

## Bayesian Joint Modeling of Interrater and Intrarater Reliability with Multilevel Data

Nour Hawila<sup>a</sup> and Arthur Berg<sup>b</sup>

<sup>a</sup>*Bristol Meyers Squibb*

<sup>b</sup>*Division of Biostatistics & Bioinformatics, Penn State University*

### ABSTRACT

We formulate three generalized Bayesian models for analyzing interrater and intrarater reliability in the presence of multilevel data. Stan implementations of these models provide new estimates of interrater and intrarater reliability. We also derive formulas for calculating marginal correlations under each of the three models. Comparisons of the kappa estimates and marginal correlations across the different models are presented from two real-world datasets. Simulations demonstrate properties of the different measures of agreement under different model assumptions.

### KEYWORDS

Reliability; Bayesian; Hierarchical; Nested; Stan

## 1. Introduction

It is often important to report on the inconsistent classification of the ratings from different raters (interrater) as well as the inconsistency of the ratings of a single rater's repeated ratings (intrarater). Although often used interchangeably, the primary distinction between reliability and agreement is that agreement is defined as the degree to which classifications are identical, whereas reliability focuses on the extent of variability and error inherent in a measurement [7]. Interrater reliability can be defined as the extent to which two or more raters agree in classifying a common observation. It is a helpful measure in assessing the consistency of the implementation of a rating system [12]. Intrarater reliability refers to the consistency of one rater's classification over multiple time points and is best determined when multiple trials are administered over a short period of time [15]. Intrarater reliability is helping in verifying the reproducibility of clinical measurements [9].

Many articles have discussed different measures of association for assessing interrater and intrarater reliability, and are mostly dependent on the type of outcome [5]. Commonly used measures for the case of binary outcomes include Cohen's Kappa statistic with a corrected standard error and some of its variants [3, 4, 6], whereas the intraclass correlation coefficient [1] is generally restricted to continuous outcomes. Generalized estimating equations (GEE) are also used for modelling outcomes where Kappa estimates are calculated as a function of the GEE parameters [20]. [8] also

---

CONTACT Author<sup>b</sup>. Email: [berg@psu.edu](mailto:berg@psu.edu)

---

### Article History

Received : 18 April 2024; Revised : 14 May 2024, Accepted : 28 May 2024; Published : 30 June 2024

### To cite this paper

Nour Hawila & Arthur Berg (2024). Bayesian Joint Modeling of Interrater and Intrarater Reliability with Multilevel Data. *Asian Journal of Statistics and Applications*. 1(1), 85-99.

proposed a weighted interrater Kappa calculation based on a GEE approach for categorical data. All of these models utilize a frequentist approach and only estimate interrater reliability and intrarater reliability separately but not jointly.

Joint estimation of interrater and intrarater reliability addresses several key issues inherent in the separate estimation approaches commonly found in the literature. First, by jointly modeling these two types of reliability, we can more accurately capture the interdependencies between the consistency of different raters and the repeatability of the same rater over time. This is particularly important in studies where the same subjects are assessed multiple times by the same raters, as the correlation between assessments can significantly influence the overall reliability of the measurements.

Moreover, joint modeling allows for the simultaneous estimation of both types of reliability, providing a more comprehensive understanding of the measurement process. This is crucial for ensuring the validity of conclusions drawn from the data, especially in clinical settings where decision-making depends heavily on the precision and repeatability of diagnostic assessments. By estimating these reliabilities together, we can also better account for variability between subjects and across time points, leading to more robust estimates of measurement error. Separately estimating inter- and intra-rater reliability ignores the dependencies in the data, whereas joint estimation may provide more accurate reliability estimates, better statistical efficiency, and less bias.

In this paper, we compare three fully Bayesian joint models of interrater and intrarater reliability of increasing complexity. The first model, which we refer to as the Bayesian independent (BIN) model, expands on the model presented in [13] by converting it to a Bayesian model and incorporating a random effect for time thus allowing intrarater reliability to also be captured. The second model, which we refer to as the Bayesian partially nested (BPN) model, expands on the model presented in [14] by converting it to a Bayesian model and incorporating multiple time points. The third model, which we refer to as the Bayesian fully nested (BFN) model, provides a Stan implementation of the WinBUGS Bayesian model presented in [10] but with more flexible priors.

In Section 2 we describe both probability-based (Section 2.1) and model-based (Section 2.2) methods in estimating the probability of a positive classification and the resulting interrater and intrarater reliability calculations. Formulas for calculating model-based marginal correlations for each of the three models are provided (Section 2.3). In Section 3, we show the results of applying the proposed Bayesian models to two real-world datasets and provide simulation results aimed at assessing the properties of the different measures of association depending on various model assumptions. We conclude the paper with an overall summary of the findings along with limitations of the proposed methods.

## 2. Methods

We consider a three-level nested data structure  $\{Y_{ijk}\}$  for subject  $i$ , evaluated by rater  $j$ , taken at time  $k$ . This type of data structure is common in observational studies in social and behavioral science, with subjects nested within clusters [18]. The quality of rating procedures is often of interest with particular emphasis on interrater and intrarater reliability. In the case of continuous outcomes, the intraclass correlation coefficient (ICC) is typically used to quantify the degree to which different raters resemble each other or the extent to which raters resemble themselves at two different

time points. For the scope of this paper, we focus on the case of dichotomous outcomes where  $J$  raters evaluate  $I$  subjects at  $K$  different time points.

### 2.1. Probability-based Methods

We first consider the simple case of  $I$  subjects evaluated by  $J$  raters, where interrater reliability is of interest. Denoting the response (classification) of the  $j^{\text{th}}$  rater on the  $i^{\text{th}}$  subject by  $Y_{ij}$ , a disagreement random variable  $Z_i$  between raters  $j$  and  $j'$  evaluating subject  $i$  could be constructed. It is assumed that  $Z_i$  is common for all pairs  $(j, j')$ . Kappa coefficients for assessing interrater reliability between raters  $j$  and  $j'$  are defined by

$$\kappa = 1 - \frac{E(Z_i)}{\bar{E}(Z_i)} \quad (1)$$

where  $E(Z_i)$  is the expectation of  $Z_i$  and  $\bar{E}(Z_i)$  is the expectation of  $Z_i$  assuming statistical independence of raters  $j$  and  $j'$ , i.e.  $P(Y_{ij} = k_1, Y_{ij'} = k_2) = P(Y_{ij} = k_1)P(Y_{ij'} = k_2)$  [19].

#### 2.1.1. Cohen's Kappa

Cohen's Kappa coefficient in particular is obtained when  $Z_i = 1 - I(Y_{ij}, Y_{ij'})$  where  $I(Y_{ij}, Y_{ij'}) = 1$  if  $Y_{ij} = Y_{ij'}$  and 0 otherwise. Assuming equal joint probabilities between raters  $j$  and  $j'$  across the subjects – i.e.  $P(Y_{ij} = k_1) = P(Y_{ij'} = k_1)$  for  $k_1 = 0, 1$  – the expectations for Cohen's Kappa are given by

$$E(Z_i) = P(Y_{ij} = 1, Y_{ij'} = 0) + P(Y_{ij} = 0, Y_{ij'} = 1) \quad (2)$$

$$\bar{E}(Z_i) = P(Y_{ij} = 1)P(Y_{ij'} = 0) + P(Y_{ij} = 0)P(Y_{ij'} = 1) \quad (3)$$

We can now define

$$p_o = 1 - E(Z_i) = \Pr(Y_{ij} = 1, Y_{ij'} = 1) + \Pr(Y_{ij} = 0, Y_{ij'} = 0) \quad (4)$$

$$p_c = 1 - \bar{E}(Z_i) = \Pr(Y_{ij} = 1)Pr(Y_{ij'} = 1) + \Pr(Y_{ij} = 0)Pr(Y_{ij'} = 0) \quad (5)$$

to be the observed and expected agreement due to chance respectively. Cohen's Kappa is more commonly known in following form

$$\kappa_c = \frac{p_o - p_c}{1 - p_c}. \quad (6)$$

Another formulation for  $p_c$ , without assuming equal joint probabilities between raters is given by

$$p_c = \Pr(Y_{ij} = 1, Y_{i'j'} = 1) + \Pr(Y_{ij} = 0, Y_{i'j'} = 0) \quad (7)$$

The true measure of chance agreement  $p_c$  is the probability that two randomly selected raters from a population make identical classifications on two different randomly chosen subjects [13]. The definitions of  $p_o$  and  $p_c$  are left relatively broad because the formulas are adjusted depending on whether interrater reliability or intrarater reliability is of interest and the nesting structure. For a three-level nested structure in

particular,

$$p_o = \begin{cases} \Pr(Y_{ijk} = Y_{ij'k}) & \text{for interrater reliability} \\ \Pr(Y_{ijk} = Y_{ijk'}) & \text{for intrarater reliability} \end{cases} \quad (8)$$

$$p_c = \begin{cases} \Pr(Y_{ijk} = Y_{i'j'k}) & \text{for interrater reliability} \\ \Pr(Y_{ijk} = Y_{i'jk'}) & \text{for intrarater reliability} \end{cases} \quad (9)$$

Like Cohen's Kappa, Scott's Pi is also applicable when evaluating categorical outcomes and therefore also binary data. Scott's Pi is calculated using the formula given in (6) but with  $p_c$  calculated slightly differently. Scott's Pi assumes the distribution of the raters is the same and therefore  $p_c$  is calculated using squared "joint proportions" which are squared arithmetic means of the marginal proportions (whereas Cohen's uses squared geometric means of them) [16]. Fleiss' Kappa is a generalization of Cohen's Kappa to multiple raters (or time points) [6]. Congers Kappa which is a correction of Fleiss' Kappa and will be used throughout this paper when computing Kappa [4]. In the following two subsections, we present closed-form formulas for computing Fleiss' and Conger's Kappa.

### 2.1.2. Fleiss' Kappa (1971)

Let  $N$  represent the total number of subjects and  $M$  the number of raters per subject. Define  $n_{ij}$  to be the number of raters assigning subject  $i$  to category  $j$  where  $j = 0, 1$ . We also define  $p_j$  to be the proportion of all ratings which were assigned to category  $j$  and is given by

$$p_j = \frac{1}{NM} \sum_{i=1}^N n_{ij}.$$

We note here that  $\sum_j n_{ij} = M$  and  $\sum_j p_j = 1$

The agreement between the  $M$  raters for the  $i^{\text{th}}$  subject,  $P_i$  is the proportion of agreeing pairs out of all  $M(M-1)$  possible pairs of assignment given by

$$P_i = \frac{1}{M(M-1)} (n_{i1}(n_{i1} - 1) + n_{i0}(n_{i0} - 1))$$

or equivalently can be written as combinations out of all  $\binom{M}{2} = \frac{M(M-1)}{2}$  possible pair combinations as

$$P_i = \frac{1}{\binom{M}{2}} \left[ \binom{n_{i1}}{2} + \binom{n_{i0}}{2} \right].$$

The overall observed extent of agreement can be measured by the mean of  $P_i$ s given by:

$$P_o = \frac{1}{N} \sum_{i=1}^N P_i$$

If the raters made their assignments by chance, the expected mean proportion of agreement would be  $(p_j p_{j'} = p_j^2)$  given by

$$P_c = p_0^2 + p_1^2 = \left[ \frac{1}{NM} \sum_{i=1}^N n_{i0} \right]^2 + \left[ \frac{1}{NM} \sum_{i=1}^N n_{i1} \right]^2.$$

The Kappa coefficient of agreement between the  $m$  raters can now be given by

$$\kappa = \frac{P_o - P_c}{1 - P_c}.$$

### 2.1.3. Conger's Kappa (1980)

Using the same notation  $N, M$ , for  $j = 0, 1$   $P_o$  and  $P_c$  are given by:

$$P_o = \frac{1}{N} \sum_{i=1}^N \frac{n_{i1}(n_{i1} - 1) + n_{i0}(n_{i0} - 1)}{M(M - 1)}$$

$$P_c = \bar{p}_{+0}^2 - s_0^2/M + \bar{p}_{+1}^2 - s_1^2/M$$

where  $s_k^2$  is given by  $s_k^2 = \frac{1}{M-1}(p_{j0} - \bar{p}_{+0})^2 + \frac{1}{M-1}(p_{j1} - \bar{p}_{+1})^2$  and represents the variance of the proportion of  $p_{j0}$  and  $p_{j1}$ . The Kappa coefficient of agreement between the  $m$  raters can now be similarly given as

$$\kappa = \frac{P_o - P_c}{1 - P_c}.$$

## 2.2. Model-based Methods

The measures typically used in (2.1) possess some unfavorable traits. Namely, the measures do not incorporate uncertainty about the estimates in addition to the lack of ability to adjust for covariates that could potentially impact the outcomes and most importantly ignore the nested aspect of the data. In this section, we propose some and generalize other Bayesian models for modelling  $p_{ijk} := \Pr(Y_{ijk} = 1)$ , the probability the  $i^{th}$  subject is classified as a 'success' ( $Y = 1$ ) by the  $j^{th}$  rater at the  $k^{th}$  time point.

### 2.2.1. Independent model

[13] proposed a generalized linear mixed model (GLMM) for the case of a single time point given in (10).

$$g(p_{ij}) = \eta + u_i + v_j, \quad i = 1, \dots, I, j = 1, \dots, J \tag{10}$$

$g(\cdot)$  is a link function (usually chosen to be the probit or logit link) and  $\eta$  is the intercept term. The random effects terms are given by  $u_i$  and  $v_j$  for the  $i^{th}$  subject and  $j^{th}$  rater which are assumed to be independent and normally distributed with

mean 0 and variances  $\sigma_u^2$  and  $\sigma_v^2$ , respectively. Under the logit and probit models, estimates of  $p_0$  and  $p_c$  for a single time point ( $k = 1$ ) are calculated as a function of  $\eta, \sigma_u^2, \sigma_v^2$  which are typically replaced by their estimates from the GLMM. This model formulation is useful for evaluating interrater reliability in the presence of a single time point.

We propose the Bayesian model presented in (11) that incorporates an additional random effects terms  $w_k$  for the  $k^{th}$  time point in addition to a fixed effects term for the  $i^{th}$  subject evaluated by the  $j^{th}$  rater at the  $k^{th}$  time point.

$$g(p_{ijk}) = X_{ijk}\beta^I + u_i + v_j + w_k, \quad i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K \quad (11)$$

Positive values of  $u_i$  indicate that the  $i^{th}$  subject was more likely to be classified as a success amongst raters across time points, positive values of  $v_j$  indicate that the  $j^{th}$  rater was more likely to classify subjects as success across time points, and positive values of  $w_k$  indicate that raters were more likely to classify subjects as success at the  $k^{th}$  time point. The underlying hierarchical Bayesian structure for model (11) is given by

$$\begin{aligned} u_i &\sim N(0, \sigma_u^2) \\ v_j &\sim N(0, \sigma_v^2) \\ w_k &\sim N(0, \sigma_w^2) \\ \sigma_u^2, \sigma_v^2, \sigma_w^2 &\sim IG(\alpha, \gamma) \\ \beta^I &\sim MVN(\mu_p^\beta, \Sigma_p^\beta) \end{aligned} \quad (12)$$

Univariate Normal priors are assumed for the random effect terms with 0 mean and independent variance terms that are assumed to have inverse-gamma hyperpriors with parameters  $\alpha$  and  $\gamma$ . The inverse-gamma parameters are typically chosen to be small and equal to reflect non-informative priors. The fixed effects coefficient term  $\beta^I$  is assumed to have a p-variate Normal prior with known mean  $\mu_p^\beta$  and unstructured covariance matrix  $\Sigma_p^\beta$ .

### 2.2.2. Fully nested model

[10] proposed a Bayesian model able to assess interrater and intrarater reliability for nested binary data structures. The model is given in (13).

$$\text{logit}(p_{ijk}) = X_{ijk}\beta^F + u_i + v_{ij} + w_{ijk}, \quad i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K \quad (13)$$

A logit link function was chosen with a flexible fixed effects term  $X_{ijk}$  of length  $p$  and accompanying fixed effects coefficients vector  $\beta^F$ . A first level of independent random effects is imposed at the subject level through  $u_i$ , followed by  $v_{ij}$  that represents a joint subject/rater random effect and finally a subject/rater/time point joint random effect  $w_{ijk}$ . Similar to the independent model,  $u_i$  are assumed to follow an independent univariate Normal distribution with mean 0 and variance  $\sigma_u^2$ . The joint subject/rater random effects  $v_i$  are assumed to be  $J$ -variate Normal with mean  $[0 \ 0 \ \dots \ 0]^T$  and covariance matrix  $\Sigma_v^F = D_v^F \Omega_v^F D_v^F$  where  $\Omega_v^F$  is the correlation matrix with unit entries on the diagonal and  $\rho_v^F$  entries on the off-diagonal, and  $D_v^F = \sqrt{\text{diag}(\Sigma_v^F)} =$

$[\sigma_{v1}, \dots, \sigma_{vJ}]^T$  so that

$$\begin{aligned} \text{Var}(v_{ij}) &= \sigma_{vj}^2 \\ \text{Cov}(v_{ij}, v_{ij'}) &= \rho_v \sigma_{vj} \sigma_{vj'} \text{ for } j = 1, 2, \dots, J \text{ and } j \neq j' \end{aligned} \quad (14)$$

We assume a  $JK$ -variate Normal distribution for  $w_{ijk}$  with similar structure to  $v_{ij}$  with 0 mean and covariance matrix  $\Sigma_w^F = D_w^F \Omega_w^F D_w^F$  where  $\Omega_w^F$  is the correlation matrix with unit entries on the diagonal and  $\rho_w^F$  entries on the off-diagonal, and  $D_w^F = \sqrt{\text{diag}(\Sigma_w^F)} = [\sigma_{w11}, \dots, \sigma_{wJK}]^T$  so that

$$\begin{aligned} \text{Var}(w_{ijk}) &= \sigma_{wjk}^2 \\ \text{Cov}(w_{ijk}, w_{ij'k'}) &= \rho_w \sigma_{wjk} \sigma_{wjk'} \text{ for } j = 1, \dots, J, k = 1, \dots, K \text{ } j \neq j', k \neq k' \end{aligned} \quad (15)$$

We note that the covariance matrices  $\Sigma_v^F$  and  $\Sigma_w^F$  are assumed to be *unstructured* for generalizability, but other assumptions could be made. A *common* covariance structure would assume  $\sigma_{vj} = \sigma_{vj'}, \forall j \neq j'$  for  $\Sigma_v^F$  and  $\sigma_{wjk} = \sigma_{wjk'}, \forall j \neq j'$  and  $k \neq k'$  for  $\Sigma_w^F$ . For  $\Sigma_w^F$ , we can also assume *separate* covariance structures for different raters (or time points), i.e.  $\sigma_{wjk} = \sigma_{wjk'}, \forall j \neq j'$  (or  $\sigma_{wjk} = \sigma_{wjk'}, \forall k \neq k'$ ). The full hierarchical Bayesian structure for model (13) is given by

$$\begin{aligned} u_i &\sim N(0, \sigma_u^2) \\ v_i &\sim MVN(0, \Sigma_v) \\ w_{i..} &\sim MVN(0, \Sigma_w) \\ \sigma_u^2, \sigma_{vj}^2, \sigma_{wjk}^2 &\sim IG(\alpha, \gamma) \\ \Omega_v^F, \Omega_w^F &\sim LKJ(\eta) \\ \beta^F &\sim MVN(\mu_p^\beta, \Sigma_p^\beta) \end{aligned} \quad (16)$$

Lewandowski-Kurowicka-Joe (LKJ) prior distributions are assumed for the correlation matrix with tuning parameter  $\eta$  to control the strength of the correlations. The fixed effects coefficient term  $\beta^F$  is assumed to have a  $p$ -variate Normal prior with known mean  $\mu_p^\beta$  and unstructured covariance matrix  $\Sigma_p^\beta$ .

We call this model the fully-nested model due to the hierarchical nature of modelling the random effects. Due to the complex structure of this model, different correlation measurements are of interest, namely, marginal correlations and correlations among random effects. Correlation among rater random effects  $\rho^R = \text{Corr}(v_{ij}, v_{ij'})$  was used to imply the strength of similarity among the underlying random effect mechanism for raters  $j$  and  $j'$  and the marginal correlation between these raters (interrater marginal correlation) was calculated as  $\text{Corr}^R = \text{Corr}(\text{logit}(p_{ijk}), \text{logit}(p_{ij'k}))$ . Similarly, for time points  $k$  and  $k'$  the correlation among time point random effects is given by  $\rho^T = \text{Corr}(w_{ijk}, w_{ij'k'})$  and the marginal correlation between these time points (intrarater marginal correlation) by  $\text{Corr}^T = \text{Corr}(\text{logit}(p_{ijk}), \text{logit}(p_{ij'k'}))$ . The correlation among rater and time point random effects is used to infer similarity between raters and time points, whereas the marginal correlation measures are used to differentiate between raters or time points as they examine all dependencies involved.

### 2.2.3. Partially nested model

Twelve years after [13] proposed the independent GLMM, [14] proposed an ordinal GLMM model with a crossed random effects structure given in (17). The ordinal GLMM model uses a probit link function to reflect an underlying continuous outcome, and creates a categorized version of the latent variable with  $C$  levels, evaluated by  $J$  raters at  $K = 2$  time points. The primary usage of this model was to evaluate intrarater reliability across two time points. The model is given by

$$g(p_{ijk}) = \alpha_c - (\beta^P x_k + u_i + v_{jk}), \quad i = 1, \dots, I, j = 1, \dots, J, k = 1, 2 \quad (17)$$

where  $\alpha_c$  represents the cutoffs defining the ordinal scaling with  $\alpha_0 = -\infty$ ,  $\alpha_C = +\infty$  and  $c = 1, \dots, C - 1$ ,  $x_k$  is an indicator variable for the  $k^{th}$  time point. A fixed effects coefficient  $\beta^P$  provides an overall adjustment for raters at different time points. The subject level random effects  $u_i$  are assumed to be independent and normally distributed with mean 0 and variances  $\sigma_u^2$ , whereas the rater/time point joint random effects  $v_j$  are assumed to be bivariate Normal (BVN) with mean  $[0 \ 0]^T$  and covariance matrix  $\Sigma_v^P = D^P \Omega^P D^P$  where  $\Omega^P$  is the correlation matrix with unit entries on the diagonal and  $\rho_v^P$  entries on the off-diagonal, and  $D^P = \sqrt{\text{diag}(\Sigma_v^P)} = [\sigma_{v1}, \sigma_{v2}]^T$  so that

$$\begin{aligned} \text{Var}(v_{jk}) &= \sigma_{vk}^2 \\ \text{Cov}(v_{jk}, v_{jk'}) &= \rho_v \sigma_{vk} \sigma_{vk'} \text{ for } k = 1, 2 \text{ and } k \neq k' \end{aligned} \quad (18)$$

[14] also developed integrals for calculating the observed intrarater association and chance intrarater association using the parameters  $(\beta, \sigma_u^2, \Sigma_v^P)$  of the model given in (17). To prevent overestimating or underestimating the strength of association between raters, and to account for chance association, [14] suggest minimizing the impact of chance association on the Kappa estimate by selecting thresholds  $\alpha_{1\min}, \dots, \alpha_{C-1\min}$  such that the chance intrarater association integrand is minimized. A form of Cohen's Kappa (6), could also be calculated based on model (17) by computing the measure for a single rater's paired classification as

$$\kappa_c^{\text{avg}} = \frac{p_o^{\text{avg}} - p_c^{\text{avg}}}{1 - p_c^{\text{avg}}}, \quad (19)$$

where  $p_o^{\text{avg}}$  and  $p_c^{\text{avg}}$  are the average observed and chance interrater (or intrarater) association taken for all pairs of raters  $j$  and  $j'$  (or time points  $k$  and  $k'$ ).

To make this model compatible the case of a binary outcome ( $C = 2$ ), and  $K$  time points, an updated version of (17) is given by

$$g(p_{ijk}) = \eta - (\beta_1^P x_k + u_i + v'_{jk}), \quad i = 1, \dots, I, j = 1, \dots, J, k = 1, \dots, K \quad (20)$$

where  $\beta_1^P$  is a  $K \times 1$  coefficients vector where the  $k^{th}$  entry represents the fixed effect of the  $k^{th}$  time point.  $v'_{jk}$  is now  $K$ -variate Normally distributed with mean  $[0 \ 0 \ \dots \ 0]^T$  and covariance matrix  $\Sigma_v^{P'} = D^{P'} \Omega^{P'} D^{P'}$  with similar formulation to



$\Sigma_v^P$  but  $K$ -dimensional so that

$$\begin{aligned} \text{Var}(v_{jk}) &= \sigma_{vk}^2 \\ \text{Cov}(v_{jk}, v_{jk'}) &= \rho_v \sigma_{vk} \sigma_{vk'} \text{ for } k = 1, 2, \dots, K \text{ and } k \neq k' \end{aligned} \quad (21)$$

Incorporating multiple time points allow for a generalized assessment of intrarater reliability. The underlying hierarchical Bayesian structure for model (20) is given by

$$\begin{aligned} u_i &\sim N(0, \sigma_u^2) \\ v_j &\sim MVN(0, \Sigma_v^P) \\ \sigma_u^2, \sigma_{vk}^2 &\sim IG(\alpha, \gamma) \\ \Omega^P &\sim LKJ(\eta) \\ \beta_1^P &\sim MVN(\mu_K^\beta, \Sigma_K^\beta) \end{aligned} \quad (22)$$

Univariate Normal priors are assumed for the subject random effect term with 0 mean and independent variance terms that are assumed to have inverse-gamma hyperpriors with parameters  $\alpha$  and  $\gamma$ . The inverse-gamma parameters are typically chosen to be small and equal to reflect non-informative priors. Multivariate Normal priors are assumed for the rater/time point joint random effect with 0 mean and covariance matrix as given in (21). The fixed effects coefficient term  $\beta_1^P$  is assumed to have a p-variate Normal prior with known mean  $\mu_K^\beta$  and unstructured covariance matrix  $\Sigma_K^\beta$ .

### 2.3. Marginal Correlations and Correlations between random effects

After fitting one of the proposed Bayesian models from the previous section, two different measures other than Kappa could be of interest. Namely, the marginal correlation between two raters (or two time points) which is defined as the correlation between the observations made by raters  $j$  and  $j'$  (or time points  $k$  and  $k'$ ).

$$\text{Corr}^R = \text{Corr}(g(p_{ijk}), g(p_{ij'k})) \quad (23)$$

$$\text{Corr}^T = \text{Corr}(g(p_{ijk}), g(p_{ijk'})) \quad (24)$$

Another measure of interest is the correlation between random rater effects  $\rho^R$  or between random time point effects  $\rho^T$ . These measures differ based on the type of model selected (independent, partially nested, or fully nested).

#### 2.3.1. Independent Model

For the independent model, the subject, rater and time point random effects are a priori univariate independent Normally distributed. The marginal correlations based on this model are given by:

$$\text{Corr}_{\text{IN}}^R = \frac{\sigma_u^2 + \sigma_w^2}{\sigma_u^2 + \sigma_v^2 + \sigma_w^2} \quad (25)$$

$$\text{Corr}_{\text{IN}}^T = \frac{\sigma_u^2 + \sigma_v^2}{\sigma_u^2 + \sigma_v^2 + \sigma_w^2} \quad (26)$$

The correlation between rater random effects and time point random effects are  $\rho^R = \rho^T = 0$  due to independence of the random effects.

### 2.3.2. Fully Nested Model

Assuming a fully nested Bayesian model as given in 13, the marginal correlations between two raters and two time points are given by

$$Corr_{FN}^R = \frac{\sigma_u^2 + \rho_R \times \sigma_v^2}{\sigma_u^2 + \sigma_v^2 + \sigma_w^2} \quad (27)$$

$$Corr_{FN}^T = \frac{\sigma_u^2 + \rho_T \times \sigma_w^2}{\sigma_u^2 + \sigma_v^2 + \sigma_w^2} \quad (28)$$

where  $\rho_R$  is the correlation between rater random effects and  $\rho_T$  is the correlation between time point random effects.

### 2.3.3. Partially Nested Model

Assuming a Bayesian partially nested model as given in 20, the marginal correlations between two raters and two time points are given by

$$Corr_{PN}^R = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_{vk}^2} \quad (29)$$

$$Corr_{PN}^T = \frac{\sigma_u^2 + \rho_T \sigma_{vk} \sigma_{vk'}}{\sigma_u^2 + \sigma_{vk}^2} \quad (30)$$

The correlation between rater random effects is  $\rho^R = 0$  and  $\rho^T$  is the correlation between time point random effects.

## 3. Simulation and Examples

### 3.1. Software and Implementation

All the developed functions that will allow users to implement the three Bayesian models proposed in Sections 2.2.1, 2.2.2 and 2.2.3 using RStan [17] in addition to easily computing interrater and intrarater reliability posterior estimates, credible intervals and simulation results are available through <https://github.com/NourHawila/IRR>.

### 3.2. Data Examples

All three Bayesian models – BIN, BPN and BFN – are fit to two real-world datasets. The resulting model fits are subsequently used to explore model properties through simulations.

#### 3.2.1. Dataset: Running gait

In a crossover study published in 2022, [11] recruited 32 cross-country, track and field, and recreational athletes with current running mileage of at least 15 km per week to compare indoor and outdoor running environments. Athletes first ran on a

treadmill in an indoor environment recorded using static video, followed by outdoor video recorded by a drone. Three judges were shown both videos, two weeks apart, and independently performed running gait analysis on each foot for each of the 32 runners, with the goal of assessing interrater and intrarater reliability in addition to indoor vs. outdoor agreement. This dataset consists of a total of 768 data points (32 runners, 3 raters, 2 time points, 2 feet, 2 locations) in a complete block design. Fleiss' Kappa was used to assess reliability in [11] after averaging across the nested structures.

### 3.2.2. Dataset: Radiograph

The quality of root canal is typically evaluated by a radiographic assessment that is known to be necessary yet subjective. A study was conducted to assess the consistency and accuracy of the radiographic evaluation of 7 endodontists who partook in a training course and evaluated radiographs from 35 participants before and after the training [2]. The quality of endodontic treatment is defined as "good" if the filling reaches an adequate length within 2 mm of the radiographic apex and if the complete obturation is in the apical one-third of the root canal [10]. This dataset consists of a total of 490 data points (35 subjects, 7 raters, 2 time points) in a complete block design. [10] used the BFN model to assess interrater and intrarater reliability using WinBUGS.

### 3.2.3. Model fits of two real-world datasets

The three Bayesian models were fit to each of the two datasets. Each of the models used 2000 iterations, 200 warmup iterations per chain across 2 chains. The functions `model_BIN`, `model_BPN`, and `model_BFN`, available at <https://github.com/NourHawila/IRR>, are used to fit the datasets to the respective Bayesian models. The syntax for fitting the three models is presented below. The `fixed_eff_intercept` argument controls whether the fixed effects design matrix includes an intercept term or not. The parameters `beta_a` and `beta_b` are the hyperpriors for the correlation terms `rho_R` and `rho_T` which are assumed to follow Beta distributions. The parameters `gamma_a` and `gamma_b` are the hyperpriors for the variance terms `sigma_S`, `sigma_R`, and `sigma_T` which are assumed to follow inverse gamma distributions. The parameters `rho_R_eta` and `rho_T_eta` are the hyperprior shape parameters for the correlation terms `rho_R` and `rho_T` respectively which are assumed to follow an LKJ correlation distribution.

```
BIN model_BIN(df, fixed_eff_intercept = TRUE, beta_a = 5, beta_b =
  5, gamma_a = 3, gamma_b = 1.5, beta_mean = 0, beta_sigma = 1/0.3,
  niters = 2000, nwarmup = 200, nchains = 2)
BPN model_BPN(df, cov_T_str = "common", fixed_eff_intercept = TRUE,
  beta_a = 5, beta_b = 5, gamma_a = 3, gamma_b = 1.5, beta_mean = 0,
  beta_sigma = 1/0.3, betak_mean = 0, betak_sigma = 1/0.3, rho_T_eta
  = 1, niters = 2000, nwarmup = 200, nchains = 2) )
BFN model_BFN(df, cov_R_str = "common", cov_T_str = "common",
  fixed_eff_intercept = TRUE, beta_a = 5, beta_b = 5, gamma_a = 3,
  gamma_b = 1.5, beta_mean = 0, beta_sigma = 1/0.3, rho_R_eta = 1,
  rho_T_eta = 1, niters = 2000, nwarmup = 200, nchains = 2) )
```

Table 1 presents fits of the three Bayesian models to the two datasets. Based on the leave-one-out information criterion (LOOIC), the best model for the running gait dataset is the BPN model with the BIN model in a close second place. The LOOIC

selected the BFN as the best model for the radiograph data, while the other two models displayed substantially larger LOOICs.

Estimates of the interrater and intrarater Kappa measures based on direct frequentist estimation and model-based methods are presented in Table 2. Using the generated quantities block in Rstan, we simulate data from the posterior predictive distribution and calculate the interrater and intrarater kappa estimates. For the running gait dataset, the LOOIC-selected BPN model produced interrater and intrarater kappa estimates (0.24 and 0.30, respectively) that substantially differ from the frequentist-based estimates (0.44 and 0.45, respectively). Similarly, for the radiograph dataset, the LOOIC-selected BFN model produced interrater and intrarater kappa estimates (0.07 and 0.33, respectively) that are even more different from the frequentist-based estimates (0.40 and 0.72, respectively). These results show that datasets with small number of raters and timepoints can lead to widely varying estimates of interrater and intrarater reliability depending on what approach is used. In the next section, we simulate datasets from different models to evaluate these differences in interrater and intrarater reliability estimation.

Table 1.: Posterior mean estimates, effective number of parameters (p\_LOO), and the LOOIC for fitting the three Bayesian models to the two datasets described in Section 3.2

Model	Parms	Running Gait	Radiograph
BIN	$\sigma_S$	0.91	1.31
	$\sigma_R$	0.79	0.77
	$\sigma_T$	0.79	0.78
	$Corr^R$	0.70	0.79
	$Corr^T$	0.70	0.78
	p_LOO	32.02	32.29
	LOOIC	720.89	396.54
BPN	$\sigma_S$	0.89	1.31
	$\sigma_T$	0.75	0.70
	$Corr^T$	0.80	0.88
	$\rho^T$	0.50	0.51
	p_LOO	33.17	35.9
	LOOIC	<b>720.44</b>	398.81
BFN	$\sigma_S$	0.70	0.63
	$\sigma_R$	0.65	0.88
	$\sigma_T$	0.64	1.47
	$Corr^R$	0.55	0.24
	$Corr^T$	0.84	0.72
	$\rho^R$	0.54	0.42
	$\rho^T$	0.47	0.51
	p_LOO	72.38	82.83
	LOOIC	741.02	<b>318.02</b>

Table 2.: Interrater and intrarater Conger Kappa estimates for the two datasets

		<b>Running Gait</b>	<b>Radiograph</b>
<b>Freq</b>	$\kappa_{\text{inter}}$	0.44	0.40
	$\kappa_{\text{intra}}$	0.45	0.72
<b>BIN</b>	$\kappa_{\text{inter}}$	0.26	0.39
	$\kappa_{\text{intra}}$	0.27	0.40
<b>BPN</b>	$\kappa_{\text{inter}}$	0.24	0.37
	$\kappa_{\text{intra}}$	0.30	0.42
<b>BFN</b>	$\kappa_{\text{inter}}$	0.18	0.07
	$\kappa_{\text{intra}}$	0.26	0.33

### 3.3. Simulations

Datasets are simulated from the three different models – independent (IN), partially-nested (PN), and fully-nested (FN) – as presented in Section 2.2. The parameter values used in these model-based simulations are the posterior estimates from the respective model fits presented in Table 1. In this way, the simulated datasets are based on the two real world datasets, and these respective simulations are labeled as ‘running gait’ and ‘radiograph’. For each dataset-based simulation (running gait and radiograph), and for each model (IN, PN, FN), a total of 148 datasets were generated. For each of the generated datasets, the three Bayesian models (BIN, BPN, BFN) were fit to the data and their respective LOOIC was recorded. The model with the smallest LOOIC is referred to as the ‘LOOIC-selected’ Bayesian model for that simulation. Over each set of 148 simulations, the relative proportions of the LOOIC-selected models are presented in Table 3.

When the true simulation is IN, the BIN model is the most common LOOIC-selected model for both dataset references. Similarly, when the true simulation is PN, the BPN model is the most common LOOIC-selected model for both dataset references. However, when the true simulation is FN, the BFN model is the most common LOOIC-selected model for the running gait dataset reference but the BPN model is the most common LOOIC-selected model for the radiograph dataset reference.

Table 3.: Relative proportions of LOOIC-selected models for 148 simulations with each row adding up to 100% and the most commonly selected model for each set of simulations is bolded

<b>Dataset Reference</b>	<b>Sim Model</b>	<b>BIN</b>	<b>BPN</b>	<b>BFN</b>
<b>Running Gait</b>	IN	<b>77.7%</b>	18.2%	4.1%
	PN	14.2%	<b>85.1%</b>	0.7%
	FN	7.4%	37.8%	<b>54.7%</b>
<b>Radio-graph</b>	IN	<b>73.6%</b>	21.6%	4.7%
	PN	12.2%	<b>82.4%</b>	5.4%
	FN	1.4%	<b>98.6%</b>	0%

For each dataset-based simulation (running gait and radiograph), and for each model (IN, PN, FN), approximate theoretical interrater and intrarater reliability measures

are generated by simulating 10,000 datasets and taking the average kappa estimates. The resulting ‘true’ kappa parameters are displayed in Table 4. The frequentist-based and model-based (BIN, BPN, BFN, LOOIC-selected) interrater and intrarater kappa estimates are calculated for each of the 148 simulated datasets and the average parameter estimates and root mean square error (RMSE) performances are presented. The overall performance of each method (frequentist, BIN, BPN, BFN, and LOOIC-selected) over all  $148 \times 3 = 444$  simulated datasets within each dataset reference is highlighted in blue. The Bayesian model consistent with the respective data simulation is highlighted in yellow. The method with the smallest RMSE is bolded in each row of the table.

The simulations show all three Bayesian models perform similarly in terms of interrater and intrarater estimation. However, the frequentist-based estimation of interrater and intrarater reliability is shown to be worse than each of the Bayesian models for every simulation scenario. All three Bayesian models tend to perform similarly with regard to interrater and intrarater reliability estimation; the Bayesian model consistent with the data simulation model was not always optimal though it is consistently close to the optimal model.

Table 4.: Kappa simulation results showing mean estimates (RMSE)

Model Parms	Sim	$\kappa$	True	Freq	BIN	BPN	BFN	LOO
Running Gait	IN	inter	0.261	0.265 (0.111)	0.266 (0.106)	0.266 (0.106)	<b>0.254 (0.102)</b>	0.265 (0.105)
		intra	0.271	0.265 (0.115)	<b>0.268 (0.108)</b>	0.263 (0.107)	<b>0.271 (0.103)</b>	0.267 (0.108)
	PN	inter	0.236	0.231 (0.06)	0.242 (0.056)	0.233 (0.056)	<b>0.226 (0.055)</b>	0.233 (0.057)
		intra	0.297	0.297 (0.086)	0.321 (0.068)	0.299 (0.071)	<b>0.325 (0.065)</b>	0.3 (0.071)
	FN	inter	0.178	0.174 (0.066)	0.201 (0.063)	0.194 (0.061)	<b>0.177 (0.056)</b>	0.182 (0.058)
Avg RMSE	intra	0.259	0.261 (0.082)	<b>0.262 (0.065)</b>	0.246 (0.07)	0.286 (0.067)	0.272 (0.068)	
Radio-graph	IN	inter	0.385	0.387 (0.088)	<b>0.381 (0.083)</b>	0.381 (0.084)	0.371 (0.085)	0.38 (0.084)
		intra	0.399	0.39 (0.113)	<b>0.386 (0.105)</b>	0.372 (0.106)	0.395 (0.107)	0.385 (0.107)
	PN	inter	0.373	0.378 (0.072)	0.373 (0.067)	0.372 (0.067)	0.365 (0.069)	<b>0.371 (0.067)</b>
		intra	0.423	0.426 (0.075)	<b>0.439 (0.061)</b>	0.415 (0.062)	0.455 (0.064)	0.421 (0.064)
	FN	inter	0.074	0.074 (0.052)	0.092 (0.049)	0.076 (0.049)	<b>0.085 (0.049)</b>	0.076 (0.049)
Avg RMSE	intra	0.331	0.327 (0.174)	0.413 (0.147)	0.308 (0.164)	<b>0.341 (0.123)</b>	0.309 (0.164)	
Avg RMSE	inter	–	0.072	0.068	0.068	0.069	<b>0.068</b>	
	intra	–	0.127	0.11	0.118	<b>0.101</b>	0.119	

#### 4. Discussion

In this paper we have presented three Bayesian GLMMs and assessed the properties of different model assumptions on the estimation of interrater and intrarater reliability. The proposed models are comprehensive in terms of jointly adjusting for subject, rater and time point effects. The models are also flexible and able to incorporate different prior beliefs and knowledge about specific model parameters through Bayesian modelling. Contrary to commonly used models, the ones presented are not restricted to two raters or time points but are applicable for the general case of  $I$  subjects,  $J$  raters and  $K$  time points.

Some drawbacks in utilizing these methods for assessing interrater and intrarater reliability include the complexity in analyzing all the different parameters of interest and ensuring proper convergence of the model. Selecting the most appropriate model among the three Bayesian models requires a good understanding of the data, though the simulations show that utilizing the LOOIC model selection method will often select

the optimal or near-optimal model. As in any Bayesian model fit, the choice of priors may have substantial impact on the posterior estimates, so a sensitivity analysis that varies the priors should be performed.

## References

- [1] Bartko, J. J. [1966], ‘The intraclass correlation coefficient as a measure of reliability’, *Psychological reports* **19**(1), 3–11.
- [2] Chueh, L.-H., Chen, S.-C., Lee, C.-M., Hsu, Y.-Y., Pai, S.-F., Kuo, M.-L., Chen, C.-S., Duh, B.-R., Yang, S.-F., Tung, Y.-L. et al. [2003], ‘Technical quality of root canal treatment in Taiwan’, *International Endodontic Journal* **36**(6), 416–422.
- [3] Cohen, J. [1960], ‘A coefficient of agreement for nominal scales’, *Educational and psychological measurement* **20**(1), 37–46.
- [4] Conger, A. J. [1980], ‘Integration and generalization of kappas for multiple raters.’, *Psychological Bulletin* **88**(2), 322.
- [5] Donner, A. [1986], ‘A review of inference procedures for the intraclass correlation coefficient in the one-way random effects model’, *International Statistical Review/Revue Internationale de Statistique* pp. 67–82.
- [6] Fleiss, J. L. [1971], ‘Measuring nominal scale agreement among many raters.’, *Psychological bulletin* **76**(5), 378.
- [7] Gisev, N., Bell, J. S. and Chen, T. F. [2013], ‘Interrater agreement and interrater reliability: key concepts, approaches, and applications’, *Research in Social and Administrative Pharmacy* **9**(3), 330–338.
- [8] Gonin, R., Lipsitz, S. R., Fitzmaurice, G. M. and Molenberghs, G. [2000], ‘Regression modelling of weighted  $\kappa$  by using generalized estimating equations’, *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **49**(1), 1–18.
- [9] Gwet, K. L. [2008], ‘Intrarater reliability’, *Wiley encyclopedia of clinical trials* **4**.
- [10] Hsiao, C. K., Chen, P.-C. and Kao, W.-H. [2011], ‘Bayesian random effects for interrater and test–retest reliability with nested clinical observations’, *Journal of Clinical Epidemiology* **64**(7), 808–814.
- [11] Lafferty, L., Wawrzyniak, J., Chambers, M., Pagliarulo, T., Berg, A., Hawila, N. and Silvis, M. [2022], ‘Clinical indoor running gait analysis may not approximate outdoor running gait based on novel drone technology’, *Sports health* **14**(5), 710–716.
- [12] Lange, R. T. [2011], ‘Inter-rater reliability’, *Encyclopedia of clinical neuropsychology* pp. 1348–1348.
- [13] Nelson, K. P. and Edwards, D. [2008], ‘On population-based measures of agreement for binary classifications’, *Canadian Journal of Statistics* **36**(3), 411–426.
- [14] Nelson, K. P., Zhou, T. J. and Edwards, D. [2020], ‘Measuring intrarater association between correlated ordinal ratings’, *Biometrical Journal* **62**(7), 1687–1701.
- [15] Scheel, C., Mecham, J., Zuccarello, V. and Mattes, R. [2018], ‘An evaluation of the inter-rater and intra-rater reliability of occupro’s functional capacity evaluation’, *Work* **60**(3), 465–473.
- [16] Scott, W. A. [1955], ‘Reliability of content analysis: The case of nominal scale coding’, *Public opinion quarterly* pp. 321–325.
- [17] Stan Development Team [2020], ‘RStan: the R interface to Stan’. R package version 2.21.2.  
**URL:** <http://mc-stan.org/>
- [18] Ten Hove, D., Jorgensen, T. D. and van der Ark, L. A. [2021], ‘Interrater reliability for multilevel data: A generalizability theory approach.’, *Psychological Methods* .
- [19] Vanbelle, S. [2017], ‘Comparing dependent kappa coefficients obtained on multilevel data’, *Biometrical Journal* **59**(5), 1016–1034.
- [20] Williamson, J. M. and Manatunga, A. K. [1997], ‘Assessing interrater agreement from dependent data’, *Biometrics* pp. 707–714.