



HAL
open science

Comparing Two Samples Through Stochastic Dominance: A Graphical Approach

Etor Arza, Josu Ceberio, Ekhiñe Irurozki, Aritz Pérez

► **To cite this version:**

Etor Arza, Josu Ceberio, Ekhiñe Irurozki, Aritz Pérez. Comparing Two Samples Through Stochastic Dominance: A Graphical Approach. *Journal of Computational and Graphical Statistics*, 2023, 32 (2), pp.551-566. 10.1080/10618600.2022.2084405 . hal-04253752

HAL Id: hal-04253752

<https://hal.science/hal-04253752>

Submitted on 23 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparing Two Samples Through Stochastic Dominance: A Graphical Approach

Etor Arza

BCAM - Basque Center for Applied Mathematics

and

Josu Ceberio

University of the Basque Country UPV/EHU

and

Ekhiñe Irurozki

Télécom Paris

and

Aritz Pérez

BCAM - Basque Center for Applied Mathematics

August 31, 2022

Abstract

Non-deterministic measurements are common in real-world scenarios: the performance of a stochastic optimization algorithm or the total reward of a reinforcement learning agent in a chaotic environment are just two examples in which unpredictable outcomes are common. These measures can be modeled as random variables and compared among each other via their expected values or more sophisticated tools such as null hypothesis statistical tests. In this paper, we propose an alternative framework to visually compare two samples according to their estimated cumulative distribution functions. First, we introduce a dominance measure for two random variables that quantifies the proportion in which the cumulative distribution function of one of the random variables stochastically dominates the other one. Then, we present a graphical method that decomposes in quantiles i) the proposed dominance measure and ii) the probability that one of the random variables takes lower values than the other. With illustrative purposes, we re-evaluate the experimentation of an already published work with the proposed methodology and we show that additional conclusions—missed by the rest of the methods—can be inferred. Additionally, the software package *RVCompare* was created as a convenient way of applying and experimenting with the proposed framework.

Keywords: Data visualization, Random variables, Cumulative distribution function, First-order stochastic dominance




Figure 1: Density estimates of the error rates produced by the optimizers *adam* and *RMSProp* in the MNIST dataset. The sci-kit learn Pedregosa et al. (2011) package was used in the estimation.

1 Introduction

The objective value obtained by an optimization algorithm may be non-deterministic. For example, in stochastic algorithms, the objective value measured depends on the seed used in the random number generator. In these kinds of scenarios, we can think that these non-deterministic measurements are observations of random variables with unknown distributions. Based on these measurements, we sometimes need to choose the random variable that takes the lowest (or largest) values. The expected values of the random variables—usually estimated as an average of several repeated observations—can be used for this purpose. However, many statisticians have claimed that summarizing data with simple statistics such as the average or the standard deviation is misleading, as very different data can still have the same statistics Matejka and Fitzmaurice (2017); Chatterjee and Firat (2007).

Motivating example 1) A real-world motivation for this work is as follows. Suppose we need to choose the best option between two stochastic gradient-based methods for optimizing the parameters of a neural network. A neural network classifier trained with a gradient-based method will produce different error rates Goodfellow et al. (2016) each time it is trained-tested, even if the same train-test dataset is used in each repeated measurement. One of the reasons is that the learned classifier depends on the initialization of its weights (before applying a gradient-based optimizer), which are often initialized randomly Glorot and Bengio (2010).

To illustrate the previous scenario, we trained and tested a neural network¹ in the MNIST dataset, and we compared two gradient-based optimizers in this data set: *adam* and *RMSProp* Goodfellow et al. (2016). The error rate in the test set depends on the seed used to train the neural network, and therefore, we can model the error rate of each of the algorithms in this problem as a random variable. An observation of each of the two random variables (the error rate of each gradient-based optimizer is modeled as a random variable) involves training the neural network in the training set and measuring its error rate in the test set: the training and test sets are the same for each trained neural network. Figure 1 shows the kernel density estimations of these random variables using the uniform kernel. As we see in the figure, the error rate is not the same in each measurement and ranges between 0.022 and 0.04. This shows that, in this context, it makes sense to model the error rate as a random variable rather than a constant: a unique value cannot represent the error rate without a significant amount of information loss.

Motivating example 2) In the following, we present another example with synthetic data. Let us consider the two random variables X_A and X_B shown in Figure 2. X_B has a lower expected value than X_A , $\mathbb{E}[X_B] < \mathbb{E}[X_A]$. If we use the expected value as the only criterion, then X_B takes lower values than X_A . However, notice that with a low but nonzero probability,

¹We follow an example in the Keras Chollet et al. (2015) library, and train the neural network for one epoch.




Figure 2: The probability density of two random variables X_A and X_B , with probability density functions $g_A = 0.925 \cdot g_{\mathcal{N}(0.210325, 0.002)} + 0.075 \cdot g_{\mathcal{N}(0.010325, 0.025)}$ and $g_B = 0.975 \cdot g_{\mathcal{N}(0.01875, 0.002)} + 0.025 \cdot g_{\mathcal{N}(0.06875, 0.001)}$ where $\mathcal{N}(\rho, \sigma)$ is the normal distribution with mean ρ and standard deviation σ . Their expected values are $\mathbb{E}[X_A] = 0.0205$ and $\mathbb{E}[X_B] = 0.02$ respectively.

X_B will take very large values that are undesirable in the context of minimization. Without loss of generality, in this paper, we assume that lower values are preferred.

An error with low variance is very important in an environment where reliability is key, even if it means a slightly worse expected value. Some examples include breast cancer detection Cruz-Roa et al. (2017), or some reinforcement learning tasks François-Lavet et al. (2018); Mnih et al. (2013) like self-driving cars Badue et al. (2021).

In other circumstances, obtaining the lowest possible error can be more important than reliability. One could argue that reliability is less important in sentiment analysis Zhang et al. (2018), or in certain real-world optimization problems Regnier-Coudert et al. (2016), where obtaining the best possible solution is key. When obtaining the best possible score is more important than reliability, it may even be worth running an optimization algorithm several times and choosing the best solution out of all the runs. In that case, X_A would also be preferred to X_B , as X_A has a higher probability of taking a value lower than 0.01 (see Figure 2).

Related work In these two examples, we have seen that summarizing and comparing random variables with only the expected value can leave important information out (such as which of the random variables can take lower values), especially when neither random variable clearly takes lower values than the other one. Many works in the literature use null hypothesis tests Mann and Whitney (1947); Conover and Conover (1980); Wilcoxon (1945) to analyze observed samples and choose one of the random variables accordingly. Nonetheless, as claimed in Benavoli et al. Benavoli et al. (2017), null hypothesis tests have their limitations too: when the null hypothesis is not rejected—this will happen often when the random variables being compared take similar values—, we get no information. Not only that but even when the null hypothesis is rejected, it does not quantify the amount of evidence in favor of the alternative hypothesis Benavoli et al. (2017).

Contribution In this paper, we propose a graphical framework that compares two random variables using their associated cumulative distribution functions, in the context of choosing the one that takes lower values. The proposed methodology can compare the scores of two stochastic optimization algorithms or the error rates of two classifiers, among other applications. To achieve this, the performances of the optimization algorithms (or the error rates of the classifiers) are modeled as random variables, and then, we compare them by measuring the *dominance*.

Specifically, we first propose 8 desirable properties for *dominance measures*: functions that compare two random variables in this context. From the measures in the literature, we find that the *probability that one of the random variables takes a lower value than the other random variable* satisfies most of these properties. In addition, we propose a new dominance measure, the *dominance rate*, that also satisfies most of the properties and is related to the first-order stochastic dominance Quirk and Saposnik (1962). Then, we propose a graphical method that involves visually comparing the random variables through these two dominance measures. The graphical method, named *cumulative difference-plot*, can also be used to compare the quantiles of the random variables, and it models the uncertainty associated with the estimate. By re-evaluating the experimentation of a recently published paper with the proposed methodology, we demonstrate that this new plot can be useful to compare two random variables, especially in the case when the random variables take similar values.

Finally, an R package named *RVCompare*, available in CRAN, is distributed alongside this paper. With this package, the *cumulative difference-plot* can be conveniently computed. The source code of the package and the supplementary material for the paper are available at github.com/EtorArza².

The rest of the paper is organized as follows: in the next section, we propose eight desirable properties for dominance measures. Then, in Section 3, we study two dominance measures that satisfy most of these properties. Section 4 introduces a graphical method to compare random variables. In Section 5, we discuss related methods in the literature and compare them to the proposed approach. Section 6, evaluates the proposed graphical method and other alternatives in an already published experimentation. In Section 7, we state the assumptions and limitations of the proposed *cumulative difference-plot*. Finally, Section 8 concludes the paper.

2 Desirable properties for dominance measures

2.1 Background

When we have two random variables and we need to choose the one that takes the lowest values, we usually take i) the random variable with the lowest expected value or ii) the random variable with the lowest median. The *median* Conover and Conover (1980) of a continuous random variable X_A , denoted as m_A , is the value that satisfies $\mathcal{P}(X_A < m_A) = \mathcal{P}(X_A > m_A)$. In other words, if m_A is the median of X_A , a sample of X_A is as likely to be lower than m_A as it is to be higher.

Interestingly enough, the median and the expected value have their strengths and weaknesses when it comes to choosing the random variable that takes the lowest values. In the following, we elaborate on this point with two particular cases of study. The first case is shown in Figure 3, with two random variables X_A and X_B . Each of the random variables is a mixture of two Gaussian distributions with the same shape and similar weight in the mixture. It is clear that X_A tends to take values lower than X_B , as the Gaussian distributions of X_A are centered in 0.05 and 0.07, while the Gaussian distributions of X_B are centered in 0.06 and 0.08. While the expected values of X_A and X_B are aligned with this intuition, the medians are not; as $\mathbb{E}[X_A] < \mathbb{E}[X_B]$ and $m_A > m_B$. However, the expected value does a poor job of summarizing the bimodal shape of X_A or X_B : both of these random variables usually take much higher or much lower values than their expected values.

The second case is shown in Figure 4. With a very high probability, X_A takes lower values than X_B , even though X_B will rarely take really low values, which might prove useful in some particular applications. In this case, $m_A < m_B$ and $\mathbb{E}[X_A] > \mathbb{E}[X_B]$, hence, the comparison of the medians are aligned with the intuition that X_A takes lower values than X_B , while the

²The source of the package *RVCompare* can be found at github.com/EtorArza/RVCompare. The code to reproduce every figure in the paper is available at github.com/EtorArza/SupplementaryPaperRVCompare.




Figure 3: Case 1. The probability density functions of X_A and X_B : $g_A = 0.489 \cdot g_{\mathcal{N}(0.05, 0.00125)} + 0.511 \cdot g_{\mathcal{N}(0.07, 0.00125)}$ and $g_B = 0.511 \cdot g_{\mathcal{N}(0.06, 0.00125)} + 0.489 \cdot g_{\mathcal{N}(0.08, 0.00125)}$ where $g_{\mathcal{N}(\rho, \sigma)}$ is the density function of the normal distribution with mean ρ and standard deviation σ .




Figure 4: Case 2. The probability density functions of X_A and X_B : $g_A = g_{\mathcal{N}(0.211325, 0.002)}$ and $g_B = 0.925 \cdot g_{\mathcal{N}(0.21875, 0.002)} + 0.075 \cdot g_{\mathcal{N}(0.04875, 0.002)}$ respectively.


expected values are not. In the presence of outliers Carreño et al. (2020), the median is considered more robust than the expected value Rousseeuw and Hubert (2011).

Notice that, in the second case, it is not trivial to choose between X_A and X_B , as X_B can take lower values, but X_A is more likely to be lower than X_B . So, when can we claim that one of them clearly takes lower values than the other? When the cumulative distribution of X_A is higher than the cumulative distribution of X_B in the entire domain of definition: in that case, X_A has a higher probability than X_B of taking values lower than x , for all x in the domain of definition. This is known Mann and Whitney (1947) as X_A being stochastically smaller than X_B . Depending on the field of study, it can also be referred to Schmid and Tiede (1996); Bennet (2013); Quirk and Saposnik (1962) as “ X_A stochastically dominates X_B ”³. The stochastic dominance can be further relaxed, obtaining what is known as *first-order stochastic dominance* in the literature Schmid and Tiede (1996); Quirk and Saposnik (1962), although, for the sake of brevity, we will call it *stochastic dominance* throughout the paper.


Definition 1. (*Stochastic dominance*) Let X_A and X_B be two continuous random variables defined in a connected subset $N \subseteq \mathbb{R}$. We say that X_A stochastically dominates X_B , denoted as $X_A \succ X_B$, when

- i) $G_A(x) \geq G_B(x)$ for all $x \in N$
and

³Without loss of generality, minimization is assumed in this paper.



(a) Case 1



(b) Case 2

Figure 5: The cumulative distributions of the two cases shown in Figures 3 and 4.

ii) *There exists an $x \in N$ such that $G_A(x) > G_B(x)$.*

where G_A and G_B are the cumulative distributions of X_A and X_B respectively.

For X_A not to stochastically dominate X_B (denoted as $X_A \not\preceq X_B$ ⁴), either condition i) or ii) must be violated. The special case that $X_A \not\preceq X_B$ and $X_B \not\preceq X_A$ at the same time is defined, it is said that X_A and X_B *cross* Benett (2013), and we denote it as $X_A \leq X_B$. In the non trivial ($X_A \neq X_B$) case that $X_A \leq X_B$, there exists two points $x_1, x_2 \in N$ such that $G_A(x_1) < G_B(x_1)$ and $G_A(x_2) > G_B(x_2)$: we cannot say, for all $x \in N$, that one of the random variables has a higher probability of taking values lower than x .

Let us now see how the cumulative distributions can be used to compare random variables in an example. In Figure 5a, the cumulative distributions of the random variables described in Figure 3 are shown. We can see that $G_A(x) > G_B(x)$ for almost all $x \in N$. But there is at least a point $x \in (0.06, 0.07)$ where $G_A(x) < G_B(x)$, hence, $X_A \leq X_B$. The same happens in the second case (Figure 5b). As in the previous case, $X_A \leq X_B$, because even though $G_A(x) > G_B(x)$ for almost all $x \in N$ (in which $g_A(x) \neq 0$ and $g_B(x) \neq 0$), for all $x \in (0.05, 0.2)$, $G_A(x) < G_B(x)$.

In the following, we will study how to quantify the difference between two random variables, emphasizing the degree to which one of the random variables stochastically dominates the other.

2.2 Desirable properties

There are many ways to compare two random variables, each with a different point of view: some aim to find how dissimilar two random variables are (disregarding which of them takes lower values), while other methods try to guess if one of the random variables stochastically dominates the other one. In the context of this paper, we are interested in measures that, given

⁴Note that $X_A \not\preceq X_B$ is not equivalent to $X_B \succ X_A$.

two random variables, quantify through the stochastic dominance how much one of the random variables tends to take lower values than the other. We use the term *dominance measure* to refer to functions that quantify the difference between two random variables following this intuition. In this section, we define eight desirable properties for these dominance measures, and we study the suitability of several measures from the literature.

Definition 2. Let X_A and X_B be two continuous random variables. We define a dominance measure between two random variables as a function \mathcal{C} that maps two random variables into a real value $\mathcal{C}(X_A, X_B)$.

It is desirable that $\mathcal{C}(X_A, X_B)$ quantifies the stochastic dominance. Hence, we want $\mathcal{C}(X_A, X_B)$ to be proportional to the portion of the support of X_A and X_B in which $G_A(x) < G_B(x)$. Formally, this intuitive idea can be represented as:

Property 1. \mathcal{C} is defined in the $[0, 1]$ interval, where:

$$i) \quad \mathcal{C}(X_A, X_B) = 1 \iff X_A \succ X_B$$

$$ii) \quad \mathcal{C}(X_A, X_B) = 0 \iff X_B \succ X_A$$

$$iii) \quad \mathcal{C}(X_A, X_B) \in (0, 1) \iff X_B \preceq X_A$$

Proposition 1. If a dominance measure \mathcal{C} satisfies Property 1 i) and ii), then it also satisfies Property 1 iii).

Proof. By definition, $X_B \preceq X_A$ iff $X_A \not\succeq X_B$ and $X_B \not\succeq X_A$. Property 1 i) and ii) implies that $X_A \not\succeq X_B$ and $X_B \not\succeq X_A$ iff $\mathcal{C}(X_A, X_B) \neq 1$ and $\mathcal{C}(X_A, X_B) \neq 0$. From Property 1 i) also $\mathcal{C}(X_A, X_B) \in [0, 1]$, thus $X_B \preceq X_A$ iff $\mathcal{C}(X_A, X_B) \in (0, 1)$. \square

Property 2. (Antisymmetry) $\mathcal{C}(X_A, X_B)$ and $\mathcal{C}(X_B, X_A)$ add up to 1.

$$\mathcal{C}(X_A, X_B) = 1 - \mathcal{C}(X_B, X_A)$$

It is noteworthy that Property 1 ii) can be inferred from Property 1 i) and Property 2.

Property 3. The inversion (under the sum) of the operands of \mathcal{C} equals the inversion of \mathcal{C} :

$$\mathcal{C}(-1 \cdot X_A, -1 \cdot X_B) = 1 - \mathcal{C}(X_A, X_B)$$

Property 4. When X_A and X_B are equal, \mathcal{C} is symmetric.

$$X_A = X_B \implies \mathcal{C}(X_A, X_B) = \mathcal{C}(X_B, X_A)$$

Assuming Property 2 holds, we can rewrite the previous property as:

$$X_A = X_B \implies \mathcal{C}(X_A, X_B) = 0.5.$$

Note that the opposite is not true:

$$\mathcal{C}(X_A, X_B) = \mathcal{C}(X_B, X_A) \not\implies X_A = X_B$$

Property 5. (Invariance to translation) Moving the domain of definition of X_A and X_B by the same amount does not change \mathcal{C} ⁵.

$$\text{for all } \lambda \in \mathbb{R}, \quad \mathcal{C}(X_A + \lambda, X_B + \lambda) = \mathcal{C}(X_A, X_B)$$

⁵We define $X_A + \lambda$ as the random variable that is sampled in two steps: first obtain an observation from X_A and then add λ to this observation. We define $\lambda \cdot X_A$ in a similar way.




Figure 6: Case 3. The probability density functions of X_A and X_B : $g_A = g_{\mathcal{U}(0.2,0.21)}$ and $g_B = 0.925 \cdot g_{\mathcal{U}(0.19,0.2)} + 0.075 \cdot g_{\mathcal{U}(0.04,0.05)}$ respectively, where $\mathcal{U}(0.2, 0.21)$ is the uniform distribution in the interval $(0.2, 0.21)$.

Property 6. (*Invariance to scaling*) *Scaling both X_A and X_B by the same positive amount does not change \mathcal{C} .*

$$\text{for all } \lambda > 0, \quad \mathcal{C}(\lambda \cdot X_A, \lambda \cdot X_B) = \mathcal{C}(X_A, X_B)$$

In the following lines, we give an intuition for Property 7. In Case 2, shown in Figure 4, we saw that for all $x \in (0.075, 0.2)$, $G_A(x) < G_B(x)$. However, notice that most of the mass of X_A and X_B is in the interval $(0.2, 0.23)$, where $G_A(x) > G_B(x)$. This means that most of the observed points of X_A and X_B will be in that interval. Therefore, it makes sense that $G_A(x) > G_B(x)$ has a higher weight than $G_A(x) < G_B(x)$ in the computation $\mathcal{C}(X_A, X_B)$. In other words, the *small* mass of X_B centered in 0.05 can only account for a *small* part of $\mathcal{C}(X_A, X_B)$. In what follows, this is formalized as X_B being a mixture of two distributions, where one of the distributions represents this small mass with a small weight in the mixture. Property 7 states that the change in the computation of \mathcal{C} produced by the distribution of small weight in the mixture can be, at most, its weight in the mixture.

Property 7. *Let $X_B = \mathcal{M}_{[1-\tau, \tau]}(X_{B1}, X_{B2})$ be the mixture⁶ distribution of X_{B1} and X_{B2} with weights $1 - \tau$ and τ respectively and let X_A be another random variable. Then,*

$$|\mathcal{C}(X_A, X_B) - \mathcal{C}(X_A, X_{B1})| \leq \tau$$

Property 8 explains that, under certain circumstances, $\mathcal{C}(X_A, X_B)$ is invariant to the translation/dilatation of one of the random variables. Specifically, it states that the distribution of one of the random variables (X_B) can change without affecting the value of $\mathcal{C}(X_A, X_B)$ as long as the changed part does not overlap with the support of the other random variable (X_A). Let us assume that the random variable X_B is defined as mixture distribution $\mathcal{M}_{[1-\rho, \rho]}(X_{B1}, X_{B2})$ where the supports of X_{B2} and X_A do not overlap, with $\rho \in (0, 1)$. Property 8 states that a translation and/or dilatation can be applied to X_{B1} , as long as: i) this transformation does not cause an overlap of the supports of X_A and X_{B2} , and ii) partial transformations will also not cause an overlap (hence the need for ξ_1 and ξ_2 in Property 8). In the following, we formalize this property:

Property 8. *Let $X_B = \mathcal{M}_{[1-\rho, \rho]}(X_{B1}, X_{B2})$ be the mixture distribution of X_{B1} and X_{B2} with weights $1 - \rho$, and ρ , respectively and let X_A be another random variable with $\rho \in (0, 1)$. Suppose that $\text{supp}(X_{B2}) \cap \text{supp}(X_A) = \emptyset$. Let $\lambda_1 \in \mathbb{R}^+$, $\lambda_2 \in \mathbb{R}$ be two numbers such that for all $\xi_1, \xi_2 \in$*

⁶The probability density function of $\mathcal{M}_{[1-\tau, \tau]}(X_{B1}, X_{B2})$ is defined as $(1 - \tau) \cdot g_{B1}(x) + \tau \cdot g_{B2}(x)$. Note that $\tau \in [0, 1]$.

$[0, 1]$, $\text{supp}((1 + (\lambda_1 - 1)\xi_1) \cdot X_{B2} + \xi_2\lambda_2) \cap \text{supp}(X_A) = \emptyset$. Then,

$$\mathcal{C}(X_A, X_B) = \mathcal{C}(X_A, \mathcal{M}_{[1-\rho, \rho]}(X_{B1}, \lambda_1 \cdot X_{B2} + \lambda_2))$$

This property can be applied to the distributions in Case 3 shown in Figure 6. For example, the probability mass in the interval $(0.04, 0.05)$ could have been centered in 0.1 or 0.15 instead of 0.045, without any changes to $\mathcal{C}(X_A, X_B)$. In addition to the position, the shape of the mass can also be altered as long as its weight in the mixture stays the same and does not overlap with X_A .

Unfortunately, it is impossible that a dominance measure satisfies Properties 1 and 7 at the same time. Intuitively, the problem is that, given the distributions X_A and $X_B = \mathcal{M}_{[1-\tau, \tau]}(X_{B1}, X_{B2})$, it is possible that $X_A \succ X_{B1}$ and at the same time $X_B \succ X_A$ with $\tau < 0.5$ ⁷. We formalize and prove this claim in the following proposition:

Proposition 2. *Let \mathcal{C} be a dominance measure.*

i) *If \mathcal{C} satisfies Property 1, then it fails to satisfy Property 7.*

ii) *If \mathcal{C} satisfies Property 7, then it fails to satisfy Property 1.*

Proof. A dominance measure only satisfies a property when that property is true for every possible random variable. Consequently, to prove this proposition, it is enough to find four random variables X_A, X_B, X_{B1} and X_{B2} where

- i) $X_B = \mathcal{M}_{[0.1, 0.9]}(X_{B1}, X_{B2})$,
- ii) $X_A \succ X_{B1}$,
- iii) $X_B \succ X_A$.

If four random variables can be found that satisfy these three statements, then with Property 1 we obtain that $\mathcal{C}(X_A, X_{B1}) = 1$ and $\mathcal{C}(X_A, X_B) = 0$. This contradicts Property 7, because $|\mathcal{C}(X_A, X_B) - \mathcal{C}(X_A, X_{B1})| \not\leq 0.1$. The same is true the other way around, Property 7 states that $|\mathcal{C}(X_A, X_B) - \mathcal{C}(X_A, X_{B1})| \leq 0.1$ and this contradicts Property 1, with $\mathcal{C}(X_A, X_{B1}) < 1$ or $\mathcal{C}(X_A, X_B) > 0$.

A simple example in which this happens is for the random variables

$$X_A = \mathcal{U}(0, 1),$$

$$X_B = \mathcal{M}_{[0.9, 0.1]}(\mathcal{U}(0.1, 1), \mathcal{U}(-0.5, 0)),$$

$$X_{B1} = \mathcal{U}(0.1, 1),$$

$$X_{B2} = \mathcal{U}(-0.5, 0).$$

The cumulative distribution functions of X_A, X_B and X_{B1} are shown in Figure 7, where it is clear that $X_B \succ X_A$ and $X_A \succ X_{B1}$. \square

In the following, we will briefly review several measures in the literature and, specifically, which of the proposed properties they satisfy. Many measures describe the difference between X_A and X_B , disregarding whether the difference in cumulative density is positive or negative. Consequently, they cannot satisfy Property 1 (see Appendix 9 for details). This is the case for *f-divergences*—including Kullback-Leibler, Jensen-Shannon, the Hellinger distance and the total variation—and for the Wasserstein distance. These measures also fail to satisfy several other properties (see a summary in Table 1).

⁷See <https://etorarza.github.io/pages/2021-interactive-comparing-RV.html> for an interactive example that illustrates the above point.




Figure 7: The cumulative distribution functions of X_A , X_B and X_{B1} .

Table 1: Which of the properties in Section 2.2 does each measure satisfy?

	1	2	3	4	5	6	7	8
Kullback-Leibler divergence				✓	✓	✓		✓
Jensen-Shannon divergence				✓	✓	✓		✓
Total-Variation				✓	✓	✓	✓	✓
Hellinger distance				✓	✓	✓	✓	✓
Wasserstein distance				✓	✓			
$\mathcal{C}_{\mathcal{P}}$: Probability of $X_A < X_B$		✓	✓	✓	✓	✓	✓	✓
$\mathcal{C}_{\mathcal{D}}$: Dominance rate of X_A over X_B	✓	✓	✓	✓	✓	✓		✓

A checkmark ✓ indicates that the measure satisfies the property.

3 Dominance measures

Most of the measures in the literature fail to satisfy the eight properties introduced in Section 2.2. However, there is a dominance measure in the literature that overcomes this limitation: the probability that $X_A < X_B$ Conover and Conover (1980).

3.1 $\mathcal{C}_{\mathcal{P}}$: the probability of $X_A < X_B$

We can compare X_A and X_B with the probability that a value sampled from X_A is smaller than a value sampled from X_B . When the random variables are exactly the same, this probability is 0.5. Formally, given two continuous random variables X_A and X_B defined in a connected set $N \subseteq \mathbb{R}$, the probability that $X_A < X_B$ is defined as:

$$\mathcal{P}(X_A < X_B) = \int_N g_B(x)G_A(x)dx. \quad (1)$$

When we consider $\mathcal{P}(X_A < X_B)$ as a dominance measure, we will denote it as $\mathcal{C}_{\mathcal{P}}(X_A, X_B)$.

One of the advantages of $\mathcal{C}_{\mathcal{P}}$ is its easy interpretation. In addition, $\mathcal{C}_{\mathcal{P}}$ is a well behaved dominance measure, as it satisfies Properties 2, 3, 4, 5, 6, 7 and 8. It also satisfies a weak version of Property 1:

$$\begin{aligned} \mathcal{C}_{\mathcal{P}}(X_A, X_B) = 1 &\implies X_A \succ X_B \implies \\ &\mathcal{C}_{\mathcal{P}}(X_A, X_B) \in (0.5, 1] \\ &\text{and} \\ \mathcal{C}_{\mathcal{P}}(X_A, X_B) = 0 &\implies X_B \succ X_A \implies \\ &\mathcal{C}_{\mathcal{P}}(X_A, X_B) \in [0, 0.5). \end{aligned}$$

Note that, when $X_A \succ X_B$, $\mathcal{C}_{\mathcal{P}}(X_A, X_B) \neq 1$ is still possible, and this is why it does not satisfy Property 1 entirely. For instance, when the probability densities of X_A and X_B are two Gaussian distributions with the same variance and the mean of X_A is lower, then $X_A \succ X_B$ but $\mathcal{C}_{\mathcal{P}}(X_A, X_B) < 1$.

So far, we have seen that $\mathcal{C}_{\mathcal{P}}$ satisfies most of the properties. Unfortunately, since it does not satisfy Property 1, not all cases of $X_A \succ X_B$ can be identified by $\mathcal{C}_{\mathcal{P}}$. We now propose a dominance measure that satisfies Property 1 and, thus, can be used to identify cases in which $X_A \succ X_B$.

3.2 $\mathcal{C}_{\mathcal{D}}$: dominance rate

Intuitively, the *dominance rate* is a dominance measure that quantifies the extent to which X_A has a lower cumulative distribution function than X_B , normalized by the portion of the probability densities with different cumulative distributions.

Definition 3. (*Dominance density function*) Let X_A and X_B be two continuous random variables defined in a connected set $N \subseteq \mathbb{R}$. We define the dominance density function as follows:

$$\mathcal{D}_{X_A, X_B}(x) = \begin{cases} g_A(x) \cdot k_A & \text{if } G_A(x) > G_B(x) \\ -g_B(x) \cdot k_B & \text{if } G_A(x) < G_B(x) \\ 0 & \text{otherwise.} \end{cases}$$

where $k_A = \left(\int_{\{x \in N \mid G_A(x) \neq G_B(x)\}} g_A(t) dt \right)^{-1}$ is the normalization constant and k_B is defined likewise.

Note that the dominance density function is not correctly defined when $\int_N |g_A(x) - g_B(x)| dx = 0$.

Definition 4. (*Dominance rate*) Let X_A and X_B be two continuous random variables defined in a connected set $N \subseteq \mathbb{R}$. The dominance rate of X_A over X_B is defined as

$$\mathcal{C}_{\mathcal{D}}(X_A, X_B) = \begin{cases} 0.5, & \text{if } \int_N |g_A(x) - g_B(x)| dx = 0 \\ 0.5 \int_N \mathcal{D}_{X_A, X_B}(t) dt + 0.5, & \text{otherwise.} \end{cases}$$

Basically, we are measuring the amount of mass of X_A in which $G_A(x) > G_B(x)$ minus the amount of mass of X_B in which $G_A(x) < G_B(x)$. This value is then normalized so that all sections in which $G_A(x) = G_B(x)$ are ignored, i.e. $\int_N \mathcal{D}_{X_A, X_B}(t) dt =$

$$\frac{\mathbb{E}_A[\mathcal{I}[G_A(x) > G_B(x)]]}{\mathbb{E}_A[\mathcal{I}[G_A(x) \neq G_B(x)]]} - \frac{\mathbb{E}_B[\mathcal{I}[G_A(x) < G_B(x)]]}{\mathbb{E}_B[\mathcal{I}[G_A(x) \neq G_B(x)]]}$$

Finally, we apply the linear transformation $l(x) = 0.5x - 0.5$ ensuring the dominance rate is defined in the interval $[0, 1]$ (instead of $[-1, 1]$), required to comply with Property 1.

From

i) $\mathcal{C}_{\mathcal{D}}(X_A, X_B) = 1 \iff X_A \succ X_B$ and

ii) $\mathcal{C}_{\mathcal{D}}(X_A, X_B) = 0 \iff X_B \succ X_A$,

we deduce that the dominance rate satisfies Property 1. Note that the previous deduction is only possible when g_A and g_B are bounded, as this implies that G_A and G_B are continuous. Specifically, it is enough to find a point in N in which $G_A(x) > G_B(x)$ to satisfy that

$\int_{x \in \{t \in N \mid G_A(t) > G_B(t)\}} g_A(x) dx > 0$, and this point is guaranteed to exist when $X_A \succ X_B$ because of the definition of the dominance. The dominance rate is also a well behaved dominance measure, as it satisfies Properties 1, 2, 3, 4, 5, 6 and 8.

We have seen that the dominance measures $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ satisfy most of the properties listed in Section 2.2. As we will see in the next section, their values are related.

3.3 The relationship between $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$

In Section 2.1 we stated that $\mathcal{C}_{\mathcal{P}} = 1$ is a stronger condition than $\mathcal{C}_{\mathcal{D}} = 1$, because $\mathcal{C}_{\mathcal{P}}(X_A, X_B) = 1$ implies that for all x in N that $G_A(x) < 1$, $G_B(x) = 0$. On the other hand, $\mathcal{C}_{\mathcal{D}} = 1$ implies that $X_A \succ X_B$ (the two conditions in Definition 1), which is weaker. In the diagram below, we show the values of $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ that imply other values of $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$. Each arrow can be interpreted as an implication. The implications are transitive: i.g. $\mathcal{C}_{\mathcal{D}}(X_A, X_B) = 1$ implies $\mathcal{C}_{\mathcal{P}}(X_A, X_B) > 0.5$.




Figure 8: Implications between the values of $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$.

3.4 Estimating $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$

In the previous sections, we have assumed that the random variables X_A and X_B are known, but usually, we only have a few observed values from each random variable. Therefore, it may be interesting to estimate $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ from the observed samples. With this purpose, we propose the following empirical estimates of $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$.

Definition 5. (*estimation of $\mathcal{C}_{\mathcal{P}}$*)

Let X_A and X_B be two continuous random variables and $A_n = \{a_1, \dots, a_n\}$ and $B_n = \{b_1, \dots, b_n\}$ their n observations respectively. We define the estimation of the probability that $X_A < X_B$ as

$$\widetilde{\mathcal{C}}_{\mathcal{P}}(A_n, B_n) = \sum_{i,k=1 \dots n} \frac{\text{sign}(b_k - a_i)}{2n^2} + \frac{1}{2}.$$

This estimator is well known in the literature because it is the U statistic of the Mann-Whitney test Mann and Whitney (1947).

Definition 6. (*estimation of $\mathcal{C}_{\mathcal{D}}$*)

Let X_A and X_B be two continuous random variables and $A_n = \{a_1, \dots, a_n\}$ and $B_n = \{b_1, \dots, b_n\}$ their n observations respectively. Let $c_{2n} = \{c_j\}_{j=1}^{2n}$ be the list of all the sorted observations of A_n and B_n where c_1 is the smallest observation and c_{2n} the largest. Suppose that $a_i \neq b_k$ for all $i, k = 1, \dots, n$. We define the estimation of the dominance rate as

$$\widetilde{\mathcal{C}}_{\mathcal{D}}(A_n, B_n) = \sum_{j=1}^{2n} \frac{\mathcal{I}(\hat{G}_A(c_j) > \hat{G}_B(c_j) \wedge c_j \in A_n)}{2n}$$

$$\sum_{j=1}^{2n} \frac{\mathcal{I}(\hat{G}_A(c_j) < \hat{G}_B(c_j) \wedge c_j \in B_n)}{2n} + \frac{1}{2}.$$

where \mathcal{I} is the indicator function and $\hat{G}_A(x)$ and $\hat{G}_B(x)$ are the empirical distributions estimated from A_n and B_n respectively.

For simplicity, this estimator of the dominance rate assumes there are no repeated samples. However, it can be extended to take into account repeated values (see Appendix 11).




Figure 9: An example of the probability density functions of Y_A and Y_B given the observed samples $A_n \cup B_n$.

4 Cumulative difference-plot

In this section, we propose a graphical method called *cumulative difference-plot* that shows the estimations of $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ decomposed by quantiles: $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ can be visually estimated from the difference plot. In addition, the proposed plot allows a comparison of quantiles of the two random variables. The proposed approach also models the uncertainty associated with the estimation of the *cumulative difference-plot* from the data.

4.1 Quantile random variables

From a practical point of view, it is unlikely that the probability densities of the compared random variables X_A and X_B are known. Usually, we only have n observations $A_n = \{a_1, \dots, a_n\}$ and $B_n = \{b_1, \dots, b_n\}$ from each random variable. The proposed *cumulative difference-plot* is based on two random variables Y_A and Y_B that are defined with these observations. Specifically, we define the densities of the two *quantile random variables* Y_A and Y_B as a mixture of several uniform distributions in the interval $[0, 1]$.

The uniform distributions in the quantile random variables are placed according to their rank in $A_n \cup B_n$. Assuming no repetitions, for each value k in $A_n \cup B_n$, its corresponding kernel is centered in $\frac{\text{rank}(k)+0.5}{2n}$ where $\text{rank}(k)$ is the rank of k in $A_n \cup B_n$. The kernels have a bandwidth of $1/4n$, ensuring that the sum of the densities of Y_A and Y_B is constant. If there are repeated values in $A_n \cup B_n$, their corresponding kernel is placed at the middle of the previous and the next rank, and the width of the kernel is increased proportionally with respect to the number of repetitions. See Figure 9 for an example. In Appendix 10.1 we show how to compute the probability densities of Y_A and Y_B step by step.

A more simple approach would be to estimate and define the quantile random variables through the empirical cumulative distribution functions of the observed samples of X_A and X_B . However, the quantile random variables defined through uniform kernels have some interesting properties: they have the same $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ as the kernel density estimations of X_A and X_B (shown in Appendix 10.2). In addition, $g_{Y_A}(x) + g_{Y_B}(x) = 2$ for all $x \in [0, 1]$. As we will later see, these properties are essential for the interpretation of the *cumulative difference-plot*.




Figure 10: The confidence bands of the cumulative distributions of the quantile random variables Y_A and Y_B corresponding to the distributions X_A and X_B in Case 2 (shown in Figure 4).

4.2 Confidence bands

The *cumulative difference-plot* is based on the cumulative distribution functions of Y_A and Y_B , which are estimated from the observed samples. This means that we need to model the uncertainty associated with the estimations. Confidence bands are a suitable choice in this scenario: a confidence band is a region in which the cumulative distribution is expected to be with a certain confidence. The size of the band is determined by the number of samples and the desired level of confidence: a high number of samples or a low level of confidence are associated with a small band size. There is an extensive literature Cheng and Iles (1983); Steck (1971); Cheng and Lies (1988); Wang et al. (2013); Faraway and Myoungshic Jhun (1990); Bickel and A. M. Krieger (1989); Hall and Horowitz (2013) on how to estimate the confidence bands of cumulative distributions, and, in this work, we use a simple bootstrap approach⁸.

To illustrate how to interpret the confidence bands of the cumulative distributions of Y_A and Y_B , we will assume that we have observed $n = 400$ samples from each random variable X_A and X_B from Case 2 (see Figure 4 in Section 2.1). We show the 95% confidence bands of the cumulative distribution functions of Y_A and Y_B in Figure 10. The estimated cumulative distribution functions of Y_A and Y_B resemble the cumulative distribution functions of X_A and X_B from Figure 5b. However, there are several relevant differences. In Figure 10, we observe that Y_A and Y_B are defined in the interval $[0, 1]$, while the cumulative distribution functions of X_A and X_B are defined in the sample space. Each of the values in this interval can be used to deduce the distribution with the lowest quantile: at $x = 0.5$, the cumulative distribution function of Y_A is larger than the cumulative distribution function of Y_B , hence, the median of X_A is lower than the median of X_B . In addition, the sum of the density function of Y_A and Y_B is constant. As a result, unlike X_A and X_B , the probability density functions of Y_A and Y_B do not have large areas where the probability density is zero.

⁸The bootstrapping Efron and Tibshirani (1993) method involves considering the observed values as a population from which random samples with replacement are drawn. These samples are then used to estimate the upper and lower pointwise confidence intervals of the cumulative distribution of Y_A and Y_B . Since a pointwise estimation of the confidence interval is used, we can expect that a portion proportional to α will fall outside the confidence band.

Note that we are interested in having an overall confidence of $1 - \alpha$, thus, we want that the cumulative distributions of Y_A and Y_B are inside their confidence bands at the same time with this level of confidence Goeman and Solari (2014); Bauer (1991). This means that we have to use a higher confidence level for each band: $\sqrt{1 - \alpha}$.

4.3 The cumulative difference-plot

In this section, we introduce a new graphical method designed to visually analyze the dominance of X_A and X_B . Without loss of generality, a minimization⁹ setting is assumed: lower values in X_A and X_B are preferred to higher values. It builds upon the difference function defined as

$$\begin{aligned} \text{diff}: [0, 1] &\longrightarrow [-1, 1] \\ x &\longmapsto G_{Y_A}(x) - G_{Y_B}(x), \end{aligned} \tag{2}$$

where $G_{Y_A}(x)$ and $G_{Y_B}(x)$ are the cumulative distribution functions of Y_A and Y_B , respectively.

The *cumulative difference-plot* is the plot of the difference function (the difference between the cumulative distributions of Y_A and Y_B), including a confidence band. A positive value in the *cumulative difference-plot* can be interpreted as a quantile in which the cumulative distribution function of X_A is larger than the cumulative distribution function of X_B . Hence, if the difference is positive at 0.5, the median of X_A is lower than the median of X_B (assuming minimization). In this sense, the best values obtained by both random variables are compared on the left side, and the worst values are compared on the right side.

4.3.1 \mathcal{C}_P and \mathcal{C}_D in the cumulative difference-plot

\mathcal{C}_P and \mathcal{C}_D can be directly obtained from the proposed plot. The integral of the difference between Y_A and Y_B is $\mathcal{C}_P - 0.5$ (we prove this in Appendix 11). Formally, $\mathcal{C}_P = 0.5 + \int_0^1 \text{diff}(x) dx$. However, in practice, \mathcal{C}_P can be visually estimated by adding 0.5 to the difference in the areas over and under 0. For the example shown in Figure 11, $\mathcal{C}_P = 0.5 - \text{Area1} + \text{Area2}$. The difference can only be in the area highlighted in blue in the cumulative difference-plot. When the probability that $X_A < X_B$ is 1, the difference is at its maximum: in the cumulative difference-plot we see a line from $(x, f(x)) = (0, 0)$ to $(0.5, 1)$ and from $(0.5, 1)$ to $(1, 0)$. Similarly, when the probability that $X_A < X_B$ is 0, the difference between Y_A and Y_B is equal to the lowest possible values inside the light blue area.

By contrast, \mathcal{C}_D is represented in the plot as the total length in which the difference is positive minus the total length in which the difference is negative. Specifically,

$$\mathcal{C}_D = \frac{\int_0^1 \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx}{\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx} + \frac{1}{2}, \tag{3}$$

where \mathcal{I} is the indicator function (we prove this in Appendix 11). As an example, \mathcal{C}_D is proportional to $\text{Length2} - \text{Length1}$ in Figure 11: it is higher than 0.5, because $\text{Length2} > \text{Length1}$. In this example, there is no need to divide by the total length in which the difference is nonzero because the difference is zero in only a limited number of points. In such cases, \mathcal{C}_D can also be estimated as the total length in which $\text{diff}(x) > 0$. In the example in Figure 11, the estimation is $\mathcal{C}_D = \text{Length2} \approx 0.75$. Note that Equation (3) is not correctly defined when Y_A and Y_B are equal, but this is an easy case to identify, as the difference is constantly 0.

⁹Note that if the random variables being compared take values in a maximization setting (higher values are preferred), then the random variables need to be redefined as the inverse with respect to the sum (this simply means the sampled values are multiplied by -1) before generating the cumulative difference-plot. With this change, the interpretation of the cumulative difference-plot is consistent and intuitive: for either minimization or maximization, on the left side of the cumulative difference-plot, the most desirable values that the random variables take are compared. If the difference is positive on the left side of the cumulative difference-plot, then the best values that X_A takes are better than the best values that X_B takes. Similarly, the worst values are compared on the right side of the cumulative difference-plot: if the difference is positive on this side, then the worst values of X_A are better than the worst values of X_B .




Figure 11: The areas and lengths in the cumulative difference-plot that can be used to deduce $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$.




Figure 12: The cumulative difference-plot for Case 2: the difference between the cumulative distribution functions of Y_A minus Y_B corresponding to the distributions X_A and X_B in Case 2 (shown in Figure 4).

4.3.2 Illustrative example

Figure 12 shows the cumulative difference-plot for the random variables X_A and X_B from Case 2 (their densities were shown in Figure 4). First, we see that the difference is both negative and positive, hence, neither random variable dominates the other. The difference is negative when $x = 0.05$ or lower. This can be interpreted as X_B having a smaller 5% quantile than X_A . The difference is positive otherwise, thus we deduce from the cumulative difference-plot that the 25%, the 50% (the median), the 75% and 95% quantiles are smaller in X_A than in X_B . In other words, the random variable X_B can take really low values with a small probability, but apart from these really low values, X_A takes lower values than X_B . This is also reflected by $\mathcal{C}_D(X_A, X_B) > 0.75$, as deduced from $\text{diff}(x) > 0$ for all $x \in (0.25, 1)$. The difference is also near its maximum value, implying that its integral is high and thus \mathcal{C}_P (the probability that $X_A < X_B$) is also near 1.

5 Related work

Statistical assessment of experimental results is a very studied research topic. In this section, we locate our proposal in the field and focus on similarities and differences with respect to other random variable comparison methods.

5.0.1 Visualizing densities

As mentioned in the introduction, it makes sense to model the performance of stochastic optimization algorithms as random variables. Therefore, statistical tools that compare random variables have become an increasingly important part of the analysis of experimental data. Among these tools, visualization techniques such as histograms or box-plots are usually applied before the rest of the methods. The advantage of these methods is their simplicity. If one of the random variables clearly takes lower values than the other, then these two methods effectively convey this message simply and naturally. Unfortunately, when both random variables have similar probability densities, these two methods might fail to represent the random variables in a way that makes it easy to compare them (example shown in Section 6.1).

The simplicity of these methods is also a drawback: for example, they give no information about the uncertainty associated with the estimates. The histogram suffers from the bin positioning bias Thas (2010); scikit-learn developers (2021). A kernel density estimation with the uniform kernel—considered to be the moving window equivalent of the histogram—overcomes this limitation Thas (2010), at the cost of using a more complex model. Similarly, the box-plot has a “non-injectivity” problem: very different data can still have the same box-plot Matejka and Fitzmaurice (2017); Chatterjee and Firat (2007). The violin-plot is an extension of the box-plot that overcomes the above limitation by combining the kernel density estimate of the random variables with the traditional box-plot Hintze and Nelson (1998). The proposed cumulative difference-plot improves on these methods because it represents the data clearly, even when the two random variables being compared are similar.

5.0.2 Null hypothesis statistical testing

Null hypothesis tests can be used to compare random variables without having to visually represent them. In a very general way, carrying out a null hypothesis test involves the following: first, a null hypothesis is proposed. Under certain assumptions, the null hypothesis implies that a given statistic obtained from the data follows a known distribution. Then, assuming the null

hypothesis is true, the probability of obtaining data with a more extreme statistic value¹⁰ than the observed Conover and Conover (1980) is computed. When the probability under the null hypothesis of the observed statistic is lower than a predefined threshold, the null hypothesis is rejected and the alternative hypothesis is accepted Greenland et al. (2016). Usually, this threshold is set at an arbitrary but well established Wasserstein and Lazar (2016) $p = 0.05$, although recently, further reducing the threshold to $p = 0.005$ has been proposed Benjamin et al. (2018); Ioannidis (2018).

In the context of comparing two random variables X_A and X_B , in general, we cannot assume that a statistic obtained from the data follows a known distribution under the null hypothesis. In this case, a non-parametric test Conover and Conover (1980) is a suitable choice. Specifically, the Mann-Whitney Mann and Whitney (1947) test is a good choice, as the samples observed from the random variables are i.i.d for each random variable¹¹. With this test, the null hypothesis is that $\mathcal{P}(X_A > X_B) = \mathcal{P}(X_B > X_A)$, and a possible alternative hypothesis is that $X_A \succ X_B$ Mann and Whitney (1947).

Null hypothesis tests have some limitations: for example, the p -value does not separate between the effect size and the sample size Benavoli et al. (2017); Calvo et al. (2019). In addition, rejecting the null hypothesis does not always mean that there is evidence in favor of the alternative hypothesis: it just means that the observed statistic (or a more extreme statistic) is very unlikely when the null hypothesis is true.

To show this, we generate 400 samples of the distributions X_A and X_B from Case 2 (density functions shown in Figure 4) and we apply the Mann-Whitney test, rejecting the null hypothesis when $p < 0.005$. If we repeat this experiment 10^4 times (with different samples each time), the null hypothesis is rejected every time¹². However, $X_A \not\prec X_B$ and $X_B \not\prec X_A$, implying that the alternative hypothesis is not true. Note that the proposed cumulative difference-plot (shown in Figure 12) avoids this problem because it correctly points out that neither random variable dominates the other one, for the same case and with the same number of samples.

5.0.3 Bayesian analysis

As an alternative Benavoli et al. (2017); Calvo et al. (2019) to the limitations of null hypothesis test, Bayesian analysis has been proposed. Bayesian analysis Gelman (2014); Bernardo and Smith (2009) estimates the probability that a hypothesis is true, conditioned to the observed data. This estimation requires the prior probabilities of the hypotheses and the data, but usually, they are assumed to follow a distribution that gives equal probability to all hypotheses and data. Recently a Bayesian version of the Wilcoxon signed-rank test Benavoli et al. (2014, 2017) has been proposed. In this paper, we will consider the simplex-plot of its posterior distribution. For convenience, in the rest of the paper, we will call it *simplex-plot*.

Once the posterior distribution is known, the probability that the difference between a sample from X_A and a sample from X_B is in the interval $(-\infty, -r)$, $[-r, r]$ or $(r, +\infty)$ can be computed. These probabilities can be interpreted as the probability that $X_A > X_B$, $X_A = X_B$ and $X_B > X_A$, where two samples x_a and x_b are considered equal when $|x_a - x_b| \leq r$. Note that the simplex-plot is just a convenient representation of the posterior distribution, where ‘rope’ or *range of practical equivalence* denotes hypothesis $X_A = X_B$ (when the difference is in the interval $[-r, r]$).

¹⁰The definition of what *data with a more extreme statistic value* is not the same for every null hypothesis test, and it depends on the test being used.

¹¹For paired data, the Wilcoxon signed-rank test Wilcoxon (1945) or the sign test Conover and Conover (1980) should be used. However, in the context of this paper, the samples observed from the random variables are not paired. In this paper, we consider the Mann-Whitney test as it is probably the most well known non-parametric test for unpaired data, although take into account that more modern alternatives have been proposed Ledwina and Wylupek (2012); Baumgartner et al. (1998); Biswas and Ghosh (2014).

¹²The source code to replicate this experiment is available in the file `mann_whitney_counter_example.R` in our Github repository.




Figure 13: The simplex-plot computed with the package *scmamp* Calvo and Santafé Rodrigo (2016) of the posterior distribution for Case 2.

We computed the simplex-plot (Figure 13) with the 400 samples of X_A and X_B from Case 2 obtained in Section 4.2. Two samples were considered equal when their difference is lower than $r = 10^{-3}$, and we used the prior proposed by Benavoli et al. (2017). We can deduce from this figure that the hypothesis $X_B > X_A$ is much more likely than $X_A = X_B$ or $X_B < X_A$.

The simplex plot summarizes the data through the probabilities of $X_B > X_A$, $X_A = X_B$ or $X_B < X_A$, but does not offer any additional information: we cannot deduce from these probabilities in which intervals the values of a random variable are lower than the other. In this sense, the cumulative difference-plot is a more detailed comparative visualization. Specifically, the observation that the 1% lowest values of X_A are lower than the 1% lowest values of X_B cannot be deduced from the simplex-plot, while it is easy to see in the cumulative difference-plot. Also, the cumulative difference-plot shows a comparison of the cumulative distributions through the dominance rate, while the simplex-plot does not.

5.0.4 Other plots in the interval $[0, 1]$

The probability-probability plot is defined as

$$PP : [0, 1] \rightarrow [0, 1]^2 : p \rightarrow (p, G_A(G_B^{-1}(p))).$$

As proposed by Schmid et al. Schmid and Tiede (1996), it can be interpreted via the integral of the non-negative part, which represents the amount of violation against the hypothesis that X_A dominates X_B .

The quantile-quantile plot Thas (2010); Wilk and Gnanadesikan (1968) is defined as

$$QQ : [0, 1] \rightarrow N^2 : p \rightarrow (G_A^{-1}(p), G_B^{-1}(p)),$$

and it is a natural way to visualize the differences in quantiles of X_A and X_B in N (the domain of definition of the random variables).

The quantile-quantile plot also allows a comparison between quantiles, just like the *cumulative difference-plot*. However, the cumulative difference-plot proposed in this paper is distinct from the two plots above in three aspects: i) the proposed cumulative difference-plot is defined directly from the observed samples. Because of its definition, it has a confidence band built-in, which allows the uncertainty associated with the estimation to be directly interpreted within the plot. ii) The proposed cumulative difference-plot contains several statistics simultaneously. Specifically, the estimated \mathcal{C}_D , \mathcal{C}_P and the comparison of the quantiles can be visually interpreted. iii) The proposed plot is just the difference of two cumulative distributions (G_{Y_A} and

G_{Y_B}), and thus, unlike in the pp-plot and qq-plot mentioned above, it can be defined without the need of the inverse function. The random variables Y_A and Y_B have the same $\mathcal{C}_D, \mathcal{C}_P$ as the kernel density estimations of the original distributions, and therefore, we can think of the cumulative difference-plot as the difference between the cumulative distribution function of two simpler versions of the original random variables.

6 Experimentation with the cumulative difference-plot

To illustrate the applicability of the proposed methodology, in the following, we re-evaluate the experimentation of a recently published work. In a recent paper, Santucci et al. (2020) introduced a gradient-based optimizer for solving problems defined in the space of permutations (from now on *PL-GS*). In their experimentation, they compared it with an estimation of distribution algorithm Larrañaga and Lozano (2001) (from now on *PL-EDA*). These two algorithms were tested in a set of 50 problem instances of the linear ordering problem Schiavinotto and Stützle (2004). The performance of each algorithm in each instance was estimated with the median relative deviation from the best-known objective value, with $n = 20$ repetitions. From now on, we call *score* to the relative deviation from the best-known objective value and note that a low score is better than a high score, as it means that the objective value found is closer to the best-known.

In the work by Santucci et al. (2020), when the score of one of the algorithms was at least 10^{-4} higher than the other, it was considered that one of the algorithms performed better than the other in that instance. Santucci et al. (2020) concluded that both algorithms performed equally in the instance *N-t70n11xx*, as the median scores were exactly the same for both algorithms in this instance.

In the following, we take a closer look at the performance of *PL-EDA* and *PL-GS* in this problem instance by comparing $n = 10^3$ measurements of the score from each algorithm. We increase the sample size from $n = 20$ to $n = 10^3$ because the difference between the performance of the algorithms is small. With a sample size of $n = 20$, the uncertainty is too high to come to any meaningful conclusion (regardless of the statistical methodology considered). With this increased sample size, we obtained more accurate estimates of the median scores—*PL-GS* = 0.00407, *EDA* = 0.00433, lower is better—and *PL-GS* obtains a better value by a difference higher than 10^{-4} .

6.1 Step 1: Visualization

Figure 14 shows the histogram of the scores. It can be deduced from the figure that neither algorithm clearly produces better scores. In particular, neither algorithm dominates the other: *PL-EDA* has a longer tail both to the right and to the left. Also, notice that the score of the algorithms is not normally distributed: *PL-GS* has a bimodal shape, and *PL-EDA* has a very long tail to the right (while the tail to the left is shorter).

Figure 15 shows the box-plot and the violin-plot of the data. Both algorithms have a similar median, but due to the high number of outliers Carreño et al. (2020), it is difficult to compare the scores of the algorithms with the box-plot. The same happens with the violin-plot.

6.2 Step 2: Comparing PL-GS with PL-EDA

Sometimes, visualization is enough to compare the performance of two algorithms: if one of the algorithms always performs better than the other, there is no need for further analysis. However,

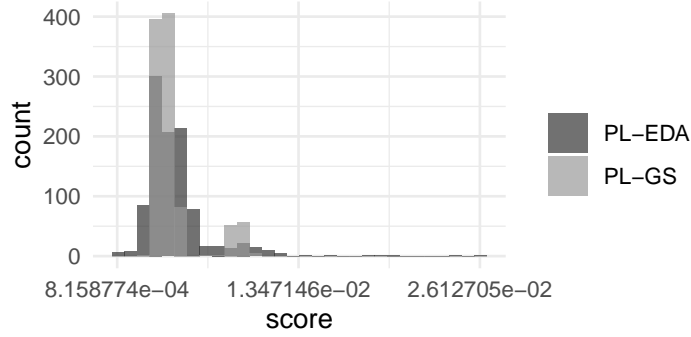


Figure 14: Histogram of the scores obtained in the instance $N-t70n11xx$. Lower is better.




Figure 15: Box-plot and violin-plot of the scores obtained in the instance $N-t70n11xx$. Lower is better.

in this case, the three visualization methods considered (histogram, box-plot, and violin-plot) have not been able to summarize the scores obtained with the algorithms in a way that enables an easy comparison. In the following, we further study the scores of the algorithms with statistical tests, the simplex-plot, and the cumulative difference-plot.

6.2.1 Mann-Whitney test

Applying the Mann-Whitney test we obtain a p -value of $p = 0.035$, lower than the usually used 0.05 threshold. With $p < 0.05$, we reject the null hypothesis and accept the alternate hypothesis: the random variable associated with the score of $PL-GS$ dominates $PL-EDA$. Note that neither rejecting the null hypothesis nor a small p -value reflect the magnitude of the difference in score of the algorithms. In addition, as stated when we studied the histogram, we know that it is unlikely that $PL-GS$ dominates $PL-EDA$.

6.2.2 Simplex-plot

We show the simplex-plot Benavoli et al. (2017) of the scores in Figure 16. Following the criterion by Santucci et al. (2020), we considered that two scores are equal when they differ by less than $r = 10^{-4}$. Unlike in the statistical test, one can deduce the probability that one of the algorithms has a better score than the other from simplex-plot: it is more likely that $PL-GS$ takes a lower value than $PL-EDA$. A closer position in the plot to $PL-EDA$ indicates a higher probability of measuring a higher score in $PL-EDA$ than in $PL-GS$. Specifically, from the simplex-plot shown in Figure 16, we can deduce that given two samples x_{gs} and x_{eda} of the scores of $PL-GS$ and $PL-EDA$ respectively,

$$\mathcal{P}(x_{eda} < x_{gs}) < \mathcal{P}(x_{gs} < x_{eda}).$$




Figure 16: Simplex-plot of *PL-GS* and *PL-EDA* in the instance *N-t70n11xx*. A closer position in the plot to *PL-EDA* indicates a higher probability of measuring a higher score in *PL-EDA* than in *PL-GS*. A low score is preferred to a high score.

However, the difference in these probabilities is small. Also, the probability that $\mathcal{P}(x_{gs} = x_{eda})$ is low (no data points near ‘rope’).

6.2.3 Cumulative difference-plot

We show the 95% confidence cumulative difference-plot in Figure 17. From this plot, we can deduce the following:

1. $\mathcal{P}(x_{eda} < x_{gs})$ and $\mathcal{P}(x_{gs} < x_{eda})$ have similar probabilities, as $\mathcal{C}_{\mathcal{P}}(PL-EDA, PL-GS) \approx 0.5$. However, The area under $\text{diff}(x) = 0$ is a little larger than the area over $\text{diff}(x) = 0$, hence $\mathcal{P}(x_{eda} < x_{gs})$ is a little smaller than $\mathcal{P}(x_{gs} < x_{eda})$.
2. Neither algorithm dominates the other one, and what is more, $\mathcal{C}_{\mathcal{D}}(PL-EDA, PL-GS) \approx 0.5$.
3. The difference is positive when $x < 0.3$, and therefore, if we only consider the best 30% values of both algorithms, *PL-EDA* dominates *PL-GS*.
4. The difference is negative when $x > 0.98$. In this case, we conclude that if we only consider the worst 2% values of *PL-EDA* and *PL-GS*, then *PL-GS* dominates *PL-EDA*.
5. These “worst” 2% values are much less likely than the “best” 30% values mentioned in 3), as the estimated probability of these “best” and “worst” values is 0.3 and 0.02 respectively.
6. The difference is negative at $x = 0.5$ and at $x = 0.75$. This can be interpreted as *PL-GS* having a better median and a better 75% quantile.

Summarizing the above points, we conclude that the performance of the algorithms is quite similar, and *PL-EDA* takes both better and worse scores than *PL-GS*. The probability that *PL-EDA* takes these better values is much higher than the probability that it takes worse values. Therefore, if we are in a setting in which repeating the execution of the algorithms is reasonable, *PL-EDA* is a much better algorithm. On the other hand, if it is critical to avoid really bad values, then *PL-GS* would be preferred. With an increased number of samples, it might be possible to better compare the algorithms (it would reduce the uncertainty associated with the size of the confidence band).




Figure 17: The cumulative difference-plot of 95% confidence of the objective values obtained by *PL-EDA* and *PL-GS* in the instance *N-t70n11xx*.

7 Assumptions and limitations

In the following, we briefly summarize the assumptions that the cumulative difference-plot requires and comment on a few caveats.

7.1 Assumptions

Correctly using the proposed cumulative difference-plot requires that the following three assumptions are satisfied. The first assumption is that all samples of both X_A and X_B are i.i.d, consequently, it should not be used with paired data. This is also an assumption made by the Mann-Whitney test.

The second assumption is that the values of the random variables represent a minimization setting: lower values are preferred to higher values. To apply the proposed method in a maximization setting, it is enough to redefine the objective function by multiplying it by -1 .

The third assumption is that X_A and X_B are continuous random variables defined in a connected subset of \mathbb{R} . This also implies that the cumulative distribution functions of X_A and X_B are continuous and that their probability density functions are bounded. Although having a bounded density means that there should never be two identical samples—the probability of observing two independent equal samples is 0 with a bounded density—, in reality, the proposed cumulative difference-plot can deal with repeated samples. To do so, when defining the kernel density estimations of Y_A and Y_B in Section 4, repeated samples were assigned the same rank. Then, the size of the uniform distributions was adjusted (with the γ function) ensuring that the sum of the estimated densities of Y_A and Y_B remains constant even in the case of repeated observations.

7.2 Limitations and future work

Just like with other methods, the number of samples determines in part the stability of the results. With a small sample size, the confidence band of the cumulative difference-plot will be larger. There are three reasons why a larger sample size increases the stability of the plot: i) we are doing a kernel density estimation, and a higher sample size Danica and power (2009) implies that the estimation is closer to the real distribution, ii) the bootstrap method also requires

several samples to be meaningful Chernick (2011); Hall (2013) and iii) the sample size needs to be reasonable with respect to the quantiles being estimated. For example, it would not make sense to use 10 samples to estimate a 1% quantile. In all of these cases, however, determining what is a *too small* sample size is a highly debated question, and is beyond the scope of this paper. To be on the safe side, we recommend using a sample size of at least $n = 100$. It is worth noting that this was arbitrarily chosen, and a suitable sample size should be chosen depending on the desired conclusions (for example, comparing small and big quantiles requires more data). With $n = 100$ we ensure that the comparison of 1% quantiles in the cumulative difference-plot is meaningful.

The most obvious limitation of the proposed approach is in its applicability: it should only be used in case of doubt between two random variables, and when none of the random variables dominates the other one. Otherwise, there are more suitable alternatives such as Bayesian analysis Benavoli et al. (2014); Calvo et al. (2019), or directly comparing box-plots. For instance, if we take 10^3 samples of X_A and X_B and all samples of X_A are lower than all samples of X_B , then there is no need for further statistical comparison, as the results speak for themselves.

The proposed approach assumes X_A and X_B are continuous random variables and that all samples of both X_A and X_B are i.i.d, and consequently, it cannot be used with paired data. As future work, the proposed methodology could be extended for paired data and ordinal random variables. Also, the bootstrap method is the slowest part of the *cumulative difference-plot*, especially as the number of samples increases. To increase the computation speed, this slow part was written in C++ (the rest of the package was written in R R Core Team (2020)). However, its speed can probably be further improved with a better implementation.

8 Conclusion

In this paper, we approached the problem of choosing between two random variables in terms of which of them takes lower values. We proposed eight desirable properties for dominance measures: functions that compare random variables in the context of quantifying the dominance. Among the measures in the literature, we found out that *the probability that one of the random variables takes lower values than the other* was the one that satisfies the most properties. However, it fails to satisfy Property 1, hence it cannot be used to determine when one of the random variables stochastically dominates the other. To overcome this limitation, we introduced a new dominance measure: the dominance rate, which quantifies how much higher one of the cumulative distribution function is than the other.

Based on the above, we proposed a cumulative difference-plot that allows two random variables to be compared in terms of which of them takes lower values. This cumulative difference-plot contains a comparison of the quantiles, in addition to allowing a graphical estimation of the dominance rate and the probability that one of the random variables takes lower values than the others. It also models the uncertainty associated with the estimate through a confidence band. Finally, in Section 6 we showed that the proposed methodology is suitable to compare two random variables, especially when they take similar values and other methods fail to give detailed and clear answers.

Supplementary Material

RVCompare: With this R package, users can compute the \mathcal{C}_P and \mathcal{C}_D of two distributions, given their probability density functions. Furthermore, it can be used to produce the

proposed cumulative difference-plot, given the observed data. (The package can be directly installed from CRAN and is also available in the GitHub repo <https://github.com/EtorArza/RVCompare>)

Reproducibility: Alongside the paper, we provide the code to generate the figures in this paper and replicate the experimentation. For instructions on how to install the dependencies and replicate the results, refer to the README.md file in the repository. (GitHub repo <https://github.com/EtorArza/SupplementaryPaperRVCompare>)

Appendices: To keep the length of the paper at a reasonable size, the appendices have been moved to another document. (The appendices are available for download at <https://doi.org/10.5281/zenodo.6528669>).

Acknowledgments

This work was funded in part by the Spanish Ministry of Science, Innovation and Universities through PID2019-106453GA-I00/AEI/10.13039/501100011033 and the BCAM Severo Ochoa excellence accreditation SEV-2017-0718; by the Basque Government through the Research Groups 2019-2021 IT1244-19, ELKARTEK Program (project code KK-2020/00049) and BERC 2018-2021 program.

Disclosure statement

The authors report that there are no competing interests to declare.

References

- Badue, C., R. Guidolini, R. V. Carneiro, P. Azevedo, V. B. Cardoso, A. Forechi, L. Jesus, R. Berriel, T. M. Paixão, F. Mutz, L. de Paula Veronese, T. Oliveira-Santos, and A. F. De Souza (2021, March). Self-driving cars: A survey. *Expert Systems with Applications* 165, 113816.
- Bauer, P. (1991, June). Multiple testing in clinical trials. *Statistics in Medicine* 10(6), 871–890.
- Baumgartner, W., P. Weiß, and H. Schindler (1998, September). A Nonparametric Test for the General Two-Sample Problem. *Biometrics* 54(3), 1129.
- Benavoli, A., G. Corani, J. Demšar, and M. Zaffalon (2017). Time for a change: A tutorial for comparing multiple classifiers through Bayesian analysis. *The Journal of Machine Learning Research* 18(1), 2653–2688.
- Benavoli, A., G. Corani, F. Mangili, M. Zaffalon, and F. Ruggeri (2014). A Bayesian Wilcoxon signed-rank test based on the Dirichlet process. In *International Conference on Machine Learning*, pp. 1026–1034. PMLR.
- Benjamin, D. J., J. O. Berger, M. Johannesson, B. A. Nosek, E.-J. Wagenmakers, R. Berk, K. A. Bollen, B. Brembs, L. Brown, C. Camerer, D. Cesarini, C. D. Chambers, M. Clyde, T. D. Cook, P. De Boeck, Z. Dienes, A. Dreber, K. Easwaran, C. Efferson, E. Fehr, F. Fidler, A. P. Field, M. Forster, E. I. George, R. Gonzalez, S. Goodman, E. Green, D. P. Green, A. G. Greenwald, J. D. Hadfield, L. V. Hedges, L. Held, T. Hua Ho, H. Hoijtink, D. J. Hruschka, K. Imai, G. Imbens, J. P. A. Ioannidis, M. Jeon, J. H. Jones, M. Kirchler, D. Laibson, J. List, R. Little, A. Lupia, E. Machery, S. E. Maxwell, M. McCarthy, D. A. Moore, S. L. Morgan, M. Munafó, S. Nakagawa, B. Nyhan, T. H. Parker, L. Pericchi, M. Perugini, J. Rouder, J. Rousseau, V. Savalei, F. D. Schönbrodt, T. Sellke, B. Sinclair, D. Tingley, T. Van Zandt, S. Vazire, D. J. Watts, C. Winship, R. L. Wolpert, Y. Xie, C. Young, J. Zinman, and V. E.

- Johnson (2018, January). Redefine statistical significance. *Nature Human Behaviour* 2(1), 6–10.
- Bennet, C. J. (2013). Inference for Dominance Relations. *International Economic Review* 54(4), 1309–1328.
- Bernardo, J. M. and A. F. Smith (2009). *Bayesian Theory*, Volume 405. John Wiley & Sons.
- Bhattacharyya, A. (1943). On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society* 35, 99–109.
- Bickel, P. J. and A. M. Krieger (1989). Confidence bands for a distribution function using the bootstrap. *Journal of the American Statistical Association* 84(405), 95–100.
- Biswas, M. and A. K. Ghosh (2014, January). A nonparametric two-sample test applicable to high dimensional data. *Journal of Multivariate Analysis* 123, 160–171.
- Calvo, B. and G. Santafé Rodrigo (2016). Scmamp: Statistical comparison of multiple algorithms in multiple problems. *The R Journal*, Vol. 8/1, Aug. 2016.
- Calvo, B., O. M. Shir, J. Ceberio, C. Doerr, H. Wang, T. Bäck, and J. A. Lozano (2019). Bayesian performance analysis for black-box optimization benchmarking. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion on - GECCO '19*, Prague, Czech Republic, pp. 1789–1797. ACM Press.
- Carreño, A., I. Inza, and J. A. Lozano (2020, June). Analyzing rare event, anomaly, novelty and outlier detection terms under the supervised classification framework. *Artificial Intelligence Review* 53(5), 3575–3594.
- Chatterjee, S. and A. Firat (2007). Generating Data with Identical Statistics but Dissimilar Graphics: A Follow up to the Anscombe Dataset. *The American Statistician* 61(3), 248–254.
- Cheng, R. C. and T. Lies (1988). One-sided confidence bands for cumulative distribution functions. *Technometrics* 30(2), 155–159.
- Cheng, R. C. H. and T. C. Iles (1983, February). Confidence Bands for Cumulative Distribution Functions of Continuous Random Variables. *Technometrics* 25(1), 77–86.
- Chernick, M. R. (2011). *Bootstrap Methods: A Guide for Practitioners and Researchers*, Volume 619. John Wiley & Sons.
- Chollet, F. et al. (2015). Keras.
- Conover, W. J. and W. J. Conover (1980). Practical nonparametric statistics.
- Cruz-Roa, A., H. Gilmore, A. Basavanthally, M. Feldman, S. Ganesan, N. N. Shih, J. Tomaszewski, F. A. González, and A. Madabhushi (2017, June). Accurate and reproducible invasive breast cancer detection in whole-slide images: A Deep Learning approach for quantifying tumor extent. *Scientific Reports* 7(1), 46450.
- Danica and power (2009). What is the minimum number of data points required for kernel density estimation?
- Devroye, L., A. Mehrabian, and T. Reddad (2020, May). The total variation distance between high-dimensional Gaussians. *arXiv:1810.08693 [math, stat]*.
- Efron, B. and R. Tibshirani (1993). *An Introduction to the Bootstrap*. Number 57 in Monographs on Statistics and Applied Probability. New York: Chapman & Hall.
- Faraway, J. J. and Myoungshic Jhun (1990). Bootstrap choice of bandwidth for density estimation. *Journal of the American Statistical Association* 85(412), 1119–1122.

- François-Lavet, V., P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau (2018). An Introduction to Deep Reinforcement Learning. *Foundations and Trends in Machine Learning* 11(3-4), 219–354.
- Gelman, A. (2014). *Bayesian Data Analysis* (Third edition ed.). Chapman & Hall/CRC Texts in Statistical Science. Boca Raton: CRC Press.
- Glorot, X. and Y. Bengio (2010, May). Understanding the difficulty of training deep feedforward neural networks. In Y. W. Teh and M. Titterton (Eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, Volume 9 of *Proceedings of Machine Learning Research*, Chia Laguna Resort, Sardinia, Italy, pp. 249–256. JMLR Workshop and Conference Proceedings.
- Goeman, J. J. and A. Solari (2014, May). Multiple hypothesis testing in genomics. *Statistics in Medicine* 33(11), 1946–1978.
- Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT press.
- Greenland, S., S. J. Senn, K. J. Rothman, J. B. Carlin, C. Poole, S. N. Goodman, and D. G. Altman (2016, April). Statistical tests, P values, confidence intervals, and power: A guide to misinterpretations. *European Journal of Epidemiology* 31(4), 337–350.
- Hall, P. (2013). *The Bootstrap and Edgeworth Expansion*. Springer Science & Business Media.
- Hall, P. and J. Horowitz (2013, August). A simple bootstrap method for constructing nonparametric confidence bands for functions. *The Annals of Statistics* 41(4).
- Hintze, J. L. and R. D. Nelson (1998). Violin plots: A box plot-density trace synergism. *The American Statistician* 52(2), 181–184.
- Ioannidis, J. P. (2018). The proposal to lower P value thresholds to .005. *Jama* 319(14), 1429–1430.
- Kailath, T. (1967, February). The Divergence and Bhattacharyya Distance Measures in Signal Selection. *IEEE Transactions on Communications* 15(1), 52–60.
- Kullback, S. and R. A. Leibler (1951, March). On Information and Sufficiency. *The Annals of Mathematical Statistics* 22(1), 79–86.
- Larrañaga, P. and J. A. Lozano (2001). *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*, Volume 2. Springer Science & Business Media.
- Ledwina, T. and G. Wyłupek (2012, December). Nonparametric tests for stochastic ordering. *TEST* 21(4), 730–756.
- Liese, F. and I. Vajda (2006, October). On Divergences and Informations in Statistics and Information Theory. *IEEE Transactions on Information Theory* 52(10), 4394–4412.
- Mann, H. B. and D. R. Whitney (1947). On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics* 18(1), 50–60.
- Matejka, J. and G. Fitzmaurice (2017). Same stats, different graphs: Generating datasets with varied appearance and identical statistics through simulated annealing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 1290–1294.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller (2013, December). Playing Atari with Deep Reinforcement Learning. *arXiv:1312.5602 [cs]*.
- Panaretos, V. M. and Y. Zemel (2019). Statistical aspects of wasserstein distances. *Annual Review of Statistics and Its Application* 6(1), 405–431.
- Papadopoulos, A. (2017). Interpretation of the Kullback-Leibler divergence.

- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Polyanskiy, Y. and Y. Wu (2012). Lecture notes on information theory. *MIT (6.441)*, *UIUC (ECE 563)*, *Yale (STAT 664)*.
- Quirk, J. P. and R. Saposnik (1962, February). Admissibility and measurable utility functions. *The Review of Economic Studies* 29(2), 140–146.
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Regnier-Coudert, O., J. McCall, M. Ayodele, and S. Anderson (2016). Truck and trailer scheduling in a real world, dynamic and heterogeneous context. *Transportation research part E: logistics and transportation review* 93, 389–408.
- Rényi, A. et al. (1961). On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California.
- Rousseeuw, P. J. and M. Hubert (2011, January). Robust statistics for outlier detection. *WIREs Data Mining and Knowledge Discovery* 1(1), 73–79.
- Santucci, V., J. Ceberio, and M. Baiocchi (2020, July). Gradient search in the space of permutations: An application for the linear ordering problem. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, Cancún Mexico, pp. 1704–1711. ACM.
- Schiavinotto, T. and T. Stützle (2004). The Linear Ordering Problem: Instances, Search Space Analysis and Algorithms. *Journal of Mathematical Modelling and Algorithms* 3(4), 367–402.
- Schmid, F. and M. Tiede (1996). Testing for First-Order Stochastic Dominance: A New Distribution-Free Test. *The Statistician* 45(3), 371.
- Schuhmacher, D. (2021). Compute The Wasserstein Distance Between Two Univariate Samples.
- scikit-learn developers (2021). Density Estimation.
- Steck, G. P. (1971). Rectangle probabilities for uniform order statistics and the probability that the empirical distribution function lies between two distribution functions. *The Annals of Mathematical Statistics* 42(1), 1–11.
- Thas, O. (2010). *Comparing Distributions*. Springer.
- Tsybakov, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. New York ; London: Springer.
- Vapnik, V. N. (1998). *Statistical Learning Theory*. Adaptive and Learning Systems for Signal Processing, Communications, and Control. New York: Wiley.
- Wang, J., F. Cheng, and L. Yang (2013, June). Smooth simultaneous confidence bands for cumulative distribution functions. *Journal of Nonparametric Statistics* 25(2), 395–407.
- Wasserstein, R. L. and N. A. Lazar (2016, April). The ASA Statement on p-Values: Context, Process, and Purpose. *The American Statistician* 70(2), 129–133.
- Wilcoxon, F. (1945, December). Individual Comparisons by Ranking Methods. *Biometrics Bulletin* 1(6), 80.
- Wilk, M. B. and R. Gnanadesikan (1968, March). Probability Plotting Methods for the Analysis of Data. *Biometrika* 55(1), 1.

Xi, a. (2017). Differences between Bhattacharyya distance and KL divergence.

Zhang, L., S. Wang, and B. Liu (2018, July). Deep learning for sentiment analysis: A survey.
Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 8(4).

Appendices

9 A literature review of measures

9.1 f -divergences

The f -divergence is a family of functions that can be used to measure the difference between two random variables. Given a strictly convex¹³ function $f : (0, +\infty) \rightarrow \mathbb{R}$ with $f(1) = 0$, and two continuous random variables X_A and X_B , the f -divergence Liese and Vajda (2006); Rényi et al. (1961) is defined as

$$D_f(X_A, X_B) = \int_{\mathbb{R}} g_B(x) f\left(\frac{g_A(x)}{g_B(x)}\right) dx \quad (4)$$

where g_A and g_B are the probability density functions of the random variables X_A and X_B respectively. Since $g_B(x)$ can be 0, we assume Polyanskiy and Wu (2012) that $0 \cdot f(0/0) = 0$ and $0 \cdot f(a/0) = \lim_{x \rightarrow 0^+} x \cdot f(a/x)$. Notice that if g_A and g_B are the same probability density functions, then $D_f(X_A, X_B) = 0$.

Kullback–Leibler divergence: The Kullback–Leibler divergence Kullback and Leibler (1951) is a particular case of the f -divergence, for $f(x) = x \cdot \ln(x)$. Given two random variables X_A and X_B , $D_{KL}(X_A, X_B)$ can be interpreted Papadopoulos (2017) as the amount of entropy increased by using g_B to model data that follows the probability density function g_A .

The Kullback–Leibler divergence is non-negative, and non symmetric $D_{KL}(X_A, X_B) \neq D_{KL}(X_B, X_A)$, and therefore, it is not actually a distance Goodfellow et al. (2016). It will not satisfy Property (2), as it is not antisymmetric either. This also makes the interpretation less intuitive. The Kullback–Leibler divergence is often used to measure the difference between two random variables Goodfellow et al. (2016), but since $D_{KL}(X_A, X_B) \neq D_{KL}(X_B, X_A)$, it may be better to interpret the Kullback–Leibler divergence as stated above Papadopoulos (2017).

In Figure 18, we show the probability density functions and cumulative distribution functions of four random variables X_A, X_B, X_C and X_D . Looking at their cumulative distributions (Figure 18b), one can clearly see that $X_A \succ X_B$, $X_B \preceq X_C$ and $X_B \succ X_D$. However, as shown in Table 9.1, $D_{KL}(X_B, X_A) = D_{KL}(X_B, X_C) = D_{KL}(X_B, X_D) = 15.4$ and $D_{KL}(X_A, X_B) = D_{KL}(X_C, X_B) = D_{KL}(X_C, X_D) = 6.2$. This means that, given any two random variables X_A and X_B , the Kullback–Leibler is not able to distinguish if $X_A \succ X_B$, $X_B \succ X_A$ or $X_A \preceq X_B$. We can interpret this as the Kullback–Leibler divergence only caring about the difference between two random variables, and not if this difference is related to one of the random variables taking lower values than the other. Hence, it cannot satisfy Property 1, even if we try to transform it to be defined in the $[0, 1]$ interval. We conclude that the Kullback–Leibler divergence is not suitable to gain information regarding which of the random variables takes lower values.

Jensen–Shannon divergence: The Jensen–Shannon divergence Polyanskiy and Wu (2012) is very similar to the Kullback–Leibler divergence, and is another the particular case of the f -divergence for $f(x) = x \cdot \ln\left(\frac{2x}{x+1}\right) + \ln\left(\frac{2}{x+1}\right)$. It is also known as the symmetrized version of the Kullback–Leibler divergence Polyanskiy and Wu (2012), because

$$D_{JS}(X_A, X_B) = D_{KL}(X_A, X_M) + D_{KL}(X_B, X_M)$$

¹³A function $f : (0, +\infty) \rightarrow \mathbb{R}$ is strictly convex if for all $t \in [0, 1]$, for all $x_1, x_2 \in (0, +\infty)$, $f(tx_1 + (1-t)x_2) < tf(x_1) + (1-t)f(x_2)$

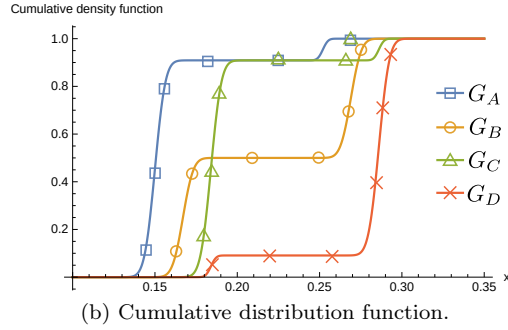



Figure 18: The probability density function and cumulative distribution of the four random variables. The distances between these random variables are listed in Table 9.1.





Figure 19: The probability density function and cumulative distribution of four other random variables. The Wasserstein distance between X_B and each of the other random variables is 0.017.

		Kullback–Leibler			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.0	6.2	28.6	88.8
	X _B	15.4	0.0	15.4	15.4
	X _C	29.4	6.2	0.0	2.6
	X _D	88.8	6.2	2.6	0.0

		Jensen-Shannon			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.0	1.2	1.4	1.4
	X _B	1.2	0.0	1.2	1.2
	X _C	1.4	1.2	0.0	0.8
	X _D	1.4	1.2	0.8	0.0

		Total variation			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.000	0.934	0.999	1.000
	X _B	0.934	0.000	0.934	0.934
	X _C	0.999	0.934	0.000	0.818
	X _D	1.000	0.934	0.818	0.000

		Hellinger			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.00	1.28	1.41	1.41
	X _B	1.28	0.00	1.28	1.28
	X _C	1.41	1.28	0.00	0.99
	X _D	1.41	1.28	0.99	0.00

		Wasserstein			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.000	0.06	0.03	0.12
	X _B	0.06	0.00	0.04	0.06
	X _C	0.03	0.04	0.000	0.083
	X _D	0.12	0.06	0.083	0.000

		C _P			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.50	0.95	0.92	0.99
	X _B	0.05	0.50	0.54	0.95
	X _C	0.08	0.46	0.50	0.91
	X _D	0.01	0.05	0.09	0.50

		C _D			
		RV ₂			
		X _A	X _B	X _C	X _D
RV ₁	X _A	0.50	1.00	1.00	1.00
	X _B	0.00	0.50	0.59	1.00
	X _C	0.00	0.41	0.50	1.00
	X _D	0.00	0.00	0.00	0.50

Table 2: $\mathcal{C}(\text{RV}_1, \text{RV}_2)$ for the random variables X_A, X_B, X_C and X_D shown in Figure 18.

where the probability density function of $X_{\mathcal{M}}$ is $g_{\mathcal{M}}(x) = 0.5(g_A(x) + g_B(x))$. Thus, we can interpret this divergence as the sum of the Kullback–Leibler divergences of g_A and g_B with respect to the average probability density function $g_{\mathcal{M}}$. The Jensen-Shannon divergence also fails to identify (see Table 9.1) the dominance relationships between X_B and the rest of the random variables in Figure 18, thus, it cannot satisfy Property 1. In addition, the Jensen-Shannon divergence also fails to satisfy Properties 2 and 3. See Table 1 for a detailed list of the properties that each measure satisfies.

Total variation: The total variation Polyanskiy and Wu (2012) is also a particular f -divergence, for $f(x) = \frac{1}{2}|x - 1|$. Unlike the Kullback–Leibler divergence, the total variation is symmetric. In fact, it is a properly defined distance Tsybakov (2009); Polyanskiy and Wu (2012). In addition, it is defined between 0 and 1.

Given two random variables X_A, X_B , the total variation can also be defined as:

$$TV(X_A, X_B) = \sup_{C \subseteq \mathbb{R}} |\mathcal{P}_A(C) - \mathcal{P}_B(C)|,$$

where \mathcal{P}_A and \mathcal{P}_B are the probability distributions¹⁴ of X_A and X_B respectively. Since the subset C that takes the supremum is $C = \{x \in \mathbb{R} \mid g_A(x) > g_B(x)\}$ Devroye et al. (2020), we can interpret the total variation as the “size” of the difference in the density functions in all points where g_A is more likely than g_B . Following this intuition, when $TV(X_A, X_B) = 1$, g_A and g_B have disjoint supports Polyanskiy and Wu (2012), and thus X_A and X_B are at their maximum difference with respect to this metric. On the other hand, when $TV(X_A, X_B) = 0$ the random variables are identical.

¹⁴Given the random variable X_A defined in \mathbb{R} , its probability distribution, noted as \mathcal{P}_A , is a mapping that, for all $U \subseteq \mathbb{R}$ that is measurable, $A(U) = \mathcal{P}(X_A \in U)$ Vapnik (1998).

The Total-Variance also fails to identify (see Table 9.1) the dominance relationships between X_B and the rest of the random variables in Figure 18.

Hellinger distance and the Bhattacharyya distance: The Hellinger distance is the square root of the f -divergence for $f(x) = (1 - \sqrt{x})^2$ Polyanskiy and Wu (2012). It is related to the Bhattacharyya coefficient, since $D_H(X_A, X_B) = 2(1 - \text{BhattCoef}(X_A, X_B))$ Xi (2017); Polyanskiy and Wu (2012), where $\text{BhattCoef}(X_A, X_B)$ is the Bhattacharyya coefficient Kailath (1967); Bhattacharyya (1943). This coefficient is defined as $\text{BhattCoef}(X_A, X_B) = \int_{\mathbb{R}} \sqrt{g_A(x)g_B(x)}dx$, and has proven useful on signal processing Kailath (1967). Given two probability density functions g_A and g_B , the Bhattacharyya coefficient can be interpreted as the integral of the geometric mean of the probability density functions. The Bhattacharyya coefficient is also related to the Bhattacharyya distance, as $D_{\text{Bhatt}}(X_A, X_B) = -\ln(\text{BhattCoef}(X_A, X_B))$.

The Hellinger distance and the Bhattacharyya distance also fail to identify (see Table 9.1) the dominance relationships between X_B and the rest of the random variables in Figure 18.

9.2 Wasserstein distance

The Wasserstein distance is another type of distance between probability random variables. Given two continuous random variables X_A, X_B , the Wasserstein distance (of order 1) is defined as Schuhmacher (2021); Panaretos and Zemel (2019)

$$D_W(X_A, X_B) = \int_{\mathbb{R}} |G_A(x) - G_B(x)|dx$$

In Figure 19, we show a different set of four random variables X_A, X_B, X_C and X_D . In this case, it is also clear that $X_A \succ X_B$, $X_B \leq X_C$ and $X_B \succ X_D$ (Figure 19b), but $D_W(X_B, X_A) = D_W(X_B, X_C) = D_W(X_B, X_D) = 0.017$. Therefore, in this case, the Wasserstein distance does not give any insights about the dominance between X_B and the rest of the random variables, thus, it cannot satisfy Property 1 even with a transformation. It also does not satisfy Properties 2, 3, 6, 7, 8.

However, with a small change, the Wasserstein distance can comply with Properties 2 and 3. This change also improves its correlation with the dominance, even though it still does not comply with Property 1. We remove the absolute value, such that the *signed Wasserstein* distance is defined as

$$D_{SW}(X_A, X_B) = \int_{\mathbb{R}} G_A(x) - G_B(x)dx.$$

For the random variables in Figure 19, the signed Wasserstein distance has different values: $D_{SW}(X_B, X_A) = 0.17$, $D_{SW}(X_B, X_C) = 0$ and $D_{SW}(X_B, X_D) = -0.017$. Notice that

$$X_A \succ X_B \implies D_{SW}(X_B, X_A) > 0 \text{ and } X_B \succ X_A \implies D_{SW}(X_B, X_A) < 0,$$

but unfortunately, when $X_A \leq X_B$, $D_{SW}(X_B, X_A)$ could be positive or negative. This implies that $D_{SW}(X_B, X_A)$ still can not determine if $X_A \succ X_B$, $X_B \succ X_A$, or $X_A \leq X_B$.

9.3 Heuristic derivation of the first-order stochastic dominance

A measure similar to the Wasserstein distance has been proposed in the literature Schmid and Trede (1996) in the context of comparing random variables. Specifically, this measure is part

of the heuristic derivation of a distribution-free statistical test for first-order stochastic dominance Schmid and Tiede (1996). Given two random variables X_A, X_B , this measure is defined as

$$\mathcal{C}_I(X_A, X_B) = \int_{\mathbb{R}} \max(0, G_A(x) - G_B(x)) dG_B(x).$$

Note that the values of \mathcal{C}_I range between 0 and 0.5. When $\mathcal{C}_I(X_A, X_B) = 0.5$, we know that $X_A \succ X_B$. Unfortunately, when $\mathcal{C}_I(X_A, X_B) \in (0, 0.5)$, it could be that $X_A \succ X_B$ or $X_A \not\succeq X_B$. Consequently, $\mathcal{C}_I(X_A, X_B)$ cannot satisfy Property 1.

10 Quantile random variables

10.1 Computing the probability density functions of Y_A and Y_B

In Section 4.1 we introduced the quantile random variables Y_A and Y_B . We now describe how to compute the probability density functions of g_{Y_A} and g_{Y_B} step by step, with the pseudocode shown in Algorithm 1. We define a function r that returns the position of an observation according to its rank in the sorted list of the observation $A_n \cup B_n$ (lines 1–4). The ranks go from 0 (for the smallest observation) to r_{max} (for the largest), where r_{max} is the number of unique observation in $A_n \cup B_n$ minus 1. Repeated observations are assigned the same rank, and no ranks are skipped: there is at least a value in $\mathbf{a} \cup \mathbf{b}$ corresponding to each rank from 0 to r_{max} . For each observation in $\{a_1, \dots, a_n\}$, a uniform distribution defined in the interval $(\frac{r(a_i) + \gamma(r(a_i) - 1)}{2n}, \frac{r(a_i) + \gamma(r(a_i))}{2n})$ is added to the mixture (lines 10–19), where $\gamma(k)$ (lines 7–9) counts the number of ranks in $A_n \cup B_n$ that are lower than or equal to k (since the lowest rank is 0, $\gamma(-1) = 0$). The kernel density estimation for Y_B is defined similarly, but with the observations $\{b_1, \dots, b_n\}$ instead.

10.2 The quantile random variables have the same $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ as the kernel density estimates of X_A and X_B .

In Section 4.1, we claimed that when a “small enough” uniform scikit-learn developers (2021) kernel is used in the kernel density estimations of X_A and X_B , these estimations will have the same $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ as the quantile random variables Y_A and Y_B . Specifically, the size of the uniform kernels needs to be smaller than $\min_{i,j \in \{1 \dots n\} | a_i \neq b_j} 2|a_i - b_j|$, where $A_n = \{a_1, \dots, a_n\}$ and $B_n = \{b_1, \dots, b_n\}$ are the n observed samples of X_A and X_B respectively. As a result, the $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ of the kernel density estimations will not change when the size of the kernels is reduced below its initial size. This can be deduced from Property 8 in Section 2.2, which both $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ satisfy.

The quantile random variables Y_A and Y_B can also be obtained by applying a sequence of transformations to the kernel density estimations (with small uniform kernels) of X_A and X_B . Three consecutive transformations are required, none of which modify the $\mathcal{C}_{\mathcal{D}}$ and $\mathcal{C}_{\mathcal{P}}$ due to Property 8. The first transformation involves further *reducing* the size of the kernels to $1/(4n)$. Secondly, each kernel k is moved into the position $r(k)/(2n) + (4n)^{-1}$, where $r(k)$ is the rank of the sample in k in $A_n \cup B_n$. In the case of ties, r assigns the same rank to all kernels and this same rank is the average of the previous and the next rank. Since each of the possible positions are at distance $1/(2n)$ from each other, this transformation will not change the $\mathcal{C}_{\mathcal{D}}$ and $\mathcal{C}_{\mathcal{P}}$. Finally, the length of the kernels is increased to $mult/(4n)$, where $mult$ is the number of

Algorithm 1: Kernel density estimation of Y_A and Y_B

Input: $A_n = \{a_1, \dots, a_n\}$: The n observed samples of X_A . $B_n = \{b_1, \dots, b_n\}$: The n observed samples of X_B .**Output:** g_{Y_A} : The probability density of Y_A . g_{Y_B} : The probability density of Y_B .

```
/* Compute the ranks of  $A_n \cup B_n$ . The lowest value has rank 0. Assign the same rank to
   ties without skipping any rank. */
1 for  $i = 1, \dots, n$  do
2   |  $r(a_i) \leftarrow$  rank of  $a_i$  in  $A_n \cup B_n$ 
3   |  $r(b_i) \leftarrow$  rank of  $b_i$  in  $A_n \cup B_n$ 
4 end
5  $R \leftarrow \{r(a_1), \dots, r(a_n), r(b_1), \dots, r(b_n)\}$ 
6  $r_{max} \leftarrow \max(R)$ 
7 for  $k = -1, 0, 1, \dots, r_{max}$  do
8   |  $\gamma(k) \leftarrow$  the number of items in  $R$  lower than or equal to  $k$ 
9 end
/* The probability density function of  $g_{Y_A}$  is represented as a mixture of  $n$  uniform
   distributions.  $g_{Y_A}[s]$  is the probability density of  $Y_A$  in the interval  $[\frac{s}{2n}, \frac{s+1}{2n})$ . */
10  $g_{Y_A} \leftarrow$  array of zeros of length  $2n$ 
11  $g_{Y_B} \leftarrow$  array of zeros of length  $2n$ 
12 for  $x_i = a_1, \dots, a_n, b_1, \dots, b_n$  do
13   |  $A_{mult} \leftarrow$  number of times that  $x_i$  is in  $A_n$ 
14   |  $B_{mult} \leftarrow$  number of times that  $x_i$  is in  $B_n$ 
15   | for  $mult = 1, \dots, (A_{mult} + B_{mult})$  do
16     |  $g_{Y_A}[\gamma(r(a_i) - 1) + mult - 1] \leftarrow (n \cdot A_{mult})^{-1}$ 
17     |  $g_{Y_B}[\gamma(r(b_i) - 1) + mult - 1] \leftarrow (n \cdot B_{mult})^{-1}$ 
18   | end
19 end
20 return  $g_{Y_A}, g_{Y_B}$ 
```

times that the sample defining the kernel is repeated in $A_n \cup B_n$. Note that this increase in the length will in no case cause an overlap of kernels.

11 $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ in the cumulative difference-plot

In this section, we mathematically prove and experimentally verify that the cumulative difference-plot can be used to deduce $\mathcal{C}_{\mathcal{D}}$ and $\mathcal{C}_{\mathcal{P}}$. First, we describe which estimators are used when these dominance measures are visually estimated from the cumulative difference-plot. Then, we show that these estimators converge to $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ as the number of samples increases.

11.1 Estimating $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ from the cumulative difference-plot

Definition 7. (*observations of random variables*)

Let X_A be a continuous random variable. We define n observations of X_A as the realizations of the i.i.d random variables $\{X_A^i\}_{i=1}^n$ that are distributed as X_A , denoted as $A_n = \{a_i\}_{i=1}^n$.

Definition 8. (*estimation of $\mathcal{C}_{\mathcal{P}}$*)

Let X_A and X_B be two continuous random variables and A_n and B_n their n observations respectively. We define the estimation of the probability that $X_A < X_B$ as

$$\widetilde{\mathcal{C}}_{\mathcal{P}}(A_n, B_n) = \sum_{i,k=1\dots n} \frac{\text{sign}(b_k - a_i)}{2n^2} + \frac{1}{2}.$$

Definition 9. (*estimation of $\mathcal{C}_{\mathcal{D}}$*)

Let X_A and X_B be two continuous random variables and A_n and B_n their n observations respectively. Let $\{c_j\}_{j=1}^{2n}$ the sorted list of all the observations of A_n and B_n where c_1 is the smallest observation and c_{2n} the largest. Let $\{c_d\}_{d=1}^{d_{max}}$ be the sorted list of unique values in $\{c_j\}_{j=1}^{2n}$. We define the estimation of the dominance rate as

$$\widetilde{\mathcal{C}}_{\mathcal{D}}(A_n, B_n) = \frac{\sum_{j=1}^{2n} \frac{\psi(c_j)}{2n} + 1}{2} \cdot k_c^{-1}$$

$k_c = \frac{\sum_{j=1}^{2n} \mathcal{I}[\psi(c_j) \neq 0]}{2n}$ is the normalization constant and ψ_j is defined as

$$\psi(c_d) = \begin{cases} 0 & \text{if } \hat{G}_A(c_{d-1}) = \hat{G}_B(c_{d-1}) \\ & \text{and } \hat{G}_A(c_d) = \hat{G}_B(c_d) \\ 1 & \text{if } \hat{G}_A(c_{d-1}) \geq \hat{G}_B(c_{d-1}) \\ & \text{and } \hat{G}_A(c_d) > \hat{G}_B(c_d) \\ 1 & \text{if } \hat{G}_A(c_{d-1}) > \hat{G}_B(c_{d-1}) \\ & \text{and } \hat{G}_A(c_d) \geq \hat{G}_B(c_d) \\ -1 & \text{if } \hat{G}_B(c_{d-1}) \geq \hat{G}_A(c_{d-1}) \\ & \text{and } \hat{G}_B(c_d) > \hat{G}_A(c_d) \\ -1 & \text{if } \hat{G}_B(c_{d-1}) > \hat{G}_A(c_{d-1}) \\ & \text{and } \hat{G}_B(c_d) \geq \hat{G}_A(c_d) \\ 1 - 2\gamma(c_d) & \text{if } \hat{G}_B(c_{d-1}) > \hat{G}_A(c_{d-1}) \\ & \text{and } \hat{G}_A(c_d) > \hat{G}_B(c_d) \\ 2\gamma(c_d) - 1 & \text{if } \hat{G}_A(c_{d-1}) > \hat{G}_B(c_{d-1}) \\ & \text{and } \hat{G}_B(c_d) > \hat{G}_A(c_d) \end{cases}$$

with $\gamma(c_d) = \frac{\hat{G}_B(c_{d-1}) - \hat{G}_A(c_{d-1})}{[B_n = c_d] - [A_n = c_d]}$. Note that $[A_n = c_d]$ counts the number of items in A_n equal to c_d and \hat{G}_A is the empirical distribution Steck (1971) estimated from A_n . To improve the readability, we abuse the notation and assume that $\hat{G}_A(c_0) = 0$.

We now show that these estimates can be directly computed from the cumulative difference plot. First, we show that the estimation of $\mathcal{C}_{\mathcal{P}}$ from the cumulative difference-plot is equivalent to the estimation in Definition 8. As mentioned in Section 4.3, the $\mathcal{C}_{\mathcal{P}}$ estimated from the cumulative difference-plot is $0.5 + \int_0^1 \text{diff}(x)dx$ where diff is the difference function introduced in Equation (2). Specifically, the difference function was defined as $\text{diff}(x) = G_{Y_A}(x) - G_{Y_B}(x)$.

Lemma 1. *Let X_A and X_B be two continuous random variables and A_n and B_n their n observations respectively. Then,*

$$\int_0^1 \text{diff}(x)dx = \sum_{j=1}^{2n} \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n})}{2n}$$

Proof. Considering that the density functions of Y_A and Y_B are constant in each interval $[\frac{j}{2n}, \frac{j+1}{2n})$ for $j = 0, \dots, (2n - 1)$, we get that

$$\begin{aligned} \int_{\frac{j}{2n}}^{\frac{j+1}{2n}} \text{diff}(x)dx &= \frac{\text{diff}(\frac{j}{2n}) + \text{diff}(\frac{j+1}{2n})}{4n} = \\ &= \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n}) + G_{Y_A}(\frac{j+1}{2n}) - G_{Y_B}(\frac{j+1}{2n})}{4n} \end{aligned}$$

Taking into account that $G_{Y_A}(0) = G_{Y_B}(0) = 0$ and $G_{Y_A}(1) = G_{Y_B}(1) = 1$,

$$\begin{aligned} \int_0^1 \text{diff}(x)dx &= \sum_{j=0}^{2n-1} \int_{\frac{j}{2n}}^{\frac{j+1}{2n}} \text{diff}(x)dx = \\ &= \frac{G_{Y_A}(\frac{0}{2n}) - G_{Y_B}(\frac{0}{2n}) + G_{Y_A}(\frac{2n}{2n}) - G_{Y_B}(\frac{2n}{2n})}{4n} + \end{aligned}$$

$$\sum_{j=1}^{2n-1} \frac{2 \cdot G_{Y_A}(\frac{j}{2n}) - 2 \cdot G_{Y_B}(\frac{j}{2n})}{4n} =$$

$$\sum_{j=1}^{2n-1} \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n})}{2n}$$

Finally, since $G_{Y_A}(1) = G_{Y_B}(1) = 1$, we have that

$$\sum_{j=1}^{2n-1} \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n})}{2n} =$$

$$\sum_{j=1}^{2n} \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n})}{2n}$$

□

Proposition 3. ($\mathcal{C}_{\mathcal{P}}$ estimated from the cumulative difference-plot)

Let X_A and X_B be two random variables and A_n and B_n their n observations respectively. Let diff be the difference function obtained from the samples A_n and B_n as defined in Equation (2). Then,

$$\widetilde{\mathcal{C}}_{\mathcal{D}}(A_n, B_n) = \int_0^1 \text{diff}(x) dx + \frac{1}{2}$$

Proof. Given the observations A_n and B_n , we need to prove that

$$\sum_{i,k=1\dots n} \frac{\text{sign}(b_k - a_i)}{2n^2} + \frac{1}{2} = \int_0^1 \text{diff}(x) dx + \frac{1}{2}$$

With Lemma 1, it is enough to prove that

$$\sum_{i,k=1\dots n} \frac{\text{sign}(b_k - a_i)}{2n^2} = \sum_{j=1}^{2n} \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n})}{2n}$$

Let $C_{2n} = \{c_j\}_{j=1}^{2n}$ be the list of all the sorted observations of A_n and B_n where c_1 is the smallest observation and c_{2n} the largest. Then, we have that

$$G_{Y_A}(\frac{j}{2n}) = \frac{[A_n < c_j] + \frac{[A_n = c_j][k \leq j | c_k = c_j]}{[C_{2n} = c_j]}}{n} \text{ and}$$

$$G_{Y_B}(\frac{j}{2n}) = \frac{[B_n < c_j] + \frac{[B_n = c_j][k \leq j | c_k = c_j]}{[C_{2n} = c_j]}}{n}$$

where $[A_n < c_j]$ counts the number of items in A_n lower than c_j , and $[k \leq j | c_k = c_j]$ counts the number of items in C_{2n} equal to c_j but with a lower or equal position in C_{2n} . Therefore, we have that

$$\sum_{j=1}^{2n} \frac{G_{Y_A}(\frac{j}{2n}) - G_{Y_B}(\frac{j}{2n})}{2n} =$$

$$\sum_{j=1}^{2n} \frac{[A_n < c_j] + \frac{[A_n = c_j][k \leq j | c_k = c_j]}{[C_{2n} = c_j]} - [B_n < c_j] - \frac{[B_n = c_j][k \leq j | c_k = c_j]}{[C_{2n} = c_j]}}{2n^2}$$

$$\sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j] + \frac{([A_n = c_j] - [B_n = c_j])[k \leq j | c_k = c_j]}{[C_{2n} = c_j]}}{2n^2} \quad (5)$$

Now we group the terms in Equation (5) into d_{max} groups such that each group contains all the terms with the same c_j , and each group d contains $[C_{2n} = c_d]$ terms, with $c_j = c_d$.

$$\sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} + \sum_{d=1}^{d_{max}} \sum_{c_j} \frac{([A_n = c_j] - [B_n = c_j])[k \leq j | c_k = c_j]}{2n^2} =$$

$$\begin{aligned}
& \sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} + \sum_{d=1}^{d_{max}} \frac{([A_n = c_d] - [B_n = c_d])((C_{2n} = c_d) + 1) \cdot [C_{2n} = c_d]/2}{[C_{2n} = c_d] 2n^2} = \\
& \sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} + \sum_{d=1}^{d_{max}} \frac{([A_n = c_d] - [B_n = c_d])((C_{2n} = c_d) + 1)/2}{2n^2} = \\
& \sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} + \sum_{d=1}^{d_{max}} \frac{([A_n = c_d] - [B_n = c_d])((C_{2n} = c_d)/2 + ([A_n = c_d] - [B_n = c_d])/2)}{2n^2} = \\
& \sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} + \sum_{j=1}^{2n} \frac{([A_n = c_j] - [B_n = c_j])/2}{2n^2} + \underbrace{\sum_{d=1}^{d_{max}} \frac{([A_n = c_d] - [B_n = c_d])/2}{2n^2}}_{=0} = \\
& \underbrace{\sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2}}_{\text{first sum}} + \underbrace{\sum_{j=1}^{2n} \frac{([A_n = c_j] - [B_n = c_j])/2}{2n^2}}_{\text{second sum}}
\end{aligned}$$

Focusing on the first sum, we have that

$$\begin{aligned}
& \sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} = \\
& \frac{\sum_{j=1}^{2n} [A_n < c_j] - \sum_{j=1}^{2n} [B_n < c_j]}{2n^2} = \\
& \frac{\sum_{j=1}^{2n} \sum_{i=1}^n [\{a_i\} < c_j] - \sum_{j=1}^{2n} \sum_{i=1}^n [\{b_i\} < c_j]}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n [\{a_i\} < a_k] + \sum_{k=1}^n \sum_{i=1}^n [\{a_i\} < b_k]}{2n^2} - \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n [\{b_i\} < a_k] + \sum_{k=1}^n \sum_{i=1}^n [\{b_i\} < b_k]}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n [\{a_i\} < a_k] + [\{a_i\} < b_k] - [\{b_i\} < a_k] - [\{b_i\} < b_k]}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n [\{a_i\} < b_k] - [\{b_i\} < a_k] + [\{a_i\} < a_k] - [\{b_i\} < b_k]}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i) + [\{a_i\} < a_k] - [\{b_i\} < b_k]}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{k=1}^n [A_n < a_k] - [B_n < b_k]}{2n^2}
\end{aligned}$$

From the second sum, we obtain

$$\sum_{j=1}^{2n} \frac{([A_n = c_j] - [B_n = c_j])/2}{2n^2} = \sum_{k=1}^n \frac{([A_n = a_k] - [B_n = a_k] + [A_n = b_k] - [B_n = b_k])/2}{2n^2}$$

Combining these summations,

$$\begin{aligned}
& \sum_{j=1}^{2n} \frac{[A_n < c_j] - [B_n < c_j]}{2n^2} + \sum_{j=1}^{2n} \frac{([A_n = c_j] - [B_n = c_j])/2}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \\
& \frac{\sum_{k=1}^n [A_n < a_k] - [B_n < b_k]}{2n^2} + \frac{\sum_{k=1}^n ([A_n = a_k] - [B_n = a_k] + [A_n = b_k] - [B_n = b_k])/2}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \\
& \frac{\sum_{k=1}^n [A_n \leq a_k] - [B_n \leq b_k]}{2n^2} + \frac{\sum_{k=1}^n (-[A_n = a_k] - [B_n = a_k] + [A_n = b_k] + [B_n = b_k])/2}{2n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{k=1}^n [A_n \leq a_k] - [B_n \leq b_k]}{2n^2} + \frac{\sum_{k=1}^n (-[C_{2n} = a_k] + [C_{2n} = b_k])}{4n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{n(n+1)/2 + \sum_{d=1}^{d_{max}} \frac{[A_n=c_d]^2 - [A_n=c_d]}{2}}{2n^2} - \\
& \frac{n(n+1)/2 + \sum_{d=1}^{d_{max}} \frac{[B_n=c_d]^2 - [B_n=c_d]}{2}}{2n^2} + \frac{\sum_{k=1}^n (-[C_{2n} = a_k] + [C_{2n} = b_k])}{4n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{d=1}^{d_{max}} \frac{[A_n=c_d]^2 - [A_n=c_d]}{2} - \sum_{d=1}^{d_{max}} \frac{[B_n=c_d]^2 - [B_n=c_d]}{2}}{2n^2} + \\
& \frac{\sum_{k=1}^n (-[C_{2n} = a_k] + [C_{2n} = b_k])}{4n^2} =
\end{aligned}$$

considering that $\sum_{d=1}^{d_{max}} \frac{[B_n=c_d] - [A_n=c_d]}{2} = 0$, we simplify the previous equation to

$$\begin{aligned}
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{d=1}^{d_{max}} \frac{[A_n=c_d]^2 - [B_n=c_d]^2}{2}}{2n^2} + \frac{\sum_{k=1}^n (-[C_{2n} = a_k] + [C_{2n} = b_k])}{4n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{d=1}^{d_{max}} [A_n = c_d]^2 - [B_n = c_d]^2}{4n^2} + \frac{\sum_{k=1}^n (-[C_{2n} = a_k] + [C_{2n} = b_k])}{4n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{d=1}^{d_{max}} [A_n = c_d]^2 - [B_n = c_d]^2}{4n^2} + \\
& \frac{\sum_{d=1}^{d_{max}} (-[C_{2n} = c_d][A_n = c_d] + [C_{2n} = c_d][B_n = c_d])}{4n^2} = \\
& \frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{d=1}^{d_{max}} [A_n = c_d]^2 - [B_n = c_d]^2}{4n^2} + \underbrace{\frac{\sum_{d=1}^{d_{max}} [C_{2n} = c_d]([B_n = c_d] - [A_n = c_d])}{4n^2}}_{\text{third sum}} =
\end{aligned}$$

We expand the third sum,

$$\begin{aligned}
& \frac{\sum_{d=1}^{d_{max}} [C_{2n} = c_d]([B_n = c_d] - [A_n = c_d])}{4n^2} = \\
& \frac{\sum_{d=1}^{d_{max}} ([B_n = c_d] + [A_n = c_d])([B_n = c_d] - [A_n = c_d])}{4n^2} = \frac{\sum_{d=1}^{d_{max}} ([B_n = c_d]^2 - [A_n = c_d]^2)}{4n^2}
\end{aligned}$$

Finally,

$$\frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2} + \frac{\sum_{d=1}^{d_{max}} [A_n = c_d]^2 - [B_n = c_d]^2}{4n^2} + \frac{\sum_{d=1}^{d_{max}} ([B_n = c_d]^2 - [A_n = c_d]^2)}{4n^2} =$$

$$\frac{\sum_{k=1}^n \sum_{i=1}^n \text{sign}(b_k - a_i)}{2n^2}$$

□

Proposition 4. *Let X_A and X_B be two random variables and A_n and B_n their n observations respectively. The $\mathcal{C}_{\mathcal{D}}$ estimated from the cumulative difference-plot is $\widetilde{\mathcal{C}}_{\mathcal{D}}$.*

Proof. In Section 4.3, we defined the $\mathcal{C}_{\mathcal{D}}$ estimated from the cumulative difference-plot as

$$\mathcal{C}_{\mathcal{D}} = \frac{\int_0^1 \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx}{2} + \frac{1}{2},$$

$$\frac{\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx}{\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx},$$

where \mathcal{I} is the indicator function. This proposition claims that

$$\frac{\int_0^1 \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx}{2} + \frac{1}{2} =$$

$$\frac{\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx}{\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx} =$$

$$\frac{\sum_{j=1}^{2n} \frac{\psi(c_j)}{2n} + 1}{2} \cdot k_c^{-1}.$$

To prove it, we show that

$$i) \int_0^1 \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx = \sum_{j=1}^{2n} \frac{\psi(c_j)}{2n}$$

and

$$ii) \int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx = k_c.$$

Let us focus our attention in *i*). We split the integral into $2n$ parts:

$$\int_0^1 \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx =$$

$$\sum_{j=1}^{2n} \int_{\frac{j-1}{2n}}^{\frac{j}{2n}} \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx \quad (6)$$

Let $C_{2n} = \{c_j\}_{j=1}^{2n}$ be the list of all the sorted observations of A_n and B_n where c_1 is the smallest observation and c_{2n} the largest and let $\{c_d\}_{d=1}^{d_{max}}$ be the sorted list of unique values in C_{2n} . We group the terms in the sum of Equation (6) into d_{max} groups such that for every j in a group, $c_j = c_d$.

$$\sum_{d=1}^{d_{max}} \sum_j \int_{\frac{j-1}{2n}}^{\frac{j}{2n}} \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx$$

Now we join the integrals for every j in each group, such that the j of the integral goes from $j_{d\downarrow} - 1$ to $j_{d\uparrow}$ (if the sample c_d is unique in C_{2n} , then $j_{d\downarrow} = j_{d\uparrow} = j$).

$$\sum_{d=1}^{d_{max}} \int_{\frac{j_{d\downarrow}-1}{2n}}^{\frac{j_{d\uparrow}}{2n}} \mathcal{I}[\text{diff}(x) > 0] - \mathcal{I}[\text{diff}(x) < 0] dx \quad (7)$$

In the interval $(\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$, diff evaluates to one of these four possibilities:

1. $\text{diff}(x) = 0$ for all $x \in (\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$
2. $\text{diff}(x) > 0$ for all $x \in (\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$
3. $\text{diff}(x) < 0$ for all $x \in (\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$
4. $\text{diff}(x) = 0$ in one point in the interval $(\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$ and $\text{diff}(x) > 0$ or $\text{diff}(x) < 0$ for every other x in the interval. However, we can safely ignore this point as the value of the integral is invariant to the value of the function in sets of zero measure.

By looking at the empirical distributions $\hat{G}_A(x)$ and $\hat{G}_B(x)$ estimated from A_n and B_n respectively, we can guess which of these possibilities corresponds to each interval.

$$\left\{ \begin{array}{l} 1) \quad \text{if } \hat{G}_A(c_{d-1}) = \hat{G}_B(c_{d-1}) \\ \quad \text{and } \hat{G}_A(c_d) = \hat{G}_B(c_d) \\ 2) \quad \text{if } \hat{G}_A(c_{d-1}) \geq \hat{G}_B(c_{d-1}) \\ \quad \text{and } \hat{G}_A(c_d) > \hat{G}_B(c_d) \\ 2) \quad \text{if } \hat{G}_A(c_{d-1}) > \hat{G}_B(c_{d-1}) \\ \quad \text{and } \hat{G}_A(c_d) \geq \hat{G}_B(c_d) \\ 3) \quad \text{if } \hat{G}_B(c_{d-1}) \geq \hat{G}_A(c_{d-1}) \\ \quad \text{and } \hat{G}_B(c_d) > \hat{G}_A(c_d) \\ 3) \quad \text{if } \hat{G}_B(c_{d-1}) > \hat{G}_A(c_{d-1}) \\ \quad \text{and } \hat{G}_B(c_d) \geq \hat{G}_A(c_d) \\ 4) \quad \text{if } \hat{G}_B(c_{d-1}) > \hat{G}_A(c_{d-1}) \\ \quad \text{and } \hat{G}_A(c_d) > \hat{G}_B(c_d) \\ 4) \quad \text{if } \hat{G}_A(c_{d-1}) > \hat{G}_B(c_{d-1}) \\ \quad \text{and } \hat{G}_B(c_d) > \hat{G}_A(c_d) \end{array} \right.$$

The value of the integral in Equation (7) corresponding to these possibilities are the following:

1. 0
2. $[C_{2n} = c_d] \cdot \frac{1}{2n}$

3. $-[C_{2n} = c_d] \cdot \frac{1}{2n}$
4. $[C_{2n} = c_d] \cdot (2 \cdot l_d - 1) \cdot \frac{1}{2n}$

where $[C_{2n} = c_d]$ counts the number of items in C_{2n} equal to c_d and l_d is the proportion in which $\text{diff}(x) > 0$ in the interval $(\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$. For example, $l_d = 0.75$ would represent that $\text{diff}(x) > 0$ in 75% of the total length of the interval, and $\text{diff}(x) < 0$ in the other 25%.

With this, we can rewrite Equation (7) as

$$\sum_{d=1}^{d_{max}} [C_{2n} = c_d] \cdot \psi(c_d) \cdot \frac{1}{2n} = \sum_{j=1}^{2n} \frac{\psi(c_j)}{2n},$$

where ψ is the function introduced in Definition 9.

Now, we only need to prove *ii*). Specifically, we need to show that

$$\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx = k_c.$$

We have that

$$\int_0^1 \mathcal{I}[\text{diff}(x) \neq 0] dx = \sum_{d=1}^{d_{max}} \int_{\frac{j_{d\downarrow}-1}{2n}}^{\frac{j_{d\uparrow}}{2n}} \mathcal{I}[\text{diff}(x) \neq 0] dx,$$

and

$$k_c = \frac{\sum_{j=1}^{2n} \mathcal{I}[\psi(c_j) \neq 0]}{2n} = \sum_{d=1}^{d_{max}} [C_{2n} = c_d] \frac{\mathcal{I}[\psi(c_d) \neq 0]}{2n}.$$

Finally, it is easy to see that

$$\int_{\frac{j_{d\downarrow}-1}{2n}}^{\frac{j_{d\uparrow}}{2n}} \mathcal{I}[\text{diff}(x) \neq 0] dx = [C_{2n} = c_d] \frac{\mathcal{I}[\psi(c_d) \neq 0]}{2n},$$

because $\text{diff}(x) = 0$ in the interval $(\frac{j_{d\downarrow}-1}{2n}, \frac{j_{d\uparrow}}{2n})$ if and only if $\psi(c_d) = 0$.

□

11.2 Convergence of the estimators

Proposition 5. *Let X_A and X_B be two continuous random variables and $\{a_i\}_{i \in \mathbb{N}}$ and $\{b_i\}_{i \in \mathbb{N}}$ be two infinite sequences of their observations respectively. Let A_n and B_n be the two finite subsequences that contain the first n elements of $\{a_i\}_{i \in \mathbb{N}}$ and $\{b_i\}_{i \in \mathbb{N}}$ respectively. Then,*

$$\mathcal{C}_{\mathcal{P}}(X_A, X_B) = \lim_{n \rightarrow \infty} \widetilde{\mathcal{C}}_{\mathcal{P}}(A_n, B_n)$$

Proof. Let $\{P_s\}_{s \in \mathbb{N}}$ be a sequence of estimators with every estimator is determined randomly with the following procedure:

- 1) generate two random permutations σ_s and τ_s of size n .
- 2) define each estimation as

$$P_s(A_n, B_n) = \sum_{i=1}^n \frac{\text{sign}(b_{\sigma_s(i)} - a_{\tau_s(i)})}{2n} + \frac{1}{2}.$$

It is easy to see that each P_s is an estimator of $\mathcal{P}(X_A < X_B)$ (since X_A, X_B are continuous, we know that $\mathcal{P}(X_A = X_B) = 0$). Now observe that the sequence $\left\{ \frac{\sum_{t=1}^s P_t(A_n, B_n)}{s} \right\}_{n \in \mathbb{N}}$ converges to $\widetilde{\mathcal{C}}_{\mathcal{P}}(A_n, B_n) = \sum_{i,k=1 \dots n} \frac{\text{sign}(b_k - a_i)}{2n^2} + \frac{1}{2}$, which means that $\widetilde{\mathcal{C}}_{\mathcal{P}}(A_n, B_n)$ is also an estimator of $\mathcal{P}(X_A < X_B)$. □

Unfortunately, the estimator $\widetilde{\mathcal{C}}_{\mathcal{D}}$ will not always converge: $\mathcal{C}_{\mathcal{D}}$ fails to satisfy Property 7, and this means that a few points can still have a big impact in the estimation of $\mathcal{C}_{\mathcal{D}}$. Specifically, given the continuous random variables X_A and X_B defined in N , $\widetilde{\mathcal{C}}_{\mathcal{D}}$ will converge iff $\int_N \mathcal{I}[G_A(x) = G_B(x)] \cdot (g_A + g_B) dx = 0$.

Luckily, this lack of convergence is not a problem when the estimation of $\mathcal{C}_{\mathcal{D}}$ is carried out visually in the cumulative difference-plot. Since the visual representation of the cumulative difference-plot involves rendering the plot with pixels, there exists a small $\delta > 0$ such that when $|\text{diff}(x)| < \delta$, the difference is displayed as 0.

In practice, we do not even need to account for the case that $\text{diff}(x) = 0$. The cumulative difference-plot models the uncertainty with a confidence band, and when $\text{diff}(x) = 0$ is inside the confidence band, then so are $\text{diff}(x) > 0$ and $\text{diff}(x) < 0$. If we assume that the difference is positive, negative or zero every time that $\text{diff}(x) = 0$ is inside the confidence band, we obtain the estimations $\widetilde{\mathcal{C}}_{\mathcal{D}}^+$, $\widetilde{\mathcal{C}}_{\mathcal{D}}^-$ and $\widetilde{\mathcal{C}}_{\mathcal{D}}^0$ respectively. Now since $\widetilde{\mathcal{C}}_{\mathcal{D}}^+ > \widetilde{\mathcal{C}}_{\mathcal{D}}^- > \widetilde{\mathcal{C}}_{\mathcal{D}}^0$, the estimation of $\mathcal{C}_{\mathcal{D}}$ with the highest part of the confidence band is an upper bound of $\mathcal{C}_{\mathcal{D}}$. The same is true for the estimation with the lowest part of the confidence band: it is a lower bound of $\mathcal{C}_{\mathcal{D}}$.

Although $\widetilde{\mathcal{C}}_{\mathcal{D}}$ does not converge to $\mathcal{C}_{\mathcal{D}}$, for any $\epsilon > 0$ we can find a δ small enough such that the difference between $\widetilde{\mathcal{C}}_{\mathcal{D}}^\delta$ and $\mathcal{C}_{\mathcal{D}}$ is smaller than ϵ . We formalize this claim in Conjecture 1, and we leave the proof for future work.

Definition 10. (δ -estimation of $\mathcal{C}_{\mathcal{D}}$)

Let X_A and X_B be two continuous random variables and A_n and B_n their n observations respectively. Let $\{c_j\}_{j=1}^{2n}$ the sorted list of all the observations of A_n and B_n where c_1 is the smallest observation and c_{2n} the largest. Let $\{c_d\}_{d=1}^{d_{max}}$ be the sorted list of unique values in $\{c_j\}_{j=1}^{2n}$.

We define the δ -estimation of $\mathcal{C}_{\mathcal{D}}$, denoted as $\widetilde{\mathcal{C}}_{\mathcal{D}}^\delta$, as the same estimation as $\widetilde{\mathcal{C}}_{\mathcal{D}}$, but assuming that the empirical distributions computed from A_n and B_n are equal when $|\hat{G}_A(x) - \hat{G}_B(x)| < \delta$.

The previous definition can also be based in the δ -difference, defined as $\text{diff}^\delta(x) = \text{diff}(x)$ when $\text{diff}(x) \geq \delta$, and $\text{diff}^\delta(x) = 0$ otherwise.

Conjecture 1. Let X_A and X_B be two continuous random variables and $\{a_i\}_{i \in \mathbb{N}}$ and $\{b_i\}_{i \in \mathbb{N}}$ be two infinite sequences of their observations respectively. Let A_n and B_n be the two finite

subsequences that contain the first n elements of $\{a_i\}_{i \in \mathbb{N}}$ and $\{b_i\}_{i \in \mathbb{N}}$ respectively. Then, for all $\epsilon > 0$, there exists a $\delta > 0$ such that

$$\left| \mathcal{C}_{\mathcal{D}}(X_A, X_B) - \lim_{n \rightarrow \infty} \widetilde{\mathcal{C}}_{\mathcal{D}}^{\delta}(A_n, B_n) \right| < \epsilon$$

11.3 Experimental verification

In the following, we experimentally verify that the cumulative difference-plot can be used to deduce $\mathcal{C}_{\mathcal{D}}$ and $\mathcal{C}_{\mathcal{P}}$. To do so, we define six pairs of example random variables and measure the $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ with three different methods: the definition of $\mathcal{C}_{\mathcal{D}}$ and $\mathcal{C}_{\mathcal{P}}$ (Equation (1) and Definition 4), the estimators in Definitions 8 and 9 and from the *cumulative difference-plot*. The *cumulative difference-plot* has a confidence band in addition to the estimation, and this confidence band allows the lower and upper bounds of $\mathcal{C}_{\mathcal{D}}$ and $\mathcal{C}_{\mathcal{P}}$ to be computed.

The probability density functions of the six examples are shown in Figures 20 through 25. The probability density of these random variables is a mix of normal distributions, the beta distribution, and the log-normal distribution.

The difference plot and the estimations were carried out with 5000 samples from each random variable. The $\mathcal{C}_{\mathcal{P}}$ and $\mathcal{C}_{\mathcal{D}}$ values computed are shown in Figures 26 and 27 respectively. In every case, the estimations with the three methods match, except for $\mathcal{C}_{\mathcal{D}}$ in Example 4 (Figure 23). This is a deceptive example because, in most of the probability mass of X_A and X_B , the cumulative distribution functions are equal. Consequently, in this example, the estimator of $\mathcal{C}_{\mathcal{D}}$ introduced in Definition 9 is unstable: it is very likely that the estimated empirical distributions are different even though the cumulative distribution functions are identical. Overcoming this limitation involves choosing a small $\delta > 0$, such that when the difference between the empirical distributions is smaller than δ , they are considered equal.

We conclude that, in most cases, the three estimation methods (from densities, using the estimators and with the cumulative difference-plot) yield a similar result, which validates the statements in the previous section.




Figure 20: Probability density functions of Example 1




Figure 21: Probability density functions of Example 2




Figure 22: Probability density functions of Example 3




Figure 23: Probability density functions of Example 4




Figure 24: Probability density functions of Example 5




Figure 25: Probability density functions of Example 6




Figure 26: The C_P values obtained in the six examples with the three methods.




Figure 27: The \mathcal{C}_D values obtained in the six examples with the three methods.