

March 12, 2019
<https://phik.rtfd.io>

A new correlation coefficient between categorical, ordinal and interval variables with Pearson characteristics

M. Baak^a, R. Koopman^a, H. Snoek^b, S. Klous^{a,c}

^aAdvanced Analytics & Big Data, KPMG Advisory N.V., Amstelveen, The Netherlands

^bNikhef National Institute for Subatomic Physics / University of Amsterdam, Amsterdam, The Netherlands

^cInformatics Institute, University of Amsterdam, Amsterdam, The Netherlands

Abstract

A prescription is presented for a new and practical correlation coefficient, ϕ_K , based on several refinements to Pearson's hypothesis test of independence of two variables. The combined features of ϕ_K form an advantage over existing coefficients. First, it works consistently between categorical, ordinal and interval variables. Second, it captures non-linear dependency. Third, it reverts to the Pearson correlation coefficient in case of a bi-variate normal input distribution. These are useful features when studying the correlation between variables with mixed types. Particular emphasis is paid to the proper evaluation of statistical significance of correlations and to the interpretation of variable relationships in a contingency table, in particular in case of low statistics samples and significant dependencies. Three practical applications are discussed. The presented algorithms are easy to use and available through a public Python library.

Contents

1	Introduction	3
2	Measures of variable association	4
3	Test of variable independence	6
3.1	Dependent frequency estimates	8
3.2	Independent frequency estimates	8
4	Definition of ϕ_K	9
4.1	Performance on benchmark samples	11
4.2	Example correlation matrix	11
4.3	Global correlation coefficients	13
4.4	Statistical noise correction	13
5	Statistical significance	14
5.1	Effective number of degrees of freedom	16
5.2	Modified χ^2 distribution	18
5.3	Evaluation of significance	21
5.4	Example significance matrix	23
5.5	Sampling approaches	24
6	Interpretation of relation between two variables	25
7	Three practical applications	27
7.1	Modeling the frequency of insurance claims	27
7.2	Finding unexpected answers in questionnaires	29
7.3	Comparison of clustering algorithms	29
8	Public implementation	29
9	Conclusion	30
	References	30

1 Introduction

The calculation of correlation coefficients between paired data variables is a standard tool of analysis for every data analyst. Pearson’s correlation coefficient [1] is a *de facto* standard in most fields, but by construction only works for interval variables. While many coefficients of association exist, each with different strengths, we have not been able to identify a correlation coefficient¹ with Pearson-like characteristics and a sound statistical interpretation that works for interval, ordinal and categorical variable types alike.

This paper describes a novel correlation coefficient, ϕ_K , with properties that – taken together – form an advantage over existing methods. Broadly, it covers three related topics typically encountered in data analysis:

1. Calculation of the correlation coefficient, ϕ_K , for each variable-pair of interest.

The correlation ϕ_K follows a uniform treatment for interval, ordinal and categorical variables. This is particularly useful in modern-day analysis when studying the dependencies between a set of variables with mixed types, where some variables are categorical. The values for levels of correlation are bound in the range $[0, 1]$, with 0 for no association and +1 for complete association. By construction, the interpretation is similar to Pearson’s correlation coefficient, and is equivalent in case of a bi-variate normal input distribution. Unlike Pearson, which describes the average linear dependency between two variables, ϕ_K also captures non-linear relations. Finally, ϕ_K is extendable to more than two variables.

2. Evaluation of the statistical significance of each correlation.

The correlation ϕ_K is derived from Pearson’s χ^2 contingency test [2], *i.e.* the hypothesis test of independence between two (or more) variables in a contingency table, henceforth called factorization assumption. In a contingency table each row is the category of one variable and each column the category of a second variable. Each cell describes the number of records occurring in both categories at the same time. The asymptotic approximation commonly advertised to evaluate the statistical significance of the hypothesis test, *e.g.* by statistics libraries such as R [3] and `scipy` [4], makes particular assumptions on the number of degrees of freedom and the shape of the χ^2 distribution. This approach is unusable for sparse data samples, which may happen for two variables with a strong correlation and for low- to medium-statistics data samples, leading to incorrect p -values. (Examples follow in Section 5.) Presented here is a robust and practical statistical prescription for the significance evaluation of the level of variable association, based on an adjustment of the χ^2 distribution when using the G -test statistic [5].

3. Insights in the correlation of each variable-pair, by studying outliers and their significances.

To help interpret any relationship found, we provide a method for the detection of significant excesses or deficits of records with respect to the expected values in a contingency table, so-called outliers, using a statistically independent evaluation for expected frequency of records.

We evaluate the significance of each outlier frequency, putting particular emphasis on the

¹The convention adopted here is that a correlation coefficient is bound, *e.g.* in the range $[0, 1]$ or $[-1, 1]$, and that a coefficient of association is not.

statistical uncertainty on the expected number of records and on the scenario of low statistics data samples.

The methods presented in this work can be applied to many analysis problems. Insights in variable dependencies serve as useful input to all forms of model building, be it classification or regression based, such as the identification of customer groups, outlier detection for predictive maintenance or fraud analytics, and decision making engines. More general, they can be used to find correlations across (big) data sets, and correlations over time (in correlograms). Three use-cases are discussed, the study of numbers of insurance claims, survey responses, and clustering compatibility.

This document is organized as follows. A brief overview of existing correlation coefficients is provided in Section 2. Section 3 describes the contingency test, which serves as input for Section 4, detailing the derivation of the correlation coefficient ϕ_K . The statistical significance evaluation of the contingency test is discussed in Section 5. In Section 6 we zoom in on the interpretation of the dependency between a specific pair of variables, where the significance evaluation of outlier frequencies in a contingency table is presented. Three practical applications of this can be found in Section 7. Section 8 describes the implementation of the presented algorithms in publicly available Python code, before concluding in Section 9.

2 Measures of variable association

A correlation coefficient quantifies the level of mutual, statistical dependence between two variables. Multiple types of correlation coefficients exist in probability theory, each with its own definition and features. Some focus on linear relationships where others are sensitive to any dependency, some are robust against outliers, etc. Typically their values range from -1 to $+1$ or 0 to $+1$, where 0 means no statistical association, $+1$ means the strongest possible association, and -1 means the strongest negative relation. In general, different correlation coefficients are used to describe dependencies between interval, ordinal, and categorical variables.

This section briefly discusses existing correlations coefficients and other measures of variable association. This is done separately for interval, ordinal, and categorical variables. In addition, several related concepts used throughout this work are presented.

An *interval variable*, sometimes called continuous or real-valued variable, has well-defined intervals between the values of the variable. Examples are distance or temperature measurements. The Pearson correlation coefficient is a *de facto* standard to quantify the level of association between two interval variables. For a sample of size N with variables x and y , it is defined as the covariance of the two variables divided by the product of their standard deviations:

$$\rho = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}}, \quad (1)$$

where \bar{x} and \bar{y} are the sample means. Notably, ρ is symmetric in x and y , and $\rho \in [-1, 1]$.

Extending this to a set of input variables, Pearson's correlation matrix C , containing the ρ values

of all variable pairs, is obtained from the covariance matrix V as:

$$C_{ij} = \frac{V_{ij}}{\sqrt{V_{ii}V_{jj}}}, \quad (2)$$

where ij are the indices of a variable pair.

The Pearson correlation coefficient measures the strength and direction of the linear relationship between two interval variables; a well-known limitation is therefore that non-linear dependencies are not (well) captured. In addition, ρ is known to be too sensitive to outlier records. Pearson's correlation coefficient, like many statistics formulas, requires interval variables as input, which can be unbinned or binned. It cannot be evaluated for categorical variables, and ordinal variables can only be used when ranked (see below).

A direct relationship exists between ρ and a bi-variate normal distribution:

$$f_{\text{b.n.}}(x, y | \bar{x}, \bar{y}, \sigma_x, \sigma_y, \rho) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}} \exp\left(-\frac{1}{2(1-\rho^2)}\left[\frac{(x-\bar{x})^2}{\sigma_x^2} + \frac{(y-\bar{y})^2}{\sigma_y^2} - \frac{2\rho(x-\bar{x})(y-\bar{y})}{\sigma_x\sigma_y}\right]\right), \quad (3)$$

where σ_x (σ_y) is the width of the probability distribution in x (y), and the correlation parameter ρ signifies the linear tilt between x and y . We use this relation in Section 4 to derive the correlation coefficient ϕ_K .

Another measure is the global correlation coefficient [6], which is a number between zero and one obtained from the covariance matrix V that gives the highest possible correlation between variable k and the linear combination of all other variables:

$$g_k = \sqrt{1 - [V_{kk} * (V^{-1})_{kk}]^{-1}}. \quad (4)$$

An **ordinal variable** has two or more categories with a clear ordering of these categories. For example, take the variable “level of education” with six categories: no education, elementary school graduate, high school graduate, college and university graduate, PhD. A rank correlation measures the statistical relationship between two variables that can be ordered. The rank of a variable is its index in the ordered sequence of values. For ordinal variables a numbering is assigned to the categories, *e.g.* 0, 1, 2, 3. Note the equidistant spacing between the categorical values.

Examples of rank correlation coefficients are Spearman's ρ [7], Kendall's τ [8], Goodman-Kruskal's γ [9–12], and the polychoric correlation [13]. The definition of Spearman's ρ is simply Eqn. 1, using the ranks of x_i and y_i as inputs, essentially treating the ranks as interval variables. This makes Spearman's ρ very robust against outliers. Noteworthy, Goodman-Kruskal's γ is dependent on the order of the two input variables, resulting in an asymmetric correlation matrix.

Although ranking is regular practice, the assumption of equidistant intervals – often made implicitly – can sometimes be difficult to justify. Adding the category of “MBA” to the above example increases the distance between “PhD” and “no education”, where one could argue that this distance should be independent of the number of educational categories.

A **categorical variable**, sometimes called a nominal or class variable, has two or more categories which have no intrinsic ordering. An example is the variable gender, with two categories: male and

female. Multiple measures of association exist that quantify the mutual dependence between two (or more) categorical variables, including Pearson’s χ^2 contingency test [2], the G -test statistic [5], mutual information [14], Fisher’s exact test [15, 16], Barnard’s test [17, 18]. For an overview see Ref. [19]. These measures determine how similar the joint distribution $p(x, y)$ is to the product of the factorized marginal distributions $p(x)p(y)$. Each measure of association consists of a sum of contributions, one from each cell of the contingency table, and therefore does not depend on the intrinsic ordering of the cells.

Though typically limited to categorical variables, these test statistics can also be applied to interval and ordinal type variables. However, their values are not bound in the range $[0, 1]$, and can become large. Moreover, their interpretation is often complex, as their values not only depend on the level of association, but also on the number of categories or intervals and the number of records.

Most comparable to this work is Cramér’s ϕ [20], a correlation coefficient meant for two categorical variables, denoted as ϕ_C , based on Pearson’s χ^2 test statistic, and with values between 0 (no association) and +1 (complete association):

$$\phi_C = \sqrt{\frac{\chi^2}{N \min(r-1, k-1)}}, \quad (5)$$

where r (k) is the number of rows (columns) in a contingency table. Notably, with a relatively small number of records, comparable with the number of cells, statistical fluctuations can result in large values of ϕ_C without strong evidence of a meaningful correlation. (An example of this follows in Fig. 4a.)

Cramér’s ϕ can also be used for ordinal and binned interval variables. Fig. 1 shows ϕ_C for a binned bi-variate normal input distribution with correlation parameter ρ . Compared to Pearson’s ρ , ϕ_C shows relatively low values for most values of ρ , and only shoots up to one for values of ρ close to one. Moreover, the value found for ϕ_C is dependent on the binning chosen per variable, as also seen in the figure. This effect makes ϕ_C difficult to interpret, and essentially unsuitable for interval variables.

One more alternative is the contingency coefficient C_p , which suffers from the disadvantage that its maximum value depends on the number of categories r and k , and does not reach a maximum of one. The recommendation [21] is not to use C_p to compare correlations in tables with variables that have different numbers of categories (*i.e.* when $r \neq k$).

To address the aforementioned issues, in this paper we define the coefficient of correlation ϕ_K , derived from Pearson’s χ^2 contingency test in Section 4, and its statistical significance, derived using the G -test in Section 5.

3 Test of variable independence

The contingency test, also called the test of variable independence, determines if a significant relationship exists between two (or more) categorical variables. Though usually performed on two categorical variables, the test can equally be applied to ordinal and binned interval variables, and can be extended to an arbitrary number of variables. Specifically, the contingency test indicates




Figure 1: Cramér’s ϕ versus Pearson’s ρ . The two curves for Cramér’s ϕ have been evaluated with different numbers of rows r and columns k : 10×10 and 5×20 bins. The value found for ϕ_C is dependent on the number of rows and columns. Smaller values for ϕ_C are found for the case of equal number of rows and columns (red) compared with different number of rows and columns (green).

how well the joint data distribution $p(x, y)$ of variables x and y is described by the product of its factorized marginal distributions $p(x)p(y)$.

Throughout this paper we employ two contingency tests, where each compares the observed frequency of each category for one variable with the expectation across the categories of the second variable:

1. Pearson’s χ^2 test:

$$\chi^2 = \sum_{i,j} \frac{(O_{ij} - E_{ij})^2}{E_{ij}}, \quad (6)$$

which is used to define the correlation coefficient ϕ_K in Section 4. Pearson’s χ^2 test is the standard test for variable independence.

2. The G -test, sometimes called log-likelihood ratio test:

$$G = 2 \sum_{i,j} O_{ij} \log(O_{ij}/E_{ij}), \quad (7)$$

which is used to evaluate the significance of the contingency test in Section 5. The sum is taken over all non-empty cells.

In both formulas, O_{ij} (E_{ij}) is the observed (expected) frequency of records for row i and column j of the contingency table. The stronger the dependency between x and y , the less well modeled is their distribution by the factorized distribution $p(x)p(y)$, and the larger each test statistic value.

Under the factorization assumption, the expected frequencies can be obtained in two ways: statistically dependent and independent.

3.1 Dependent frequency estimates

The default method of frequency estimation for row i and column j includes O_{ij} , so E_{ij} is statistically dependent on the observed frequency of its bin.

The expected value of the two nominal variables is calculated as:

$$E_{ij} = N p_r(i) p_k(j) = \frac{(\sum_{n=1}^k O_{in})(\sum_{m=1}^r O_{mj})}{N}, \quad (8)$$

where $p_r(i)$ ($p_k(j)$) is the i^{th} (j^{th}) bin of the row-projected (column-projected) marginal probability mass function (p.m.f.) and N is the number of records. The statistical dependency between E_{ij} and O_{ij} arises as the expectation E_{ij} for cell ij includes the cell's observation O_{ij} in both the sum over columns and rows, and as part of N . The formula can be easily extended to an arbitrary number of variables.

We use Eqn. 8 for the definition of ϕ_K in Section 4 and for the calculation of its significance in Section 5, as this distribution matches the observed frequencies most closely.

3.2 Independent frequency estimates

The second method of estimation of E_{ij} excludes O_{ij} , *i.e.* is statistically independent of the observed frequency of records for row i and column j . This estimate, known in high energy physics as the ABCD formula [22], is given by:

$$E_{ij} = \frac{B_{ij} C_{ij}}{D_{ij}} = \frac{(\sum_{n \neq j} O_{in})(\sum_{m \neq i} O_{mj})}{\sum_{m \neq i} \sum_{n \neq j} O_{mn}}, \quad (9)$$

where by construction O_{ij} is not part of E_{ij} , which allows for an objective comparison between observed and expected frequencies per bin. This formula can also be extended to more variables, except that the denominator of Eqn. 9, which is different for each pair of indices, can easily become zero for low statistics samples.

Note that B_{ij} , C_{ij} , and D_{ij} are sums of frequencies, each obeying Poisson statistics, and are statistically independent. Consequently, the statistical uncertainty on E_{ij} is evaluated with straightforward error propagation [23] as:

$$\sigma_{E_{ij}}^2 = \frac{\sigma_{B_{ij}}^2 C_{ij}^2}{D_{ij}^2} + \frac{\sigma_{C_{ij}}^2 B_{ij}^2}{D_{ij}^2} + \frac{\sigma_{D_{ij}}^2 E_{ij}^2}{D_{ij}^2}. \quad (10)$$

For an observed frequency of Q records, $\sigma_Q = \sqrt{Q}$, except when $Q = 0$, in which case we set $\sigma_Q = 1$. By doing so, when B_{ij} or C_{ij} is zero, and thus $E_{ij} = 0$, this approach results in a non-zero error on E_{ij} . The statistical uncertainty on the expected frequency, $\sigma_{E_{ij}}$, is only zero when both B_{ij} and C_{ij} are zero.

The expectation from Eqn. 9 is built with fewer statistics than Eqn. 8 and thus is slightly less accurate. Another difference is that the ABCD formula is not a true product of two (or more) factorized marginal distributions, *i.e.* the relative predictions for one row are not identical to those for another row, as is the case for dependent frequency estimates.

We use the independent frequency estimates of Eqn. 9 for the detection of significant excesses or deficits of records over expected values in a contingency table in Section 6, for reasons described there.

4 Definition of ϕ_K

The correlation coefficient ϕ_K is obtained by inverting the χ^2 contingency test statistic through the steps outlined below. Although the procedure can be extended to more variables, the method is described with two variables for simplicity.

We define the bi-variate normal distribution of Eqn. 3 with correlation parameter ρ and unit widths, centered around the origin, and in the range $[-5, 5]$ for both variables. Using uniform binning for the two interval variables, with r rows and k columns, results in a corresponding bi-variate p.m.f.. With N records, the observed frequencies, O_{ij} , are set equal to the probability per bin multiplied by N . The expected frequencies E_{ij} , are set to the predictions from the bi-variate normal distribution with $\rho = 0$, with N records and the same binning. We then evaluate the χ^2 value of Eqn. 6.

Let us define this function explicitly. First, we perform the integral of the bi-variate normal distribution over the area of bin ij

$$F_{ij}(\rho) = \int_{\text{area}_{ij}} f_{\text{b.n.}}(x, y | \rho) dx dy, \quad (11)$$

leading to the sum over bins:

$$\chi_{\text{b.n.}}^2(\rho, N, r, k) = N \sum_{i,j}^{k,r} \frac{(F_{ij}(\rho = \rho) - F_{ij}(\rho = 0))^2}{F_{ij}(\rho = 0)}. \quad (12)$$

This χ^2 value explicitly ignores statistical fluctuations in observed frequencies, and is a function of the numbers of rows and columns, N , and the value of ρ .

To account for statistical noise, we introduce a sample-specific pedestal related to a simple estimate of the effective number of degrees of freedom of the bi-variate sample, n_{sdof} :

$$n_{\text{sdof}} = (r - 1)(k - 1) - n_{\text{empty}}(\text{expected}), \quad (13)$$

with number of rows r and columns k , and where $n_{\text{empty}}(\text{expected})$ is the number of empty bins of the dependent frequency estimates of the sample. The pedestal is defined as:

$$\chi_{\text{ped}}^2 = n_{\text{sdof}} + c \cdot \sqrt{2n_{\text{sdof}}}. \quad (14)$$

The noise pedestal is configurable through parameter c , and by default $c = 0$. See Section 4.4 for the impact of the noise pedestal on ϕ_K and Section 5 for a discussion on the effective number of degrees of freedom.

The maximum possible χ^2 value [20] of the contingency test is:

$$\chi_{\text{max}}^2(N, r, k) = N \min(r - 1, k - 1), \quad (15)$$

which depends only the number of records N , rows r , and columns k , and is reached when there is a one-on-one dependency between the two variables. Specifically note that χ_{\max}^2 is independent of the shape of distribution² $p(x, y)$.

We scale Eqn. 12 to ensure it equals χ_{ped}^2 for $\rho = 0$ and χ_{\max}^2 for $\rho = 1$.

$$X_{\text{b.n.}}^2(\rho, N, r, k) = \chi_{\text{ped}}^2 + \left\{ \frac{\chi_{\max}^2(N, r, k) - \chi_{\text{ped}}^2}{\chi_{\text{b.n.}}^2(1, N, r, k)} \right\} \cdot \chi_{\text{b.n.}}^2(\rho, N, r, k). \quad (16)$$

This function is symmetric in ρ , and increases monotonically from χ_{ped}^2 to χ_{\max}^2 as ρ goes from zero to one.

We can now perform the necessary steps to obtain the correlation coefficient ϕ_K :

1. In case of unbinned interval variables, apply a binning to each one. A reasonable binning is generally use-case specific. As a default setting we take 10 uniform bins per variable.
2. Fill the contingency table for a chosen variable pair, which contains N records, has r rows and k columns.
3. Evaluate the χ^2 contingency test using the Pearson's χ^2 test statistic (Eqn 6) and the statistically dependent frequency estimates, as detailed in Section 3.1.
4. Interpret the χ^2 value as coming from a bi-variate normal distribution without statistical fluctuations, using Eqn. 16.
 - If $\chi^2 < \chi_{\text{ped}}^2$, set ρ to zero.
 - Else, with fixed N, r, k , invert the $X_{\text{b.n.}}^2$ function, *e.g.* using Brent's method [24], and numerically solve for ρ in the range $[0, 1]$.
 - The solution for ρ defines the correlation coefficient ϕ_K .

The procedure can be extended to more variables by using a multi-variate Gaussian instead of a bi-variate one.

In summary, we interpret the χ^2 value found in data as coming from a bi-variate normal distribution with a fixed amount of statistical noise and with correlation parameter ϕ_K . Non-linear relations are captured by ϕ_K through the χ^2 test of variable independence. The correlation ϕ_K reverts to the Pearson correlation coefficient in case of a bi-variate normal input distribution, with uniformly binned interval variables. Unlike Cramér's ϕ , the value of ϕ_K is stable against the number of bins chosen per interval variable, making it unambiguous to interpret. (In Fig. 1, overlaying the ϕ_K values evaluated with (a)symmetric binning gives a line indistinguishable from Pearson's ρ .) Like Cramér's ϕ , ϕ_K is affected by statistical fluctuations, which is relevant when the number of records is comparable with the number of cells (or lower); however, unlike Cramér's ϕ , ϕ_K has a correction for the statistical noise. Note that ϕ_K is independent of the order of the two input variables, and that the procedure can be extended to more than two variables³.

²Note that the G -test does not have this useful feature, making the G -test unsuitable for the calculation of ϕ_K .

³For more than two variables, follow the same procedure and assume a common correlation for each variable pair of the multivariate normal input distribution.

car color	driver age	area	mileage	car size
blue	60.4	suburbs	3339	XS
blue	30.9	suburbs	53370	XL
blue	18.5	suburbs	112557	XL
green	40.9	downtown	29605	L
gray	23.7	downtown	15506	M
multicolor	60.3	downtown	33148	L
white	66.7	suburbs	91132	XL
red	69.2	downtown	152445	XXL
metallic	43.5	hills	147275	S
...

Table 1: A synthetic data set with car insurance data. The data set consists of 2000 records and is used to illustrate the calculations of ϕ_K , statistical significance (in Section 5) and outlier significance (in Section 6).

All coefficients presented in Section 2 are computationally inexpensive to evaluate. The calculation of ϕ_K is computationally expensive because of the integrals of correlated bi-variate normal distributions evaluated in Eqn. 16, but is well-doable on any modern laptop, typically taking only a fraction of a second per ϕ_K calculation.

4.1 Performance on benchmark samples

A comparison with alternative correlation coefficients based on benchmark samples is given in Fig. 2. By construction, the interpretation of ϕ_K is similar to that of Pearson’s correlation coefficient, in particular for the bi-variate normal input distributions and the linear shapes, shown in the left and middle columns. Unlike Pearson, however, ϕ_K also captures non-linear relations as shown in the right column. Moreover, ϕ_K can be determined for categorical, ordinal, and interval variables alike. Note that Cramér ϕ gives relatively low values for all samples.

4.2 Example correlation matrix

When studying the dependencies of a set of variables with a mixture of types, one can now calculate the correlation matrix for all variable pairs, filled with ϕ_K values, which is a useful overview to have for a data analyst.

For illustration purposes a synthetic data set with car insurance data has been created. The data set consists of 2000 records. Each record contains 5 (correlated) variables of mixed variable types, see Table 1. These data are used throughout the paper to provide insights in the practical application of the methods introduced in this work. The ϕ_K correlation matrix measured on the car insurance data set is shown in Fig. 3.




Figure 2: Benchmark sample results for ϕ_K . Each synthetic data set contains 2000 data points. For the left column, from top to bottom, the bi-variate normal distributions have been generated with true correlations: $\{0.9, 0.7, 0.4, 0, -0.4, -0.7, -0.9\}$. For the middle column a linear data set is generated which is rotated around the origin. In the right column various data sets are generated with non-linear correlations. Note that these non-linear correlations are well-captured by ϕ_K , while Pearson's ρ is close to zero for all cases.

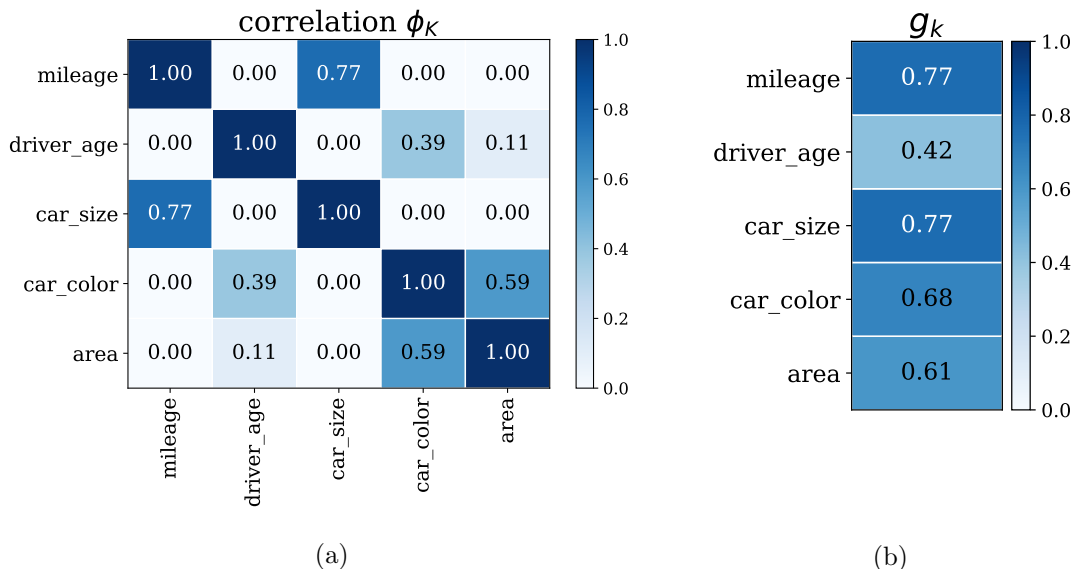


Figure 3: Correlation coefficients calculated on the synthetic car insurance data set (Table 1) containing mixed variables types. a) The ϕ_K correlation matrix. b) The global correlations g_k .

4.3 Global correlation coefficients

Besides the variable-pair information available from the correlation matrix C in Fig. 3a, it is also interesting to evaluate per variable the global correlation coefficient, g_k , of Eqn. 4. Strictly speaking, g_k is only defined for interval variables, as it requires a covariance matrix V . Here, we set the variances of all variable types to one (anyhow undefined for categorical variables⁴) and use $V = C$. Example global correlations measured in the car insurance data are shown in Fig. 3b. They give a tenable estimate of how well each variable can be modeled in terms of all other variables, irrespective of variable type.

4.4 Statistical noise correction

The calculation of ϕ_K contains a correction for statistical fluctuations: for any χ^2 value below the sample-specific noise threshold χ_{ped}^2 of Eqn. 14, indicating that no meaningful correlation can be determined, ϕ_K is set to 0 by construction.

The impact of the noise correction is seen in Fig. 4a, showing the absolute value of Pearson's ρ , Cramér's ϕ , and ϕ_K measured for 1000 synthetic data sets with only 500 records each, simulated from a bi-variate normal distribution with no correlation, and each binned in a 10x10 grid. Without absolute function applied, the distribution of Pearson's ρ values would be centered around zero, as expected. The calculation of Cramér's ϕ results in a seemingly significant bump at 0.2. This cannot be interpreted as a meaningful correlation, but results from the statistical noise contributing to each sample's χ^2 value.

⁴Interval variable can always be re-scaled to have unit variance.

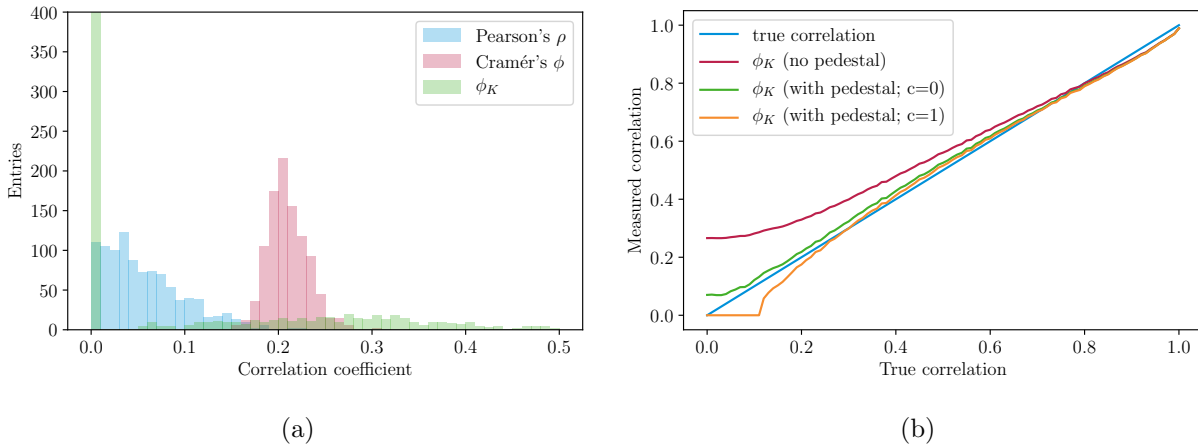


Figure 4: a) The correlation coefficients Pearson’s ρ , Cramér’s ϕ and ϕ_K , measured for 1000 synthetic data sets with 500 data points each, which are simulated using a bi-variate normal with parameter $\rho = 0$. The absolute value of the Pearson’s ρ is taken as the measured ρ can also take on negative values. b) The median ϕ_K value measured in 1000 synthetic data sets containing 500 data points each, simulated using a bi-variate normal distribution, as a function of true correlation ρ . The value of ϕ_K is evaluated using three different configurations of the noise pedestal parameter c (see Eqn. 14).

For ϕ_K , only when $\chi^2 > \chi_{\text{ped}}^2$ does the calculation of ϕ_K kick into action. The noise threshold is set such that about 50% of the simulated samples gets assigned $\phi_K = 0$. The remaining samples result in a wide distribution of ϕ_K values⁵.

Fig. 4b shows ϕ_K as a function of true correlation, where ϕ_K is obtained from the median χ^2 value of 1000 synthetic data sets with 500 data points each. The median gives the most representative, single synthetic data sample. In the calculation of ϕ_K three configurations for the noise pedestal of Eqn. 14 are tested: no pedestal, and $c \in \{0, 1\}$. No pedestal gives ϕ_K values that overshoot the true correlation significantly at low values. Configuration $c = 1$ undershoots: the calculation of ϕ_K turns on too late. Configuration $c = 0$ follows the true correlation most closely. The residual differences disappear for larger data samples, and we deem this acceptable for this level of statistics (with on average only 5 records per bin).

The sample-specific noise threshold χ_{ped}^2 depends mostly on the number of filled cells, and stabilizes for larger sample sizes. Consequently, its impact is rather limited for large samples with a meaningful, non-zero correlation, typically having $\chi^2 \gg \chi_{\text{ped}}^2$. For small sample sizes, as is also obvious from Fig. 4b, any correlation coefficient value should first be held up against the significance of the hypothesis test of variable independence – the topic of Section 5 – before further interpretation.

5 Statistical significance

In practice, when exploring a data set for variable dependencies, the studies of correlations and their significances are equally relevant: a large correlation may be statistically insignificant, and

⁵Without noise correction, the ϕ_K distribution shows a similar peak as Cramér’s ϕ , at value 0.5.

vice versa a small correlation may be very significant.

Both Pearson's χ^2 test and the G -test asymptotically approach the χ^2 distribution [2]. For samples of a reasonable size (Cochran's rule on what defines "reasonable size" follows below), the default approach to obtain the p -value for the hypothesis test of variable independence is to integrate the χ^2 probability density function $g(x|k)$ over all values⁶ equal to or greater than the observed test statistic value t_{obs} :

$$p = \int_{t_{\text{obs}}}^{\infty} g(x|k) dx, \quad (17)$$

with the p.d.f. of the χ^2 distribution:

$$g(x|k) = \frac{1}{2^\mu \Gamma(\mu)} \cdot x^{\mu-1} \cdot e^{-x/2}, \quad (18)$$

where $\mu = k/2$, $\Gamma(\mu)$ is the gamma function, and k is set to the number of degrees of freedom n_{dof} . The solution of this integral is expressed as the regularized gamma function. This approach holds for samples of a reasonable size, and when using the χ^2 test statistic or G -test.

For the independence test of n_{dim} variables, the number of degrees of freedom is normally presented [25] as the difference between the number of bins n_{bins} and model parameters n_{pars}

$$\begin{aligned} n_{\text{dof}} &= n_{\text{bins}} - n_{\text{pars}} \\ &= \left[\prod_{i=1}^{n_{\text{dim}}} n_i \right] - \left[\sum_{i=1}^{n_{\text{dim}}} (n_i - 1) + 1 \right]. \end{aligned} \quad (19)$$

where n_i is the number of categories of variable i . Explained using Eqn. 8, each dimension requires $(n_i - 1)$ parameters to model its p.m.f., which is normalized to one, and the p.m.f. product is scaled to the total number of events, which requires one more parameter. For just two variables this reduces to:

$$n_{\text{dof}} = (r - 1)(k - 1). \quad (20)$$

In practice Eqn. 19 does not hold for many data sets, in particular for distributions with unevenly filled or unfilled bins. For example, in the case of two (binned) interval variables with a strong dependency. The *effective* number of degrees of freedom, n_{edof} , is often smaller than the advocated value, n_{dof} , and can even take on floating point values, because the number of available bins is effectively reduced.

The asymptotic approximation, Eqns. 17-18, breaks down for sparse data sets, for example for two (interval) variables with a strong correlation, and for low-statistics data sets. The literature on evaluating the quality of this approximation is extensive; for an overview see Refs. [26, 27]. Cochran's rule of thumb is that at least 80% of the expected cell frequencies is 5 counts or more, and that no expected cell frequency is less than 1 count. For a 2x2 contingency table, Cochran recommends [28, 29] that the test should be used only if the expected frequency in each cell is at least 5 counts.

How to properly evaluate the p -value if the test statistic does not follow the χ^2 distribution and hence Eqn. 17 cannot be safely applied. A reasonable approach is to evaluate Eqn. 17 directly

⁶The integral runs up to infinity, even though the contingency test has a maximum test statistic value. In practice the difference is negligible.

with Monte Carlo data sets, sampled randomly from the distribution of expected frequencies. However, this approach quickly becomes cumbersome for p -values smaller than 0.1%, *i.e.* once more than 1000 simulations are needed for a decent p -value estimate, and practically impossible when needing at least a million simulations. Given that variable dependencies can be very significant, we prefer a common approach that works for both strong and weak dependencies and both low- and high-statistics samples.

In this section we propose another option: a hybrid approach where a limited number of Monte Carlo simulations is used to fit an analytical, empirical description of the χ^2 distribution. Specifically, we describe two corrections to Eqn. 17:

1. A procedure to evaluate the effective number of degrees of freedom for a contingency test;
2. A correction to Eqn. 18 for low statistics samples, when using the G -test statistic.

We conclude the section with a prescription to evaluate the statistical significance of the hypothesis test of variable independence, and a brief overview of sampling methods to help evaluate the p -value.

5.1 Effective number of degrees of freedom

To obtain the effective number of degrees of freedom of any sample, we use the property of Eqn. 18 that, for a test statistic distribution obeying $g(x|k)$, to good approximation the average value of $g(x|k)$ equals k .

The effective number of degrees of freedom for any sample is obtained as follows:

1. For the two variables under study, the dependent frequency estimates form the factorized distribution most accurately describing the observed data. Using Monte Carlo sampling techniques, this distribution is used to randomly generate 500 independent synthetic data sets with the same number of records as in the observed data set.
Optionally, sampling with fixed row and/or column totals may be chosen. A short discussion of sampling methods is held in Section 5.5.
2. For each synthetic data set, evaluate the G -test statistic using the statistically dependent frequency estimates, as detailed in Section 3.
3. The effective number of degrees of freedom, n_{edof} , is taken as the average value of the G -test distribution of all generated Monte Carlo samples.

Fig. 5 shows a “smiley” data set of two interval variables, consisting of two blobs and a wide parabola, which are binned into a 20x20 histogram, for which we can generate an arbitrary number of records.

The bottom two curves in Fig. 6 show n_{edof} obtained for this sample, as a function of the number of records in the data set, N , and evaluated using the G -test and χ^2 test statistic. Using Eqn. 20, the advocated number of degrees of freedom of this sample equals 361. For both test statistics




Figure 5: Example “smiley” data set of two interval variables binned in 20 bins in the x and y direction.




Figure 6: The effective number of degrees of freedom as a function of the number of data points in the input data set (Figure 5). The theoretical number of degrees of freedom, $n_{dof} = 361$, is indicated with the dashed line.

this number is only reached for very large sample sizes ($\geq 10^6$), and drops significantly for smaller values of N , where the drop is slightly steeper for the G -test statistic. The top two curves show the same data set on top of a uniform background of 1 record per cell, ensuring that each is always filled, again evaluated using the G -test or χ^2 test statistic. Now the G -test overshoots, and the χ^2 test statistic happens to level out at the expected value.

To understand the behavior of under- and overshooting, realize that n_{edof} relates directly to the distribution of dependent frequency estimates. By construction, the dependent frequency estimates of Eqn. 8 make non-zero predictions for each bin in the distribution, as long as the input data set contains at least one record per row and column. Under the assumption of variable independence, each bin in the distribution is expected to be filled.

First consider the bottom two curves of Fig. 6. For an uneven input distribution, for example two strongly correlated interval variables, one may expect many bins with low frequency estimates. A data set sampled randomly from a distribution with very low frequency estimates, such as the data set in Fig. 5, is likely to contain empty bins. On average, high-statistics bins contribute $n_{\text{dof}}/n_{\text{bins}}$ ($\lesssim 1$) to the G -test or χ^2 test statistic, but the low-statistics bins do not obey this regime. As an example, let us focus on the empty bins. By construction their contribution to the G -test is zero. The contribution to the χ^2 test statistic is non-zero: $\sum_i E_i$, where the sum runs over all empty bins. It is clear however, when $E_i \ll 1$, that this sum is relatively small and contributes only marginally. Taken over many randomly sampled data sets, this effect reduces the average value of the G -test or χ^2 test statistic distribution to lower values, and likewise decreases n_{edof} compared with n_{dof} .

For the top two curves, by construction $E_i > 1$ for each bin, bringing them closer to the nominal regime and increasing the G -test and χ^2 test statistics. For a discussion of the contribution of low-statistics contingency table cells to the χ^2 test statistic, see Ref. [30].

In summary, depending on the shape and statistics of the input data set, the effective number of degrees of freedom of a contingency table can differ from the advocated value of n_{dof} (Eqn. 19). To be certain of the effective value to use, this is best derived as the average value of the test statistic distribution, which is obtained with Monte Carlo simulations of the expected frequency distribution.

5.2 Modified χ^2 distribution

Given a large enough data sample, and given the hypothesis that the observed frequencies result from a random sampling from the distribution of expected frequencies, the G -test statistic can be approximated⁷ by Pearson's χ^2 . In this scenario both the G -test and χ^2 value are described by the χ^2 distribution of Eqn. 18, with the same number of degrees of freedom, and applying any one test leads to the same conclusions.

For low statistics samples – to be more specific, samples with many bins of low expected and observed frequencies – the distributions of G and χ^2 start to differ, and both distributions diverge from the nominal χ^2 distribution. This can be seen in Fig. 7, which uses the smiley data set of Fig. 5 as input. The simulated distribution of test statistics is wider than the χ^2 -distribution in

⁷The approximation is obtained with a second-order Taylor expansion of the logarithm around 1.




Figure 7: The simulated distribution of χ^2 and G -test statistics using the smiley data set as input (Figure 5). a) The χ^2 distribution (green) is wider than the expected distribution (dashed lines), and b) the G -test distribution (green) is narrower. The corresponding p -value distributions are shown in panels c) and d), both when using n_{dof} (blue) and n_{edof} (red).

case of the Pearson χ^2 -test statistic (Fig. 7a) and narrower than the χ^2 -distribution in case of the G -test statistic (Fig. 7b). This results in p -value distributions with elevated frequencies around zero and one for the Pearson χ^2 -test statistic (Fig. 7c) and lower frequencies near zero and one for the G -test statistic (Fig. 7d). Note that the effective number of degrees of freedom is much lower than the theoretical value; using n_{dof} in the p -value calculation results in an uneven distributions peaked towards one.

This section addresses the question whether the test statistic distribution for the contingency test can be modeled for all sample sizes, knowing that Eqn. 18 cannot be safely used for low statistics data sets. In particular we are interested in assessing the p -values of large test statistic values, coming from possibly strong variable dependencies. To evaluate these correctly, it is important to properly model the high-end tail of the test statistic distribution.

We observe empirically that for low-statistics samples the G -test statistic distribution converges towards a Gaussian distribution $G(x|\mu, \sigma)$, with mean $\mu = n_{\text{edof}}$ and width $\sigma = \sqrt{n_{\text{edof}}}$. For high-statistics samples the distribution is modeled by $g(x|k)$, with $k = n_{\text{edof}}$ degrees of freedom. Experimentally we find that, for any sample size, the G -test statistic distribution can be well described by the combined probability density function $h(x|f)$:

$$h(x|f) = f \cdot g(x|n_{\text{edof}}) + (1 - f) \cdot G(x|n_{\text{edof}}, \sqrt{n_{\text{edof}}}), \quad (21)$$

where the parameters of $g(x|k)$ and $G(x|\mu, \sigma)$ are fixed as above, and f is a floating fraction parameter between $[0, 1]$.

Below we use $h(x|f)$ as the modified χ^2 p.d.f. to model the G -test statistic distribution for any data set.

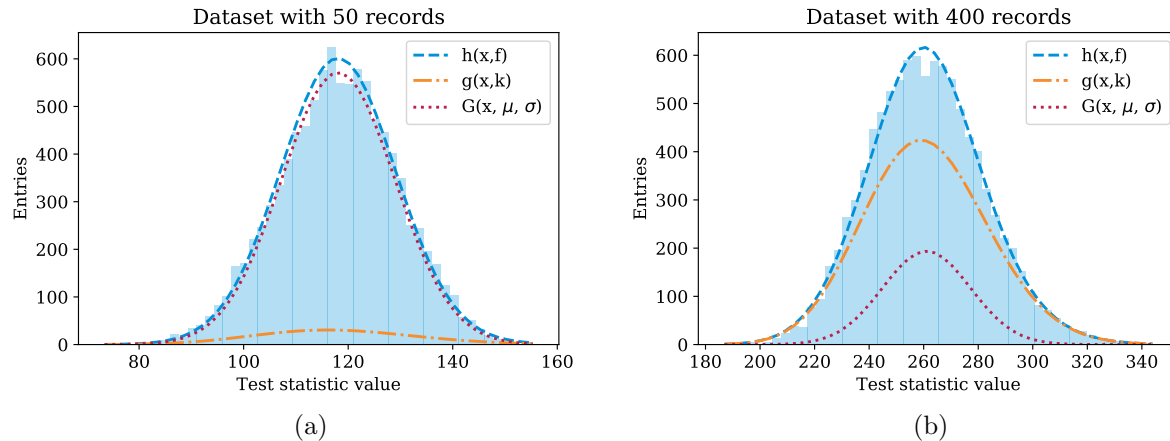


Figure 8: The G -test statistic distribution for two smiley data sets containing a) 50 and b) 400 data points. The distribution is modeled with the $h(x|f)$ distribution.

Fig. 8 shows the results of binned log-likelihood fits of $h(x|f)$ to two G -test statistic distributions, each with 10k entries generated with the procedure of Section 5.1, using the smiley data set with 20×20 bins with: a) $N = 50$ and b) $N = 400$ records for the simulated data sets. Clearly, these distributions are not well modeled using $g(x|n_{\text{edof}})$ or $G(x|n_{\text{edof}}, \sqrt{n_{\text{edof}}})$ alone. The fit of $h(x|f)$ can separate the two component p.d.f.'s given that the RMS-value of $g(x|n_{\text{edof}})$ is $\sqrt{2n_{\text{edof}}}$ and

the width of the Gaussian is fixed to $\sqrt{n_{\text{edof}}}$. For $N = 50$, the distribution is dominated by the Gaussian, and for $N = 400$ by the theoretical χ^2 distribution. Note that $G(x|n_{\text{edof}}, \sqrt{n_{\text{edof}}})$, when present, contributes to the core of the distribution while $g(x|n_{\text{edof}})$ dominates in the tails.




Figure 9: a) The fit fraction f as a function of the number of records per simulated data set, N . b) The same data points, but here f is shown as a function of the average number of records per bin.

Fig. 9 uses a similar setup, with 20x20 or 50x50 bins, where the fit fraction f is shown as a function of a) the number of records per simulated data set, N , and b) the average number of records per cell, \bar{n} . The fraction f rises as a function of sample size, such that $h(x|f)$ turns into $g(x|n_{\text{edof}})$ for large enough data sets. With 20x20 bins, for a fraction of 0.50 (0.99) the approximately sample size equals 175 (700), and the average number of entries per cell equals 0.4 (1.8). Note that the fraction reaches 1 well before n_{edof} reaches the advocated value of n_{dof} in Fig. 6.

In summary, to assess the p -value for the hypothesis test of variable independence, in this work we choose to work with the G -test statistic, and not Pearson's χ^2 , for two reasons:

1. We manage to describe the G -test statistic distribution most successfully for any sample size.
2. As seen from Fig. 7b, for a large observed test statistic value, corresponding to a large significance of variable dependency, applying the naive formula of Eqn. 17 over-covers, *i.e.* gives a conservative p -value (the green distribution is narrower than expected).

We use the distribution $h(x|f)$ of Eqn. 21 as modified χ^2 distribution in Eqn. 17 to assess the p -value for the hypothesis test.

5.3 Evaluation of significance

The statistical significance of the hypothesis test of any variable independence is obtained with the following procedure:

1. Calculate the average number of entries per cell, \bar{n} . If $\bar{n} < 4$, set $n_{\text{sim}} = 2000$, else $n_{\text{sim}} = 500$ samples.

2. Follow the procedure of Section 5.1 to generate n_{sim} synthetic data sets based on the dependent frequency estimates of the input data set. For each synthetic data set evaluate its G -test value. Take the average of the G -test distribution to obtain n_{edof} .
3. If $\bar{n} < 4$, to obtain f fit the probability density function $h(x|f)$ to the G -test distribution, with n_{edof} fixed. Else, skip the fit and set $f = 1$.
4. With this fraction, use Eqn. 17 with $h(x|f)$ as modified χ^2 distribution to obtain the p -value for the hypothesis test, using the G -test value from data as input.
5. The p -value is converted to a normal Z -score:

$$Z = \Phi^{-1}(1 - p); \quad \Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z e^{-t^2/2} dt, \quad (22)$$

where Φ^{-1} is the quantile (inverse of the cumulative distribution) of the standard Gaussian, *e.g.* Z is the significance in 1-sided Gaussian standard deviations. For example, the threshold p -value of 0.05 (95% confidence level) corresponds to $Z = 1.64$.

When the p -value is too small to evaluate Eqn. 22 numerically, at $p \lesssim 10^{-310}$, anyhow a very strong variable dependency, Z is estimated using Chernoff's bound [31] to ensure a finite value. Let $z \equiv G/n_{\text{edof}}$, Chernoff states when $z > 1$:

$$p \leq f \cdot (ze^{1-z})^{n_{\text{edof}}/2}, \quad (23)$$

where we safely ignore the contribution from the narrow Gaussian in $h(x|f)$. This is converted to Z with the approximation (valid for large $Z > 1.5$):

$$Z = \sqrt{u - \log u}; \quad u = -2 \log(p\sqrt{2\pi}). \quad (24)$$




Figure 10: The p -value distributions corresponding of the two G -test distributions of Fig. 8, with a) $N = 50$ and b) $N = 400$ records per sample. See the text for a description of the two p -value calculations performed.

The significance procedure is illustrated in Fig. 10, which shows the p -value distributions of the two G -test distributions of Fig. 8, with $N = 50$ and $N = 400$ records per sample. The two p -value distributions in each figure have been calculated in two ways.

1. Using the original χ^2 distribution $g(x|k)$ of Eqn. 17, with the effective number of degrees of freedom, n_{edof} . This results in the blue distributions.
2. Fitting each test statistic distribution with $h(x|f)$ of Eqn. 21, and using that to calculate the p -values, resulting in the red distributions.

The blue distributions drop around zero and one, in particular for the low statistics sample ($N = 50$). This is because the G -test distribution is more narrow than the χ^2 distribution, as shown in Fig. 7. The red p -value distributions, evaluated with $h(x|f)$, are uniform, as desired in both setups.

Let us apply the statistical procedure to a low-statistics data sample. A smiley data set with 100 entries, in a histogram with 20x20 bins, has correlation value $\phi_K = 0.73$ and test statistic value $G = 227.4$. The Z calculation is done in three consecutively more refined ways:

1. *The asymptotic approximation*: using $n_{\text{dof}} = 361$ and the asymptotic χ^2 distribution $g(x|k)$ gives: $Z = -5.7$;
2. *Effective number of degrees of freedom*: using $n_{\text{edof}} = 189.3$ and the asymptotic χ^2 distribution $g(x|k)$ results in: $Z = 1.9$;
3. *Modified χ^2 distribution*: with $n_{\text{edof}} = 189.3$, the modified χ^2 distribution $h(x|f)$, and fit fraction $f = 0.10$ one finds: $Z = 2.5$.

In this example, between the three approaches the Z -value increases with more than 8 units! Typically, using the effective number of degrees of freedom gives the largest correction to Z , and the modified χ^2 distribution only gives a small correction on top of that.

The choice of 2000 synthetic data sets for the fit of $h(x|f)$ is a compromise between accuracy and speed. With this number, Z typically varies at the level of 0.04, and is calculated in just a fraction of a second.

Based on our findings, for any sample size we recommend the p -value to be calculated with the modified χ^2 distribution $h(x|f)$, using n_{edof} degrees of freedom. If not, the p -value may over-cover for strong variable dependencies and at low-statistics, resulting in a Z -value that is too small, possibly by multiple units. This is important to know, as it can lead to rather incorrect conclusions regarding the studied variable dependency.

5.4 Example significance matrix

In practice a correlation value may be small but its statistical significance can still be large, and vice versa. For this reason, when exploring a data set, the levels of correlation and significance should always be studied together.

Fig. 11 shows the significance matrix determined for the car insurance data set of Table 1. Compared with the correlation matrix of Fig. 3, the low ϕ_K values happen to be statistically insignificant, but the higher values are very significant.




Figure 11: The significance matrix, showing the statistical significances of correlated and uncorrelated variable pairs. The color scale indicates the level of significance, and saturates at ± 5 standard deviations.

5.5 Sampling approaches

Based on the statistically dependent frequency estimates, three sampling approaches are offered to generate synthetic data sets for testing the hypothesis of no variable association:

- *Multinomial sampling*: with only the total number of records fixed. The hypothesis of no association is independent of the row and column variables.
- *Product-multinomial sampling*: with the row or column totals fixed in the sampling. The hypothesis of no association is also called homogeneity of proportions. This approach is commonly used in cohort and case-control studies.
- *Hypergeometric sampling*: both the row or column totals are fixed in the sampling. This approach is also known as Fisher's exact test. We use Patefield's algorithm [32] to generate the samples.

There is an ongoing debate about sampling design for tests of variable independence. Although in practice most people are not too worried about the sampling approach, at least not in the high-statistics regime, because asymptotically the different approaches lead to the same result. The default approach used in this paper is multinomial sampling. For a discussion and further references see Ref. [33].

6 Interpretation of relation between two variables

After the evaluation of ϕ_K and its significance, the specific relationship between two variables is typically inspected. To facilitate the interpretation of any dependency found, the significance of observed excesses or deficits of records with respect to expected values in the contingency table is discussed here.

The statistical significance for each cell in the table is obtained from an hypothesis test between a background-only and signal-plus-background hypothesis for a Poisson process. Such hypothesis tests, *i.e.* for the presence of new sources of (Poisson) counts on top of known “background” processes, are frequently performed in many branches of science, for example gamma ray astronomy and high energy physics, and have been discussed extensively in the literature [34].

We employ a measure of statistical significance commonly used in both fields, one that accounts for the mean background rate having a non-negligible uncertainty. The background estimate and its uncertainty have been derived from an auxiliary or side-band measurement, typically assumed to be a Poisson counting setup, as in the case of the ABCD estimate of Section 3.2. Here we use as background estimate the statistically independent frequency estimate (and related uncertainty) of Eqn. 9 (10).

The hybrid Bayesian-Frequentist method from Linneman [35] is used to evaluate the probability of the hypothesis test (p -value). Per cell, Linneman’s probability calculation requires the observed count n_o , the expected count n_e , and the uncertainty on the expectation σ_e :

$$p_B = B(1/(1 + \tau), n_o, n_e \tau + 1), \quad (25)$$

where B is the incomplete Beta function, and $\tau = n_e/\sigma_e^2$.

We apply four corrections on top of this calculation:

1. The incomplete Beta function returns no number for $n_o = 0$, when by construction the p -value should be 1.
2. The incomplete Beta function is undefined when $\sigma_e = 0$, in which case we simply revert to the standard Poisson distribution.
3. The incomplete Beta function always returns 1 when $n_e = 0$, irrespective of n_o and σ_e . The scenarios $n_o = 0$ and $\sigma_e = 0$ are captured by the previous two fixes. In all other cases we set $n_e = \sigma_e$ before evaluating Eqn. 25. In particular, this procedure prevents (minus) infinite significances for low statistics cells where uncertainty-wise these are not expected⁸.
4. As we combine an integer-valued measurement (namely the observed frequency) with a continuous expectation frequency and uncertainty, resulting in a continuous (combined) test statistic, we correct p_B to Lancaster’s mid- P value [36], which is the null probability of more

⁸When $n_e = 0$, B or C is zero in Eqn. 9, so Eqn. 10 typically gives $\sigma_e < 1$. For example, for $n_e = 0$, $\sigma_e = 0.14$, and $n_o = 0$ (1), correction three to p_B results in $Z = -0.29$ (1.10). Varying n_e between $\sigma_e/2$ and $3\sigma_e/2$ gives a maximum absolute shift in Z of 0.05 (0.12). To do outlier detection, for this procedure we deem this level of systematic error acceptable.

extreme results plus only half the probability of the observed result⁹:

$$p = P(s = \text{observed} \mid \text{background})/2 + P(s > \text{observed} \mid \text{background}), \quad (26)$$

with s the integrated-over number of cell counts. This p -value is then translated into the Z -value using Eqn. 22. When observing the expected frequency by construction Lancaster's mid- P value (Z -value) is close to 0.5 (0), even at low statistics. Likewise, for background-only samples the Lancaster's mid- P correction centers the Z distribution around zero.

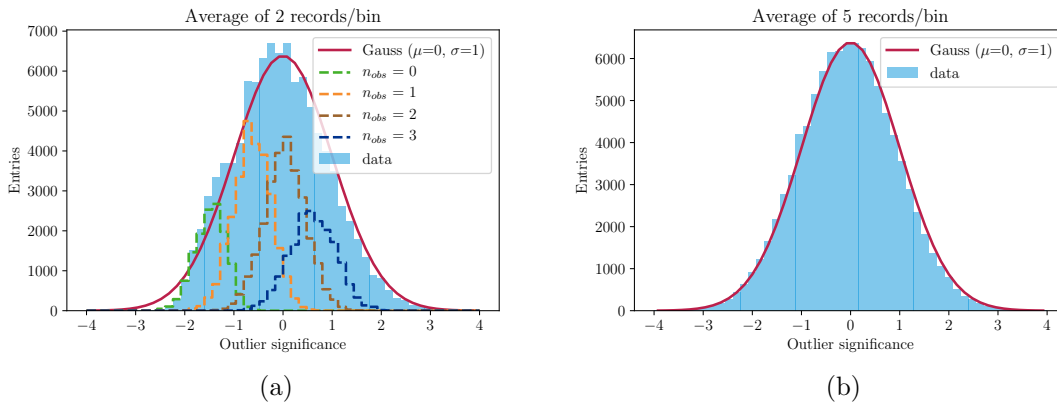


Figure 12: The distribution of outlier significances measured in 1000 randomly generated data sets of two variables obeying a uniform probability mass distribution for a data set containing a) 200 and b) 500 records, collected in a 10×10 contingency table. Normal distributions have been overlaid. In plot a) the Z distributions from 0, 1, 2, and 3 observed entries per cell are shown as well.

Fig. 12a shows the Z distribution from 1000 randomly generated samples of two variables obeying a uniform probability mass distribution, *i.e.* the samples have no variable dependency. Each sample contains only 200 records collected in a 10×10 contingency table, so on average each cell contains 2.0 records. As can be seen from the Gaussian curve, even for such low statistics samples the Z distribution is fairly consistent with a normal distribution, albeit slightly shifted towards negative values. Fig. 12b shows a similar distribution, built from samples with on average 5.0 records per contingency table cell. Clearly, with more statistics the distribution converges to the normal distribution relatively quickly¹⁰.

To filter out significant excesses or deficits of records over expected values in the contingency table, one simply demands $|Z|$ to be greater than a specified value, *e.g.* 5 standard deviations. For two variables with a dependency, note that excesses and deficits always show up together, since the frequency estimates of Section 3.2 smooth the input distribution.

Two example contingency tables are shown in Fig. 13, one for a combination of categorical variables, and one for the combination of an interval and ordinal variable, both based on the synthetic car

⁹The standard p -value definition is: $p = P(s \geq \text{observed} \mid \text{background})$.

¹⁰With an average of less than 1.0 records per bin, the Z distribution gets more distorted, and breaks up into individual peaks of 0, 1, 2, etc. observed entries per cell. The distribution peaks at negative Z values, corresponding to no observations, and the tail at negative Z gets truncated. Relevant here: the mean of the distribution remains close to zero, its width is (slightly less than) one, and the positive tail is similar to that of a normal distribution.




Figure 13: Significances of excesses or deficits of records over the expected values in a contingency table for a) the categorical variables “car color” and “area” and b) the ordinal variable “car size” and the interval variable “mileage”, measured on the synthetic data of Table 1.

insurance data of Table 1. Per cell each figure shows the evaluated Z -value. For example, black-colored cars occur significantly more in suburbs and significantly less down-town, and XXL-sized cars have significantly higher mileage.

In practice these turn out to be valuable plots to help interpret correlations, in particular between categorical variables. In essence, for a data sample with a dependency, the contingency table cells with large $|Z|$ values show the variable dependency.

7 Three practical applications

Given a set of mixed-type variables and using the methods described in this work, one can:

- Find variable pairs that have (un)expected correlations;
- Evaluate the statistical significance of each correlation;
- Interpret the dependency between each pair of variables.

The methods presented in this work can be applied to many analysis problems, and in particular they are useful for model building purposes. Three interesting applications using the methods presented in this paper are briefly discussed below.

7.1 Modeling the frequency of insurance claims

One interesting application is the modeling of numbers of expected insurance claims, *e.g.* car damage claims as a function of car type, type of residential area, mileage, age of driver, etc. – a set of variables with a mixture of types.

The aggregate loss incurred by an insurer S is the total amount paid out in claims over a fixed time period: $S = \sum_{n=1}^N s_n$, where s_n is an individual claim amount, known as the severity, and N is the total number of claims paid out in the time period. Traditionally it is assumed that the individual claim amounts are mutually independent, and that N does not depend on the values of the claims. The total expected severity is then expressed as a product of the expected number of claims times the average claim amount: $E(S) = E(N) \cdot E(s)$, where each term can be estimated separately. When a vector of variables \vec{x} is available at the individual claim level, this information is incorporated through two independent generalized linear models (GLMs): one for the claim frequency N , and the other for the severity s . See Ref. [37] for more information.

Here we focus on the GLM for modeling the frequency of insurance claims. Suppose that claims data are available for m different classes of policy holders, and that class i has N_i claims. Assume the claim frequency for each class is Poisson distributed, $N_i \sim P(\nu_i)$, where ν_i is the expectation for N_i . Let $\vec{x}_i = (x_{i0}, \dots, x_{ik})$ be the vector of variables for class i at claim level, with the baseline convention that $x_{i0} \equiv 1$. One writes:

$$\nu_i = E(N_i|\vec{x}_i) = g^{-1}(\vec{\alpha} \cdot \vec{x}_i), \quad (27)$$

where $\vec{\alpha} = (\alpha_{i0}, \dots, \alpha_{ik})$ is a vector of regression constants¹¹. In GLM terminology g is the link function. When the frequency GLM uses a logarithmic link function, Eqn. 27 simplifies to:

$$\nu_i = e^{\vec{\alpha} \cdot \vec{x}_i}, \quad (28)$$

yielding a simple rating structure which ensures that $\nu_i > 0$. The logarithmic function reflects the common practice that each variable alters the baseline claim rate by a multiplicative factor.

In initial GLM risk models, no relations are typically assumed between the input variables, and each variable category j (or interval bin) is one-hot-encoded, $x_{ij} \in \{0, 1\}$, and assigned one model parameter. The number of regression parameters per variable equals its number of categories or interval bins. Note that it is common practice to merge low-statistics categories until they contain sufficient records.

Take the example variables of residential area and car type, each with multiple categories. Three classes of policy holders could be: “city, small car”, “city, SUV”, and “countryside, SUV”, where the first two share the regression parameter α_{city} , and the last two the regression parameter α_{SUV} . The predicted, factorized number of claims for class “city, SUV” simply reads: $N_0 e^{\alpha_{\text{city}} + \alpha_{\text{SUV}}}$, where $N_0 \equiv e^{\alpha_0}$ is the nominal number of claims shared between all classes, and $x_{\text{city}} = x_{\text{SUV}} = 1$.

In a refinement modeling step, to improve the factorized estimates, cross-terms between categories of variable pairs can be added to the linear sum in the power of Eqn. 28. However, there is an exponentially large number of cross-terms to choose from. Practical modeling questions are: which are the most relevant terms to add? And can they be picked in an effective way that limits their number?

To help answer these, realize that the shape of Eqn. 28 and the assumption of variable independence are identical to the factorization assumption of Eqn. 8. A practical approach can then be:

1. Using the ϕ_K values and their significances, select the variable pairs with the strongest correlations.

¹¹Sometimes the ratio of claims to no claims per class of policy holders is modeled instead.

2. The most relevant model cross-terms for each variable pair pq , having the largest impact in the model's likelihood, can be identified by studying the outliers in the correlation plots of Section 6.
3. Cross-terms can also be included in a manner that limits the number of extra regression parameters. For example, for a given variable pair pq , introduce one cross-term parameter β_{pq} that affects only the contingency table cells with a Z value greater than a predefined value (and one for those smaller). To model those outlier cells, use Eqn. 10: the cross term for each selected cell ij should scale with the uncertainty on the statistically independent estimate for that cell, $\sigma_{E_{ij}}\beta_{pq}x_{p,i}x_{q,j}$.

7.2 Finding unexpected answers in questionnaires

When interpreting questionnaires one is often interested in finding all “unexpected” correlations between ordinal or categorical answers given to a set of survey questions (the definition of what constitutes an unexpected correlation is typically survey specific). The methods presented in this paper can help to do so:

1. By selecting question-pairs that have an interesting (“unexpected”) ϕ_K correlation and significance on the one hand;
2. And selecting those with relatively high $|Z|$ values in the contingency tables of their respective answers on the other hand.

This allows one to compile a list with all answer-pairs significantly deviating from the norm of no correlation, of which the unexpected pairs are a subset.

7.3 Comparison of clustering algorithms

When looking for groups of similar data records, a typical approach is to run multiple unsupervised clustering algorithms to cluster the data, and study the results. In trying to understand the compatibility in clusters created by the various algorithms, the methods presented in this work come in useful.

For each data record, store the cluster-ID assigned by each clustering algorithm. Using this information, one can now:

1. Calculate the correlation matrix between the various clustering algorithms;
2. For two specific algorithms, study where the two sets of predicted clusters overlap and deviate.

8 Public implementation

The ϕ_K correlation analyzer code is publicly available as a Python library through the PyPi server, and from GitHub at <https://github.com/KaveIO/PhiK>. Install it with the command:

```
pip install phik
```

The web-page <https://phik.readthedocs.io> contains a description of the source code, a tutorial on how to set up an analysis, and working examples of how to use and run the code.

9 Conclusion

We have presented a new correlation coefficient, ϕ_K , based on the χ^2 contingency test, with Pearson-like behavior and the practical feature that it applies to all variable types alike. Compared to Cramér's ϕ , the calculation of ϕ_K is stable against the binning per interval variable, making it easy to interpret, and contains a noise correction against statistical fluctuations.

The asymptotic approximation breaks down for sparse and low-statistics data sets. To evaluate the statistical significance of the hypothesis test of variable independence, a hybrid approach is proposed where, using the G -test statistic, a number of Monte Carlo simulations is used to determine the effective number of degrees of freedom and to fit an analytical, empirical description of the χ^2 distribution. We have evaluated the statistical significance of outlier frequencies with respect to the factorization assumption, which is a helpful technique for interpreting any dependency found, *e.g.* between categorical variables.

Three practical use-cases are discussed, studying the numbers of insurance claims, survey responses, and clustering compatibility, but plenty of other applications exist. The methods described are easy to apply through a Python analysis library that is publicly available.

Acknowledgments

We thank our colleagues of KPMG's Advanced Analytics & Big Data team for fruitful discussions and in particular Ralph de Wit for carefully reading the manuscript. This work is supported by KPMG Advisory N.V.

References

- [1] K. Pearson, Proceedings of the Royal Society of London (1854-1905) **58**, 240242 (1895).
 - [2] G. A. Barnard, Springer Series in Statistics Breakthroughs in Statistics , 110 (1992).
 - [3] `r.stats.chisq.test`.
 - [4] `scipy.stats.chi2_contingency`.
 - [5] R. R. Sokal and F. J. Rohlf, *Biometry: the principles and practice of statistics in biological research* (Palgrave Macmillan, 2012).
 - [6] F. James and M. Roos, Computer Physics Communications **10**, 343367 (1975).
 - [7] C. Spearman, American Journal of Psychology **15**, 72 (1904).
-

-
- [8] M. G. Kendall, *Biometrika* **30**, 81 (1938).
- [9] L. Goodman and W. Krustal, *Journal of the American Statistical Association* **49**, 732 (1954).
- [10] L. Goodman and W. Krustal, *Journal of the American Statistical Association* **54**, 123 (1959).
- [11] L. Goodman and W. Krustal, *Journal of the American Statistical Association* **58**, 310 (1963).
- [12] L. Goodman and W. Krustal, *Journal of the American Statistical Association* **67**, 415 (1972).
- [13] F. Drasgow, *Encyclopedia of Statistical Sciences* (2006).
- [14] T. M. Cover and J. A. Thomas, *Elements of information theory* (John Wiley & Sons, 2006).
- [15] R. A. Fisher, *Journal of the Royal Statistical Society* **85**, 87 (1922).
- [16] R. A. Fisher, *Statistical methods for research workers* (Oliver and Boyd, 1970).
- [17] G. Barnard, *Nature* **156**, 177 (1945).
- [18] G. A. Barnard, *Biometrika* **34**, 123 (1947).
- [19] A. Agresti, *Statistical Science* **7**, 131153 (1992).
- [20] C. Harald, *Mathematical methods of statistics* (Princeton University Press, 1999).
- [21] S. M. Smith, *Fundamentals of marketing research* (Sage Publications, 2009).
- [22] ATLAS, M. Aaboud *et al.*, *Eur. Phys. J.* **C78**, 154 (2018), [1708.07875].
- [23] H. H. Ku, (1965).
- [24] R. P. Brent, *Algorithms for minimization without derivatives* (Prentice-Hall, 1973).
- [25] D. E. Bock, P. F. d. Velleman and D. V. R. D., *Stats: modeling the world* (Pearson/Addison-Wesley, 2007).
- [26] A. Agresti, *Statistics in Medicine* **20**, 27092722 (2001).
- [27] P. M. Kroonenberg and A. Verbeek, *The American Statistician* **72**, 175183 (2018).
- [28] W. Cochran, *Annals of Mathematical Statistics* **23**, 315 (1952).
- [29] W. G. Cochran, *Biometrics* **10**, 417 (1954).
- [30] F. Yates, *Supplement to the Journal of the Royal Statistical Society* **1**, 217 (1934).
- [31] X. Lin *et al.*, *Past, present, and future of statistical science* (CRC Press, 2014).
- [32] W. M. Patefield, *Applied Statistics* **30**, 91 (1981).
- [33] D. Kim and A. Agresti, *Computational Statistics & Data Analysis* **24**, 89104 (1997).
- [34] R. D. Cousins, J. T. Linnemann and J. Tucker, *Nucl. Instrum. Meth.* **A595**, 480 (2008), [physics/0702156].
-

-
- [35] J. T. Linnemann, eConf **C030908**, MOBT001 (2003), [physics/0312059].
- [36] H. O. Lancaster, *Journal of the American Statistical Association* **56**, 223 (1961).
- [37] J. Garrido, C. Genest and J. Schulz, *Insurance: Mathematics and Economics* **70**, 205215 (2016).
-