of \mathcal{C} lies. The value of \mathcal{C}^* is the relative location of \mathcal{C} in the bar shown with circles. These values are overlain because the range of \mathcal{C} and \mathcal{C}^* are the same but the range of the raw and standardized interestingness measures are not coincident in general unless the range of the raw interestingness measure is also the unit interval. The rules are sorted by ascending values of \mathcal{C}^* and show that standardizing the interestingness measure is not simply a transformation of the raw interestingness measure.

When comparing the standardized values against raw values, a variety of behaviours are observed (cf. Figure 3). Standardizing \mathcal{L} results in two very noticeable trends in these rules. The values of \mathcal{L} and \mathcal{L}^* are, generally, positively correlated when the value of \mathcal{L} is not near one. Rules that have values of \mathcal{L} near one obtain many values of \mathcal{L}^* in the unit interval with a concentration near one. Even though the cosine similarity is defined in a similar form as lift, no constant value of \mathcal{C} indicates independence and so the trend observed in the plot with lift is not identifiable. In general, high values of \mathcal{C} identify positively correlated items and low values of \mathcal{C} identify negatively correlated items.

The plot of Yule’s Q reveals several features. There is a high concentration of values with \mathcal{Q} near zero, indicating the same independent rules as for when \mathcal{L} had a value near one. There are also no rules with low values of \mathcal{Q} while achieving high values of \mathcal{Q}^* . This is not particularly surprising because the upperbound of \mathcal{Q} is constant. The plot of the Gini index is similar to the others. The majority of rules have values of \mathcal{G} very near zero. No




Figure 2: Cosine similarity values for randomly selected rules. Vertical lines represent the range of possible cosine similarity values for the rule, with the horizontal dash indicating the achieved value of the cosine similarity. The circle indicates the relative value of the standardized cosine similarity within the range of possible values.

values of \mathcal{G} that are high obtain a low value of \mathcal{G}^* . In contrast to the other standardized interestingness measures, \mathcal{G}^* maps rules indicating independence to a single value, namely zero, as \mathcal{G} would already be very close to zero.

In Figure 3, red points highlight rules that obtained values of \mathcal{L}^* more than 0.9 and green points highlight rules that obtained values of \mathcal{C}^* greater than 0.95. We choose these values to observe how high values of a standardized interestingness measure correspond to other interestingness measures, including the other standardized interestingness measures. In this figure, as well as similar figures in following sections, the raw interestingness measure is on the horizontal axis and the standardized on the vertical axis. In this example, lift and the cosine similarity are specifically selected together because we know that the similar form of \mathcal{L} and \mathcal{C} results in similar raw interestingness measures. Not surprisingly, the relationship between high values of \mathcal{L}^* and high values of \mathcal{C}^* is strong. These points coincide with relatively high values of \mathcal{Q} , \mathcal{Q}^* , and even \mathcal{G}^* but not \mathcal{G} . The correspondence of all four standardized interestingness measures speaks to a common type of measurement. This is in contrast to the raw interestingness measures: mostly moderate values of \mathcal{L} , a wide range of values of \mathcal{C} , mostly high values of q , and mostly low values of \mathcal{G} . These are reflected more

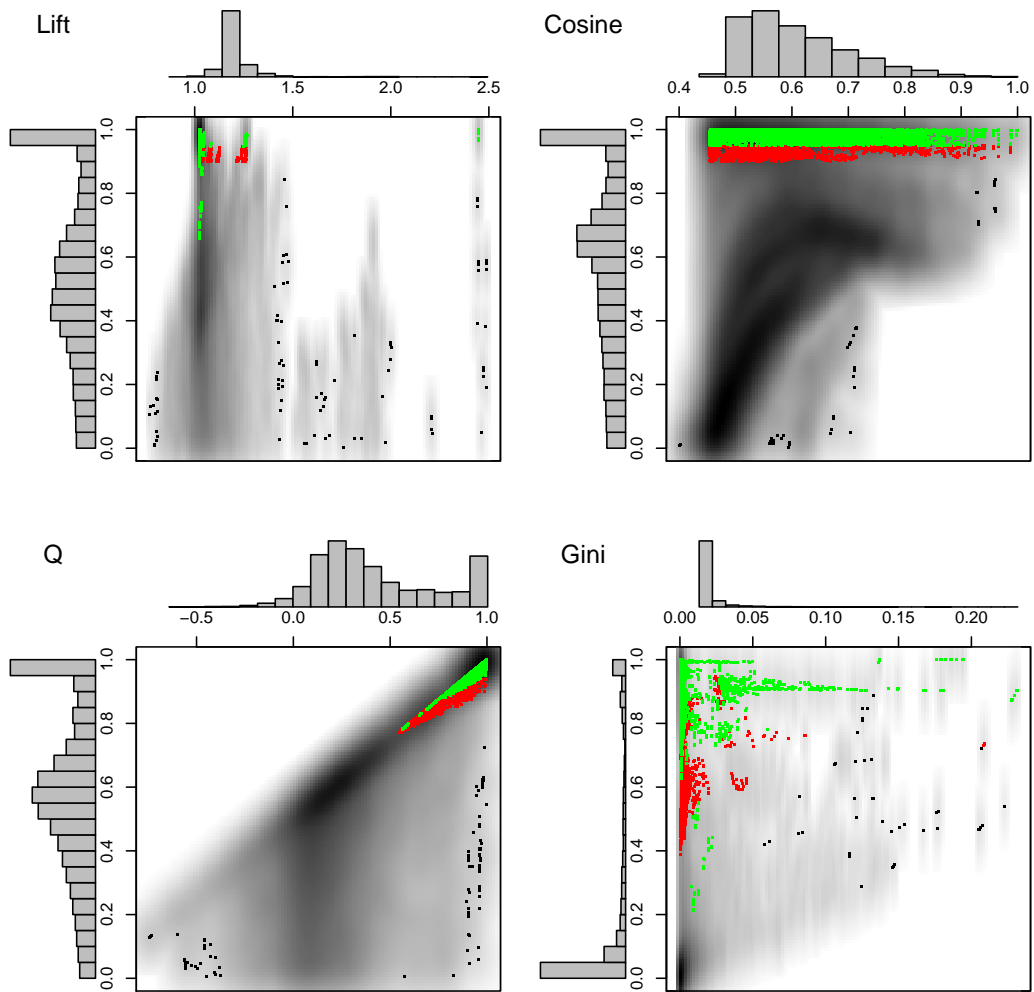


Figure 3: Plots of interestingness measures against their standardized counterparts for the Belgian accident data. Dark regions indicate a higher density of rules.

generally in Appendix A.4.

Table 2 shows rules with the consequent {35}. These rules show similar values for an interestingness measure as well as similar values for a standardized interestingness measure. There is, generally, greater disparity in the standardized interestingness measure than in the raw interestingness measure. For example, all of the rules obtained values of \mathcal{L} of approximately 1.95, but \mathcal{L}^* ranged from 0.32 to 0.39 with the lowest values occurring when item {18} is in the antecedent. The standardized cosine similarity and standardized Yule’s Q exhibit the same behaviour and, to a lesser extent, we observe this with \mathcal{Q}^* . Table 3 contains Kendall’s tau-b for these data. Standardizing lift reorders the rankings of rules the most, while standardizing Yule’s Q reorders rules the least. The cosine similarity and the Gini index maintain a more moderate relationship when each is standardized. These trends are shown by deciles in Figure 9(a) of Appendix A.5. The values of tau-b are fairly close to zero, suggesting a significant reordering of rules for all interestingness measures and all deciles, except for high deciles of Yule’s Q that obtain higher values of tau-b. The maintained ordering is most evident in higher deciles of Yule’s Q, which may be due to the interestingness measure being the only of the four discussed that always has a constant for an upper bound.

Table 2: Rules for the Belgian accident data with the consequent {35} sorted by support and confidence.

antecedent	supp	conf	raw lift	s. lift	raw cos	s. cos	raw Q	s. Q	raw Gini	s. Gini
{12,18,23}	0.203	0.411	1.951	0.317	0.630	0.317	0.959	0.334	0.079	0.868
{17,18,23}	0.203	0.411	1.951	0.325	0.630	0.325	0.960	0.342	0.079	0.870
{18,23}	0.203	0.411	1.951	0.327	0.630	0.327	0.960	0.343	0.079	0.870
{12,17,23}	0.204	0.411	1.952	0.368	0.631	0.368	0.962	0.386	0.079	0.878
{12,23}	0.204	0.411	1.952	0.370	0.631	0.37	0.962	0.388	0.079	0.878
{17,23}	0.204	0.411	1.952	0.385	0.631	0.385	0.963	0.403	0.079	0.881
{23}	0.204	0.411	1.952	0.387	0.631	0.387	0.963	0.405	0.079	0.881

Table 3: Rule orders for the Belgian accident data.

Interestingness Measure	tau-b
\mathcal{L} and \mathcal{L}^*	-0.033
\mathcal{C} and \mathcal{C}^*	0.380
\mathcal{Q} and \mathcal{Q}^*	0.657
\mathcal{G} and \mathcal{G}^*	0.250

5 Reuters

The Reuters-21578 Distribution 1.0 data (Lewis, 1996) are well-known data in text categorization. They are a collection of documents appearing on the Reuters newswire in 1987, made available in 1990, and reformatted in 1996. The reformatted version of the data is analyzed here. Words in documents are also reduced to word stems using the `Rstem` package (Lang, 2010) for `R` (R Development Core Team, 2012) and stop words are removed. There were over 1,100 key word stems and 17,000 documents containing at least one of these key word stems, with an average of 38 unique key word stems.

Support and confidence thresholds, both set to 0.005, provide over 500,000 rules with the Apriori algorithm. Figure 4 shows the plots of raw interestingness measures against their standardized counterparts. The highest values of lift are higher than the previous data, indicating greater dependence. Similar to the previous data, most rules have values of \mathcal{L} near one and these rules have values of \mathcal{L}^* spread throughout the unit interval with concentrations near zero and one. Most rules obtain positive values of \mathcal{Q} with these data. When standardized, the range of values of \mathcal{Q}^* between zero and one is well represented. One concentration occurs at the maximum value of \mathcal{Q} near one, which has a large proportion of values of \mathcal{Q}^* near its maximum of one. The plot of \mathcal{G}^* against \mathcal{G} shows that most values of \mathcal{G} are close to zero and correspond to a variety of \mathcal{G}^* values, not just those near zero. However, most of the rules with \mathcal{G} near zero also have values of \mathcal{G}^* near zero.

Rules represented in Figure 4 are highlighted in red when \mathcal{G} is above 0.05 and in green when \mathcal{G}^* is above 0.8. High values of \mathcal{G} correspond to high values of \mathcal{C} and \mathcal{Q} , while rules with high values of \mathcal{G}^* correspond to high values of \mathcal{L} , \mathcal{L}^* , \mathcal{C} , and \mathcal{C}^* as well as very high values of \mathcal{Q} and hence \mathcal{Q}^* . We also observe a few rules that exhibit the opposite behaviour among the negatively correlated rules. This most interestingly manifests in the Yule’s Q subplot, which shows the few negatively correlated rules with values of \mathcal{Q} below zero. Similar to the previous data, these data indicate high values of \mathcal{G} related to highly dependent rules. In this second data set, we again observe that values of raw interestingness measures indicating independence have standardized counterparts spread throughout the unit interval rather than being centred around a single value. This shows that while the value of a raw interestingness measure may suggest independence for several rules (Prop. 1 from Section 2), these rules may have different standardized interestingness measures.

Table 5 shows values of tau-b for the Reuters-21578 data. Unlike the previous data, lift and the cosine similarity are reordered substantially. Figure 9(c) shows how most rules have been reordered across deciles and interestingness measures. An example is illustrated in Table 4 for rules containing the consequent {gold}. Unlike the previous data, we observe greater differences in raw interestingness measures and even greater differences in the standardized interestingness measures. We also again see that the standardized interestingness measures share some similarities, including generally monotonic relationships with each other (see Figure 7 in Appendix A.4). We also observe the near-boundary case for the standardized Gini index. The raw values for the Gini index are all very, very close to zero but the standardized

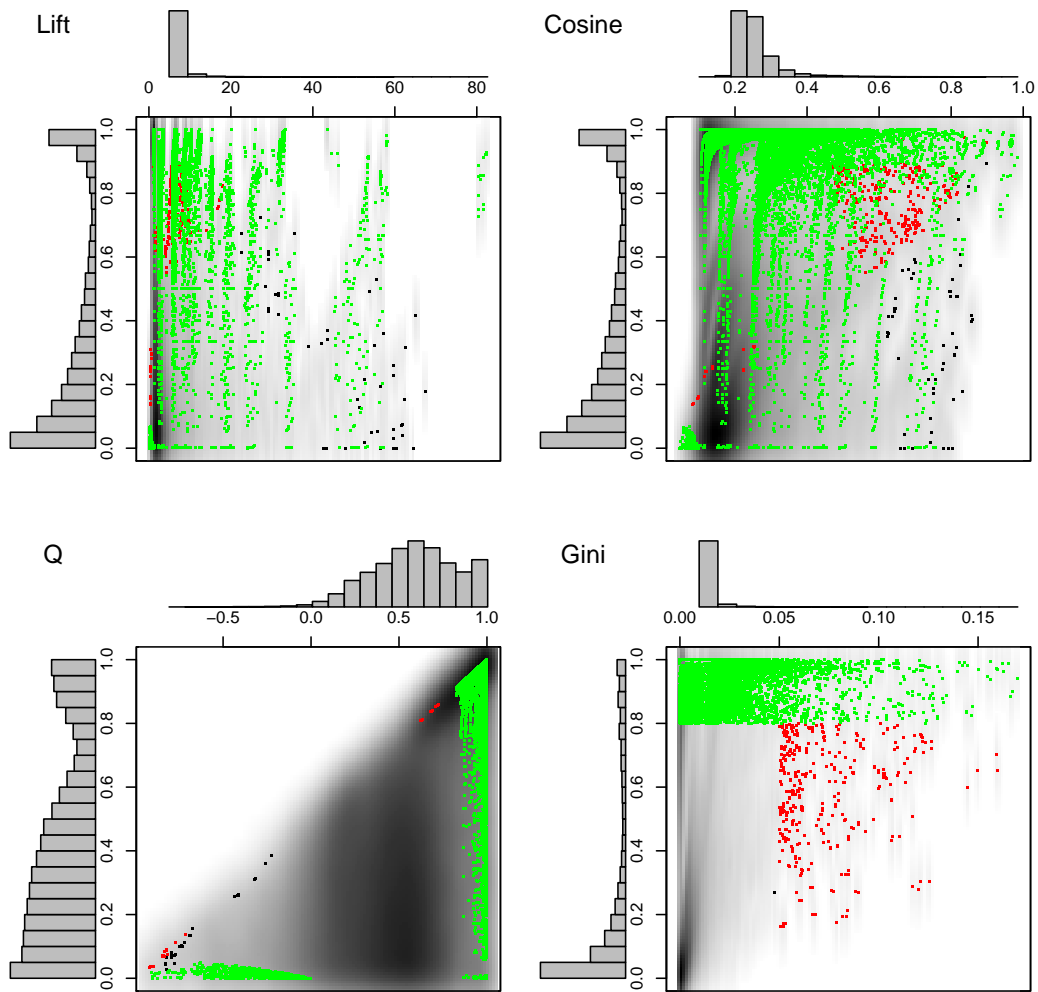


Figure 4: Plots of interestingness measures against their standardized counterparts for the Reuters data. Dark regions indicate a higher density of rules.

interestingness measures clearly differ from zero in two rules while remaining essentially zero in another rule. Despite being so close to the raw interestingness measure’s global lower bound, we see that the standardized interestingness measure for the rule $\{\text{reuter}\} \Rightarrow \{\text{gold}\}$ is relatively high; this shows that this indication of nearly independent rules may not be so close to independence after all. A somewhat similar phenomenon is observed in the previous data in Table 2, where Q^* values were moderate even though q values were near the global maximum of one.

Table 4: Rules for the Reuters Data with the consequent $\{\text{gold}\}$ sorted by support and confidence.

antecedent	supp	conf	raw lift	s. lift	raw cos	s. cos	raw Q	s. Q	raw Gini	s. Gini
$\{\text{march}\}$	0.011	0.659	1.001	0.198	0.107	0.198	0.002	0.153	3×10^{-8}	7×10^{-6}
$\{\text{said}\}$	0.016	0.895	1.183	0.754	0.136	0.754	0.472	0.623	0.00062	0.325
$\{\text{reuter}\}$	0.017	0.983	1.074	0.960	0.136	0.96	0.691	0.828	0.00016	0.641

Table 5: Rule orders for the Reuters data.

Interestingness Measure	tau-b
\mathcal{L} and \mathcal{L}^*	-0.264
\mathcal{C} and \mathcal{C}^*	0.023
\mathcal{Q} and \mathcal{Q}^*	0.353
\mathcal{G} and \mathcal{G}^*	0.234

6 Random Transactions

The presence of some associations was a reasonable assumption for the two real data sets investigated thus far. We now consider a simulated data set with completely random transactions of independent items. This investigation is necessary to understand how randomness, such as noise, affects standardizing interestingness measures. We expect that many of the raw interestingness measures that fulfil Prop. 1 of Section 2 will achieve values very near the constant, thus implying independence. For this reason, we expect no pattern for the values of \mathcal{C} and, by extension, \mathcal{C}^* . As in special cases of previous examples, we expect that the standardized counterparts will have no pattern for these rules.

We generate 100,000 random transactions of up to 10,000 items using `random.transactions` from the `arules` package in R. The probability of items appearing in transactions was left at the default value of 0.01. We then mine the rules with support and confidence thresholds each set to 0.0001 to obtain over 630,000 rules.

Figure 5 shows standardized interestingness measures plotted against their raw counterparts with these random transactions. The scale of each of the plots is important to note:




Figure 5: Plots of interestingness measures against their standardized counterparts for the random transactions data. Dark regions indicate a higher density of rules

some only show a small range near zero. The ‘shape’ of the plots are similar for the interestingness measures. In three cases, the transformation is nearly linear and the maximum standardized values of \mathcal{L}^* and \mathcal{C}^* are small. The relationship between \mathcal{Q} and \mathcal{Q}^* is also linear but a larger range is observed, in part because of the constant upperbound, but many of the rules are very close to zero. Contrary to this trend, \mathcal{G} is very near zero but \mathcal{G}^* polarizes to either zero or one. We investigated this further and the values that tended to one are the entire set of negatively correlated rules. All of the positively correlated rules tended to zero. Though we are observing a dichotomy, whether positively or negatively correlated, we see independent transactions tending to a particular standardized value, depending on the case.

Values of tau-b are shown in Table 6. Though one may have expected that rule orders would appear random with random transactions, this was not the case. The standardized rule orders are surprisingly much more similar to the original interestingness measure rankings compared to previous data. The Gini index was reordered the most, likely because negatively correlated rules score higher than independent rules. Also of interest is the generally monotonically increasing relationship between the decile and value of tau-b for all interestingness measures (see Figure 9(c) in Appendix A.5).

Table 6: Rule orders for randomly generated transactions.

Interestingness Measure	tau-b
\mathcal{L} and \mathcal{L}^*	0.678
\mathcal{C} and \mathcal{C}^*	0.777
\mathcal{Q} and \mathcal{Q}^*	0.809
\mathcal{G} and \mathcal{G}^*	0.284

With previous data, we highlighted high values of an interestingness measure. This is not necessary with these data as the transformations appear to be generally monotonic and most rules achieve standardized values very near zero.

7 Conclusion

As we have shown through several examples, standardizing interestingness measures serves several purposes not satisfied by raw interestingness measures alone. We have also investigated the behaviour of standardized interestingness measures over different data sets. In the Belgian accidents data, we observed that high values of one standardized interestingness measure correspond to high values of another standardized interestingness measure. Comparing the relative order of rules of a raw against standardized interestingness measure resulted in rules being reordered to various degrees, depending on the interestingness measures; some interestingness measures reordered rules considerably while others maintained much of the original order.

In the Reuters-21578 data, high values of both the Gini index and standardized Gini index related to relatively high values of the standardized interestingness measures and some of the raw interestingness measures. Exceptions coincide with the natural Gini index scoring high for very negatively correlated data as well as very positively correlated data. The relative order of rules also changed more dramatically with these data, such as when standardizing lift and the cosine similarity. We once again observed greater similarities between standardized interestingness measures than between raw interestingness measures. We also highlighted an example where an interestingness measure took on one of the smallest possible values in its range. The value of the interestingness measure indicated that the antecedent and consequent were very close to being independent, but the standardized interestingness measure was quite high indicating that the value obtained was nearly as far from independence as bounds permit.

In the last example, we observed how standardized interestingness measures perform on data designed to have no pattern. Although three interestingness measures indicated some form of relationship, standardizing interestingness measures always resulted in values near zero. Only in the case of \mathcal{Q}^* were values above 0.1 observed. The order of rules by raw interestingness measures was generally similar to the corresponding standardized interestingness measure. Standardizing interestingness measures showed how data with no patterns do not result in an interestingness measure that is misleadingly high whereas the raw interestingness measure alone may have indicated patterns. Standardized interestingness measures often appear to capture the same aspect of a rule by providing similar values; plots in Appendix A.4 show how standardized interestingness measures are more similar to each other than raw interestingness measures. In the cases of \mathcal{L}^* and \mathcal{Q}^* , independent rules do not evaluate to a single value but \mathcal{G}^* maps independent rules to zero akin to its raw counterpart; this is not surprising as this is the lower bound of the interestingness measure. Random transactions, on the other hand, showed standardized interestingness measures of rules near zero in most cases.

One may further study the effects of standardizing interestingness measures on specific properties by considering two families of interestingness measures: one containing interestingness measures possessing the property and the other containing interestingness measures that do not possess the property. A thorough evaluation of several data sets with known patterns and features with members from each of the two families can address more specific aspects of the impact of standardizing interestingness rules.

Acknowledgements

This work was supported by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada and by an Early Researcher Award from the Ontario Ministry of Research and Innovation (McNicholas).

References

- Agrawal R, Srikant R (1994) Fast algorithms for mining association rules in large databases. In: 20th International Conference on Very Large Data Bases, Santiago, Chile, pp 487–489
- Agrawal R, Imielinski T, Swami A (1993) Mining association rules between sets of items in large databases. In: ACM SIGMOD International Conference on Management of Data, Washington DC, USA, pp 207–216
- Agresti A (2002) *Categorical Data Analysis*, 2nd edn. John Wiley & Sons, New York, USA
- Blaikie N (2003) *Analyzing Quantitative Data: From Description to Explanation*. Sage Publications Limited, London
- Brin S, Motwani R, Ullman JD, Tsur S (1997) Dynamic itemset counting and implication rules for market basket data. In: SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data, Tucson, AZ, pp 255–264
- Freitas AA (1999) On rule interestingness measures. *Knowledge-Based Systems* 12(5-6):309–315
- Geng L, Hamilton H (2006) Interestingness measures for data mining: A survey. *ACM Computing Surveys (CSUR)* 38(3):9
- Geurts K, Wets G, Brijs T, Vanhoof K (2003) Profiling high frequency accident locations using association rules. In: Proceedings of the 82nd Annual Transportation Research Board, Washington D.C., pp 123–130
- Goodman LA, Kruskal WH (1954) Measures of association for cross classifications. *Journal of the American Statistical Association* 49(268):734–764
- Han J, Pei J, Yin Y (2000) Mining frequent patterns without candidate generation. *SIGMOD Rec* 29:1–12
- Houtsma M, Swami A (1995) Set-oriented mining for association rules in relational databases. In: Proceedings of the Eleventh International Conference on Data Engineering, Taipei, pp 25–33
- Lang DT (2010) Rstem: Interface to Snowball implementation of Porter’s word stemming algorithm. R package version 0.4-0
- Lewis DD (1996) Reuters-21578 Distribution 1.0. URL www.daviddlewis.com/resources/testcollections/reuters21578, README file version 1.3 dated May 2004
- McNicholas PD, Murphy TB, O’Regan M (2008) Standardising the lift of an association rule. *Computational Statistics and Data Analysis* 52:4712–4721

- Park JS, Chen MS, Yu PS (1995) An effective hash-based algorithm for mining association rules. *SIGMOD Rec* 24:175–186
- Piatetsky-Shapiro G (1991) Discovery, analysis, and presentation of strong rules. *Knowledge Discovery in Databases* pp 229–238
- R Development Core Team (2012) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0
- Tan PN, Kumar V, Srivastava J (2002) Selecting the right interestingness measure for association patterns. In: *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Edmonton, Alberta, pp 32–41
- Yule GU (1903) Notes on the theory of association of attributes in statistics. *Biometrika* 2(2):121–134
- Zaki MJ (2000) Scalable algorithms for association mining. *IEEE Transactions on Knowledge and Data Engineering* 12:372–390

A Appendix

Before deriving any inequalities, recall the following notation:

- σ is the minimum support threshold. When not explicitly given, $\sigma = 1/n$.
- κ is the minimum confidence threshold. When not explicitly given, $\kappa = 1/n$.
- $m = \min \{P(A), P(B)\}$
- $M = \max \{P(A), P(B)\}$

A.1 Cosine Similarity

The lower bounds of \mathcal{C} are obtained by considering several cases. First, because $P(A, B) \geq \sigma$,

$$\mathcal{C} \geq \frac{\sigma}{\sqrt{P(A)P(B)}}.$$

Because $P(A, B) \leq m$,

$$\mathcal{C} = \frac{P(A, B)}{\sqrt{P(A)P(B)}} = \frac{P(A, B)}{\sqrt{mM}} \leq \frac{m}{\sqrt{mM}} = \sqrt{\frac{m}{M}}.$$

We obtain another by considering that $P(A \cup B) = P(A) + P(B) - P(A, B) \leq 1$, and so $P(A, B) \geq P(A) + P(B) - 1$, giving

$$\mathcal{C} \geq \frac{P(A) + P(B) - 1}{\sqrt{P(A)P(B)}}.$$

By considering support threshold σ , if $P(A, B) \geq \sigma$, the cosine similarity is smallest if $P(A) = P(B) = (1 - \sigma)/2$, giving

$$\mathcal{C} \geq \frac{2\sigma}{1 + \sigma}.$$

If a confidence threshold κ is employed, then $P(A, B)/P(A) \geq \kappa$ and it follows that

$$\mathcal{C} \geq \frac{P(A, B)}{\sqrt{\frac{P(A, B)}{\kappa} P(B)}} = \sqrt{\kappa \frac{P(A, B)}{P(B)}},$$

and because $P(A, B) \geq \sigma$, we have

$$\mathcal{C} \geq \sqrt{\kappa \frac{\sigma}{P(B)}}.$$

Also, because $P(A, B) \geq \kappa P(A)$,

$$\mathcal{C} \geq \frac{\kappa P(A)}{\sqrt{P(A)P(B)}} = \kappa \sqrt{\frac{P(A)}{P(B)}}.$$

Upper bounds of \mathcal{C} come from considering that $P(A, B) \leq m$ and $P(A)P(B) = Mm$, giving

$$\mathcal{C} \leq \sqrt{\frac{m}{M}}.$$

A.2 Yule's Q

Yule's Q is given by

$$\begin{aligned} \mathcal{Q} &= \frac{P(A, B)P(\bar{A}, \bar{B}) - P(A\bar{B})P(\bar{A}B)}{P(A, B)P(\bar{A}, \bar{B}) + P(A\bar{B})P(\bar{A}B)} \\ &= \frac{P(A, B) [1 - P(A) - P(B) + P(A, B)] - [P(A) - P(A, B)] [P(B) - P(A, B)]}{P(A, B) [1 - P(A) - P(B) + P(A, B)] + [P(A) - P(A, B)] [P(B) - P(A, B)]} \\ &= \frac{P(A, B) - P(A)P(B)}{P(A, B) + P(A)P(B) - 2P(A, B)(P(A) + P(B) - P(A, B))}. \end{aligned} \tag{6}$$

Yule's Q achieves the maximum value of 1 iff $P(A, B) = \min\{P(A), P(B)\}$ and achieves the minimum value of -1 only when $P(A, B) = 0$ or when $P(\overline{A}\overline{B}) = 0$. It is a strictly monotonically increasing function of $P(A, B)$ so only these endpoints will be considered.

Given the minimum support threshold σ , $P(A, B) \geq \sigma$ so the minimum value of -1 is not attained. If, however, $P(A, B) = \sigma$, then Yule's Q attains the value

$$\frac{\sigma - P(A)P(B)}{\sigma + P(A)P(B) - 2\sigma(P(A) + P(B) - \sigma)}. \quad (7)$$

Given the minimum confidence threshold κ , because $P(A, B) = P(B | A)P(A)$, we know that $P(A, B) > \kappa P(A)$ and so Yule's Q is greater than the value

$$\begin{aligned} & \frac{\kappa P(A) - P(A)P(B)}{\kappa P(A) + P(A)P(B) - 2\kappa P(A)(P(A) + P(B) - \kappa P(A))} \\ &= \frac{\kappa - P(B)}{\kappa + P(B) - 2\kappa(P(A) + P(B) - \kappa P(A))}. \end{aligned} \quad (8)$$

From (7) and (8), the lower bound of Yule's Q is given by

$$\lambda = \max \left\{ -1, \frac{\sigma - P(A)P(B)}{\sigma + P(A)P(B) - 2\sigma(P(A) + P(B) - \sigma)}, \frac{\kappa - P(B)}{\kappa + P(B) - 2\kappa(P(A) + P(B) - \kappa P(A))} \right\}.$$

A.3 Gini

Derivation of an alternative form of the Gini index.

$$\begin{aligned} \mathcal{G} &= P(A) [P(B | A)^2 + P(\overline{B} | A)^2] + P(\overline{A}) [P(B | \overline{A})^2 + P(\overline{B} | \overline{A})^2] - P(B)^2 - P(\overline{B})^2 \\ &= P(A) \left[\frac{P(A, B)^2}{P(A)^2} + \frac{P(A, \overline{B})^2}{P(A)^2} \right] + P(\overline{A}) \left[\frac{P(\overline{A}, B)^2}{P(\overline{A})^2} + \frac{P(\overline{A}, \overline{B})^2}{P(\overline{A})^2} \right] - P(B)^2 - (1 - P(B))^2 \\ &= \left[\frac{P(A, B)^2 + P(A, \overline{B})^2}{P(A)} \right] + \left[\frac{P(\overline{A}, B)^2 + P(\overline{A}, \overline{B})^2}{P(\overline{A})} \right] - 2P(B)^2 + 2P(B) - 1 \\ &= -2P(B)^2 + 2P(B) - 2 \frac{P(A, B)[P(A) - P(A, B)]}{P(A)} - 2 \frac{P(\overline{A}, B)[P(\overline{A}) - P(\overline{A}, B)]}{P(\overline{A})} \\ &= -2P(B)^2 + 2P(B) - 2(P(B)) + 2 \left[\frac{P(A, B)^2}{P(A)} + \frac{[P(B) - P(A, B)]^2}{1 - P(A)} \right] \\ &= \frac{-2P(B)^2 P(A)(1 - P(A))}{P(A)(1 - P(A))} + 2 \left[\frac{P(A, B)^2 + P(A)P(B)^2 - 2P(A)P(B)P(A, B)}{P(A)(1 - P(A))} \right] \\ &= 2 \frac{[P(A, B) - P(A)P(B)]^2}{P(A)(1 - P(A))} \end{aligned} \quad (9)$$

$$= \frac{2}{1 - P(A)} [P(B | A) - P(B)][P(A, B) - P(A)P(B)] \quad (10)$$

Using (9), we know immediately that the (global) lowerbound of zero is achieved iff A and B are independent. When A and B are not independent, the Gini index is non-zero.

Using Expression (9), we can write the Gini index as the product

$$\mathcal{G} = \frac{2}{P(A)(1 - P(A))} \times \left(P(A, B) - P(A)P(B) \right)^2.$$

The first factor is largest when $P(A) = \frac{1}{2}$, making it smaller than 8. The second term can be no larger than $\frac{1}{16}$, precisely when A and B are most positively correlated. Let $z = P(A) = P(B) = P(A, B)$ in this instance,

$$(P(A, B) - P(A)P(B))^2 \leq (z - z^2)^2 \leq \left(\frac{1}{2} - \left(\frac{1}{2} \right)^2 \right)^2 = \frac{1}{16}.$$

From these deductions, we know that

$$\mathcal{G} \leq 8 \times \frac{1}{16} = \frac{1}{2}. \quad (11)$$

We may obtain other inequalities by considering the highest and lowest values of $P(A, B)$ influences the numerator of (9). Let $l = \max\{\sigma, P(A) + P(B) - 1, \kappa P(A)\}$. From Section A.1, $l \leq P(A, B)$ and $P(A, B) \leq m$, as before.

Due to the quadratic nature of the numerator, we must be slightly more cautious. When $P(A, B) - P(A)P(B) \geq 0$,

$$2 \frac{(l - P(A)P(B))^2}{P(A)(1 - P(A))} \leq \mathcal{G}. \quad (12)$$

When $P(A, B) - P(A)P(B) > 0$, the above inequalities are reversed to obtain

$$2 \frac{(m - P(A)P(B))^2}{P(A)(1 - P(A))} \leq \mathcal{G}. \quad (13)$$


When $P(A, B) < P(A)P(B)$, we can also employ $P(A, B) \geq \kappa P(A)$ to obtain

$$\mathcal{G} \geq 2 \frac{(P(A, B) - P(A)P(B))(\kappa - P(B))}{1 - P(A)},$$


and when $P(A, B) > P(A)P(B)$, we find that

$$\mathcal{G} \geq 2 \frac{(P(A, B) - P(A)P(B))(1 - P(B))}{1 - P(A)}.$$

A.4 Scatterplot Matrices




(a) Raw Interestingness Measures

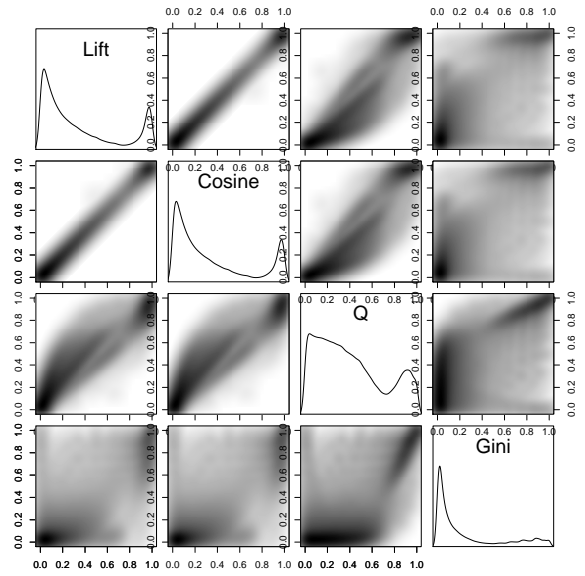


(b) Standardized Interestingness Measures

Figure 6: Scatterplot matrices of interestingness measures for the Belgian accident data.



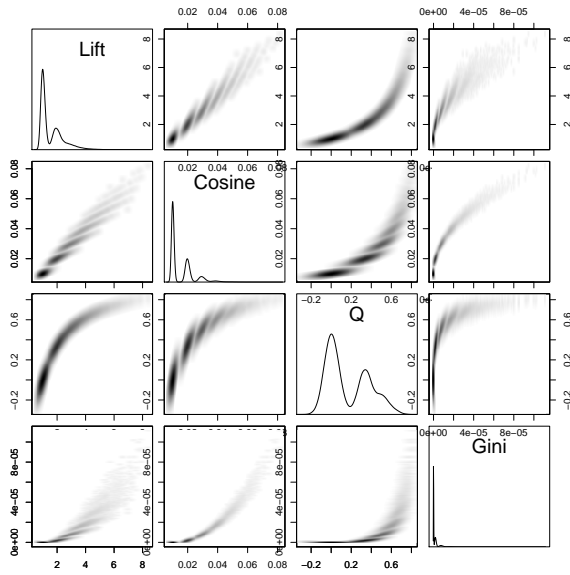
(a) Raw Interestingness Measures




(b) Standardized Interestingness Measures

Figure 7: Scatterplot matrices of interestingness measures for the Reuters data.

A.5 Tau-b by Deciles

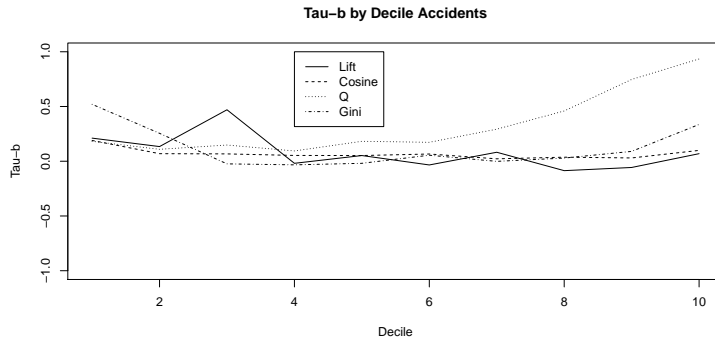


(a) Raw Interestingness Measures

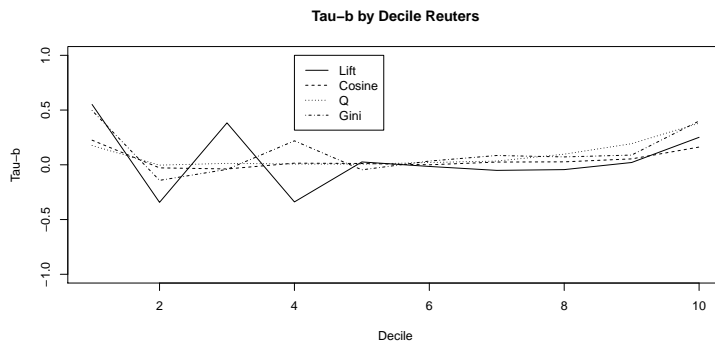


(b) Standardized Interestingness Measures

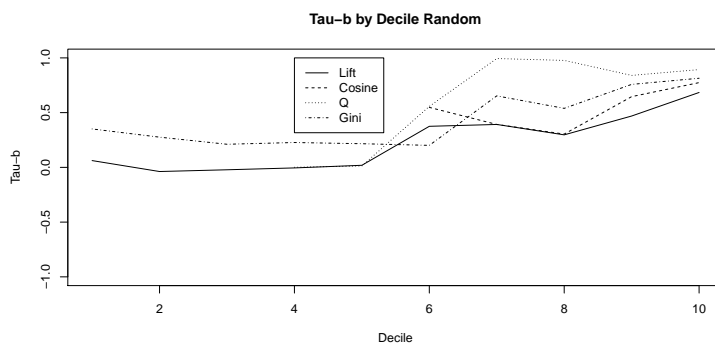
Figure 8: Scatterplot matrices of interestingness measures for random transactions.



(a) Accidents Data



(b) Reuters-21578 Data



(c) Random Data

Figure 9: Tau-b by deciles for each data set.