Knowledge and Information Systems (2020) 62:239–258
https://doi.org/10.1007/s10115-019-01352-3

REGULAR PAPER

Rate of change analysis for interestingness measures

Nandan Sudarsanam[1] · Nishanth Kumar[2] · Abhishek Sharma[3] · Balaraman Ravindran[4]

Received: 18 April 2018 / Revised: 28 February 2019 / Accepted: 8 March 2019 / Published online: 20 March 2019
© Springer-Verlag London Ltd., part of Springer Nature 2019

Abstract
The use of association rule mining techniques in diverse contexts and domains has resulted in the creation of numerous interestingness measures. This, in turn, has motivated researchers to come up with various classification schemes for these measures. One popular approach to classify the objective measures is to assess the set of mathematical properties they satisfy in order to help practitioners select the right measure for a given problem. In this research, we discuss the insufficiency of the existing properties in the literature to capture certain behaviors of interestingness measures. This motivates us to adopt an approach where a measure is described by how it varies if there is a unit change in the frequency count $(f_{11}, f_{10}, f_{01}, f_{00})$, at different preexisting states of the counts. This rate of change analysis is formally defined as the first partial derivative of the measure with respect to the various frequency counts. We use this analysis to define two novel properties, unit-null asymptotic invariance (UNAI) and unit-null zero rate (UNZR). UNAI looks at the asymptotic effect of adding frequency patterns, while UNZR looks at the initial effect of adding frequency patterns when they do not preexist in the dataset. We present a comprehensive analysis of 50 interestingness measures and classify them in accordance with the two properties. We also present multiple empirical studies, involving both synthetic and real-world datasets, which are used to cluster various measures according to the rule ranking patterns of the measures. The study concludes with the observation that classification of measures using the empirical clusters shares significant similarities to the classification of measures done through the properties presented in this research.

Keywords Association rule mining · Objective measures · Properties of measures · Rate of change analysis

1 Introduction

Association rule mining (ARM) has emerged as a powerful and specialized tool to identify patterns in large datasets. It can be used in applications or business operations where instances of some occurrence, typical spatial or temporal, are represented in tabular format across a set of common attributes. An ARM study results in rules of the form $A \Rightarrow B$, which would mean that based on evidence from the data, the presence of itemset A is likely to indicate

Extended author information available on the last page of the article

Table 1 Standard 2 × 2 contingency table representing the frequency counts of A and B

	B	B^c
A	f_{11}	f_{10}
A^c	f_{01}	f_{00}

the presence of itemset B. There are two major challenges to an ARM implementation: (i) Candidate generation: This involves the process of filtering all the possible combinations of items that satisfy a given condition for selection. Given the exponentially large possibilities of rules, this condition focuses on the use of frequency-based thresholds to remove potentially uninteresting rules [1]. The second major challenge is (ii) Candidate evaluation: This involves the use of an appropriate metric (*interestingness measure*) to evaluate all the different rules that can be defined from the selected itemsets [7,19].

While several works have dealt with the first challenge [1,14,16,21,26,27], this research concerns itself with the latter challenge. Candidate evaluation can be challenging because there are different ways of describing interestingness of rules. A recent study [22] showed that even among *objective measures*, there exist more than 61 that are defined in the literature [1,15,17,21,22]. Also, the information derived from these different interestingness measures (IM) may not always be consistent [19].

The properties are typically defined using a contingency table (see Table 1), a simplified adaptation from [19]. Here, two states, present and absent, are defined for two itemsets, A (rows) and B (columns). The frequency counts f_{11} and f_{00} define the co-presence and co-absence of A and B, respectively. While the term f_{10} would represent the presence of A and absence of B, f_{01} the opposite.

In this research, we posit that the popularly used set of 8 properties covered in [20] do not fully capture some important aspects of interestingness measures, and this motivates us to define a more relevant and new property-based analysis of IMs. Specifically, our motivation is built on the observations of [22], who state that the empirical classification of measures based on how they rank rules has little to do with the property-based classification. A deeper study on this mismatch leads us to believe that preexisting mathematical properties are only useful in specific environmental contexts. These observations lead us to devise simpler, more generic property definitions which can be applied to different environmental contexts and bear a stronger affiliation to rule ranking patterns exhibited by the measures on empirical datasets. The major contributions of this research are listed as follows:

- Introduction of two novel properties to classify interestingness measures, UNAI and UNZR, based on the rate of change analysis (RCA) approach. Note that we do not attempt to define a new interestingness measure. Our objective is to meaningfully classify the existing measures to facilitate an appropriate selection of a measure for a given problem.
- An in-depth analysis of the performance of these properties in classifying wide set of popular interestingness measures, as well as a comparison with other properties presented in [19].
- Presenting empirical case studies that provide validation for the findings and also demonstrate the usefulness of the properties using real-world and synthetic datasets.

1.1 An illustrative example for the intuition behind the properties UNAI and UNZR

While we formally introduce the properties in Sect. 3, here we provide a jargon-free motivation for our approach through a simple example. Most of the current properties are defined

Table 2 Contingency tables for different scenarios

	Table				Probabilities		Measure		
	f_{11}	f_{10}	f_{01}	f_{00}	$P(A)$	$P(B)$	Cosine	Lift	PS
Sparse									
S1	1	100	100	10,000	0.010	0.010	0.01	1.00	0.000
S2	1	100	100	11,000	0.009	0.009	0.01	1.10	0.000
S3	1	100	100	12,000	0.008	0.008	0.01	1.20	0.000
Dense									
D1	10,000	100	100	1	0.99	0.99	0.99	1.00	0.000
D2	10,000	100	100	1001	0.90	0.90	0.99	1.10	0.080
D3	10,000	100	100	2001	0.83	0.83	0.99	1.20	0.134

on the basis of how the IMs change when f_{ij}s are increased or decreased, with respect to each other. The idea being that similar behavior on perturbing the f_{ij} count suggests that the IMs possess similar properties. However, the properties in extant literature do not account for the preexisting states of f_{ij}s when the perturbations are made, whereas UNAI and UNZR (the two properties presented in this study) achieve this.

Lets us take the example of two environments. One is a sparse matrix which is typical in a point-of-sale system often discussed in the market basket analysis. Here, the system is marked by pairs of itemsets with high f_{00} count compared to f_{11}, implying co-absence is more likely than co-presence. Second, we take a dense matrix, which could occur in cases where the attributes are easily possessed by the instances. This is common when the number of attributes is small. Here, the system is marked by high f_{11} count compared to f_{00}. In both these environments, we explore the effect of increasing f_{00} in two steps. The effect of these increases in both environments on three popular measures–Cosine [18], Lift [5] and Piatetsky-Shapiro (PS) [15]—is captured in Table 2. The base case and the two increases in the sparse setting are shown as S1, S2 and S3, with equivalents for the dense case.

It can be seen that Cosine and PS behave similarly in the sparse setting, whereas the Lift and PS behave similarly in the dense setting. No preexisting property captures this binning. We can see that a common property, *null invariance* [19], which captures whether a measure is invariant in co-absence of items, can be used to discern between Cosine and the other two measures. While this could be useful in the dense environment, it is misleading in the sparse environment. In the subsequent sections, and Table 3 in particular, we demonstrate that the property $UNAI_{f_{00}}$ captures the difference between Lift and the other two measures that we see in the sparse environment, while $UNZR_{f_{00}}$ classifies all three measures differently (while both Lift and PS change with the increase in density levels, their rate of the change is significantly different and therefore all three measures should be classified separately), which reflects the difference we seen in the dense environment.

The property *unit-null asymptotic invariance* ($UNAI$) studies the effect that an increase of a frequency count f_{ij} has on a measure when the frequency count is already very large (asymptotic effect as the frequency count tends to $+\infty$). The property *unit-null zero rate* ($UNZR$) looks at the effect of increasing the frequency count on the IM when it is currently nonexistent in the dataset (effect when the frequency count is or tends to 0).

The rest of this document is structured as follows: Sect. 2 presents a brief overview of the literature in interestingness measures and properties used in ARM. Section 3 formally defines the mathematical requirements for UNAI and UNZR and also presents the different states that these two properties can take. In Sect. 4, we analyze a set of 50 IMs and compare our properties with 8 other properties. In Sect. 5, we present empirical studies on synthetic and real datasets. We finally summarize and conclude the study in Sect. 6.

2 Related work

Given the abundance of measures and difficulty in choosing the appropriate IM, researchers have suggested various classification schemes (of the IMs) to help identify the appropriate measure for a given application [6,15,19,20,22]. There are two different types of classification that exist in the literature: classification based on the properties of IMs (e.g., [6,15,19,20]) and classification based on empirical results of IMs on different datasets (e.g., [22]).

Research conducted by [15] formalized a framework consisting of three properties that an IM should satisfy, namely: it must take value 0 if the occurrences of itemsets are independent (P1); the measure must be monotonically increasing with the co-presence of itemsets (P2); and the measure should be monotonically decreasing with the occurrences of either itemsets (P3).

The authors in [19] proposed the following 5 properties in addition to the 3 proposed by [15]: symmetry under variable permutation (O1), row/column scaling invariance (O2), antisymmetry under row/column permutation (O3), inversion invariance (O4) and null invariance (O5). They conducted a comparative study, testing 21 different IMs against the resulting 8 properties. The authors further proposed that the optimal way of finding a suitable IM would be to let the user define a property vector indicating the properties that would be ideally required for the given application. This property vector would then be compared to the property vectors of the different objective measure to pick out the ideal IM for that particular case. For instance, the null-invariance property is considered to be important for IMs used in the context of small probability events in a large dataset [25]. While there has been further work in introducing new properties (e.g., [2,4–6,8,11]), these have not been as commonly used or cited as the work of [15,19]. The use of partial derivatives of frequency counts to understand the behavior of IMs is proposed in the literature [10,12,23]. These properties look at the impact on IM with the addition of counterexamples to rules. The first property is satisfied if IMs should be decreasing functions of the number of counterexamples of the rules, while the second property, and the more rigorous one, addresses the shape of the IM curve when the first counterexamples appear. The properties presented in [11] also categorize IMs based on their sensitivity to marginal frequency counts of counterexamples. These results are used to propose a multi-criteria aid approach assessing the issue of selecting an IM adapted to the users context.

We provide a significant extension of these properties in our paper. The properties presented in this paper provide a more detailed assessment of the behavior of IMs with respect to changes to as they differentiate between the types of counterexamples and behavior of IM at very low frequency counts.

In comparison with property-based classifications, there has been limited work on classification of IMs based on empirical results from different datasets. Research by [9] proposed the classification of 35 IMs based on their empirical performance on 2 different datasets by studying the correlation of IMs. These were classified using a graph-based clustering approach to create high-correlation and low-correlation graphs.

The work of [22] performed a comprehensive classification of 61 different objective IMs based on empirical results with 110 datasets. It suggested that there exist 21 clusters of measures which are distinct and each of these were studied in detail. The work by [13,24] follows a similar approach to this paper, where 20 IMs are empirically classified using 8 different properties across 10 data-sets. The study identifies three main groups of IMs in the two approaches, which may be refined in five smaller classes.

3 Mathematical definitions for properties UNAI and UNZR

An interestingness measure (IM) can be represented as a function of the frequency counts (see Eq. 1). RCA analysis seeks to assess the relative change in the interestingness measure per unit change of the frequency counts. This is essentially the first partial derivative of the interestingness measure with respect to the variables representing the counts, as shown in Eq. 2. The set of formulas representing the first partial derivative of the interestingness measure with respect to each of the four state variables f_{11}, f_{00}, f_{10} and f_{01} represents the RCA analysis as shown in Eq. 3.

$$\mathrm{IM} = \phi(f_{11}, f_{10}, f_{01}, f_{00}) \tag{1}$$

$$\phi'_{f_{ij}} = \frac{\partial(\mathrm{IM})}{\partial f_{ij}} \tag{2}$$

$$\mathrm{RCA}\ (\mathrm{IM}) = \left\{\phi'_{f_{11}}, \phi'_{f_{10}}, \phi'_{f_{01}}, \phi'_{f_{00}}\right\} \tag{3}$$

$$\mathrm{UNAI}_{ij} = \lim_{f_{ij} \longrightarrow +\infty} (\phi'_{f_{ij}}) \tag{4}$$

$$\mathrm{UNZR}_{ij} = \lim_{f_{ij} \longrightarrow 0} (\phi'_{f_{ij}}) \tag{5}$$

We use the RCA analysis to define two novel properties: the *unit-null asymptotic invariance* (UNAI) and the *unit-null zero rate* (UNZR). Mathematically, both these properties are the *derivative at a point* or the *instantaneous rate of change*, at two specific points. We can define the property unit-null asymptotic invariance (UNAI) as the derivative of the interestingness measure (IM) with respect to f_{ij} as $f_{ij} \rightarrow \infty$, and this instantaneous rate of change can be written as shown in Eq. 4. UNAI can be defined for each of the four frequency count variables by substituting ij with the count of interest. Similar to UNAI, UNZR can be captured by looking at the instantaneous rate of change at 0. Formally, this would be the derivative of the interestingness measure (IM) with respect to f_{ij} as $f_{ij} \rightarrow 0$, and this instantaneous rate of change can be written as shown in Eq. 5. To compute, UNAIs and UNZRs, in some cases we can simply take the first partial derivative and directly substitute the point of interest and in other scenarios, we use the limit notation for derivative at a point (also shown in Eqs. 4 and 5). Having defined the framework for computing the satisfaction of UNAIs and UNZRs, in the subsequent sections we define the conditions where an interestingness measure can be said to satisfy these properties. These sections present a classification scheme for the properties UNAI and UNZR which are presented at the individual f_{ij} level as well as the metric as a whole.

3.1 UNAI property definition

We create a two-pronged classification scheme for UNAI. We define $UNAI_{f_{ij}}$ which is $UNAI$ defined for each frequency count (f_{11}, f_{10}, f_{01}, f_{00}). We do this explicitly for f_{11}

which can then be extended to the other frequency counts. The intuition behind these definitions is discussed in 1.1.

1. $UNAI_{f_{11}}$ is satisfied when: $\lim_{f_{11}\to+\infty}(\phi'_{f_{11}}) = 0$, for all feasible combination of values of f_{00}, f_{10}, and f_{01}. We define a *feasible combination* of values as ones which enable the calculation of the metric in deterministic forms for a database with nonzero rows.
   By extension, we can say that the $UNAI_{f_{11}}$ condition is not met when $\lim_{f_{11}\to+\infty}(\phi'_{f_{11}})$ $\neq 0$, for any feasible combination of values of f_{00}, f_{10}, and f_{01}.
   Similarly, we can define $UNAI_{f_{ij}}$ for the other three frequency counts by swapping the variables accordingly.
2. $UNAI$ is satisfied when $UNAI_{f_{ij}}$ is satisfied $\forall(ij)$. This is essentially an extension of the classification from $UNAI_{f_{ij}}$ to a general property for the metric as a whole.

In our example in Sect. 1.1, the $UNAI_{f00}$ would have been satisfied for Cosine and PS, whereas it would have not been satisfied for Lift. This would have been the useful classification in the sparse setting.

3.2 UNZR property definition

The classification scheme we adopt for UNZR is more complex than UNAI. Similar to UNAI, we adopt a two-pronged approach of defining $UNZR$ at the f_{ij} level as well as a defining it for the metric as a whole. However, we differ from $UNAI$ in that $UNZR$ states are not binary, but have three states that correspond to the property being satisfied, partially satisfied and not satisfied. Another aspect of the difference is that the definitions at the f_{ij} level are different for $\{f_{11}, f_{00}\}$ and $\{f_{10}, f_{01}\}$. As shown below, they are symmetrical opposites in terms of the inequality conditions. We formally defined the property for f_{11} and f_{10} below and extend it to the other frequency counts f_{00} and f_{01}, respectively. The intuition behind these definitions is discussed in 1.1.

1. $UNZR_{f_{11}}$ is satisfied when $\lim_{f_{11}\to 0}(\phi'_{f_{11}}) > 0$ for all feasible combinations of f_{00}, f_{10}, and f_{01}. Again, a *feasible combination* is one that enables the computation of the metric in deterministic forms. This formulation can be extended to $UNZR_{f_{00}}$ by swapping the variables accordingly.
   $UNZR_{f_{10}}$ is satisfied when $\lim_{f_{10}\to 0}(\phi'_{f_{10}}) < 0$ for all feasible combinations of f_{11}, f_{00}, and f_{01}. This formulation can be extended to $UNZR_{f_{01}}$ by swapping the variables accordingly.
2. $UNZR_{f_{11}}$ is partially satisfied when two conditions are met. These are: (i) $\lim_{f_{11}\to 0}(\phi'_{f_{11}})$ ≥ 0 for all feasible combinations of f_{00}, f_{10}, and f_{01}, and (ii) $\lim_{f_{11}\to 0}(\phi'_{f_{11}}) > 0$ for at least one or more feasible combinations of f_{00}, f_{10}, and f_{01}. This formulation can be extended to $UNZR_{f_{00}}$ by swapping the variables accordingly.
   Similarly, $UNZR_{f_{10}}$ is partially satisfied when two conditions are met. These are: (i) $\lim_{f_{10}\to 0}(\phi'_{f_{10}}) \leq 0$ for all feasible combinations of f_{11}, f_{00}, and f_{01}, and (ii) $\lim_{f_{10}\to 0}(\phi'_{f_{10}}) < 0$ for at least one or more feasible combinations of f_{11}, f_{00}, and f_{01}. This formulation can be extended to $UNZR_{f_{01}}$ by swapping the variables accordingly.
3. Finally, by extension, we can say that $UNZR_{f_{11}}$ is not satisfied when either of these two conditions are met: (i) $\lim_{f_{11}\to 0}(\phi'_{f_{11}}) < 0$ for any feasible combination of f_{00}, f_{10}, and f_{01} or, (ii) $\lim_{f_{11}\to 0}(\phi'_{f_{11}}) = 0$ for all feasible combinations of f_{00}, f_{10}, and f_{01}. This formulation can be extended to $UNZR_{f_{00}}$ by swapping the variables accordingly.

Similarly, we can say that $UNZR_{f_{10}}$ is not satisfied when either of these two conditions are met: (i) $\lim_{f_{10}\to 0}(\phi'_{f_{10}}) > 0$ for any feasible combination of f_{11}, f_{00}, and f_{01} or, (ii) $\lim_{f_{10}\to 0}(\phi'_{f_{10}}) = 0$ for all feasible combinations of f_{11}, f_{00}, and f_{01}. This formulation can be extended to $UNZR_{f_{01}}$ by swapping the variables accordingly.

4. At the overall metric level, we say that $UNZR$ property is satisfied for a metric if the $UNZR_{f_{ij}}$ is satisfied $\forall (ij)$. We say that UNZR property is partially satisfied for a metric if $UNZR_{f_{ij}}$ is at least partially satisfied for all f_{ij}s. Finally, a metric fails to satisfy the UNZR property if one or more $UNZR_{f_{ij}}$s do not satisfy the property.

In our example in Sect. 1.1, the $UNZR_{f00}$ would have been fully satisfied for PS, partially satisfied for Lift, and not been satisfied for Cosine. This is a meaningful classification for the dense setting, since PS changes at a significant rate (given that the range for this measure is -0.25 to $+0.25$), whereas Lift reflects smaller increases, and Cosine is unaffected.

4 Mapping UNAI and UNZR to commonly used measures and other properties

This section is divided in two parts. The first part performs a detailed analysis that uses the proposed properties to classify commonly used measures. The second part then compares these classifications to the classification done by other popular properties in the literature [20]. This twofold approach is used because it is important to show that a property can actually differentiate between measures (Sect. 4.1), and that it classifies measures in a way that is different from other properties (Sect. 4.2).

4.1 Classification of existing measures using UNAI and UNZR

In this section, we classify 50 common measures across the two properties $UNAI$ and $UNZR$, at both the f_{ij} level as well as the metric level. We use all 21 metrics from [20] and also borrow popular metrics from [22]. We consciously avoid metrics which are mathematically identical as suggested by [22], but choose to have metrics which could still be rank-wise indistinguishable. We do this because practitioners might make sense of an absolute score and the rate at which it increases or decreases. We also avoid metrics which need us to make any a priori assumptions on probability distributions or cannot be abstracted as a function of f_{ij}s. The analysis is carried out in accordance with the definitions in Sect. 3, and findings are summarized in Table 3.

The results on the classification of these measures provide two important insights. First, that $UNAI$ property for the metrics as a whole is satisfied by a majority of the measures (37 of the 50). These numbers are even higher for the individual $UNAI_{f_{ij}}$ (ranging from 45 for f_{11}, 44 for f_{00}, 46 for f_{10} and 45 for f_{01} out of the 50 measures). This suggests that UNAI would be less useful as a tool to eliminate measures that nullify the unstable effect of one frequency count being particularly large. Instead, this property can be useful when due importance needs to be given when a frequency count is expected to be high and continues to grow. A classic scenario would be Lift. In certain contexts, an increase in co-absence in a sparse database should continue to increase the metric value since it makes co-presence even less probabilistic through random chance.

The second insight from the case of $UNZR$ is of a different nature. At the overall metric level, there are only 3 measures that fully satisfy the UNZR property: They are *Novelty*,

Table 3 The UNAI and UNZR properties exhibited by 50 interestingness measures

Measure	$\text{UNAI}_{f_{11}}$	$\text{UNAI}_{f_{00}}$	$\text{UNAI}_{f_{10}}$	$\text{UNAI}_{f_{01}}$	UNAI	$\text{UNZR}_{f_{11}}$	$\text{UNZR}_{f_{00}}$	$\text{UNZR}_{f_{10}}$	$\text{UNZR}_{f_{01}}$	UNZR
Lift	Y	N	Y	Y	N	Y	P	P	P	P
Jaccard	Y	Y	Y	Y	Y	Y	N	P	P	N
Confidence	Y	Y	Y	Y	Y	Y	N	Y	N	N
Recall	Y	Y	Y	Y	Y	Y	N	N	Y	N
Specificity	Y	Y	Y	Y	Y	N	Y	N	Y	N
Precision	Y	Y	Y	Y	Y	Y	N	Y	N	N
Ganascia	Y	Y	Y	Y	Y	Y	N	Y	N	N
Kulczynski-1	N	Y	Y	Y	N	Y	N	P	P	N
F-Measure	Y	Y	Y	Y	Y	Y	N	P	P	N
Causal confidence	Y	Y	Y	Y	Y	Y	Y	Y	N	N
Odd’s ratio	N	N	Y	Y	N	P	P	P	P	P
Negative reliability	Y	Y	Y	Y	Y	N	Y	N	Y	N
Sebag Schoenauer	N	Y	Y	Y	N	Y	N	P	N	N
Accuracy	Y	Y	Y	Y	Y	P	P	P	P	P
Support	Y	Y	Y	Y	Y	Y	N	P	P	N
Coverage	Y	Y	Y	Y	Y	P	N	N	P	N
Prevalence	Y	Y	Y	Y	Y	P	N	P	N	N
Relative risk	Y	N	Y	Y	N	Y	P	Y	P	P
Novelty	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Yule’s Q	Y	Y	Y	Y	Y	P	P	P	P	P
Yule’s Y	Y	Y	Y	Y	Y	P	P	P	P	P
Cosine	Y	Y	Y	Y	Y	Y	N	Y	Y	N
Least contradiction	Y	Y	N	Y	N	Y	N	Y	N	N
Odd multiplier	Y	N	Y	Y	N	Y	P	P	Y	P
Descriptive confirm	Y	Y	Y	Y	Y	Y	N	Y	N	N
Causal confirm	Y	Y	Y	Y	Y	Y	Y	Y	N	N

Table 3 continued

Measure	$\mathrm{UNAI}_{f_{11}}$	$\mathrm{UNAI}_{f_{00}}$	$\mathrm{UNAI}_{f_{10}}$	$\mathrm{UNAI}_{f_{01}}$	UNAI	$\mathrm{UNZR}_{f_{11}}$	$\mathrm{UNZR}_{f_{00}}$	$\mathrm{UNZR}_{f_{10}}$	$\mathrm{UNZR}_{f_{01}}$	UNZR
Certainty factor	Y	Y	Y	N	N	P	P	N	Y	N
Conviction	Y	Y	Y	Y	Y	P	P	P	Y	P
Informational gain	Y	Y	Y	Y	Y	Y	Y	P	P	P
Laplace	Y	Y	Y	Y	Y	Y	N	Y	N	N
Klosgen	Y	Y	Y	Y	Y	P	N	N	N	N
Piatetsky-Shapiro	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Zhang	Y	N	Y	N	N	Y	P	Y	P	P
Y and L's 1-way support[a]	Y	Y	Y	Y	Y	N	P	N	P	N
Y and L's 2-way support[a]	Y	Y	Y	Y	Y	N	P	Y	Y	N
Implication Index	Y	Y	Y	Y	Y	N	N	N	N	N
Leverage	Y	Y	Y	Y	Y	Y	P	Y	N	N
Kappa	Y	Y	Y	Y	Y	P	P	Y	Y	P
Causal confirm confidence	Y	Y	Y	Y	Y	Y	Y	Y	N	N
Examples and counterexamples	Y	Y	N	Y	N	P	N	Y	N	N
Putative casual dependency	Y	Y	Y	Y	Y	P	P	Y	Y	P
Dependency	Y	Y	Y	Y	Y	P	P	P	P	P
J-measure	Y	Y	Y	Y	Y	N	N	Y	N	N
Collective strength	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Gini Index	Y	Y	Y	Y	Y	N	N	P	P	N
Goodman–Kruskal	N	N	N	N	N	N	N	N	N	N
Mutual information	Y	Y	Y	Y	Y	N	N	Y	Y	N
Normalized mutual information	Y	Y	Y	Y	Y	N	N	N	N	N
Loevinger	Y	Y	Y	N	N	P	P	N	Y	N
Added value	N	Y	N	N	N	P	P	P	P	P

Y: indicates that the property is satisfied, P: indicates that the property is partially satisfied, and N: indicates that the property is not satisfied
[a]These metric names are shortened to fit into the table: Y and L's stand for Yao and Liu's for both the shortened names

Piatetsky-Shapiro and *Collective Strength*. Of the remaining, 14 measures partially satisfy the property and 33 fail to satisfy the property. For each f_{ij}, the UNZR measures are more discerning. In the case of f_{11}, 25 satisfy the property, 9 for f_{00}, 22 for f_{10} and 15 for f_{01}. These suggest that UNZR at the f_{ij} level could be more meaningfully used to pick metrics, especially for the case of f_{00}, which is satisfied by only nine measures. A particular case could be when the practitioner expects an f_{ij} to be low or close to zero and would like to see the metric impacted when presented with evidence of it. The use of $UNZR$ at the overall metric level could also be useful if the practitioner suspects that any of the frequency values can be close to zero but would like to see its presence or absence to have a meaningful impact on the metric.

4.2 Comparing the UNAI and UNZR mapping with other properties

In this section, we compare the classification of measures done through $UNZR$ and $UNAI$, with the classification done through other properties in the literature [20]. This is important because, in addition to fulfilling other criteria, it is necessary that a property classifies measures differently from other preexisting properties. Otherwise, there is a redundancy, and one could question the need for the new property in question. We conduct our comparison on the properties proposed by [20]. This includes five new properties proposed in that study, as well as three previous properties from [15]. In order to perform the analysis, we take all the 50 measures analyzed in Table 3 which include the 21 measures analyzed by [20]. We conduct an analysis that compares the classification of these measures across the two states of $UNAI$ and three states of $UNZR$ and compare it to the two states (satisfied or not satisfied) across the 8 properties presented in [20]. This leads us to create the contingency Table 4.

The findings from Table 4 suggest that the classification of measures through $UNAI$ and $UNZR$ is more or less independent of the classification done through all of the eight preexisting properties. The few cases where we see low overlaps are also easily explainable by the low membership to a certain class and not a relationship between properties (for instance, observe that only 3 of the 50 measures satisfy the 'row and column scaling invariance' or fully satisfy UNZR).

5 Empirical studies

The work of [22] has established that empirical clustering of measures bears no meaningful relationship to properties presented in [20] (which also cover three properties originally presented in [15]). While the properties UNAI and UNZR have been constructed to intuitively convey a certain mathematical aspect of the measure, an important motivation and therefore requirement in design were that they have a meaningful map to the actual behavior of measures, empirically. Our studies across a wide range of datasets, both synthetic and real, suggest that these two properties bear strong relationships with the empirical clusters. More interestingly, we find that the results are substantially more pronounced in certain environmental conditions. Specifically, we find that $UNZR_{f_{11}}$ and $UNAI_{f_{00}}$ are valuable in sparse datasets, and correspondingly, $UNZR_{f_{00}}$ and $UNAI_{f_{11}}$ are better properties to consider in dense data. In the following sections, we do an illustrative analysis showing how the $UNZR_{f_{11}}$ classification of measures is useful in sparse datasets and $UNZR_{f_{00}}$ is useful in dense datasets. The motivation to choose the $UNZR$ properties over the $UNAI$ is the fact that the $UNZR$ creates groups of more or less equal sizes. For instance, $UNZR_{f_{11}}$ splits the

Table 4 Contingency table: relationship between classification of UNAI and UNZR and the classification of other prominent properties

		UNAI		UNZR		
		Satisfied	Not satisfied	Satisfied	Partially satisfied	Not satisfied
P1: Statistical independence	Satisfied	15	4	2	8	9
	Not satisfied	22	9	1	6	24
P2: (Refer [15])	Satisfied	34	13	3	14	30
	Not satisfied	3	0	0	0	3
P3: (Refer [15])	Satisfied	27	11	3	14	21
	Not satisfied	10	2	0	0	12
O1: Symmetry under variable permutation	Satisfied	13	4	3	7	7
	Not satisfied	24	9	0	7	26
O2: Row and column scaling invariance	Satisfied	2	1	0	3	0
	Not satisfied	35	12	3	11	33
O3: Anti-symmetry row or column permutation	Satisfied	4	0	2	2	0
	Not satisfied	33	13	1	12	33
O3’: Inversion invariance	Satisfied	10	1	3	5	3
	Not satisfied	27	12	0	9	30
O4: Null invariance	Satisfied	8	4	0	0	12
	Not satisfied	29	9	3	14	21

measures with 25 of them satisfying the property, 15 of them partially satisfying it, and 10 of them failing to satisfy the property, whereas with $UNAI_{f_{00}}$ we see that 44 of the 50 measures satisfy this property. A similar comparison exists between $UNZR_{f_{00}}$ and $UNAI_{f_{11}}$.

We conduct our empirical studies by first considering synthetic contingency tables that mimic sparse and dense datasets, and we explore further by choosing 8 real-world datasets to validate our findings. We choose 4 of these datasets with a high number of sparse columns and extract these for analysis in the sparse setting. Similarly, the remaining 4 are used for the dense analysis. Based on the rule ranking of the measures in the two environmental conditions, we then cluster the measures into sets and see how they correlate with the property of interest. The entire analysis works with 50 measures obtained from the literature, and no new measure is introduced.

5.1 Sparse datasets

Sparse datasets are characterized by having a relatively high f_{00} count with respect to f_{11}, primarily, and to a lesser extent f_{10}, and f_{01}. As discussed in the previous section, we choose to analyze the effect of the $UNZR_{f_{11}}$ property in this setting.

We mimic the rules from a synthetic dataset using artificially created sets of rules in form of contingency tables. We do this specifically for the sparse settings. We achieve these environments by assigning low values to f_{11}, high values for f_{00}, while f_{10}, f_{01} fall in between the two extremes. The f_{11}, f_{00}, f_{10} and f_{01} cells of the tables took the values {0, 1, 10, 11}, {1000, 5000, 10,000, 25,000, 50,000, 75,000, 100,000}, {10, 100, 250, 500, 600, 800, 1000} and {10, 100, 250, 500, 600, 800, 1000}, respectively. This resulted in 1372 unique contingency tables, each representing a rule in a sparse dataset.

For the real-world dataset, we chose the fairly popular datasets such as 'Adult,' 'Breast Cancer (wisconsin),' 'Glass Classification' and 'Chess(King Rook vs King Pawn)' from the UCI Machine Learning archive [3].

A detailed discretization and binarization of variables were carried out in conformance to the best practices suggested in [21]. These help us create the transactional tables for each dataset. We confine the analysis to one-to-one rules. We use a basic support-based pruning with a threshold close to 0, in order to get a full enumeration of all one-to-one rules but avoid a variable mapping to itself.

Similar to [22], we choose a subset of the rules to compare. However, given the unique nature of our problem, unlike [22], we do not randomly select the rules. Instead, we choose a subset of rules that are typically encountered in sparse datasets, by selecting top 10% of rules with the highest $\frac{f_{00}}{f_{11}}$, and capping this to 5000 in large datasets.

In the next steps, we follow the same procedure as [22]. Each rule is evaluated using each measure, and a rank ordering of rules is done for each measure. Using Spearman's rank correlation, we create a matrix of pairwise similarities between measures which acts as the adjacency matrix for a complete graph. We perform hierarchical agglomerative clustering on the measures using the similarities. This process naturally creates groups of measures. We explore clusters of size 2, 3 and 4 and showcase the formation that points to the most meaningful results. While there are various other graph clustering algorithms that can be implemented, the simplicity of this approach is appealing.

Our study finds that there is a significant match between the three property states and the clusters that are formed for both the synthetic and real datasets. We split the measures into clusters as described above. The cluster memberships can be found in "Appendix."

Table 5 Empirical analysis—sparse dataset

Dataset	$\mathrm{UNZR}_{f_{11}}$	Synthetic			Adult		Breast cancer			Kr versus KP		Glass			
		A	*B*	*C*	*A*	*B*	*A*	*B*	*C*	*A*	*B*	*A*	*B*	*C*	*D*
Total	50	21	20	9	36	14	6	40	4	46	4	4	13	4	19
N	10	0	4	6	2	8	3	3	4	6	4	2	0	4	4
P	15	4	9	2	12	3	1	14	0	15	0	1	2	0	12
Y	25	17	7	1	22	3	2	23	0	25	0	1	11	0	13
ARI		0.22			0.21		0.19			0.12		0.11			
ARI_{P+Y}		0.18			0.42		0.51			0.43		0.18			

Table 6 Empirical analysis—dense dataset

Dataset	$\mathrm{UNZR}_{f_{00}}$	Synthetic			Mushroom				Spambase			Soybean				Lung cancer			
		A	*B*	*C*	*A*	*B*	*C*	*D*	*A*	*B*	*C*	*A*	*B*	*C*	*D*	*A*	*B*	*C*	*D*
Total	50	24	19	7	23	12	12	3	22	4	24	3	17	28	2	3	19	4	24
N	23	3	15	6	2	7	11	3	17	4	2	1	15	6	1	2	15	4	2
P	18	14	2	3	15	3	0	0	2	0	16	0	1	16	1	1	1	0	16
Y	9	6	1	0	6	2	1	0	3	0	6	2	1	6	0	0	3	0	6
ARI		0.22			0.23				0.33			0.24				0.30			
ARI_{P+Y}		0.39			0.36				0.43			0.32				0.41			

The relationship between empirical cluster memberships and property affiliations is summarized in Table 5. We summarize the cluster relationships with the Adjusted Rand index (ARI) for the three states. We also present ARI_{P+Y} where the partial and complete conformance to the property is grouped as a single class. In certain environments, this shows promising results. In the synthetic dataset, all of the 21 measures of cluster A satisfy $UNZR_{f_{11}}$, either completely of partially. The split is rather more even in cluster *B*, but cluster C is dominated by measures which do not satisfy $UNZR_{f_{11}}$. In the 'Adult' dataset, cluster *A* again overwhelmingly consists of measures which satisfy $UNZR_{f_{11}}$, either partially or completely (34 out of 36), whereas the properties that do not satisfy $UNZR_{f_{11}}$ tend to exist more in cluster *B*. A similar distribution can be seen in the other datasets as well where at least one cluster is a clear partition of measures that do not satisfy $UNZR_{f_{11}}$.

5.2 Dense datasets

We characterize dense dataset as one which has relatively higher f_{11} count compared to f_{00} count, primarily, and to a lesser extent f_{10}, and f_{01}. As discussed earlier, we choose to study the effect of $UNZR_{f_{00}}$ property in this environment. The motivation for using synthetic tables is the same as in the sparse case, and the values are identical but inverted (between f_{11} and f_{00}).

For the real-world dataset, we chose the following datasets: 'Mushroom,' 'Soybean,' 'Lung Cancer' and the 'Spambase' from the UCI Machine Learning archive. The methodology of rule generation was identical to that of the sparse real-world datasets, with the focus to create rules from a dense environment.

The results from this analysis are summarized in Table 6 along with ARI scores as in the sparse setting. In the synthetic dataset, cluster A is populated by measures which satisfy the $UNZR_{f_{00}}$ (21 out of 24), either partially or completely. Clusters B (14 out of 19) and C (6 out of 7) are dominated by measures that do not satisfy $UNZR_{f_{00}}$. In the 'Mushroom' dataset, cluster A again consisted of measures which satisfy $UNZR_{f_{00}}$, either partially or completely (21 out of 23). Cluster B is split between the measures that satisfy $UNZR_{f_{00}}$ and measure that do not (7 N's vs 3 P's and 2 Y's). Clusters C and D overwhelmingly consisted of measures which do not satisfy $UNZR_{f_{00}}$, with only 1 measure satisfying the property among the 15 in both clusters combined. In similar fashion, we see clusters A and C of 'Spambase' dataset representing measures that negate and satisfy $UNZR_{f_{00}}$. In general, it is evident that the clustering holds a clear mapping to the $UNZR_{f_{00}}$ property for the selected rules in a dense setting.

6 Conclusions and future work

This study presents two new properties, UNAI and UNZR, which are based on taking the partial derivative of an IM with respect to a frequency count. UNAI corresponds to the derivative at infinity and UNZR at zero. The study then showcases the classification of a broad set of measures in accordance with these properties and also compares them to the classification done by other properties in the literature. The classifications through these properties are fairly independent of those done by other preexisting properties, suggesting that something new is being captured. Finally, the study showcases the utility of classification through the new properties by conducting empirical analyses on both synthetic and real-world datasets, which relate the rule ranking behavior of the measures with two of the properties proposed. The findings suggest that the rule ranking behavior holds a clear relationship to the classification done by the property.

The promising results shown by these two new properties in aligning with empirical behavior of measures could motivate further extensions in the development of properties that build on this idea. One possible extension is to study the shape of the partial derivative curve (linear, polynomial, etc). Finally, the authors in this study agree with the view put forth in [22] that meaningful classification of measures needs to, also, be driven by similarity (or dissimilarity) in rule ranking that can be seen on empirical datasets. We would like to extend this argument by stating that the value of mathematical properties, derived from principled arguments, can be benchmarked across-the-board in this fashion. (This study performs such an analysis exclusively for the two properties proposed in this study.) This can also be extended beyond interestingness measures in ARM. Binary classification metrics (some of which are included in this analysis like accuracy, recall, etc.) can also be defined by the same contingency table and could therefore lend themselves to a representation and segmentation using a similar analysis.

Acknowledgements This work was supported by a funding from the Robert Bosch Center for Data Science and Artificial Intelligence (RBC-DSAI) at IIT Madras.

Appendix A

Illustrative example of the UNAI and UNZR framework using Lift

In this sections, we consider the behavior of the popular interestingness measure, Lift, under the UNAI and UNZR properties defined in the previous section. Lift is defined as follows:

$$\text{Lift}(L) = \frac{P(A \cap B)}{P(A) \cdot P(B)} = \frac{f_{11}(f_{11} + f_{01} + f_{10} + f_{00})}{(f_{10} + f_{11})(f_{01} + f_{11})} \tag{6}$$

Differentiating w.r.t to f_{11} and simplifying, we get

$$\frac{\partial(L)}{\partial f_{11}} = \frac{2 f_{10} f_{11} f_{01} + f_{10} f_{01}(f_{10} + f_{00} + f_{01}) - f_{11}^2 f_{00}}{(f_{10} + f_{11})^2 (f_{01} + f_{11})^2} \tag{7}$$

We check the UNAI property for Lift by considering the derivative as $f_{11} \rightarrow \infty$

$$L_{f_{11}}(\infty) = \lim_{f_{11} \longrightarrow \infty} \frac{\partial L}{\partial f_{11}} = \lim_{f_{11} \longrightarrow \infty} \frac{2 f_{10} f_{11} f_{01} + f_{10} f_{01}(f_{10} + f_{00} + f_{01}) - f_{11}^2 f_{00}}{(f_{10} + f_{11})^2 (f_{01} + f_{11})^2} \tag{8}$$

After algebraic simplification, we can say that the above function is equal to zero for all feasible combinations of f_{00}, f_{10} and f_{01}. Hence, we can say that Lift satisfies UNAI with respect to f_{11}. Similarly, we check for UNAI property with respect to f_{00}, f_{10}, f_{01}.

$$L_{f_{00}}(\infty) = \lim_{f_{00} \longrightarrow \infty} \frac{\partial L}{\partial f_{00}} = \frac{f_{11}}{(f_{01} + f_{11})(f_{10} + f_{11})} \tag{9}$$

$$L_{f_{10}}(\infty) = \lim_{f_{10} \longrightarrow \infty} \frac{\partial L}{\partial f_{10}} = 0 \tag{10}$$

$$L_{f_{01}}(\infty) = \lim_{f_{01} \longrightarrow \infty} \frac{\partial L}{\partial f_{01}} = 0 \tag{11}$$

Here, it is evident that this function is not equal to 0 for all possible values of f_{11}, f_{10}, f_{01}. Hence, we say that $UNAI_{f_{00}}$ is not satisfied but I w.r.t to $UNAI_{f_{11}}, UNAI_{f_{01}}, UNAI_{f_{10}}$ is satisfied.

We check for the UNZR property for f_{11} by taking the partial derivative at $f_{11} = 0$, we get,

$$L_{f_{11}}(0) = \frac{\partial \text{L}}{\partial f_{11}} |_{f_{11}=0} = \frac{f_{10} + f_{00} + f_{01}}{f_{10} f_{01}} \tag{12}$$

Similarly, taking the derivative with respect to f_{00}, f_{10}, f_{01} at 0, we get

$$L_{f_{00}}(0) = \frac{\partial \text{L}}{\partial f_{00}}\bigg|_{f_{00}=0} = \frac{f_{11}}{(f_{11} + f_{10})(f_{11} + f_{01})} \tag{13}$$

$$L_{f_{10}}(0) = \frac{\partial \text{L}}{\partial f_{10}}\bigg|_{f_{10}=0} = -\frac{(f_{01} + f_{00})}{(f_{11} + f_{01}) f_{11}} \tag{14}$$

$$L_{f_{01}}(0) = \frac{\partial \text{L}}{\partial f_{01}}\bigg|_{f_{01}=0} = -\frac{(f_{10} + f_{00})}{(f_{11} + f_{10}) f_{11}} \tag{15}$$

We see that for all feasible combinations $UNZR_{f_{11}}$, $UNZR_{f_{10}}$ and $UNZR_{f_{01}}$ are satisfied. However, $UNZR_{f_{00}}$ is only partially satisfied. From Eq. 23, we can see that the following conditions are met: (i) For all feasible combinations of $f_{11}, f_{10}, f_{01}, L_{f_{00}}(0) \geq 0$.

This passes the definition of partial satisfaction for UNZR as defined in the paper. At the same time, this does not fully satisfy the $UNZR_{f_{00}}$ property since there are values where it can be 0.[1]

Illustrative example of the UNAI and UNZR framework using Piatetsky-Shapiro

In this section, we consider the behavior of another popular interestingness measure, Piatetsky-Shapiro (PS), under the UNAI and UNZR properties in the same way as in the previous section. PS is defined as follows:

$$\text{Piatetsky-Shapiro (PS)} = N(P(A \cap B) - P(A)P(B)) = f_{11} - \frac{(f_{11}+f_{01})(f_{11}+f_{10})}{f_{11}+f_{01}+f_{10}+f_{00}} \quad (16)$$

Differentiating w.r.t to f_{11} and simplifying, we get

$$\frac{\partial(\text{PS})}{\partial f_{11}} = \frac{(f_{01}+f_{00})(f_{10}+f_{00})}{(f_{11}+f_{01}+f_{10}+f_{00})^2} \quad (17)$$

We check the UNAI property for PS by considering the derivative as $f_{11} \to \infty$

$$\text{PS}_{f_{11}}(\infty) = \lim_{f_{11}\longrightarrow\infty} \frac{\partial \text{PS}}{\partial f_{11}} = \lim_{f_{11}\longrightarrow\infty} \frac{(f_{01}+f_{00})(f_{10}+f_{00})}{(f_{11}+f_{01}+f_{10}+f_{00})^2} \quad (18)$$

After algebraic simplification, we can say that the above function is equal to zero for all feasible combinations of f_{00}, f_{10} and f_{01}. Hence, we can say that Piatetsky-Shapiro satisfies UNAI with respect to f_{11}. Similarly, we check for UNAI property with respect to f_{00}, f_{10}, f_{01}. Hence, we can say that Piatetsky-Shapiro satisfies UNAI with respect to f_{11}. Similarly, we check for UNAI property with respect to f_{00}, f_{10}, f_{01}.

$$\text{PS}_{f_{00}}(\infty) = \lim_{f_{00}\longrightarrow\infty} \frac{\partial \text{PS}}{\partial f_{00}} = 0 \quad (19)$$

$$\text{PS}_{f_{10}}(\infty) = \lim_{f_{10}\longrightarrow\infty} \frac{\partial \text{PS}}{\partial f_{10}} = 0 \quad (20)$$

$$\text{PS}_{f_{01}}(\infty) = \lim_{f_{01}\longrightarrow\infty} \frac{\partial \text{PS}}{\partial f_{01}} = 0 \quad (21)$$

Since Piatetsky-Shapiro satisfies $UNAI_{f_{11}}$, $UNAI_{f_{01}}$, $UNAI_{f_{10}}$ and $UNAI_{f_{00}}$, we say it satisfied $UNAI$.

Now we check for the UNZR property for f_{11} by taking the partial derivative at $f_{11} = 0$, we get,

$$\text{PS}_{f_{11}}(0) = \frac{\partial \text{L}}{\partial f_{11}}|_{f_{11}=0} = \frac{(f_{01}+f_{00})(f_{10}+f_{00})}{(f_{01}+f_{10}+f_{00})^2} \quad (22)$$

Similarly, taking the derivative with respect to f_{00}, f_{10}, f_{01} at 0, we get

$$\text{PS}_{f_{00}}(0) = \frac{\partial \text{PS}}{\partial f_{00}}|_{f_{00}=0} = \frac{(f_{01}+f_{11})(f_{10}+f_{11})}{(f_{11}+f_{01}+f_{10})^2} \quad (23)$$

$$\text{PS}_{f_{10}}(0) = \frac{\partial \text{L}}{\partial f_{10}}|_{f_{10}=0} = -\frac{(f_{11}+f_{01})(f_{01}+f_{00})}{(f_{11}+f_{01}+f_{00})^2} \quad (24)$$

$$\text{PS}_{f_{01}}(0) = \frac{\partial \text{PS}}{\partial f_{01}}|_{f_{01}=0} = -\frac{(f_{11}+f_{10})(f_{10}+f_{00})}{(f_{11}+f_{10}+f_{00})^2} \quad (25)$$

[1] Substitute $f_{11} = 0$, while giving the others positive values.

We see that for all feasible combinations $UNZR_{f_{11}}$, $UNZR_{f_{10}}$ and $UNZR_{f_{01}}$ and $UNZR_{f_{00}}$ are satisfied. Therefore, we can say that Piatetsky-Shapiro satisfies the property $UNZR$ as defined in the paper.

Appendix B: Clustering results from Sect. 5

The cluster memberships resulting from empirical analysis on the synthetic datasets are provided as an example of the details of cluster formation:

Sparse datasets

Synthetic dataset: **Cluster A**: {Recall, Precision, Confidence, Jaccard, F-Measure, Odd's Ratio, Sebag Schoenauer, Support, Lift, Ganascia, Kulczynski-1, Relative Risk, Yule's Q, Yule's Y, Cosine, Odd Multiplier, Information Gain, Laplace, Zhang, Leverage, Examples and Counter Examples}, **Cluster B**: {Specificity, Negative Reliability, Accuracy, Descriptive Confirm, Causal Confirm, Piatetsky-Shapiro, Novelty, Causal Confidence, Certainty Factor, Loevinger, Conviction, Klosgen, 1-Way Support, 2-Way Support, Kappa, Putative Causal Dependency, Causal Confirm Confidence, Added Value, Collective Strength, Dependency}, **Cluster C**: {Mutual Information, Coverage, Prevalence, Least Contradiction, Normalized Mutual Information, Implication Index, Gini Index, Goodman Kruskal, J-Measure}.

Dense datasets

Synthetic dataset: Cluster A: {Recall, Odd's Ratio, Specificity, Negative Reliability, Lift, Coverage, Piatetsky-Shapiro, Novelty, Yule's Q, Yule's Y, Odd Multiplier, Certainty Factor, Loevinger, Conviction, Information Gain, Klosgen, Zhang, 1-Way Support, 2-Way Support, Kappa, Putative Causal Dependency, Added Value, Collective Strength, Dependency}. **Cluster B:** {Precision, Confidence, Jaccard, F-Measure, Sebag Schoenauer, Support, Accuracy, Causal Confidence, Ganascia, Kulczynski-1, Prevalence, Relative Risk, Cosine, Least Contradiction, Descriptive Confirm, Causal Confirm, Laplace, Examples and Counter Examples, Causal Confirm Confidence}. **Cluster C:** {Mutual Information, Normalized Mutual Information, Implication Index, Gini Index, Goodman Kruskal, Leverage, J-Measure}.

References

1. Agrawal R, Imieliński T, Swami A (1993) Mining association rules between sets of items in large databases. In: ACM SIGMOD international conference on management of data, pp 207–216
2. Belohlavek R, Grissa D, Guillaume S, Nguifo EM, Outrata J (2014) Boolean factors as a means of clustering of interestingness measures of association rules. Ann Math Artif Intell 70(1–2):151–184
3. Dua D, Karra Taniskidou E (2017) Uci machine learning repository. University of California, Irvine
4. Freitas AA (1999) On rule interestingness measures. Knowl Based Syst 12(5):309–315
5. Geng L, Hamilton H (2007) Choosing the right lens: finding what is interesting in data mining. In: Quality measures in data mining, pp 3–24
6. Guillaume S, Grissa D, Nguifo EM (2012) Categorization of interestingness measures for knowledge extraction. arXiv preprint arXiv:1206.6741
7. Hájek P, Havel I, Chytil M (1966) The guha method of automatic hypotheses determination. Computing 1(4):293–308
8. Hébert C, Crémilleux B (2007) A unified view of objective interestingness measures. In: Machine learning and data mining in pattern recognition, Springer, Berlin, pp 533–547

9. Huynh HX, Guillet F, Blanchard J, Kuntz P, Briand H, Gras R (2007) A graph-based clustering approach to evaluate interestingness measures: a tool and a comparative study. Qual Meas Data Min 43:25–50
10. Le Bras Y, Meyer P, Lenca P, Lallich S (2010) A robustness measure of association rules. In: Joint European conference on machine learning and knowledge discovery in databases, Springer, pp 227–242
11. Lenca P, Meyer P, Vaillant B, Lallich S (2008) On selecting interestingness measures for association rules: user oriented description and multiple criteria decision aid. Eur J Oper Res 184(2):610–626
12. Lenca P, Vaillant B, Lallich S (2006) On the robustness of association rules. In: 2006 IEEE conference on cybernetics and intelligent systems, IEEE, pp 1–6
13. Lenca P, Vaillant B, Meyer P, Lallich S (2007) Association rule interestingness measures: experimental and theoretical studies. In: Quality measures in data mining, Springer, pp 51–76
14. Liu H, Lin Y, Han J (2011) Methods for mining frequent items in data streams: an overview. Knowl Inf Syst 26(1):1–30
15. Piatetsky-Shapiro G (1991) Discovery, analysis, and presentation of strong rules. In: Piatetsky-Shapiro G, Frawley W (eds) Knowledge discovery in databases. AAAI/MIT, Menlo Park, pp 229–248
16. Salam A, Khayal MSH (2012) Mining top k frequent patterns without minimum support threshold. Knowl Inf Syst 30(1):57–86
17. Steinbach M, Kumar V (2007) Generalizing the notion of confidence. Knowl Inf Syst 12(3):279–299
18. Tan P-N, Kumar V (2000) Interestingness measures for association patterns: a perspective. In: Proceedings of workshop on postprocessing in machine learning and data mining
19. Tan P-N, Kumar V, Srivastava J (2002) Selecting the right interestingness measure for association patterns. In: ACM SIGKDD international conference on knowledge discovery and data mining
20. Tan P-N, Kumar V, Srivastava J (2004) Selecting the right objective measure for association analysis. Inf Syst 29(4):293–313
21. Tan P-N, Steinbach M, Kumar V (2005) Introduction to data mining, 1st edn. Addison-Wesley Longman Publishing Co., Inc, Boston
22. Tew C, Giraud-Carrier C, Tanner K, Burton S (2014) Behavior-based clustering and analysis of interestingness measures for association rule mining. Data Min Knowl Discov 28(4):1004–1045
23. Vaillant B, Lallich S, Lenca P (2006) Modeling of the counter-examples and association rules interestingness measures behavior. In: DMIN, pp 132–137
24. Vaillant B, Lenca P, Lallich S (2004) A clustering of interestingness measures. In: International conference on discovery science, Springer, pp 290–297
25. Wu T, Chen Y, Han J (2007) Association mining in large databases: a re-examination of its measures. In: European conference on principles of data mining and knowledge discovery, pp 621–628
26. Wu X, Kumar V, Quinlan JR, Ghosh J, Yang Q, Motoda H, McLachlan GJ, Ng AFM, Liu B, Yu PS, Zhou Z-H, Steinbach M, Hand DJ, Steinberg D (2007) Top 10 algorithms in data mining. Knowl Inf Syst 14(1):1–37
27. Zhang S, Wu X, Zhang C, Lu J (2008) Computing the minimum-support for mining frequent patterns. Knowl Inf Syst 15(2):233–257

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

Nandan Sudarsanam is a faculty member in the Department of Management Studies and a core member of the Robert Bosch Center for Data Science and Artificial Intelligence (RBC-DSAI) at IIT Madras. He earned his Ph.D. from the Engineering Systems Division at MIT, following which, he worked as a quantitative researcher for five years at a high-frequency algorithmic trading firm in New York. His research and work experience focus on applications of experimentation, machine learning and the abstraction of data to models and algorithms. This spans data and problems across different domains, including but not limited to finance, urban mobility, public policy and criminology. He publishes in premier machine learning conferences like IJCAI and AAAI, as well as various peer-reviewed journals in engineering and applied statistics.

Nishanth Kumar graduated with a master's in applied mathematical sciences from Heriot-Watt University, Edinburgh, and an undergraduate degree in Mathematics from the University of Madras. He works in the field of public policy in India specifically around household finance, social security, data protection and financial inclusion and has published articles in periodicals and working papers in these areas. Nishanth is currently the head of the Social Protection Initiative at Dvara Research, a not-for-profit research and policy advocacy institution based in Chennai, India. His research is primarily related to poverty and the role of state in delivering social protection to individuals and households in India. His research interests also include the designing of better tools for exploratory data analysis especially in the field of social sciences.

Abhishek Sharma is a master's student at College of Information and Computer Sciences, University of Massachusetts, Amherst, and worked as a research assistant at Indian Institute of Technology (IIT), Madras, in 2017. His current research covers the modeling of time series data using deep state space models, and integration of stochastic optimization and reinforcement learning.

Balaraman Ravindran is the head of the Robert Bosch Centre for Data Science and Artificial Intelligence (RBC-DSAI) at IIT Madras and a professor in the Department of Computer Science and Engineering. He received his PhD from the University of Massachusetts, Amherst, and his master's in research degree from Indian Institute of Science, Bangalore. He has nearly two decades of research experience in machine learning and specifically reinforcement learning. He has held visiting positions at the Indian Institute of Science, India, and University of Technology, Sydney. He has published nearly 100 papers in journals and conferences, including premier venues such as ICML, AAAI, IJCAI and AAMAS. He has also co-authored the chapter on reinforcement learning in the Handbook of Neural Computation published by Oxford University Press. He is currently serving on the editorial boards of PLOS One and Frontiers in Big Data as well as the program committees of several premier conferences.

Affiliations

Nandan Sudarsanam[1] · **Nishanth Kumar**[2] · **Abhishek Sharma**[3] · **Balaraman Ravindran**[4]

✉ Nandan Sudarsanam
nandan@iitm.ac.in

Nishanth Kumar
Nishanth.K@IFMR.co.in

Abhishek Sharma
shaabhishek@gmail.com

Balaraman Ravindran
ravi@cse.iitm.ac.in

1 Department of Management Studies, Robert Bosch Center for Data Science and Artificial Intelligence (RBC-DSAI), Indian Institute of Technology Madras, Chennai, India

2 IFMR Finance Foundation, Chennai, India

3 Department of Metallurgical and Materials Engineering, Indian Institute of Technology Madras, Chennai, India

4 Department of Computer Science and Engineering, Robert Bosch Center for Data Science and Artificial Intelligence (RBC-DSAI), Indian Institute of Technology Madras, Chennai, India