



UvA-DARE (Digital Academic Repository)

Interrater reliability for incomplete and dependent data

ten Hove, D.

Publication date
2023

[Link to publication](#)

Citation for published version (APA):

ten Hove, D. (2023). *Interrater reliability for incomplete and dependent data*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 2

On the Usefulness of Interrater Reliability Coefficients

Abstract

For four data sets of different measurement levels, we computed 20 coefficients that estimate interrater reliability. The results show that the coefficients provide very different numerical values when applied to the same data. We discuss possible explanations for the differences among coefficients and suggest further research that is needed to clarify which coefficient a researcher should use to estimate interrater reliability.

Chapter 2 is published as: Ten Hove, D., Jorgensen, T. D., & Van der Ark, L. A. (2018). On the usefulness of interrater reliability coefficients. In M. Wiberg, S. A. Culpepper, R. Janssen, J. González, & D. Molenaar (Eds.), *Quantitative psychology: The 82nd annual meeting of the Psychometric Society, Zurich, Switzerland, 2017* (pp. 67–75). Springer. http://doi.org/10.1007/978-3-319-77249-3_6

2.1 Introduction

Interrater reliability (IRR) entails the degree of agreement, consistency, or shared variance among two or more raters assessing the same subjects, expressed as a number between 0 (no agreement) and 1 (perfect agreement). On September 27, 2017, the term “inter-rater reliability”—including quotation marks—returned 173,000 hits on Google Scholar¹, which illustrates its academic importance. IRR also has societal relevance. For example, in the Netherlands an officer of Child Protection Services (Raad voor de Kinderbescherming) assesses the recidivism risk, risk factors, and protective factors of each juvenile delinquent (Van der Put et al., 2011). For the juvenile delinquent, the stakes are high because the assessment by the officer of Child Protection Services determines the district attorney’s sentencing recommendation. If the IRR of the assessment procedure were low, the sentencing recommendation would largely depend on the officer who did the assessment, which is highly undesirable.

In our experience, most researchers associate IRR with Cohen’s (1960) kappa, but there is an abundance of coefficients available. Just for nominal data, Popping (1988) identified over 38 coefficients. Zhao et al. (2013) discussed 22 of these coefficients and found several were mathematically equivalent, resulting in 11 unique coefficients. The R package `irr` (Gamer et al., 2012) contains 17 different coefficients for various types of data that estimate the IRR. Some coefficients have different versions, which increases the number of coefficients even further. For example, the intraclass correlation coefficient (ICC) can be calculated using a one-way or two-way model, to estimate the consistency or agreement of either a single rating or the average rating across raters. Due to the abundance of coefficients, we found that preferring a particular coefficient to estimate IRR is hard to justify.

Despite review articles on IRR (e.g., Gwet, 2014; Hallgren, 2012), it is unknown to what degree the estimated IRR depends on the coefficient. It would be desirable if coefficients that can be applied to data with the same measurement level (e.g., nominal data) produce similar results. Therefore, this paper investigates to what degree the choice of coefficient affects the estimated IRR. In the discussion, we attempt to explain some of the differences among coefficients, and suggest research that is needed to answer the question: “Which coefficient should a researcher use to estimate interrater reliability?”.

2.2 Methods

2.2.1 Data

We selected four datasets that are freely available from the R package `irr` (see Table 2.1; Gamer et al., 2012). Each dataset contained the ratings of R raters observing S subjects. The dataset `Diagnoses` (Fleiss, 1971) consists of ratings by six psychiatrists classifying 30

¹In December 2022 (when finalizing this dissertation) 329,000 hits.

patients into one of five nominal diagnostic categories: depression, personality disorder, schizophrenia, neurosis, or other. The dataset **Vision** (Stuart, 1953) consists of the distance-vision performance of 7477 subjects using their left eye and their right eye. The two eyes are considered the two instruments (i.e., two raters). The ratings were measured on a scale from 1 (low performance) to 4 (high performance), which we treat as ordinal. The dataset **Video** is an artificial dataset consisting of four raters rating the credibility of 20 videotaped testimonies. Ratings could vary from 1 (not credible) to 6 (highly credible), though observed scores only ranged from 2 to 5. Technically, rating scales cannot yield interval-level data unless it can be known that the distance between adjacent integers is equivalent for any pair of adjacent integers across the range of the scale; however, unbiased results may be obtained by treating Likert-type rating scales containing at least five points as interval-level rather than ordinal-level data (Rhemtulla et al., 2012). Therefore, we treated the ratings in the **Video** dataset as interval-level data. The dataset **Anxiety** is also an artificial dataset, in which three raters rated the anxiety of 20 subjects on a scale from 1 (not anxious at all) to 6 (extremely anxious). The measurement level of these ratings was also treated as interval.

Table 2.1: Characteristics of the Four Datasets used in this Study

Dataset	S	R	NR	Min	Max	Level
Diagnoses	30	6	180	1	5	Nominal
Vision	7477	2	14954	1	3	Ordinal
Video	20	4	80	2	5	Interval
Anxiety	20	3	60	1	6	Interval

Note. S = number of subjects; R = number of raters; NR = number of ratings. Min = minimum value; Max = maximum value; Level = Measurement Level.

2.2.2 IRR Coefficients

We considered 20 IRR coefficients from the R package `irr` (version 0.84; Gamer et al., 2012). We considered nine coefficients for nominal ratings (Table 2.2, top panel). Cohen’s kappa (κ ; Cohen, 1960) can be used only for nominal ratings with two raters. Weighted versions of κ have been derived that can be used only for nominal ratings with two raters (Cohen, 1968). We calculated two weighted κ versions: κ with equal weights (κ_W) and κ with squared weights (κ_W^2). Three generalizations of κ were available to assess nominal data with more than two raters: Fleiss’ kappa (κ_{Fleiss} ; 1971), Conger’s exact kappa (κ_{Exact} ; 1980), and Light’s kappa (κ_{Light} ; 1971). The percent agreement, Krippendorff’s (1980) alpha, and coefficient iota (Janson & Olsson, 2001) each have a version for several measurement levels, including nominal-level ratings. Their coefficients for nominal ratings are denoted PA_N , α_N , and ι_N , respectively.

We considered four coefficients for ordinal ratings (Table 2.2, central panel). Kendall’s

(1948) W and the mean of Spearman's rank-order correlation ($\bar{\rho}$; Spearman, 1904) have been designed specifically for ordinal data, whereas the percent agreement and Krippendorff's (1980) alpha have a version for ordinal ratings. The latter two coefficients are denoted PA_O and α_O , respectively.

We considered seven coefficients for interval-level ratings (Table 2.2, bottom panel). Each coefficient can also be applied to ratio-level ratings. The percent agreement, Krippendorff's (1980) alpha, and coefficient iota (Janson & Olsson, 2001) have a version for interval ratings. These coefficients are denoted PA_I , α_I , and ι_I , respectively. For the Finn (1970) coefficient and the ICC (Shrout & Fleiss, 1979), we specified two-way models to treat both raters and subjects as each being randomly drawn from a population, which is often the case in social and behavioral research. In addition, for the ICC we computed the level of consistency rather than the level of absolute agreement. Furthermore, we computed the mean of Pearson's product-moment correlation coefficients (\bar{r} ; Pearson, 1985) and Robinson's measure of agreement (A ; Robinson, 1957).

We excluded three coefficients of the R package `irr` from our analyses, because they clearly measured something different than the IRR: the Stuart-Maxwell coefficient (Maxwell, 1970) and the Bhapkar (1966) coefficient assess homogeneity in marginal distributions, and the coefficient of Eliasziw et al. (1994) estimates intrarater reliability (i.e., consistency of repeated ratings from the same rater).

2.2.3 Analyses

For the nominal dataset (`Diagnoses`), we applied only nominal IRR coefficients. For the ordinal dataset (`Vision`), we applied all ordinal, nominal, and interval-level IRR coefficients, with the exception of α_N and α_I . The results of interval-level coefficients are interesting because researchers frequently treat Likert-type scales as though they are continuous. The results of nominal IRR coefficients are interesting when the ordering is not of primary interest in the application at hand. Therefore, for the interval-level datasets (`Video` and `Anxiety`), we also computed all nominal, ordinal, and interval-level IRR coefficients, with the exception of PA_N , PA_O , α_N , α_O , and ι_N .

We investigated the range of values obtained by these coefficients. We also investigated whether the choice of coefficient affects the conclusion about the IRR using the heuristic labels suggested by Landis and Koch (1977) for the use of κ : negative values indicate a poor IRR, values between 0.00 and 0.20 indicate a slight IRR; values between 0.21 and 0.40 indicate a fair IRR; values between 0.41 and 0.60 indicate a moderate IRR; values between 0.61 and 0.80 indicate a substantial IRR, and values between 0.81 and 1.00 indicate an almost perfect IRR. Furthermore, we investigated the following aspects of the IRR coefficients in Table 2.2, by checking the literature and the functions of the package `irr`: Are standard errors available? Is it possible to conduct null-hypothesis significance testing? Are missing data allowed? And if so, how can missing data be handled? How many raters are allowed?

Table 2.2: Characteristics of the 20 IRR Coefficients used in this Study.

Symbol	Name	SE	NHST	Miss	$R > 2$
Nominal level					
κ	Cohen's kappa	•	•		
κ_W	Weighted kappa (equal weights)	•	•		
κ_W^2	Weighted kappa (squared weights)	•	•		
κ_{Fleiss}	Fleiss' kappa	•	•		•
κ_{Exact}	Conger's exact kappa	•	•		•
κ_{Light}	Light's kappa	•	•		•
PA_N	Percent agreement	•			•
α_N	Krippendorff's alpha	•		•	•
ι_N	Coefficient iota				•
Ordinal level					
W	Kendall's W		•		•
$\bar{\rho}$	Mean Spearman's rank correlation				•
PA_O	percent agreement	•			•
α_O	Krippendorff's alpha	•		•	•
Interval level					
PA_I	Percent agreement	•			•
α_I	Krippendorff's alpha	•		•	•
ι_I	coefficient iota				•
$Finn_2$	Finn's coefficient (two-way)		•		•
ICC_2	Intraclass correlation coefficient (two-way)	•	•		•
\bar{r}	Mean Pearson's correlation				•
A	Robinson's A				•

Note. SE = standard errors are available; NHST = null-hypothesis significance test is available; Miss = missing data can be handled by other methods than listwise deletion; $R > 2$ = the method can handle more than two raters.

2.3 Results

Table 2.3 shows the variability of the evaluated IRR coefficients as estimated for the four datasets. For the nominal-level dataset **Diagnoses**, the six available IRR coefficients ranged from 0.17 (PA) to 0.46 (κ_{Light} ; $M = 0.40$, $SD = 0.11$). For the ordinal-level dataset **Vision**, the IRR coefficients ranged from 0.60 (several coefficients) to 0.85 (W ; $M = 0.69$, $SD = 0.09$), but from 0.71 (several coefficients) to 0.85 (W) if only ordinal IRR coefficients are considered. For the interval-level dataset **Video**, the IRR coefficients ranged from 0.04 (κ_{Fleiss}) to 0.92 ($Finn_2$; $M = 0.26$, $SD = 0.24$), but from 0.10 (α_I) to 0.92 ($Finn_2$) if only interval-level IRR coefficients are considered. For the interval-level dataset **Anxiety**, the IRR coefficients ranged from -0.04 (κ_{Fleiss}) to 0.54 (W ; $M = 0.22$, $SD = 0.21$), but from 0.00 (PA_I) to 0.50 ($Finn_2$) if only interval-level IRR coefficients are considered.

Table 2.3 (cf. the stars next to the values) also shows that the interpretation of the IRR of a dataset by means of the benchmarks of Landis and Koch (1977) depends on the choice of coefficient. For the dataset **Diagnoses**, the IRR could be labelled either slight,

Table 2.3: IRR Estimates of 20 Coefficients for 4 Datasets.

Coefficient	Diagnoses	Vision	Video	Anxiety
Nominal level				
κ	<i>a</i>	0.60*	<i>a</i>	<i>a</i>
κ_W	<i>a</i>	0.65**	<i>a</i>	<i>a</i>
κ_W^2	<i>a</i>	0.60*	<i>a</i>	<i>a</i>
κ_{Fleiss}	0.43*	0.60*	0.04	-0.04
κ_{Exact}	0.44*	0.60*	0.10	-0.02
κ_{Light}	0.46*	0.60*	0.07	-0.02
PA_N	0.17	<i>b</i>	<i>b</i>	<i>b</i>
α_N	0.43*	<i>b</i>	<i>b</i>	<i>b</i>
ι_N	0.44*	<i>b</i>	<i>b</i>	<i>b</i>
Ordinal level				
W	<i>c</i>	0.85***	0.39	0.54*
$\bar{\rho}$	<i>c</i>	0.71**	0.24	0.34
PA_O	<i>c</i>	0.71**	<i>b</i>	<i>b</i>
α_O	<i>c</i>	0.71**	<i>b</i>	<i>b</i>
Interval level				
PA_I	<i>c</i>	<i>b</i>	0.35	0.00
α_I	<i>c</i>	<i>b</i>	0.10	0.16
ι_I	<i>c</i>	0.60*	0.15	0.19
Finn ₂	<i>c</i>	0.78**	0.92***	0.50*
ICC ₂	<i>c</i>	0.70**	0.16	0.20
\bar{r}	<i>c</i>	0.70**	0.24	0.28
A	<i>c</i>	0.85***	0.40	0.48*
Ranges of values				
Range ^d	0.17 – 0.46	0.71 – 0.85	0.10 – 0.92	0.00 – 0.50
Range ^e	0.17 – 0.46	0.60 – 0.85	0.04 – 0.92	-0.04 – 0.54

Note. *a* = coefficient cannot be computed because the number of raters is greater than 2; *b* = coefficient was not computed because a version of the coefficient that applies to another measurement level was computed; *c* = coefficient was not computed because the measurement level of data is nominal; *d* = range of all IRR coefficients that match the measurement level of the ratings; *e* = range of all IRR coefficients. * = coefficient greater than 0.40 (moderate IRR); ** = coefficient greater than 0.60 (substantial IRR); *** = coefficient greater than 0.80 (almost perfect IRR). Estimates that correspond to the correct measurement level are printed in **boldface**.

fair, or moderate; for the dataset `Vision`, the IRR could be labelled either moderate, substantial, or almost perfect; for the dataset `Video`, the IRR could be labeled anywhere from slight to almost perfect; and for dataset `Anxiety`, the IRR could be labelled either poor, slight, fair, or moderate.

For 13 of the 20 coefficients, standard errors were provided (Table 2.2). To the best of our knowledge, for the other coefficients, standard errors are not available. For nine coefficients, a test statistic is available that tests whether the coefficient equals zero. Although no dataset contained missing values, it is worth noting that the package `irr` handles missing data differently for different coefficients. Coefficients α_N , α_O , and α_I use all available data by counting disagreements among any observed pair of ratings on the same subject (i.e., pairwise deletion). Coefficients ι_N and ι_I do not allow missing ratings (i.e., the software will return a missing value for the coefficient when any ratings are missing), whereas all other coefficients handle missing data by listwise deletion.

2.4 Discussion

The results showed that the coefficients provide very different numerical values when applied to the same dataset. Depending on the choice of the coefficient, the IRR label for a single dataset can range from poor to almost perfect. This seriously questions the usefulness of IRR coefficients. We limited ourselves to coefficients available in the R package `irr` (Gamer et al., 2012), so the ranges may be even wider if more coefficients were included. This problem should be investigated further.

The usefulness of the coefficients in this paper can be investigated only if IRR has a sound definition; however, a clear definition seems to be absent. Some coefficients (e.g., the ICC) are based on variance decomposition, which is compatible with the framework of generalizability theory (e.g., Vangeneugden et al., 2005), whereas other coefficients (e.g., PA) are derived from the concept of literal agreement. Coefficients that stem from different conceptualizations of IRR cannot all measure the same thing. In a recent discussion with Feng (2015), Krippendorff (2016) wrote: “ I contend Feng discusses reliability measures with seriously mistaken conceptions of what reliability is to assure us of” (p. 139). We need to distinguish the different theories behind the IRR coefficients and come up with a more accurate terminology to identify competing conceptualizations of IRR. Only if the theories and models behind IRR are sorted out, we can start investigating why some IRR coefficients produce higher values than others, and we can separate the wheat from the chaff. In that respect, we believe the work of Zhao et al. (2013) is a valuable contribution. They explain, for example, the flaws of chance-corrected coefficients such as κ . Once we have selected estimates for different conceptualizations of IRR, we can deal with other issues identified in this study.

Another major problem is that few coefficients can handle missing data. This is problematic because ratings in the social and behavioral sciences can be expensive. For example, an assessment of a juvenile delinquent by an officer of Child Protection Services

in The Netherlands (see our Introduction) takes approximately 6 to 8 hours. A study investigating the IRR must allow for planned missingness because it is financially and practically impossible to have all officers assess all juvenile delinquents. Hence, a useful coefficient must be estimable with missing data. We also found that for some coefficients, standard errors and confidence intervals cannot be computed and null-hypothesis testing is impossible. These standard errors, confidence intervals, and hypothesis tests should first be derived. Then the bias of all standard errors, the coverage of all confidence intervals, and the Type I error rate of all hypothesis tests should be investigated. Finally, we used the benchmarks of Landis and Koch (1977). These benchmarks are considered to be the single most often used benchmarks (e.g., Gwet, 2014, p. 164). The 42,000+ citations² of the Landis and Koch paper on Google Scholar indicate at least their widespread use. A relevant question may be whether these benchmarks, which were designed for κ , can be used for coefficients stemming from different conceptualizations of IRR. In future research, it should be investigated whether different sets of heuristic rules should be provided for different types of coefficients.

²In December 2022 (when finalizing this dissertation) 76,000+ citations