

EUR Research Information Portal

Unbiased mixed variables distance

Publication status and date:

Published: 04/11/2024

Document Version

Early version, also known as preprint

Document License/Available under:

CC BY

Citation for the published version (APA):

van de Velden, M., Iodice D'Enza, A., Markos, A., & Cavicchia, C. (2024). *Unbiased mixed variables distance*. arXiv. <https://arxiv.org/abs/2411.00429>

[Link to publication on the EUR Research Information Portal](#)

Terms and Conditions of Use

Except as permitted by the applicable copyright law, you may not reproduce or make this material available to any third party without the prior written permission from the copyright holder(s). Copyright law allows the following uses of this material without prior permission:

- you may download, save and print a copy of this material for your personal use only;
- you may share the EUR portal link to this material.

In case the material is published with an open access license (e.g. a Creative Commons (CC) license), other uses may be allowed. Please check the terms and conditions of the specific license.

Take-down policy

If you believe that this material infringes your copyright and/or any other intellectual property rights, you may request its removal by contacting us at the following email address: openaccess.library@eur.nl. Please provide us with all the relevant information, including the reasons why you believe any of your rights have been infringed. In case of a legitimate complaint, we will make the material inaccessible and/or remove it from the website.

Unbiased mixed variables distance

Michel van de Velden

Econometric Institute, Erasmus University Rotterdam

and

Alfonso Iodice D’Enza

Department of Political Sciences, University of Naples Federico II

and

Angelos Markos

Department of Primary Education, Democritus University of Thrace

and

Carlo Cavicchia

Econometric Institute, Erasmus University Rotterdam

November 4, 2024

Abstract

Defining a distance in a mixed setting requires the quantification of observed differences of variables of different types and of variables that are measured on different scales. There exist several proposals for mixed variable distances, however, such distances tend to be “biased” towards specific variable types and measurement units. That is, the variable types and scales influence the contribution of individual variables to the overall distance. In this paper, we define “unbiased” mixed variable distances for which the contributions of individual variables to the overall distance are not influenced by measurement types or scales. We define the relevant concepts to quantify such biases and we provide a general formulation that can be used to construct unbiased mixed variable distances.

Keywords: Data visualization, distance-based methods, cluster analysis

1 Introduction

In many fields of study, it is necessary to measure the distance or similarity between observations or data points. These measurements are crucial for various tasks, including data visualization, clustering and classification. A significant challenge arises when a dataset consists of mixed variables, that is, when the observed variables differ in type or scale (Blanco-Mallo et al. 2023). Mixed variable datasets can consist of both numerical variables (continuous or discrete) and categorical variables (binary, ordinal, or nominal). These datasets are often encountered in machine learning and scientific computing tasks (Luo et al. 2024; Hermes, Heerwaarden, and Behrouzi 2024).

In a purely numerical context distance can be based on the sizes of the observed differences. For categorical variables, the situation is more complex as one only observes whether categories are equivalent or not. However, it may be appropriate to quantify differences between categories of the same variable differently. For example, one could argue that the dissimilarity between the colors purple and blue is smaller than the dissimilarity between yellow and blue. Similarly, the number of categories could play a role in quantifying differences. If there are many categories, there are more ways to differ. Many proposals exist to quantify distance for categorical variables; see, e.g. Boriah, Chandola, and Kumar (2008), Šulc and Řezanková (2019), and van de Velden, Iodice D’Enza, Markos, and Cavicchia (2024). The complexity arising from quantifying differences and the impact that these choices have in a clustering context when working with categorical data is thoroughly demonstrated by Amiri, B. S. Clarke, and J. L. Clarke (2018).

For mixed variables, the situation is even more complex, as this requires a method of calculating distance that can handle the diversity in variable types effectively, i.e., the aggregation over the variables should not be influenced by types or scales. One of the most recognized and widely used methods for calculating mixed variable distances is based on Gower’s general coefficient of similarity (Gower 1971). This measure, which we refer to as Gower’s distance, provides an elegant and simple way to measure the distance between observations based on measurement on numerical

and categorical variables. Although alternative proposals that allow distance calculations in mixed variable contexts exist (see, e.g., Ahmad and Dey 2007; Huang 1998; van de Velden, Iodice D’Enza, and Markos 2019), Gower’s proposal – implemented in the **gower** R package (Loo 2017) and in the `daisy` function available through the **cluster** R package (Maechler et al. 2023) – remains a popular choice.

It has been observed that biases can arise when computing distances in a multivariate context. For example, in numerical data, measurement scales can influence distances. Similarly, for categorical variables, the number of categories can affect the calculation of distances. Additionally, in settings with mixed variable types, it is important to control the balance between these types. For example, the Gower’s distance is known to be biased toward the contribution of categorical variables (Hennig and Liao 2013; Foss et al. 2016). One way to overcome potential biases is to weigh the distances resulting from the numerical and categorical variables differently (Huang 1998; Bello et al. 2021). Alternatively, different ways to scale the numerical and categorical variables may be introduced (Hennig and Liao 2013).

In this paper, we consider the problem of biased mixed variables distance. In particular, we propose a general formulation for multivariate mixed variable distances that can be used to construct unbiased mixed variable distances. To achieve this, we consider two properties: 1) *multivariate additivity* and 2) *commensurability*. Multivariate additivity means that the multivariate distance is the sum of univariate distances. The commensurability requirement ensures that the univariate contributions to the overall distance are commensurable (i.e., corresponding in size, extent, amount or degree). We show that existing mixed variable distances lead to biased distances, and we derive the magnitude of such biases. We evaluated several specific mixed variable distances and show that our new unbiased mixed variable distances provide a more objective starting point in an exploratory distance-based data analysis.

This paper is organized as follows. After introducing the two desirable properties for multivariate mixed variable distances in Section 2, we propose a general multivariate mixed variable

distance formulation that can be used to construct mixed variable distances that satisfy these properties. In Sections 4 and 5, we show how several common distance and scaling options – for numerical and categorical variables – violate the properties of Section 2 and hence may lead to biased mixed variable distances. In Section 6 we show, using our general distance formulation, how unbiased mixed variable distances can be constructed. Using a simulation study and an empirical application, we illustrate our findings in Sections 7 and 8, and end the article with our conclusions.

2 Properties of a multivariate mixed variable distance

A multivariate mixed variable distance requires the quantification and aggregation of distances for different variables. Aggregation concerns “joining” or the accumulation of variable-specific distances. For this aggregation, we only consider distances that are “additive”. We define multivariate additivity as the property that in a p –dimensional setting, the distance can be expressed as the sum of p distances. That is,

Definition 2.1 (Multivariate Additivity) *Let $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$ denote a p –dimensional vector. A distance function $d(\mathbf{x}_i, \mathbf{x}_l)$ between observations i and l is multivariate additive if*

$$d(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j=1}^p d_j(\mathbf{x}_i, \mathbf{x}_l),$$

where $d_j(\mathbf{x}_i, \mathbf{x}_l)$ denotes the j –th variable specific distance.

Note that the variable-specific distances d_j are functions of the p –dimensional vectors. These functions can be used to take into account associations between the p variables.

The Manhattan distance for $d(\mathbf{x}_i, \mathbf{x}_l)$ satisfies the additivity property, whereas the Euclidean distance, which takes the square root of the sum of squared differences, does not. Euclidean distance implicitly assumes that the variables jointly span a space. It is then defined by considering the length of a straight line connecting the points in this high-dimensional space. In a multivariate

context, the contribution per variable for the Euclidean distance decreases as more variables are included.

By defining the overall multivariate distance $d(\mathbf{x}_i, \mathbf{x}_l)$ as the sum of variable-specific distances, each variable-specific distance $d_j(\mathbf{x}_i, \mathbf{x}_l)$ is treated equivalently. This requires these distances to be on similar or equivalent scales. To achieve this we introduce an additional property: commensurability. In particular, we define distances as commensurable if the mean variable-specific distance between observations is constant across variables. That is,

Definition 2.2 (Commensurability) *Let $\mathbf{X}_i = (X_{i1}, \dots, X_{ip})$ denote a p -dimensional random variable corresponding to observation i . Furthermore, let d_j denote the distance function corresponding to the j -th variable. We have commensurability if, for all j , and $i \neq l$,*

$$E[d_j(\mathbf{X}_i, \mathbf{X}_l)] = c,$$

where c is some constant of which the actual value is of no importance.

In Definitions 2.1 and 2.2, we allow for variable-type-specific distance measures. In a mixed variable type setting, such flexibility is essential as the numerical and categorical variables require different distance functions.

3 A general multivariate mixed variable distance

To construct a mixed distance that satisfies multivariate additivity and commensurability, as defined in Definitions 2.1 and 2.2, we introduce a general multivariate mixed variable distance formulation.

Let $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$ denote a p -dimensional vector. We define the distance between observations i and l as

$$d(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j=1}^p d_j(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j=1}^p w_j \delta_j(\mathbf{x}_i, \mathbf{x}_l), \quad (1)$$

where δ_j is a function quantifying the dissimilarity between observations i and l based on the p -dimensional variables, and w_j is a weight corresponding to each of these functions.

To distinguish numerical from categorical variables, we decompose the overall distance with respect to the variable types. For a numerical variable j_n , the distance is given by

$$d_{j_n}(\mathbf{x}_i, \mathbf{x}_l) = w_{j_n} \delta_{j_n}^n(\mathbf{x}_i^n, \mathbf{x}_l^n), \quad (2)$$

where \mathbf{x}_i^n denotes the vector of p_n numerical elements of \mathbf{x}_i , $\delta_{j_n}^n$ is a function quantifying the dissimilarity between observations on the j_n -th numerical variable, and w_{j_n} is a scale value, or weight, for the j_n -th variable. Similarly, for a categorical variable j_c , the distance is

$$d_{j_c}(\mathbf{x}_i, \mathbf{x}_l) = w_{j_c} \delta_{j_c}^c(\mathbf{x}_i^c, \mathbf{x}_l^c), \quad (3)$$

where \mathbf{x}_i^c denotes the vector of p_c categorical elements of \mathbf{x}_i , $\delta_{j_c}^c$ corresponds to the category dissimilarity between the categories chosen by subjects i and l for categorical variable j_c , and w_{j_c} is a scale value, or weight, for the j_c -th variable. Hence, the category dissimilarities quantify the differences between categories. This decomposition avoids using the full mixed vector of observations as input for distance calculations. As a result, the distances for the numerical and categorical variables depend only on the values of the variables of the same type. Although a more general formulation, where interaction between the variables of different types is theoretically appealing, defining and quantifying such interactions is challenging and we do not pursue this here.

Combining Equations (1),(2) and (3), leads to the following formulation for mixed distance between observations i and l :

$$d(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j=1}^p w_j \delta_j(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j_n=1}^{p_n} w_{j_n} \delta_{j_n}^n(\mathbf{x}_i^n, \mathbf{x}_l^n) + \sum_{j_c=1}^{p_c} w_{j_c} \delta_{j_c}^c(\mathbf{x}_i^c, \mathbf{x}_l^c), \quad (4)$$

where p_n and p_c are the number of numerical and categorical variables, respectively.

This general formulation allows for a very flexible implementation of mixed distances. Gower's distance, for example, is a special case of this general distance. By definition, this general formulation satisfies the multivariate additivity property. Furthermore, although commensurability is not guaranteed, it can be achieved by applying appropriate variable specific weights. In particular,

if in (1) – and analogously in (4) – we set

$$w_j = 1/E[d_j(\mathbf{X}_i, \mathbf{X}_l)], \quad (5)$$

for all j and d_j , we obtain a distance that satisfies the multivariate additivity as well as the commensurability properties. Hence, using these weights we obtain an unbiased mixed distance where measurement types or units do not trivially impact the overall distance.

The expected values in (5) depend on the functions d_j and the distributions underlying the random variables. One may either consider appropriate assumptions concerning the underlying distributions, or follow a completely distribution-free approach to estimate the expected value based on the observed data. In Sections 4 and 5, we consider some properties of common distance and scaling options for numerical and categorical variables. In particular, we show how, depending on the underlying distributions, several popular choices for δ_j and w_j can lead to biased distances.

4 Numerical variables distances and scaling

There are many distance definitions for numerical variables. Typically, these distances are based directly on the observed numerical differences. For example, consider the Minkowski distance, defined as:

$$d_{mink}(\mathbf{x}_i^n, \mathbf{x}_l^n) = \left(\sum_{j_n=1}^{p_n} |x_{ij_n}^n - x_{lj_n}^n|^{1/m} \right)^m,$$

where $m \geq 1$. It is easily verified that this distance includes the Euclidean ($m = 2$) and Manhattan distance ($m = 1$) as special case. However, the multivariate additivity property is only satisfied for Manhattan distance. In this paper, we therefore only consider distances that apply Manhattan distances to the – possibly transformed – numerical data. That is, dropping the super- and sub-scripted n 's for ease of notation, we write

$$d_j(\mathbf{x}_i, \mathbf{x}_l) = w_j \delta_j(\mathbf{x}_i, \mathbf{x}_l) = w_j |f_j(\mathbf{x}_i) - f_j(\mathbf{x}_l)|,$$

where $f_j(\cdot)$ is a function that transforms the original values (e.g., some type of scaling). In our multivariate additive setting, inserting this distance for all p_n numerical variables gives, as total

distance due to the numerical variables

$$d_{num}(\mathbf{x}_i, \mathbf{x}_l) = \sum_{j=1}^{p_n} w_j |f_j(\mathbf{x}_i) - f_j(\mathbf{x}_l)|. \quad (6)$$

Differences between observations for different variables are aggregated through the weights w_j and functions $f_j(\cdot)$ that can also be used to account for differences in measurement scales and distributions. In multivariate contexts, common solutions for dealing with different measurement scales is to rescale the data using either the standard deviation, the range, or the robust (or interquartile) range. We refer to these options as independent scaling, as they do not involve the values of other variables. Alternatively, one could utilize the association between numerical variables to define distances. In the next section, we review common variants for scaling numerical variables and propose a new association-based scaling method that can be used in a multivariate additive context.

4.1 Scaling of numerical variables

Standard deviation scaling. We define standard deviation scaling as the procedure that transforms, for each variable separately, the measurements to standard deviations from the mean. That is,

$$f(\mathbf{x}_i) = f(x_{ij}) = \frac{x_{ij} - \bar{x}_j}{s_j},$$

where \bar{x}_j and s_j denote, respectively, the sample mean and standard deviation for variable j . This type of scaling converts all observations to so-called z -scores, and is a common procedure in many statistical analyses. Note that, regardless of the underlying distribution, by scaling in this way, we have: $E[X_i - X_l]^2 = 2$. Moreover, when the data are from a normal distribution, this conversion to z -scores converges to standard normal data.

Range scaling. For range scaling, which is used in Gower's distance, we subtract the minimum and divide by the range of a variable as follows

$$f(\mathbf{x}_i) = f(x_{ij}) = \frac{x_{ij} - x_j^{\min}}{x_j^{\max} - x_j^{\min}},$$

where x_j^{min} and x_j^{max} denote, respectively, the minimum and maximum for variable j . Consequently, all dissimilarities are between 0 and 1 and there is at least one dissimilarity equal to 1. This scaling retains the original shape of the distribution, however, it is very sensitive to outliers and may lead to a loss of information. In particular, if the data are spread out over a large interval but highly condensed, the majority of distances becomes small, making differentiation difficult.

Robust range scaling. For robust range scaling, we subtract the median and divide by the interquartile range. That is,

$$f(\mathbf{x}_i) = f(x_{ij}) = \frac{x_{ij} - x_j^{0.50}}{x_j^{0.75} - x_j^{0.25}},$$

where x_j^p denotes the p -th percentile for variable j . This method overcomes some of the problems of range scaling since it is less susceptible to the presence of outliers.

Principal component scaling. The multivariate additivity property implicitly assumes independence of the variables. However, in some cases one could argue that relationships between the numerical variables should be taken into account. Our formulation accommodates this by allowing the functions δ_j^n to take the complete (numerical) observation vector as their argument. Here, we propose a method to incorporate variable association in a multivariate additive setting.

When two variables are (linearly) related, it could be argued that distances orthogonal to the direction of the relationship should receive larger weights than distances along the direction of the relationship. To achieve this, we first rotate the original data to their principal coordinates, and then standardize them using standard deviations.

Let \mathbf{X} denote an $n \times p$ centered data matrix, of rank p , with \mathbf{S} the corresponding sample covariance matrix. Consider the eigen-decomposition

$$\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top,$$

where \mathbf{V} is orthonormal and $\mathbf{\Lambda}$ is a $p \times p$ diagonal matrix of eigenvalues. Then

$$\mathbf{X}^* = \mathbf{X}\mathbf{V}\mathbf{\Lambda}^{-1/2}$$

gives the data in standardized principal coordinates (Gower 1966). The columns of $\mathbf{V}\mathbf{\Lambda}^{-1/2}$ define p -dimensional functions $f_j(\mathbf{x}_i)$ that can be implemented in our distance formulation (6).

Furthermore, if the eigenvalues are organized in non-increasing order, the first rotated variable, i.e., the first column of \mathbf{X}^* , corresponds to the linear combination of the original variables that captures the most variation. The second column accounts for the most of the remaining variance, and so on. Note that we do not reduce the number of variables, as is typically the case when applying principal component analysis, but only rotate the data. By considering the principal coordinate orientation, the dimensions are independent.

Principal component scaling is closely related to the Mahalanobis distance. However, with Mahalanobis distance, the standardized principal coordinates are rotated back to the original orientation. That is,

$$\mathbf{X}^{*m} = \mathbf{X}\mathbf{S}^{-1/2} = \mathbf{X}\mathbf{V}\mathbf{\Lambda}^{-1/2}\mathbf{V}^\top = \mathbf{X}^*\mathbf{V}^\top.$$

When Euclidean distances are used, this rotation does not affect the distances. Hence, Euclidean distances based on \mathbf{X}^{*m} are equivalent to those based on \mathbf{X}^* . However, our multivariate additive distance is sensitive to rotation. We therefore retain the principal component orientation, \mathbf{X}^* . Consequently, in contrast to the other types of scaling considered, principal component scaling changes the orientation of the original variables, and our multivariate distance is additive with respect to the rotated rather than the original variables.

4.2 Commensurability of numerical variable distances

The scaling options described above can be used to transform variables to similar scales. Nevertheless, this type of scaling does not necessarily lead to commensurability as defined in Definition 2.2 since the underlying distributions also impact the expected values. The commensurability property for the numerical variables can be expressed as follows

$$E[d_j(\mathbf{X}_i^n, \mathbf{X}_l^n)] = w_j E[|f_j(\mathbf{X}_i^n) - f_j(\mathbf{X}_l^n)|] = c, \quad (7)$$

for all $i \neq l$ and $j \in \{1 \dots p_n\}$, and where c is some constant that depends on the chosen scaling as well as on the distribution of \mathbf{X}_l^n . To illustrate the effect of different distributions and scaling

options, we present simulation results for mean distances for various combinations of scaling and distributions in Table 1. Moreover, in Appendix A.1 and B we provide exact results concerning the mean values for Uniform distribution with range scaling, and for Normal distribution with standard deviation standardization. It can be verified that these exact results coincide with the results for $n = 500$ in the Table 1.

Table 1: Impact of distributions and scaling on mean distance $E[d_j(\mathbf{X}_i, \mathbf{X}_l)]$. Skewed refers to a $\chi_{1/2}^2$ distribution. For bimodal we considered $n/2$ draws from $\chi_{1/2}^2$ (censored at 10), and $n/2$ draws from $10 - \chi_{1/2}^2$ (censored at 0)

Distribution	Sample size	SD scaling	Range scaling	Rob. Range scaling
Normal	50	1.16	0.25	0.76
Normal	500	1.13	0.19	0.84
Uniform	50	1.16	0.35	0.69
Uniform	500	1.15	0.33	0.67
Skewed	50	0.82	0.17	1.57
Skewed	500	0.77	0.09	1.49
Bimodal	50	1.07	0.50	0.51
Bimodal	500	1.06	0.49	0.50

Table 1 illustrates the impact of the underlying distribution and the chosen scaling variant on the mean distances. In a multivariate setting, if all variables are scaled in the same way, and they are from the same distribution (possibly with different parameters; e.g., all Uniform or all Normal), the resulting values are on commensurable scales. However, when we have variables from different distributions, this is no longer the case. In particular, we see that for both range and robust range normalization the mean values vary considerably depending on the distribution. Standardization on the other hand, appears to be less sensitive to the underlying distributions. Still, even when standardization is used, distance calculations based on a combination of skewed and

uniform or normal distributed variables clearly biases the results towards the uniform or normal variables.

5 Categorical variable distances and scaling

Several distance definitions for categorical variable distance have been presented in the literature. Crucial for such distances is the quantification of differences between categories. Specifically, for two observations on one categorical variable, they either fall into the same category or different categories. In the case of differing categories, one straightforward approach is to assign a distance of 1 (with no difference being assigned a value of zero). However, there may be reasons to differentiate distances depending on the different categories. For instance, the color green might be considered closer to blue than to red, necessitating a more nuanced distance measure. This challenge in quantifying distances for categorical data has led to a multitude of different proposals. For an overview of several such proposals, refer to, for example, Boriah, Chandola, and Kumar (2008) and Šulc and Řezanková (2019).

In van de Velden, Iodice D’Enza, Markos, and Cavicchia (2024), a framework that allows a general and flexible implementation of categorical variable distances was introduced. Distances obtained using the framework are by construction multivariate additive. Here we use this framework to appraise properties of a selection of the categorical distances reviewed in Boriah, Chandola, and Kumar (2008) and van de Velden, Iodice D’Enza, Markos, and Cavicchia (2024). Additionally, we consider some new distance variants that are based on the indicator matrix representation of categorical variables.

5.1 Category dissimilarity

Crucial for the definition of categorical variable distances is a so-called category dissimilarity matrix Δ_j that quantifies, for each variable, the differences between categories. Simple matrix multiplication of an indicator matrix \mathbf{Z}_j (i.e., a binary matrix where each column represents a

category, and observed categories are indicated by ones) and Δ_j results in the categorical variable distance matrix. That is, for variable j , distances between observations can be calculated as

$$\mathbf{D}_j = \mathbf{Z}_j \Delta_j \mathbf{Z}_j^\top \quad (8)$$

In a multivariate setting, the distance simply becomes the sum of the variable specific distances over the variables. Hence, the distances satisfy the multivariate additivity property. From (8) it follows that different Δ_j 's correspond to different definitions of categorical variable distances.

In van de Velden, Iodice D'Enza, Markos, and Cavicchia (2024), a distinction is made between “independent” and “association-based” category dissimilarities. For independent category dissimilarities, one variable is treated at the time. The distribution over the categories for that variable can be taken into account to construct the dissimilarities but any possible association between the different variables is ignored. However, to better differentiate between categories, associations between categorical variables may be employed as well. The resulting distances can be referred to as association-based.

5.1.1 Independent category dissimilarities

Table 2 provides an overview of category dissimilarity matrices Δ_j . The independent dissimilarity matrices in the top panel are a subset of the distances reviewed in Boriah, Chandola, and Kumar (2008) whereas those in the bottom panel, which we label as “indicator” variants, concern new proposals based on the application of numerical variable distances to possibly scaled, indicator matrices.

The idea of the indicator-based category dissimilarities is to treat the binary coded data as “numerical”. In practice, this type of coding, in combination with Euclidean distance, appears to be not uncommon, (see, e.g., Boehmke and Greenwell 2020; Hennig and Liao 2013). As Euclidean distance is not multivariate additive, we replace it by the Manhattan distance. Distances resulting from applying Manhattan distances to differently scaled indicator matrices can then be implemented using (8) by appropriately defining the category dissimilarity matrices. Below

Table 2: Distances and corresponding category dissimilarity matrices or category dissimilarities δ_{ab} , where \odot indicates the Hadamard (i.e., elementwise) matrix product and $\log(\cdot)$ takes the logarithm of the parenthesized object and collects them in an object of the same size

Distance	Category dissimilarity matrix Δ_j (or its typical element δ_{ab} , for $a \neq b$)
Matching	$\Delta_{m_j} = \mathbf{1}\mathbf{1}^\top - \mathbf{I}$
Eskin	$\Delta_{e_j} = 2/q_j^2 \Delta_{m_j}$
Occurrence frequency (OF)	$\Delta_{OF_j} = \log(\mathbf{p}_j) \log(\mathbf{p}_j)^\top \odot \Delta_{m_j}$
Inverse OF	$\Delta_{IOF_j} = \log(n\mathbf{p}_j) \log(n\mathbf{p}_j)^\top \odot \Delta_{m_j}$
Indicator: No scaling	$\Delta_{d_j} = 2\Delta_{m_j}$
Indicator: Hennig-Liao scaling	$\Delta_{HL_j} = 2\eta_j \Delta_{m_j}$
Indicator: Standard deviation scaling	$\delta_{abj}^s = \sqrt{\frac{1}{q_j}} (s_a^{-1/2} + s_b^{-1/2})$
Indicator: Cat. dissimilarity scaling	$\Delta_{cds_j} = \frac{1}{q_j} \mathbf{S}_{j_d}^{-1/2} \Delta_{m_j} \mathbf{S}_{j_d}^{-1/2}$

we briefly show how different choices for scaling lead to the different indicator-based category dissimilarity matrices presented in Table 2.

Indicator: No scaling Applying Manhattan distance to a one-variable indicator matrix results in distances of either zero (same category was chosen) or 2, different categories are chosen. Hence, in a multivariate setting one can implement this distance by defining, for each variable j as category dissimilarity matrix $\Delta_j = 2\Delta_{m_j}$, where Δ_{m_j} is the matching category dissimilarity matrix of appropriate order.

Indicator: Hennig and Liao scaling Hennig and Liao (2013) propose to tackle the differences caused by different variable types by standardizing variances across variables. For a categorical variable with q_j categories, coded using a q_j -dimensional indicator matrix, it is proposed to scale

the indicator matrix in such a way that

$$\sum_{a \in \{1 \dots q_j\}} E[Z_{ia}^* - Z_{ia}^*]^2 = 2\phi, \quad (9)$$

where Z_{ia}^* denotes the scaled indicator variable for the i -th observation on the a -th category, and ϕ needs to be supplied by the researcher. Hennig and Liao (2013) suggest to use $\phi = 1/2$ for categorical variables and $\phi = (q_j - 1)/q_j$ for ordinal ones. Consequently, for the ordinal case, the expected value in (9) converges to 2, similar to the situation for numerical variables.

To satisfy (9), a scaling factor, say η_j , can be applied to each indicator matrix. That is, for the j -th categorical variable, $\mathbf{Z}_j^* = \eta_j \mathbf{Z}_j$. Distances are then calculated between the rows of the scaled indicator matrix. The Hennig and Liao scaling factor η_j is implemented in the function `distancefactor` in the R package `fpc` (Hennig 2024). In the Appendix we show how this η_j is calculated and derive limiting results.

Hennig and Liao (2013) use Euclidean distance between combined standardized numerical and rescaled indicator matrices, but this does not meet the multivariate additivity constraint. Switching to Manhattan distance solves this issue. For categorical variables, the Manhattan distance between rows of rescaled indicator matrices is either 0 if the categories match, or $2\eta_j$ if they differ. Thus, Hennig and Liao's scaling with Manhattan distance can be integrated into our mixed variable distance by constructing category dissimilarity matrices for each variable j as follows

$$\Delta_{HL_j} = 2\eta_j \Delta_{m_j}, \quad (10)$$

where Δ_{m_j} denotes the matching category dissimilarity matrix of appropriate order.

Indicator: standard deviation scaling By considering the columns of each indicator matrix, representing the categories of the categorical variable, as numerical variables that, together, represent one categorical variable, we can also apply the standard deviation scaling procedure for numerical variables to these indicator matrices. The sample covariance matrix corresponding to the indicator matrix \mathbf{Z}_j corresponding to a categorical variable with q_j categories, can be defined

as

$$\mathbf{S}_j = \frac{1}{n} \mathbf{Z}_j^\top \left(\mathbf{I} - (1/n) \mathbf{1} \mathbf{1}^\top \right) \mathbf{Z}_j = \mathbf{D}_{p_j} - \mathbf{p}_j \mathbf{p}_j^\top, \quad (11)$$

where \mathbf{p}_j denotes the vector of observed proportions and \mathbf{D}_{p_j} is a diagonal matrix with \mathbf{p}_j as its diagonal. We define a scaled indicator matrix as

$$\mathbf{Z}_j^* = \frac{1}{\sqrt{q_j}} \mathbf{Z}_j \mathbf{S}_{j_d}^{-1/2}, \quad (12)$$

where \mathbf{S}_{j_d} is a diagonal matrix with on its diagonal the diagonal elements of \mathbf{S}_j . Note that \mathbf{Z}_j^* is scaled such that the variance per category equals $1/q_j$. Consequently, the variance per variable is standardized to 1.

Standardizing all variables in this fashion, arranging them next to each other and considering the Manhattan distance between the rows of the resulting standardized “super”-indicator matrix, corresponds to distances obtained by applying (8), where the category dissimilarity matrices $\mathbf{\Delta}_j$ have as its elements the absolute values of the differences between the rows of the standardized indicator matrix \mathbf{Z}_j^* corresponding to different categories of variable j . That is,

$$\delta_{ab_j} = \frac{1}{\sqrt{q_j}} \left(s_a^{-1/2} + s_b^{-1/2} \right),$$

where, $a, b \in 1, \dots, q_j$ denote categories of variable j , and, generically, $s_a = p_{j_a}(1 - p_{j_a})$, denotes the a -th diagonal element of \mathbf{S}_{j_d} .

Indicator: Category dissimilarity scaling An alternative way that combines the categorical distance framework and scaling of the indicator matrix, is to directly substitute the scaled indicator matrix into (8) and apply the scaling to the category dissimilarity matrix rather than to the indicator matrices. Using standard deviation scaling and simple matching, we get

$$\mathbf{D}_j = \mathbf{Z}_j^* \mathbf{\Delta}_{m_j} \mathbf{Z}_j^{*\top} = \frac{1}{q_j} \mathbf{Z}_j \mathbf{S}_{j_d}^{-1/2} \mathbf{\Delta}_{m_j} \mathbf{S}_{j_d}^{-1/2} \mathbf{Z}_j^\top = \mathbf{Z}_j \mathbf{\Delta}_{cds_j} \mathbf{Z}_j^\top,$$

where

$$\mathbf{\Delta}_{cds_j} = \frac{1}{q_j} \mathbf{S}_{j_d}^{-1/2} \mathbf{\Delta}_{m_j} \mathbf{S}_{j_d}^{-1/2}. \quad (13)$$

Using this definition, individual pair-wise category dissimilarities become

$$\delta_{ab} = \frac{1}{q_j} [s_{asb}]^{-1/2}.$$

5.1.2 Association-based category dissimilarity

The categorical dissimilarities in Table 2 define category dissimilarities for each variable independent of the other variables. Alternatively, several authors, (e.g., Le and Ho 2005; Ahmad and Dey 2007; Jia, Cheung, and Liu 2014; Ring et al. 2015) have proposed distance measures for categorical variables that take into account the association between categorical variables. Crucial in the construction of such “association-based” category dissimilarities are the joint bivariate distributions. In van de Velden, Iodice D’Enza, Markos, and Cavicchia (2024), it was shown that extant association-based distances correspond to different ways of quantifying the differences between the conditional distributions. That is, the distributions for categories of one variable over the categories of another variable.

Let \mathbf{P}^{jk} denote the (empirical) joint distribution for categorical variables j and k . The conditional distribution for the row variable j can be obtained by dividing the elements of \mathbf{P}^{jk} by the marginal distribution for the column variable k . Using the indicator matrices as introduced before, we can derive the conditional distribution for the categories of variable j over the categories of variable k as

$$\mathbf{R}^{j|k} = \left(\mathbf{Z}_j^\top \mathbf{Z}_j \right)^{-1} \mathbf{Z}_j^\top \mathbf{Z}_k.$$

We can define the category dissimilarities for the categories of variable i , by quantifying the differences between the rows of $\mathbf{R}^{j|k}$.

For illustration purposes, we briefly review, in a bi-variate setting, two variants: 1) Kullback-Leibler variant (Le and Ho 2005) and 2) Total variance distance (Ahmad and Dey 2007; van de Velden, Iodice D’Enza, Markos, and Cavicchia 2024).

Kullback-Leibler Le and Ho (2005) propose to use a symmetric variant of Kullback-Leibler

divergence to define the category dissimilarities. That is, the category dissimilarity for categories a and b of variable j , is defined as the sum of the Kullback-Leibler divergences between rows \mathbf{r}_a and \mathbf{r}_b :

$$\begin{aligned}\delta_{ab} &= \text{KL}(\mathbf{r}_a, \mathbf{r}_b) + \text{KL}(\mathbf{r}_b, \mathbf{r}_a) \\ &= \sum_{l=1}^{q_k} \left[r_{al} \log_2 \left(\frac{r_{al}}{r_{bl}} \right) + r_{bl} \log_2 \left(\frac{r_{bl}}{r_{al}} \right) \right]\end{aligned}$$

where we dropped the superscribed $j|k$'s, KL denotes the Kullback-Leibler divergence between two discrete distributions and \log_2 denotes the base 2 logarithm.

Total variance distance Ahmad and Dey (2007) proposed a mixed variable distance that combined Euclidean distances for numerical variables with a proposal for the categorical variable distances that, as shown by van de Velden, Iodice D'Enza, Markos, and Cavicchia (2024), is equivalent to 1/2 times the Manhattan distances between rows of \mathbf{R} :

$$\delta_{ab} = 1/2 \|\mathbf{r}_a - \mathbf{r}_b\|_1,$$

where $\|\cdot\|_1$ denotes the L_1 norm. This is also known as the total variance distance between two discrete probability distributions.

Note that for these association-based measures, the category dissimilarities become zero if the two variables are independent. For a dependent scenario where the marginal probabilities coincide (i.e., we have the same number of categories and the joint distribution can be represented by a diagonal matrix), the category dissimilarity values using total variance distance is equivalent to the matching case: Δ_m . For Kullback-Leibler, perfect dependence leads to category dissimilarity matrix where all non-diagonal values require assessment of $r_{al} \log_2 \left(\frac{r_{al}}{r_{bl}} \right)$ where $r_{al} = 1$ and $r_{bl} = 0$. We can resolve this by replacing the zero conditional probability by a small value. The Kullback-Leibler implementation in the R package **philentropy** uses $r_{bl} = (1/10)^5$.

In a multivariate context, bivariate associations need to be aggregated. For instance, an average category dissimilarity can be calculated by considering all pairs of categorical variables or only

pairs with associations above a certain threshold may be considered. van de Velden, Iodice D’Enza, Markos, and Cavicchia (2024) describe several association-based methods, with more available in an R package on GitHub at https://github.com/alfonsoIodiceDE/catdist_package. It is important to note that different methods produce varying category dissimilarities, so the choice of method affects the overall distances. Care is needed when combining different variants.

5.2 Commensurability of categorical variable distances

The general framework for constructing categorical variable distances allows for a straightforward assessment of commensurability. In particular, let X_{ij} denote the i -th observation for categorical variable X_j . Let q_j denote the number of categories for this j -th variable and let \mathbf{p}_j denote the $q_j \times 1$ vector of category probabilities. Using (8) as distance function, the mean absolute distances can be expressed as follows

$$E[d_j(\mathbf{X}_i, \mathbf{X}_l)] = \sum_{a,b \in \{1 \dots q_j\}} p_{ja} p_{jb} \delta_{jab} = \mathbf{p}_j^\top \mathbf{\Delta}_j \mathbf{p}_j, \quad (14)$$

where $\mathbf{\Delta}_j$ is the matrix of category dissimilarities for variable j whose elements, δ_{jab} quantify the dissimilarity between categories a and b . Hence, the mean distance is a function of variable specific properties (the number of categories), distributional aspects (\mathbf{p}_j) and the choice of distance (the category dissimilarity matrix $\mathbf{\Delta}_j$). We illustrate this below by considering different combinations of distances and distributions. For instance, Table 3 shows category dissimilarities and expected values for one categorical variable where the distribution over categories is uniform. Moreover, for the association-based variants, the category dissimilarity matrices are presented corresponding to perfect bivariate dependence where the two variables have the same number of categories.

It is important to note that for uniformly distributed variables, all category dissimilarity matrices in Table 3 can be expressed as a function of the simple matching variant, that is, all category dissimilarities are constant. Furthermore, we see that, except for the scaled category dissimilarity matrix, the number of categories influences the mean absolute distances. Consequently, when

Table 3: Distances and corresponding category dissimilarity matrices when category probabilities are equal. That is, $p_j = 1/q$ for $j = 1 \dots q$. For Hennig-Liao we used $\phi = 1/2$ and the limiting result ($n \rightarrow \infty$) for η . For Inverse OF we used $n = 160$. For Kullback-Leibler, $\kappa = 5 \log_2(10) \approx 16.610$. We dropped all subscripted j 's for ease of notation

Distance	Cat. dissimilarity Δ_j	$E[d(X_i, X_l)]$	$q = 2$	$q = 5$
Matching	$\Delta_m = \mathbf{11}^\top - \mathbf{I}$	$\frac{q-1}{q}$	0.5	0.8
Eskin	$\Delta_e = 2/q^2 \Delta_m$	$\frac{2(q-1)}{q^3}$	0.250	0.064
Occurrence frequency (OF)	$\Delta_{OF} = \log^2(q) \Delta_m$	$\log^2(q) \frac{q-1}{q}$	0.240	2.072
Inverse OF	$\Delta_{IOF} = \log^2(n/q) \Delta_m$	$\log^2(n/q) \frac{q-1}{q}$	9.601	9.610
Indicator (no scaling)	$\Delta_d = 2\Delta_m$	$\frac{2(q-1)}{q}$	1	1.6
Indicator (Hennig-Liao scaling)	$\Delta_{HL} = \sqrt{\frac{2q}{q-1}} \Delta_m$	$\sqrt{\frac{2(q-1)}{q}}$	1	1.265
Indicator (St. dev. scaling)	$\Delta_s = 2\sqrt{\frac{q}{q-1}} \Delta_m$	$2\sqrt{\frac{(q-1)}{q}}$	1.414	1.789
Indicator (Cat. dissim. scaling)	$\Delta_{cds} = \frac{q}{q-1} \Delta_m$	1	1	1
Total Variance	$\Delta_{tvd} = \Delta_m$	$\frac{q-1}{q}$	0.5	0.8
Kullback-Leibler (Le & Ho)	$\Delta_{KL} = \kappa \Delta_m$	$\kappa \frac{q-1}{q}$	8.305	13.288

categorical variables with different numbers of categories are used, the overall distance is directly influenced by the number of categories. For non-uniform distributions, analytical results can also be obtained using (14) if a distribution is provided. In Appendix D we illustrate the effect of different distributions on the distances in Table 2 by considering, for different numbers of categories, different stylized distributions exhibiting various degrees of skewness.

6 Unbiased mixed variable distances

In Sections 4 and 5, we showed that depending on distributions, scaling and choice of distance function, multivariate distances typically do not satisfy the commensurability property even in

a single variable type setting. Furthermore, from the values in Tables 1 and 3 it is clear that combining numerical and categorical variable distances leads to variable type dependent biases.

To resolve these issues, we can apply weights based on the mean distance as defined in Equation (5). For numerical data, one can estimate the mean distances based on the observed data. That is, for given f_j , we estimate the mean distances as follows

$$m_{j_n} = \frac{1}{n} \sum_{i \neq l}^n d(\mathbf{x}_i^n, \mathbf{x}_l^n) = \frac{1}{n} \sum_{i \neq l}^n |f_j(\mathbf{x}_i^n) - f_j(\mathbf{x}_l^n)|.$$

For n sufficiently large, this estimate converges to the population mean. Hence, by defining weights as

$$w_{j_n} = 1/m_{j_n} \tag{15}$$

we obtain commensurability. We can therefore interpret the scaled measurement units of the variables as “mean distance” units.

Table 3 along with Figures 6 through 8 in Appendix D, show that for categorical variables the impact of the number of categories, underlying distributions and category dissimilarity definitions on mean distances can be significant. Consequently, in a multivariate setting, choosing the same category variable distance for all variables may not lead to commensurability. Similar to the case for numerical variables, we can resolve this by defining weights based on the mean distance. In particular, for a selected distance (i.e., Δ_j) and known distributions \mathbf{p}_j , we achieve commensurability by using

$$w_{j_c} = 1 / \left(\mathbf{p}_j^\top \Delta_j \mathbf{p}_j \right). \tag{16}$$

If, as is typically the case, \mathbf{p}_j is unknown, one can either use estimates based on theory, or estimate \mathbf{p}_j from the data. Note that through the different definitions of Δ_j , it is still possible to consider different, yet commensurable, distances.

Inserting the weights (15) and (16) into the general mixed variable distance (4) results in a family of unbiased distances. That is, regardless of the choice of distance for the numerical and categorical variables, the resulting distance is multivariate additive and commensurable.

7 Simulation study

As shown in Sections 4 and 5, individual variables may have varying impacts on overall distance, with observed differences being more or less influential depending on the data distribution, scaling method, or data type.

To illustrate the impact of different choices, we conduct a simulation study generating high-dimensional numerical data based on an underlying low-dimensional structure. We then discretize a subset of the variables to create mixed data. Next, we analyze how various mixed-distances and levels of discretization (i.e., number of categories) affect the results.

For the distances, we consider the following variants: 1) *Numerical data* Manhattan distance on the original generated data (without discretizing any variables) 2) *Naive approach*: Euclidean distance on the complete standardized matrix where categorical variables are transformed to indicator matrices (one-hot-encoding); 3) *Hennig-Liao*: Euclidean distance on the matrix with numerical variables standardized and categorical variables transformed to indicator matrices (one-hot-encoding) and scaled using the Hennig-Liao scaling factor; 4) *Additive Hennig-Liao*: Manhattan distance with numerical variables standardized and Hennig-Liao scaling of the indicator matrices (i.e., Δ_{HL}); 5) *Gower*: range normalized numerical variables and simple matching for categorical variables; 6) *Unbiased independent*: commensurable distance using simple matching for the categorical variables. 7) *Unbiased standardized*: commensurable distance using category dissimilarity scaling for the categorical variables; 8) *Unbiased dependent*: commensurable association-based mixed distance using PCA scaling of the numerical, and total variance distance for the categorical variables. Note that the variants 2) and 3) are not multivariate additive. We include them here as benchmark of methods used in practice.

7.1 Data generating process

For each simulation instance, we first generate a 2-dimensional configuration, \mathbf{Y} , by sampling $n \times 2$ times from a uniform distribution on $[-2, 2]$ and constructing an orthogonal basis. Next, we

post-multiply \mathbf{Y} by a random $2 \times p$ matrix, \mathbf{N} , with elements drawn from a uniform distribution on $[-2, 2]$. The resulting $n \times p$ matrix, $\mathbf{X}_o = \mathbf{Y}\mathbf{N}$, gives a numerical, p -dimensional “observed” matrix corresponding to the 2-dimensional configuration \mathbf{Y} . We add random noise to these observations using a normal distribution with $\sigma = 0.03$ (approximately half the standard deviation of the generated values). Finally, we discretize p_{cat} columns of \mathbf{X}_o by splitting their ranges into the desired number of intervals. This setting allows us to appraise the effect of the type of variable *within* a chosen mixed-distance variant, and with respect to retrieval of the underlying configuration.

7.2 Variable specific effects

We generate 100 instances, each with $n = 500$ observations on $p = 6$ variables. For the mixed datasets, two variables remain numerical, the other 4 variables are converted into 2, 3, 5 and 9 categories, respectively. To assess the effect of individual variables we consider the effect of leaving one-variable-out on the distance matrix. In particular, for each instance, we consider the absolute values of the differences between the leave-one-out distance and the distance using all variables. Furthermore, as the different mixed distance variants are on different scales, we also calculate a relative measure by dividing the absolute difference for each variable, by the sum of these differences over all variables.

The results are summarized in Figure 1. The left panel shows that, for the additive variants, the observed values align with the theoretical ones in Tables 1 and 3. For example, for the numerical variables, we see that the mean values for the non-commensurable additive variants (Hennig-Liao, numerical and Gower¹) are between the corresponding values in Table 1 for the normal and uniform distributions. (The generated data are a mixture of uniform and normal). For the categorical variables we see a similar correspondence to the values in Table 3. For the Euclidean variants, we did not derive theoretical mean values for individual variables as these depend on

¹The Gower implementation in `daisy` adds an additional scaling factor by dividing Gower distance by the number of variables. Here we did not apply this scaling

the number of variables. In general, for Euclidean distance, individual mean values decrease as the number of variables increases. We can see this effect by comparing the two Hennig-Liao implementations. For both variants, having more categories increases the distances due to that variable.

For the two independent unbiased variants (*Uind* and *Ustd*) the values are, by construction, 1. For the dependent variant, however, leaving a variable out also impacts the distances of the other categorical variables. Hence, for this variant the differences compared to the full distance are close to, but not exactly, 1.

To assess the contributions of variables of different types to the overall distance, we consider the relative variable contributions depicted in the right panel of Figure 1. We see that for the numerical data, as well as for the unbiased variants, the relative effect is either exactly $1/6$ or close to this value. For the other variants, distances are affected by type (i.e., numerical or categorical) as well as by the number of categories. For *Naive*, and, to a lesser extent, *Gower*, categorical variables contribute more to the overall distance. *Hennig-Liao* reduces the influence of categorical variables, increasing the relative contribution of numerical variables to the distance.

To see how these differences impact subsequent analyses, we consider the effect of individual variables on a 2-dimensional classical multidimensional scaling (MDS) solution. That is, for all distance variants, we compare the MDS solutions when leaving one-variable-out, with the solution using all variables. To measure the difference, we consider the alienation coefficient (Borg and Leutner 1985), which lies between zero and one and can be seen as a measure of unexplained variance. Hence, higher values indicate more influence on the solution. The results are visualized in Figure 2. Note that the effects on mean distances do not carry-over one-to-one to the effect on the low-dimensional configurations. In particular, although the categorical variables impact the MDS solution more than the numerical ones for *Gower* and *Naive*, the effect of the number of categories is reversed. That is, variables with fewer categories have more impact on the MDS solutions. In general, low category variables allow for little differentiation between observations




Figure 1: Effect of leaving variables out on distance. The first 4 variables (blue background) are categorical with indicated numbers of categories respectively. The last two are numerical. For the variants: G = Gower, HL = Hennig-Liao, HLa = Additive Hennig-Liao, Naive = Naive approach, Num = original data (all variables numerical), Udep = Unbiased dependent, Uind = Unbiased independent, Ustd = Unbiased standardized

and projecting to a low-dimensional space suffers from this for all distances.

7.3 Retrieval of the true configuration

To assess the effect of the different mixed-distances on retrieving the underlying configuration, we generate 100 instances, each with $n = 500$ observations on $p = 6$ variables. Three variables remain numerical, and three variables are all discretized into the same number of categories, i.e., either 2, 3, 5 or 9. For each dataset, we create a distance matrix using the same mixed distances as before. Next, we apply classical MDS and calculate the alienation coefficient of this solution with the true 2-dimensional configuration. The results are presented in Figure 3.

We see that different distances perform similarly when only binary variables are added. However, as the number of categories increases, the variation in results (as indicated by the sizes of the boxes in Figure 3) gets larger. Results for the unbiased variants appear to improve when




Figure 2: Effect of leaving variables out on the MDS solution. The first 4 variables (blue background) are categorical with indicated numbers of categories respectively. The last two are numerical. Variant abbreviations same as in Fig 1

more categories are used. *Unbiased dependent* performs better when dealing with more than 2 categories, as it effectively distinguishes categories based on their associations. Recall that the high-dimensional data ($p = 6$) were generated by expanding 2-dimensional data, therefore, meaningful associations are present. The association-based approach utilizes these dependencies.

The results in this plot are in line with the findings of the variable specific effects in Figure 1. For example, comparing Figures 1 and 2 we see how distances using *Hennig-Liao* over-emphasize the numerical variables. Similarly, by using range normalization, distance due to the numerical variables is undervalued in *Gower* distance. Consequently, when the true configuration relies more on numerical variables, *Gower* distance performs worse.

8 Illustration: FIFA player data

The simulation study confirmed, in a controlled setting, the theoretical findings of Sections 4 and 5. To further illustrate these findings, we apply a selection of the mixed variable distances to an empirical dataset. In particular, we consider a sample from the FIFA player data 2021

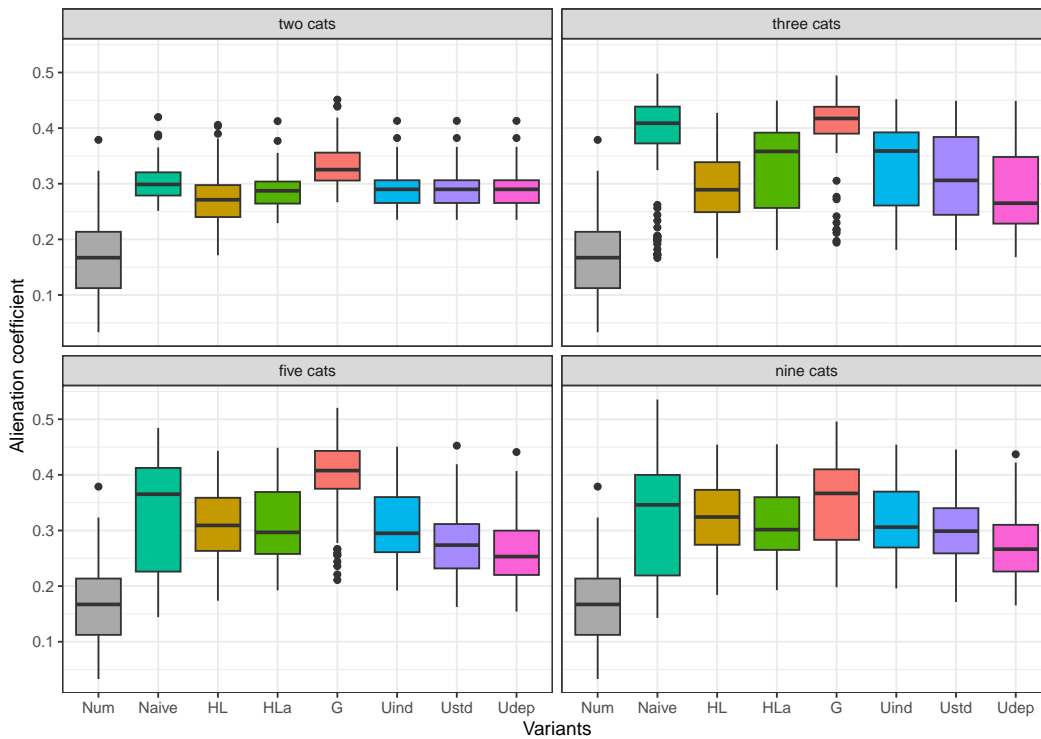


Figure 3: Distribution of alienation coefficients (with respect to the true configuration) of 2–dimensional classical MDS solutions using different mixed distances, where 3 out of 6 numerical variables have been discretized into the indicated number of categories

that can be found at [kaggle.com](https://www.kaggle.com)². In particular, we select 408 players from the Dutch league. Furthermore, we select 7 categorical and 7 numerical variables. Our selection of variables is based on our aim to illustrate how different mixed distances are affected by different measurement types and distributions. The categorical variables have categories ranging from 2 (“preferred foot”) to 25 (“position”). With the exception of the variable “club”, the distributions over categories are typically uneven. For the numerical variables, the wage and, particularly, transfer value variables are skewed. See Figure 9 in Appendix E for an overview of the distributions of the variables.

For this dataset, there is no known underlying “truth” concerning the distances or subsequent distance-based analysis. We can, however, illustrate how different choices for distances lead

²The complete link is <https://www.kaggle.com/datasets/stefanoleone992/fifa-21-complete-player-dataset>

to trivial, scale and measurement type-related, biases. In fact, based on the results presented in Sections 4 and 5, we expect that for all non-commensurable variants, distances due to the categorical variables increase with the number of categories. Furthermore, we expect that for Gower, the distance due to categorical variables exceeds that of the numerical variables. The Hennig-Liao scaling is designed to remove this bias towards categorical variables.

Similar to our analysis of simulated data, we first consider the effects of leaving individual variables out on the overall distance. Figure 4 presents the results where the left panel gives the absolute, and the right panel gives the relative contributions. We see that our expectations concerning the additive variants are confirmed. Furthermore, similar to our findings in the simulation study, we see that for the non-additive Euclidean variants, where variable specific effects depend on the number of variables, the scaling of categorical variables as proposed by Hennig-Liao, leads to an overemphasis of distances due to the numerical variables.




Figure 4: Variable importance with respect to full distance matrix. Left panel, absolute values, right panel, relative values

Similar to our analysis of variable specific effects in Section 7.2, we consider variable importance with respect to the distances as well as to the 2-dimensional MDS solutions. Note,




Figure 5: Variable importances for alienation coefficient with respect to full distance matrix

however, that there is no true underlying, numerical, configuration. Hence, we can only consider the effect of individual variables with respect to the MDS solution using all variables. The results for variable importance and alienation coefficients are presented in Figures 4 and 5 where, with the exception of the numerical variant– which is not possible in this truly mixed setting– we consider the same variants as in Section 7.

In Figure 4 we see that for the *Naive* variant, distances from categorical variables are much larger than from numerical ones, as expected. Similarly, *Gower* distance, shows a similar bias. In contrast, *Hennig-Liao* scaling with Euclidean distance over-corrects, making numerical variables overly dominant. We can also see, that a drop in mean values for the last two numerical values. This is due to the skewness of the distributions of these variables.

Important to note is that although for the unbiased variants the mean distances per variable are equivalent, when it comes to impact on the MDS solutions, there is variability. That is, commensurability does not mean that the variables play the same role in determining an underlying configuration. Or, more generally, that this is the case for any subsequent analysis. Furthermore, the biased variants not only lead to predictable mean variable distances, they also bias the MDS, or any other subsequent, analysis (Figure 5).

9 Conclusion

In this paper, we considered the problem of biased distance. That is, distances that by construction under- or over-represent contributions of variables based on variable type (e.g., numerical or categorical), distribution (e.g., symmetric or skewed) or scale (e.g., measurement scale or the number of categories). We introduced a general multivariate mixed variable distance that allows for an easy, yet flexible, way to correct for biased distances. Our proposed general distance allows for many different implementations. In particular, concerning distances due to categorical variables, such flexibility is desirable as there is no unique way to define distance for such variables.

Defining mixed-variable distances has recently received significant attention (Mousavi and Sehhati 2023; Ghashti and Thompson 2024; Boj and Grané 2024). Most methods propose a variable weighting system (Ahmad and Dey 2007; Hennig and Liao 2013; Mousavi and Sehhati 2023; Ghashti and Thompson 2024), or emphasize leveraging both numerical and categorical variable types (Huang 1998; Foss et al. 2016). Instead of adding another variant, we focus on the theoretical properties of distances for different variable types to construct a mixed distance formulation. This approach aims to address these issues in a general, rather than case- or application-specific, manner. Finally, it is important to emphasize that we do not claim optimality of any kind for our proposed method. In fact, the idea of an unbiased distance is that we treat all variables equivalently, that is, variable types, scales or measurement levels should not trivially impact the distance. If there is reason to believe that for a certain analysis, certain variables are more (less) important than others, this can and should be taken into account. However, in many, especially unsupervised, distance-based methods, there is no objective way of deciding, beforehand, whether this is the case.

Source code

The simulation study and the illustration example presented in this article can be reproduced following the instructions at https://alfonsoiodicede.github.io/blogposts_archive/Unbiased_mixed%20variables_distance_supplementary_mat.html.

References

- Ahmad, A. and L. Dey (2007). “A k-mean clustering algorithm for mixed numeric and categorical data”. In: *Data & Knowledge Engineering* 63.2, pp. 503–527.
- Amiri, Saeid, Bertrand S. Clarke, and Jennifer L. Clarke (2018). “Clustering Categorical Data via Ensembling Dissimilarity Matrices”. In: *Journal of Computational and Graphical Statistics* 27.1, pp. 195–208.
- Bello, Francesco de et al. (2021). “Towards a more balanced combination of multiple traits when computing functional differences between species”. In: *Methods in Ecology and Evolution* 12.3, pp. 443–448.
- Blanco-Mallo, E. et al. (2023). “Do all roads lead to Rome? Studying distance measures in the context of machine learning”. In: *Pattern Recognition* 141, p. 109646.
- Boehmke, Brad and Brandon M Greenwell (2020). *Hands-on machine learning with R*. Chapman and Hall/CRC.
- Boj, Eva and Aurea Grané (2024). “The robustification of distance-based linear models: Some proposals”. In: *Socio-Economic Planning Sciences* 95, p. 101992.
- Borg, I and D Leutner (1985). “Measuring the Similarity of MDS Configurations”. In: *Multivariate Behavioral Research* 20, pp. 325–334.
- Boriah, S., V. Chandola, and V. Kumar (2008). “Similarity measures for categorical data: A comparative evaluation”. In: *Proceedings of the 2008 SIAM international conference on data mining*. SIAM, pp. 243–254.

- Foss, Alex et al. (2016). “A semiparametric method for clustering mixed data”. In: *Machine Learning* 105, pp. 419–458.
- Ghashti, J.S. and J.R.J. Thompson (2024). “Mixed-Type Distance Shrinkage and Selection for Clustering via Kernel Metric Learning”. In: *Journal of Classification*. DOI: <https://doi.org/10.1007/s00357-024-09493-z>.
- Gower, John C. (1966). “Some Distance Properties of Latent Root and Vector Methods Used in Multivariate Analysis”. In: *Biometrika* 53.3/4, pp. 325–338.
- (1971). “A General Coefficient of Similarity and Some of Its Properties”. In: *Biometrics* 27.4, pp. 857–871.
- Hennig, Christian (2024). *fpc: Flexible procedures for clustering*. R package version 2.2-12. URL: <https://CRAN.R-project.org/package=fpc>.
- Hennig, Christian and Tim F. Liao (2013). “How to Find an Appropriate Clustering for Mixed-Type Variables with Application to Socio-Economic Stratification”. In: *Journal of the Royal Statistical Society Series C: Applied Statistics* 62.3, pp. 309–369.
- Hermes, Sjoerd, Joost van Heerwaarden, and Pariya Behrouzi (2024). “Copula Graphical Models for Heterogeneous Mixed Data”. In: *Journal of Computational and Graphical Statistics* 33.3, pp. 991–1005.
- Huang, Joshua Zhexue (1998). “Extensions to the k-Means Algorithm for Clustering Large Data Sets with Categorical Values”. In: *Data Mining and Knowledge Discovery* 2, pp. 283–304.
- Jia, H., Y. Cheung, and J. Liu (2014). “A New Distance Metric for Unsupervised Learning of Categorical Data”. In: *IEEE Transactions on Neural Networks and Learning Systems* 27, pp. 1065–1079.
- Le, S.Q. and T.B. Ho (2005). “An association-based dissimilarity measure for categorical data”. In: *Pattern Recognition Letters* 26.16, pp. 2549–2557. ISSN: 0167-8655.
- Loo, Mark van der (2017). *gower: Gower’s Distance*. R package version 0.1.2. URL: <https://CRAN.R-project.org/package=gower>.

- Luo, Hengrui et al. (2024). “Hybrid Parameter Search and Dynamic Model Selection for Mixed-Variable Bayesian Optimization”. In: *Journal of Computational and Graphical Statistics* 33.3, pp. 855–868.
- Maechler, Martin et al. (2023). *cluster: Cluster Analysis Basics and Extensions*. R package version 2.1.6 — For new features, see the ‘NEWS’ and the ‘Changelog’ file in the package source). URL: <https://CRAN.R-project.org/package=cluster>.
- Mousavi, E and M Sehhati (2023). “A generalized multi-aspect distance metric for mixed-type data clustering”. In: *Pattern Recognition* 138, p. 109353.
- Ring, M. et al. (2015). “ConDist: A Context-Driven Categorical Distance Measure”. In: *Machine Learning and Knowledge Discovery in Databases*. Ed. by A. Appice et al. Cham: Springer International Publishing, pp. 251–266.
- Šulc, Z. and H. Řezanková (2019). “Comparison of similarity measures for categorical data in hierarchical clustering”. In: *Journal of Classification* 36.1, pp. 58–72.
- van de Velden, Michel, Alfonso Iodice D’Enza, and Angelos Markos (2019). “Distance-based clustering of mixed data”. In: *Wiley Interdisciplinary Reviews: Computational Statistics* 11.3, e1456.
- van de Velden, Michel, Alfonso Iodice D’Enza, Angelos Markos, and Carlo Cavicchia (2024). “A general framework for implementing distances for categorical variables”. In: *Pattern Recognition* 153, p. 110547.

A Appendix

A.1 Mean distance results: Uniform distribution

Let $X \sim U(a, b)$ denote a random variable from a uniform distribution on the interval $[a, b]$, where a and b are known. Furthermore, let X^s , X^m , and X^{rr} denote the standardized, range normalized and robust range normalized versions of X . Then:

Theorem A.1 $E[|X_i^s - X_l^s|] = c = \frac{1}{3}\sqrt{12}(\approx 1.15)$

Theorem A.2 $E[|X_i^{rn} - X_l^{rn}|] = c = \frac{1}{3}$

Theorem A.3 $E[|X_i^{rr} - X_l^{rr}|] = c = \frac{2}{3}$

Proof A.1 *Standardization gives*

$$X^s = \frac{X^s - \mu}{\sigma_x} = \frac{X^s - (a + (b-a)/2)}{(b-a)/\sqrt{12}} \sim U\left(-\frac{\sqrt{12}}{2}, \frac{\sqrt{12}}{2}\right)$$

The expected value of the absolute difference between two independent uniform random variables on $[a, b]$ is $(b-a)/3$. Hence

$$E[|X_i^s - X_l^s|] = c = \frac{1}{3}\sqrt{12}.$$

■

Proof A.2 *Range normalization gives*

$$X^{rn} = \frac{X^{rn} - a}{(b-a)} \sim U(0, 1),$$

so that

$$E[|X_i^{rn} - X_l^{rn}|] = c = \frac{1}{3}.$$

■

Proof A.3 *Subtracting the median and dividing by the interquartile range for a $U(a, b)$ random variable gives*

$$X^{rr} = \frac{X^{rr} - (a + (b-a)/2)}{(b-a)/2} \sim U\left(\frac{2a}{b-a}, \frac{2b}{b-a}\right).$$

Resulting in

$$E[|X_i^{rr} - X_l^{rr}|] = c = \frac{2}{3}.$$

■

B Mean distance results: Normal distribution

Let $X \sim N(\mu, \sigma)$ denote a random variable from a normal distribution with known mean μ and standard deviation

sigma. Furthermore, let X^s, X^{rn} , and X^{rr} denote the standardized, range normalized and robust range normalized versions of X . We have

Theorem B.1 $E[|X_i^s - X_j^s|] = 2\sqrt{(1/\pi)} (\approx 1.13)$

Theorem B.2 $E[|X_i^{rn} - X_j^{rn}|] = \frac{1}{x^{0.75}} \sqrt{(1/\pi)} (\approx 0.84)$

Proof B.1 *The distribution of the difference between two identically distributed normal random variables is also normal with mean zero and standard deviation $\sqrt{2}\sigma$. Furthermore, the absolute value of a centered normal random variable, say Y is “folded normal” with as its expected value*

$$E[|Y|] = \sigma \sqrt{(2/\pi)}. \tag{17}$$

Using standard deviation standardization (assuming known σ) we have

$$X^s = \frac{X - \mu}{\sigma} \sim N(0, 1).$$

Hence,

$$E[|X_i^s - X_j^s|] = 2\sqrt{(1/\pi)}.$$

■

Proof B.2 *For the difference between normal random variables using robust range scaling, we have*

$$X_i^{rn} - X_j^{rn} \sim N\left(0, \sqrt{2} \frac{\sigma}{\text{IQR}}\right).$$

As the interquartile range IQR is linearly related to σ , we can set, without loss of generality, $\sigma = 1$ and use the interquartile range of a standard normal distribution. That is, $\text{IQR} = 2x^{0.75}$, where $P(X < x^{0.75}) = 0.75$.

As before, taking absolute values leads to a folded normal distribution. Inserting the appropriate expressions into 17 yields

$$E[|X_i^n - X_j^n|] = \frac{1}{x^{0.75}} \sqrt{(1/\pi)}.$$

■

C Hennig-Liao limiting results

Let T denote the total potential pairwise differences among all observations. In the function `distancefactor` from the R package `fpc`, this is calculated as $T = n(n + 1)/2$. Let W denote the cumulative “within-category” difference. This accounts for the number of ways two observations from the same category can be chosen, and is defined as

$$W = \sum_{a \in \{1 \dots q\}} \frac{n_a(n_a + 1)}{2},$$

where n_a denotes the number of observations with category a . Define B as the difference between the total and within category differences. That is,

$$B = T - W.$$

The normalization or scaling factor η is defined as

$$\eta = \sqrt{\phi \frac{T}{B}}, \tag{18}$$

where ϕ is user-supplied with as suggested values $1/2$ for categorical and $(q - 1)/q$ for ordinal variables. The normalization factor η depends on the sample distribution, the user-supplied ϕ , the number of categories q and the number of observations n . If all categories are equally likely, we can derive limiting results with respect to the sample size. In particular:

Theorem C.1

$$\lim_{n \rightarrow \infty} \eta(n, q, \phi) = \sqrt{\phi \frac{q}{(q - 1)}},$$

$$\lim_{n \rightarrow q} \eta(n, q, \phi) = \sqrt{\phi \frac{(q + 1)}{(q - 1)}}.$$

Proof C.1 *As each category is equally likely, the number of observations per category converges to $n_a = n/q$, for all $a \in \{1 \dots q\}$. Furthermore, for n large, $n_a + 1 \rightarrow n_a$ and $n + 1 \rightarrow n$. Consequently, $T = n(n + 1)/2 \rightarrow n^2/2$ and $W = qn_a(n_a + 1)/2 \rightarrow qn_a^2/2$. Hence, $B \rightarrow n_a^2q(q - 1)/2$. Inserting the limiting results for T and B into 18 gives the first limiting result of Theorem C.1. The second limit result can be verified directly by inserting $n = q$ and $n_a = 1$ into 18.*

■

When we have equal probabilities and n goes to infinity, the variance per variable equals

$$V(Z^*) = qV(Z_a^*) = q\eta^2V(Z_a) = q\eta^2\frac{1}{q}\left(1 - \frac{1}{q}\right) = \eta^2\left(\frac{q-1}{q}\right) = \phi.$$

Hence, equation (9) is satisfied.

D Categorical mean distances: skewed distributions

The expected mean distances for the categorical variables distances depend on the chosen distance as well as on the distributions \mathbf{p} . Table 3 gives results when all categories are equally likely. Here we consider distributions where $p_1 \in \{0.05, 0.1, 0.2, 0.33, 0.5, 0.66, 0.8, 0.9, 0.95\}$ and $p_j = (1 - p_1)/(q - 1)$, with $j = 2, \dots, q$, for $q \in \{2, 3, 5, 10\}$. For the association-based variants we again consider the perfect bivariate dependence case where the marginal probabilities coincide and the number of categories for the two variables is equal.

Figures 6 through 8 summarize the results for different numbers of categories and different levels of skewness for the distances from Table 3. Note that the scales for the different distances vary significantly. For example, the values for the OF and IOF variants are much larger than other values. Moreover, for the IOF distance, the number of observations is also of importance. Here we selected $n = 160$ for all cases. Increasing this would lead to larger values. Decreasing the value can lead to negative results for some scenarios as np becomes smaller than 1.

We see that for all variants the distribution as well as the number of categories influences the mean distances.




Figure 6: Mean values for independent categorical variable distances

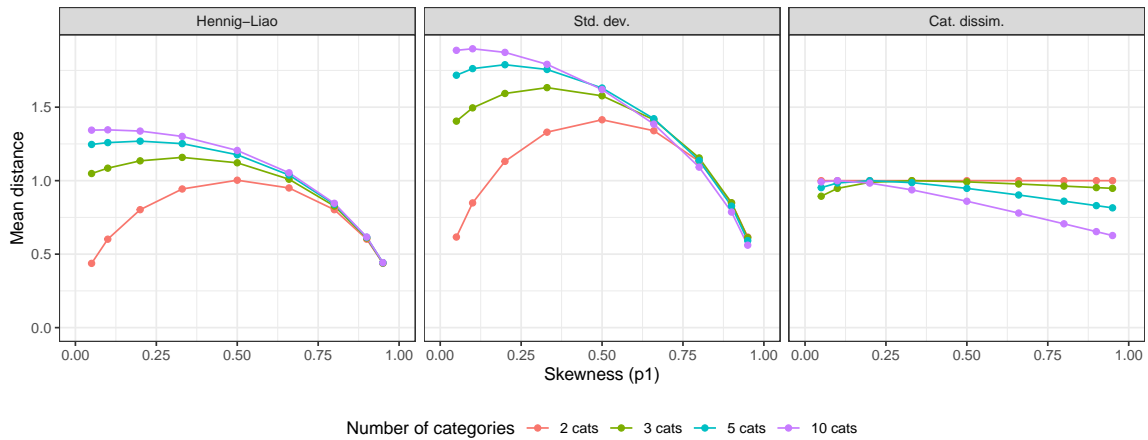


Figure 7: Mean values for independent categorical variable distances

Figure 8 illustrates the effect of the underlying marginal distributions for the dependent scenario. We see that the absolute size of the mean values differs considerably between methods. In particular, the values for the Kullback-Leibler variant are much larger. The pattern, however, is quite similar and for both variants the maximum mean distance occurs when all categories are equally likely, indicating maximum variance.

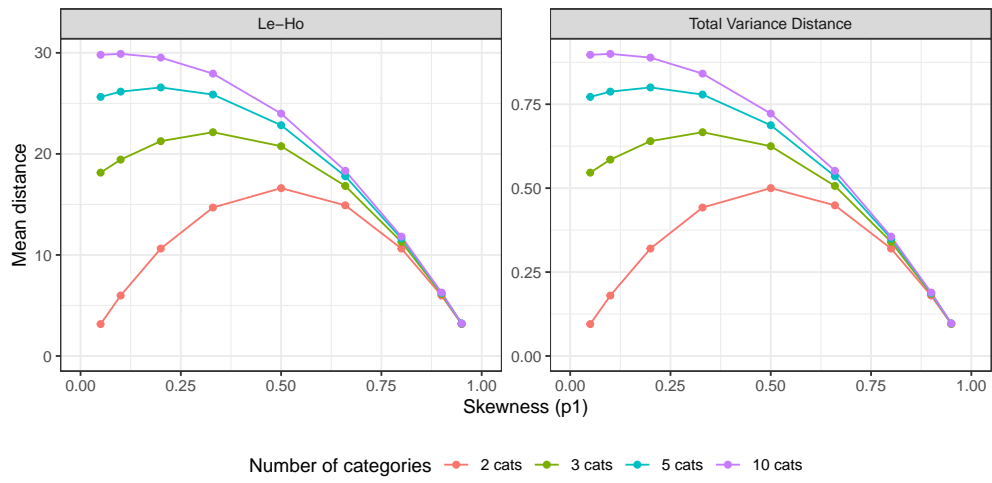


Figure 8: Mean values for association-based distances based on dependent bi-variate data

E FIFA variable distributions




Figure 9: Bar plots for the FIFA categorical variables (left), histograms for numerical variables (right)