Published in final edited form as:

Stat Med. 2020 February 10; 39(3): 279–293. doi:10.1002/sim.8398.

Assessing method agreement for paired repeated binary measurements administered by multiple raters

Wei Wang¹, Nan Lin^{1,2}, Jordan D. Oberhaus³, Michael S. Avidan³

¹Department of Mathematics and Statistics, Washington University in St. Louis, St. Louis, Missouri

²Division of Biostatistics, Washington University School of Medicine in St. Louis, St. Louis, Missouri

³Department of Anesthesiology, Washington University School of Medicine in St. Louis, St. Louis, Missouri

Abstract

Method comparison studies are essential for development in medical and clinical fields. These studies often compare a cheaper, faster, or less invasive measuring method with a widely used one to see if they have sufficient agreement for interchangeable use. Moreover, unlike simply reading measurements from devices, eg, reading body temperature from a thermometer, the response measurement in many clinical and medical assessments is impacted not only by the measuring device but also by the rater. For example, widespread inconsistencies are commonly observed among raters in psychological or cognitive assessment studies due to different characteristics such as rater training and experience, especially in large-scale assessment studies when many raters are employed. This paper proposes a model-based approach to assess agreement of two measuring methods for paired repeated binary measurements under the scenario where the agreement between two measuring methods and the agreement among raters are required to be studied simultaneously. Based upon the generalized linear mixed models (GLMMs), the decision on the adequacy of interchangeable use is made by testing the equality of fixed effects of methods. Approaches for assessing method agreement, such as the Bland-Altman diagram and Cohen's kappa, are also developed for repeated binary measurements based upon the latent variables in GLMMs. We assess our novel model-based approach by simulation studies and a real clinical application, in which patients are evaluated repeatedly for delirium with two validated screening methods. Both the simulation studies and the real data analyses demonstrate that our proposed approach can effectively assess method agreement.

Correspondence Nan Lin, Department of Mathematics and Statistics, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130. nlin@wustl.edu.

DATA AVAILABILITY STATEMENT

The CAM and 3D-CAM data in the real data section are not publicly available. Restrictions apply to the availability of these data. Data requests should be made to Michael S. Avidan. Readers can generate the simulated dataset following the R code in the Supplementary Materials.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Keywords

Bland-Altman diagram; generalized linear mixed model; interrater reliability; method agreement; paired repeated binary measurement

1 INTRODUCTION

Method comparison studies are designed to compare two measuring methods used to measure the same quantity. One is typically a new measuring method, and the other is often an existing and widely used one. The primary goal is to assess the extent of agreement between the two measuring methods and decide if they have sufficient agreement so that they can be used interchangeably. That is, it does not matter which method is being used to take the measurement as both give practically the same value. If two measuring methods agree well enough to be used interchangeably, we may prefer the one that is cheaper, faster, less invasive, or easier to use. This is the primary motivation behind method comparison studies.¹ For example, in the diagnosis of delirium, the confusion assessment method (CAM)² is the most widely used diagnostic questionnaire battery, and the 3-minute diagnostic interview for confusion assessment method (3D-CAM)³ is a 3-minute delirium assessment based upon the CAM. The goal of the comparison is to assess the agreement between the time-consuming CAM and the time-efficient 3D-CAM, and hence determine whether they can be used interchangeably.

In a broad sense, the term “measuring method” may refer to a medical device, an instrument, a questionnaire battery, or a human judge. In this paper, we use the term “rater” specifically for human judges. In the medical context, a rater may refer to a doctor in a diagnostic process, or a clinical observer in a clinical trial. There are occasions when the response measurement does not necessarily require the rater to provide subjective assessment, eg, when reading the patient’s blood pressure from a monitor. We use the term “recorder” specifically for the human judge whose subjective assessment is not required. The “interrater reliability” is the term for agreement among raters, while the “method agreement” is the term for agreement between measuring methods. Throughout the literature, reviews of assessing agreement usually discuss method agreement and interrater reliability together in the sense that agreement indices can often be used for both method agreement and interrater reliability, though entities on which agreement is assessed are different. Detailed review of assessing agreement can be found in the literature.^{4–7} Major approaches are summarized as follows for both continuous and categorical measurements.

- Continuous measurements
 - The limits of agreement (LOA) approach introduced by Bland and Altman is a widely used technique for assessing agreement between two measuring methods or two raters of interest.^{8–11} The LOA approach is accompanied by the Bland-Altman diagram to visually show the difference between two measuring methods or two raters.
 - The intraclass correlation coefficient (ICC) is used for assessing agreement among multiple measuring methods or multiple raters based

on ANOVA-type models.^{12–16} These models mainly differ in assuming random effects or fixed effects on measuring methods or raters.

- The concordance correlation coefficient (CCC)¹⁷ was originally developed to assess agreement between two measuring methods or two raters for paired measurements without replications. The CCC modifies Pearson's correlation coefficient by additionally assessing how far the best fitting line of the data is from the 45-degree line through the origin, ie, deviation from perfect agreement. It was later extended to multiple measuring methods or multiple raters for data with replications.^{18,19} In general, the CCC reduces to the ICC under the ANOVA models used to define the ICC.⁵
- Categorical measurements
 - The agreement between two measuring methods or two raters can be measured by Cohen's kappa²⁰ for binary data or its variants, eg, weighted kappa²¹ for ordinal data. Fleiss' kappa²² can deal with multiple measuring methods or multiple raters. These kappa statistics are popular chance-corrected measures of agreement for categorical data.
 - The technique of model-based analysis is also widely used in measuring agreement for categorical data.^{23–26} By assuming that the categorical data follow a generalized linear mixed model (GLMM), approaches for continuous measurements can be applied to assessing agreement for the observed data on the scale of a latent variable.

Barnhart et al⁵ classified measures of agreement by unscaled agreement indices and scaled agreement indices. Measures of agreement are scaled agreement indices if the values have regularized magnitudes, eg, the ICC, CCC, and kappa statistics, which all range from –1 to 1, and otherwise are unscaled, eg, LOA.

The proposed methodology of assessing method agreement in this paper is motivated by a fairly common situation in psychological or cognitive assessment studies when the response measurement is affected by both the measuring method and the rater. For example, in the diagnosis of mental health disorder, neuropsychological tests are used to determine the presence of cognitive strengths and weaknesses that may be the result of a psychological disorder. There are a wealth of test batteries that combine a range of neuropsychological tests to provide an overview of cognitive skills on patients, eg, the Neurobehavioral Cognitive Status Examination (NCSE)²⁷ and the Mini-Mental State Examination (MMSE).²⁸ These tests are usually designed questionnaires or interviews administered by neuropsychologists, so test scores are finally determined by neuropsychologists. Consequently, the test score is simultaneously affected by both the test battery and the neuropsychologist. This is different from reading measurements from a medical device, eg, reading body temperature from a thermometer or reading blood pressure from a monitor. When reading body temperature from a thermometer, the rater is simply a recorder whose subjective assessment is not required and hence could rarely impact the response

measurement. However, in psychological or cognitive assessment, widespread inconsistencies are commonly observed among raters due to various experience and training background, especially in large-scale assessment studies when plenty of raters are involved in administration. A special feature of method comparison studies in such settings is that method agreement and interrater reliability shall be investigated simultaneously. To our knowledge, no literature has so far focused on assessing method agreement under the scenario where method agreement and interrater reliability are required to be studied simultaneously. For example, in a recent review of the use and psychometric properties of the NCSE,²⁹ the study of comparing the NCSE with the MMSE only discussed the sensitivity, which evaluates the ability of a method to discern small changes by a relative measure of precision, but no further discussion was carried out for assessing method agreement between these two test batteries, and neither included considerations on rater's effect.

Our approach of assessing method agreement is based on a GLMM framework. The usage of GLMMs is motivated by the need of addressing several challenges simultaneously: (1) incorporating the influence of raters into assessing method agreement; (2) categorical outcomes; (3) longitudinal data; (4) missing or unbalanced data. We will illustrate our approach using a clinical research study³⁰ in which surgical patients are evaluated for postoperative delirium using two assessment tools, ie, the CAM and the 3D-CAM. In this illustration, all the aforementioned challenges are encountered. The GLMM allows modeling of multiple sources of fixed and random effects, diverse response distributions, and covariance structures, and thus is an analysis framework that can accommodate all those complexities. Throughout this paper, we consider the situation where the rater's effect is assumed random, but the two measuring methods for comparison are assumed fixed effects. This is motivated by the intention of generalizing the method agreement results to a large population of raters because our primary goal is assessing method agreement. In the GLMM-based framework, we used hypothesis testing of equal fixed effects of two measuring methods to determine whether the two methods agree. Furthermore, our proposed methodology provides straightforward graphical technique and summary statistics to measure the extent of method agreement for paired repeated binary measurements. Although the classic Bland-Altman diagram is generalized to assessing method agreement for data with multiple continuous measurements per subject,¹¹ it is still an open question how to plot the Bland-Altman diagram for longitudinal binary measurements with correlation within the subject. Gao et al²⁴ pointed out that the kappa statistic for analyzing repeated measurements is limited because it is intended for data with a single observation made by each rater on each subject. To address these limitations, we propose novel ideas of plotting the Bland-Altman diagram and calculating Cohen's kappa for repeated binary measurements based on the latent variables in the GLMMs. We prove that the newly developed diagram maintains the independence between the vertical-axis and horizontal-axis variables as the classic Bland-Altman diagram does. In addition to assessing method agreement, we provide a way to simultaneously evaluate the interrater reliability by the ICC. In summary, our new methodology fills a gap in the current agreement literature to provide a flexible modeling approach to assess the method agreement and interrater reliability simultaneously for paired repeated binary measurements. Our methodology is versatile based on the GLMM

framework in the sense that it can be extended to various data structures in more complicated cases. Nelson et al^{25,26} also employed the GLMM framework to assess interrater reliability for ordinal ratings. Their approach incorporated the rater and patient characteristics that may impact interrater reliability, but it is not intended for assessing method agreement. Roy et al^{31,32} used hypothesis testing to assess method agreement for repeated continuous measurements based on the linear mixed model (LMM). Like the clinical application of comparing imaging measurement devices therein,³² their approach is intended for the situation with the rater simply serving as a recorder. It is different from the case in this paper where the response measurement requires the rater's subjective assessment. Gao et al²⁴ provided a way to assess agreement between two measuring methods for paired repeated binary measurements impacted by a fixed set of raters as long as neither of the two methods is a gold standard. In their model, only the subject's effect is assumed random, while both the method's effect and the rater's effect are assumed fixed, which is different from the case in this paper where we consider the rater's effect assumed random. With the rater's effect assumed fixed, their approach assessed method agreement for each rater, and thus it is ideal for the case with a small set of raters.

The rest of this paper is organized as follows. In Section 2, we introduce our approach to determine whether two measuring methods agree by hypothesis testing based upon the GLMM. Measures of method agreement and interrater reliability are further provided in this section. In Section 3, simulation studies demonstrate the performance of our approach. It is further illustrated in Section 4 using real data with simultaneous CAM and 3D-CAM assessments. The conclusion and discussion are given in Section 5.

2 METHODOLOGY

2.1 The GLMM-based framework

Consider a study of method agreement between two measuring methods with I subjects, J raters, and T_i time points for subject $i = 1, \dots, I$. The term "subject" is used here to refer to the entity on which the measurements are taken, eg, a patient in a clinical trial. Each subject is measured repeatedly over time. At each time point, a pair of measurements are taken with both methods administered by two different raters. The data are usually not balanced in the medical context because patients may not be measured at every time point due to involuntary absence or rejection to continuing participation. We label the two methods as "1" and "2," and let $(y_{ij1}, y_{ij'2})$ denote the pair of binary measurements from Methods 1 and 2 on a randomly selected subject i at time point t recorded at the same time by two different raters j and j' randomly selected from a population of raters. To model the binary measurements (either 0 or 1), we use a GLMM with a probit link function, given by

$$y_{ijtm} | \gamma_i, \alpha_{jm} \sim \text{Bernoulli}(\pi_{ijtm}),$$

and

$$\pi_{ijtm} = \Phi(\mu_{ijtm}).$$

where $m = 1, 2, i = 1, 2, \dots, I, j = 1, 2, \dots, J, t = 1, \dots, T_i$ and $\Phi(\cdot)$ is the cdf of the standard normal. The linear predictor μ_{ijtm} is given as

$$\mu_{ijtm} = \beta_m + g(x_t) + \gamma_i + \alpha_{jm}. \quad (1)$$

Terms in (1) are assumed as follows.

- β_m is the fixed effect of Method m , $m \in \{1, 2\}$.
- $g(x_t)$ is a regression function that describes the dependence on a time-dependent covariate x_t . This allows to incorporate longitudinal measurements on subjects when the mean response value changes over time.
- γ_i is the random effect of subjects, and $\gamma_i \stackrel{iid}{\sim} N(0, \sigma_\gamma^2)$.
- α_{jm} is the random effect of raters within Method m , and $\alpha_{jm} \stackrel{iid}{\sim} N(0, \sigma_{am}^2)$.

Notice the dependence of the variance components on m , which means that we allow them to vary across different methods.

The probit link function $\Phi(\cdot)$ allows assessing method agreement and interrater reliability on the latent scale. Define the latent variable $\tilde{y}_{ijtm}$ as

$$\tilde{y}_{ijtm} = \mu_{ijtm} + \tilde{\varepsilon}_{ijtm}, \quad (2)$$

where $\tilde{\varepsilon}_{ijtm} \sim N(0, 1)$ denotes the random error on the latent scale. The relationship between the latent variable $\tilde{y}_{ijtm}$ and the observed y_{ijtm} is

$$y_{ijtm} = \begin{cases} 1, & \text{if } \tilde{y}_{ijtm} > 0; \\ 0, & \text{otherwise.} \end{cases}$$

Since we consider measurements on each subject collected over time, we allow dependence in the within-subject errors of each method while letting the errors arising from different methods or from different subjects remain independent. That is, we assume that

$$\text{Corr}(\tilde{\varepsilon}_{ijtm}, \tilde{\varepsilon}_{ij't'm'}) = \begin{cases} h(|t - t'|, \mathbf{w}), & \text{if } m = m'; \\ 0, & \text{if } m \neq m', \end{cases}$$

where h is a specified correlation function, and $\mathbf{w}$ is a vector of covariance parameters. For example, if the correlation matrix is AR(1) structured, the correlation function is

$$h(|t - t'|, \rho) = \rho^{|t - t'|}. \quad (3)$$

From (1), the difference in the latent variable between Methods 1 and 2 for subject i at time point t is

$$D_{it} = \tilde{y}_{ijt1} - \tilde{y}_{ijt2} = (\beta_1 - \beta_2) + (\alpha_{j1} - \alpha_{j'2}) + (\tilde{\varepsilon}_{ijt1} - \tilde{\varepsilon}_{ijt2}).$$

with marginal mean difference

$$E(D_{it}) = \beta_1 - \beta_2.$$

Therefore, to decide whether the two methods agree, it is equivalent to testing the following hypothesis:

$$H_0: \beta_1 = \beta_2 \quad \text{vs} \quad H_1: \beta_1 \neq \beta_2. \quad (4)$$

If H_0 is rejected, we conclude that there is a significant difference between the two measuring methods. Otherwise, we fail to reject the hypothesis that the two measuring methods agree, and then measures of method agreement introduced in the next section would further assess the extent of agreement.

2.2 Measures of method agreement 2.2.1 Limits of agreement

The LOA approach is popular for assessing agreement in the medical context. Consider a set of n subjects, each measured by both measuring methods, which results in n pairs of measurements with one pair per subject. Assume that the underlying difference between the two paired measurements is D and $D \sim N(\xi, \nu^2)$. The LOA approach takes the interval $(\xi - 1.96\nu, \xi + 1.96\nu)$ covering the middle 95% of the population of D as the measures of agreement. The two bounds of this interval are estimated by $\hat{\xi} \pm 1.96\hat{\nu}$, which are called the 95% LOA. In the case of paired measurements data, the estimators $\hat{\xi}$ and $\hat{\nu}^2$ are taken as the sample mean and sample variance of the observed differences, respectively. If the two LOA fall within prespecified margins $\pm\delta$, one can conclude that the two measuring methods have sufficient agreement for interchangeable use. For example, the prespecified margins $\pm\delta$ may refer to clinically acceptable difference in the medical context. Although δ is recommended to be specified in advance, it is rarely done so in practice. Instead, agreement of methods is often evaluated by judging whether the bounds of the interval $(\hat{\xi} - 1.96\hat{\nu}, \hat{\xi} + 1.96\hat{\nu})$ are unacceptably large.¹

The Bland-Altman diagram, as part of the LOA approach, is a popular graphical tool to evaluate method agreement. It plots the difference between the two paired measurements on the vertical axis against the average of the two measurements on the horizontal axis to display the data. Furthermore, three horizontal lines are added: one for the mean difference $\hat{\xi}$ and one each for the two LOA $\hat{\xi} \pm 1.96\hat{\nu}$. When two methods agree, the points in the Bland-Altman diagram scatter around zero in a random manner, and 95% of the differences are expected to lie within the two lines corresponding to the two LOA.

Barnhart et al³³ concluded that among the existing agreement indices good for both continuous and categorical data, the coverage probability is the preferred agreement index on the basis of its consistent evaluation of data quality across multiple reviewers, populations, and continuous/categorical data. However, the classic Bland-Altman diagram is used for only continuous data. Compared to other agreement indices, the most appealing feature of the Bland-Altman diagram is that it visually shows the difference between two

measurements, which is friendly to nonstatisticians in the medical and clinical fields. Although the Bland-Altman diagram is later generalized to assessing method agreement for data with multiple continuous measurements per subject,¹¹ the implicit assumption for the model therein is that there is no time effect and correlation in the difference between paired measurements. Thus, it is still an open question how to plot the Bland-Altman diagram for longitudinal binary measurements. In other words, there may be time effect and correlation among repeated measurements within the same subject. We shall show our novel idea of plotting the Bland-Altman diagram with repeated measurements over time in our model-based setting. We aim to have one point per subject as the individual-level summary measure in the diagram, and also maintain the independence between the vertical-axis and horizontal-axis variables as the classic Bland-Altman diagram does.

One natural idea of the individual-level summary measure for each method is

$$\beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \gamma_i, \quad (5)$$

which averages fixed time effect over time points and removes rater's effect. From the Bayesian perspective, we have the prior information that for any rater $j \in \{1, \dots, J\}$,

$$\mu_{ijm} \left| \beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \gamma_i \right. \stackrel{iid}{\sim} N \left(\beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \gamma_i, \sigma_{am}^2 \right).$$

Then, the posterior distribution of (5) is given by the normal distribution

$$N \left(\bar{\mu}_{im}, \left(\frac{1}{\sigma_\gamma^2} + \frac{J}{\sigma_{am}^2} \right)^{-1} \right), \quad (6)$$

where

$$\bar{\mu}_{im} = \beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \frac{J\sigma_\gamma^2}{J\sigma_\gamma^2 + \sigma_{am}^2} \left(\gamma_i + \frac{1}{J} \sum_{j=1}^J \alpha_{jm} \right), \quad m = 1, 2. \quad (7)$$

See Appendix A for detailed derivation of the posterior distribution (6).

Note that, when J is large, the variance of the above normal distribution in (6) goes to 0. The quantity $\beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \gamma_i$ on subject i is then measured by $\bar{\mu}_{im}$ accurately. Therefore, in order to measure the difference between two measuring methods for subject i after eliminating the rater's effect, one natural idea is to use the quantity $\bar{\mu}_{i1} - \bar{\mu}_{i2}$ when the number of raters, J , is large.

Remark 1. In practice, when a dataset consists of only a few raters, we could consider another scenario with a large value of $J\sigma_\gamma^2/\sigma_{am}^2$. If the rater's variance σ_{am}^2 is relatively small, the value of $J\sigma_\gamma^2/\sigma_{am}^2$ could still be large even with only a few raters in the dataset. A

relatively small value of rater's variance usually implies good agreement among raters. For example, in the medical context, training is required to guarantee that doctors give consistent diagnoses for the same patient. In this case, in spite of only a few raters, the quantity $\bar{\mu}_{i1} - \bar{\mu}_{i2}$ is still proper to measure the difference between two measuring methods for each subject. Simulation results for this scenario are given in the Supplementary Materials.

In practice, $\bar{\mu}_{im}$ in (7) can be evaluated based on the empirical best linear unbiased predictor (EBLUP)³⁴ $\hat{\mu}_{im}$. The EBLUPs can be provided by PROC GLIMMIX in SAS 9.4.³⁵ The following Theorem 1 indicates that we can plot a Bland-Altman diagram based on the paired EBLUPs $(\hat{\mu}_{i1}, \hat{\mu}_{i2})$ for $i = 1, 2, \dots, I$, on the latent scale to visually show the difference between two methods. It is worth noting that each subject is represented by one point in the Bland-Altman diagram.

Theorem 1. *Suppose that the fixed effects of methods satisfy $\beta_1 = \beta_2$. If the number of raters $J \rightarrow \infty$, the points $((\hat{\mu}_{i1} + \hat{\mu}_{i2})/2, \hat{\mu}_{i1} - \hat{\mu}_{i2})$ in the Bland-Altman diagram based on the EBLUPs $\hat{\mu}_{i1}$ and $\hat{\mu}_{i2}$ for all patients $i = 1, 2, \dots, I$, scatter around zero in a random manner. That is, for any $i \in \{1, \dots, I\}$, as $J \rightarrow \infty$,*

- a. $\hat{\mu}_{i1} - \hat{\mu}_{i2}$ is uncorrelated with $(\hat{\mu}_{i1} + \hat{\mu}_{i2})/2$;
- b. $E(\hat{\mu}_{i1}) - E(\hat{\mu}_{i2}) \rightarrow 0$.

In practice, we can further plot paired measurements $(\Phi(\hat{\mu}_{i1}), \Phi(\hat{\mu}_{i2}))$ to obtain the Bland-Altman diagram on the probability scale. If the distribution of the values on the probability scale is severely skewed, we can do transformation and then plot the Bland-Altman diagram of the transformed data

$$(\log\{\Phi(\hat{\mu}_{i1})\}, \log\{\Phi(\hat{\mu}_{i2})\}). \quad (8)$$

Alternatively, Cohen's kappa introduced in the following section could also provide a way to measure the extent of method agreement.

2.2.2 Cohen's kappa—We have so far used the unscaled agreement index LOA to evaluate the method agreement based upon the fitted GLMM. In this section, we will show how to calculate the scaled agreement index Cohen's kappa based upon the fitted GLMM. For each $i \in \{1, 2, \dots, I\}$ and $m \in \{1, 2\}$, the predicted 0–1 binary score of the i th subject by Method m is defined as

$$\hat{y}_{im} = \begin{cases} 1, & \text{if } \hat{\mu}_{im} > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Here, $\hat{\mu}_{im}$ is the EBLUP that we used for the LOA approach in the previous section.

Let Y_1 and Y_2 denote the outcomes $\hat{y}_{im}$ by Methods 1 and 2, respectively. The outcomes can then be presented as a contingency table described in Table 1.

Here, a , b , c , and d denote the numbers of patients with four possible combinations of outcomes measured by Methods 1 and 2. For example, a denotes the number of patients whose predicted outcomes from GLMM are 0 from both measuring methods.

Cohen's kappa κ is given by

$$\kappa = \frac{p_o - p_e}{1 - p_e}, \quad (9)$$

where

$$p_o = \frac{a + d}{a + b + c + d} \quad \text{and} \quad p_e = \frac{a + b}{a + b + c + d} \cdot \frac{a + c}{a + b + c + d} + \frac{c + d}{a + b + c + d} \cdot \frac{b + d}{a + b + c + d}.$$

A confidence interval for Cohen's kappa is given based on the variance estimates discussed by Fleiss et al.³⁶ The calculation of Cohen's kappa together with its confidence interval is implemented using the R package "psych."³⁷

A limitation of the original Cohen's kappa is that its magnitude is sensitive to the underlying disease prevalence, which is described as the prevalence effect.³⁸ For example, older patients may experience an increased prevalence of some disease and are more likely to be diagnosed as "Positive." For rare diseases, subjects are more likely to be diagnosed as "Negative." In either case, Cohen's kappa would decrease due to the imbalance between a and d in Table 1. It is an overcorrection for chance agreement. If one applies Cohen's kappa directly to repeated measurements, the imbalance between a and d may significantly increase due to the correlation among repeated measurements. For example, the prevalence effect for rare diseases would become severe if each subject is measured at a number of time points, ie, large T , because the number of "Negative" a may keep increasing while the number of "Positive" d remains unchanged as more observations are collected for each subject over time. It is worth noting that our predicted binary score $\hat{y}_{im}$ is a summarized predicted binary value for each subject based upon the latent variable model, in the sense that each subject has only one binary value of $\hat{y}_{im}$ rather than repeated binary measurements over time in the observed dataset. Therefore, our approach is able to correct the prevalence effect in repeated measurements. Furthermore, our GLMM-based approach could incorporate subject's characteristics related to the underlying disease prevalence, eg, patient's age, and hence could remove these effects to avoid overcorrection for chance agreement.

2.3 Measures of interrater reliability

Our framework can also measure the interrater reliability within each measuring method. Note that the effect of raters in (1) is assumed random. As mentioned in Section 1, the ICC is appropriate to deal with raters randomly selected from a large population. Therefore, agreement among raters within the m th method ($m = 1, 2$) can be measured by the ICC based upon the latent variable model (2). The ICC within the m th method is given by

$$ICC_m = \frac{\text{Cov}(\tilde{y}_{ijtm}, \tilde{y}_{ij'tm})}{\sqrt{\text{Var}(\tilde{y}_{ijtm})\text{Var}(\tilde{y}_{ij'tm})}} = \frac{\sigma_\gamma^2 + 1}{\sigma_\gamma^2 + \sigma_{\alpha m}^2 + 1}, \quad m = 1, 2. \quad (10)$$

The estimated ICC is given by simply plugging in the estimators of variance components.

3 SIMULATION STUDY

In the simulation, we set the numbers of subjects, raters, and time points as $I = 100$, $J = 30$, and $T_i = 5$ for any $i = 1, \dots, I$. We will first demonstrate our approach in Section 3.1 based on one simulated dataset for each setup. We will also present the averages of parameter estimates based on 1000 Monte Carlo replicates. In Section 3.2, we will compare our approach with approaches that do not account for the rater's effect. The GLMM fitting is implemented by PROC GLIMMIX in SAS 9.4.³⁵

3.1 Illustrative examples

We shall consider two setups of (1) and (3), one for the case where two measuring methods agree, and the other for the case where they disagree. We assign true values of parameters in (1) and (3) as follows.

- (Model 1) The two measuring methods agree.
Set $\beta_1 = \beta_2 = 1.6$. Hence, $\beta_1 - \beta_2 = 0$.
(Model 2) The two measuring methods disagree.
Set $\beta_1 = 2.2$ and $\beta_2 = 1.6$. Hence, $\beta_1 - \beta_2 = 0.6$.
- We consider the following setup of variance components.
Set $\sigma_\gamma^2 = 0.8$, $\sigma_{\alpha 1}^2 = 0.2$, $\sigma_{\alpha 2}^2 = 0.4$, $\rho = 0.1$.
- Set $g(x_t) = -0.5x_t$ and $x_t = t$, for $t = 1, \dots, T_i$.

Section 3.1.1 gives the results of our approach for one simulated dataset for each setup, particularly to demonstrate the Bland-Altman plot. Section 3.1.2 further presents the performance of our proposed method based on 1000 simulated datasets.

3.1.1 Results for method agreement—Tables 2 and 3 present the results at significance level 0.05. It shows that the estimates of $\beta_1 - \beta_2$ are 0.0276 under Model 1; 0.5706 under Model 2, which are close to the true values 0 under Model 1 and 0.6 under Model 2, respectively. Under Model 1, we do not reject the null hypothesis $\beta_1 = \beta_2$ with p -value 0.8739. Under Model 2, the p -value for testing (4) is 0.0010, which shows a significant difference between β_1 and β_2 . All these testing results agree with the underlying truth.

The Bland-Altman diagrams on both the latent scale and the probability scale are shown in Figures 1A to 1D. As we can see in all diagrams, the points are around zero in a random manner, so there is no systematic pattern. The dashed line shows the mean difference, and the dotted lines indicate the 95% LOA. On both the latent scale and the probability scale, about 95 out of 100 points lie within the 95% LOA. Under Model 1, the mean difference is

0.04, which is close to the true difference 0. Under Model 2, the mean difference is 0.57, close to the true difference 0.6 and also nearly the same as the difference 0.5706 estimated by the GLMM shown in Table 2.

The ICCs for two measuring methods are calculated based on (10) and estimates of variance components from Table 3. Table 4 summarizes the results of the ICCs. The true ICCs for Methods 1 and 2 are given by 0.9 and 0.8182, respectively. Therefore, the simulated results in Table 4 are close to the true values of the ICCs for two measuring methods.

We also investigate the performance of our implementation of Cohen's kappa compared to naive Cohen's kappa, that is, directly applying Cohen's kappa to the observed data with correlated repeated measurements. For $\beta_1 = \beta_2$, our approach gives a kappa value 0.82 with the 95% confidence interval (0.7, 0.93). On the other hand, for $\beta_1 - \beta_2 = 0.6$, our approach gives a kappa value 0.53 with the 95% confidence interval (0.36, 0.69). However, the naive Cohen's kappas have similar values in these two examples, with values 0.3 and a 95% confidence interval (0.22, 0.38) for $\beta_1 = \beta_2$, and 0.29 with a 95% confidence interval (0.21, 0.37) for $\beta_1 - \beta_2 = 0.6$. Despite the change in the difference between β_1 and β_2 , ie, 0 vs 0.6, the two naive kappa values for observed data do not change much because we find that the counts of disagreement $b + c$ are very similar due to correlated repeated measurements over time, ie, both around 170 out of the total 500 observations. It is expected that the case with $\beta_1 = \beta_2$ should show a higher agreement between two methods. Our approach gives a high kappa value 0.82, which suggests almost perfect agreement. However, naive Cohen's kappa gives a very low value 0.3, which significantly deviates the underlying truth of method agreement.

3.1.2 Parameter estimates—Table 5 presents averaged estimates of β_1 , β_2 , and ICCs over 1000 simulated datasets. It shows that the averaged estimates of (β_1, β_2) are all very close to the true values, ie, (1.6, 1.6) for Model 1, and (2.2, 1.6) for Model 2. The averaged ICCs are all very close to the true ICCs for Methods 1 and 2, ie, 0.9 and 0.8182.

3.2 Power comparison

We compare our approach with approaches without taking the rater's effect into account, ie,

$$\mu_{ijm} = \beta_m + g(x_t) + \gamma_i. \quad (11)$$

The only difference is whether the linear predictor incorporates the term α_{jm} that accounts for the random effect of raters. We carry out a simulation study to explore how the presence of unobserved rater heterogeneity would affect the size and power of test for method agreement in (4). We fix $\beta_2 = 1.6$ and let β_1 range from 1.6 to 2.8 by 0.1. Except for $\beta_{1,2}$, the setup for other parameters is exactly the same as that in Section 3.1. We can explore the size and also how the power changes with the difference between β_1 and β_2 .

Figure 2 shows the power of testing method agreement in (4) at significance level 0.05. The solid line represents the model (11) not including the rater's effect, while the dashed line represents the model (1) including the rater's effect. The simulation results are based on 1000 replicates at each value of β_1 . The simulated sizes at significance level 0.05 are 0.056

and 0.270 for models fitted with and without rater's effect, respectively. Therefore, the model without rater's random effect cannot control the Type I error, though its power is higher than the power of the model with rater's random effect.

4 REAL DATA

The CAM² is the most widely used tool for delirium screening, but it is time consuming for clinical staff to administer on a routine basis. The 3D-CAM³ is a 3-minute delirium assessment based upon the CAM algorithm. The appeal of the 3D-CAM is that it takes less time than the CAM to screen for delirium. Therefore, it is meaningful to determine whether the agreement is sufficient between the two assessment tools so that CAM and 3D-CAM could be used interchangeably. In this section, we shall apply our approach of assessing method agreement to a real dataset for CAM and 3D-CAM comparison.

Our data were collected from patients who underwent major elective surgery and were enrolled in the Electroencephalography Guidance of Anesthesia to Alleviate Geriatric Syndromes (ENGAGES)³⁰ trial, which looked at the effectiveness of electroencephalogram guidance of anesthesia at preventing postoperative delirium at Barnes-Jewish Hospital in St. Louis, Missouri, USA. Patients enrolled in ENGAGES were 60 or older with at least a two-day hospital stay. We interviewed patients on postoperative days 0–5 using both the CAM and 3D-CAM at the same time but scored them independently in order to remain blind to the outcome of the other assessment. The dataset contains 42 pairs of readings (either “Positive” or “Negative”) from 20 patients with 6 raters at 6 time points. The dataset is unbalanced since some patients were interviewed at only a couple of the time points.

Table 6 shows the result of testing the difference. Since the p -value is 0.523, the testing result indicates that the two methods may be used interchangeably.

The ICCs for the CAM and 3D-CAM are 0.99 and 0.85, respectively, which indicates that raters have an excellent degree of agreement for both measuring methods. Next, we use the Bland-Altman diagram to evaluate the degree of agreement. Figure 3A shows the Bland-Altman diagram on the latent scale based upon 20 pairs of readings (one per patient), and Figure 3B is drawn on the log-transformed probability scale (8).

As we can see from Figure 3, the mean difference on the latent scale is -0.35 . It might be hard for nonstatisticians to interpret and set a prespecified margin on the latent scale. Since the difference on the log-transformed scale $\log\{\Phi(\hat{\mu}_{i1})\} - \log\{\Phi(\hat{\mu}_{i2})\} = \log\{\Phi(\hat{\mu}_{i1})/\Phi(\hat{\mu}_{i2})\}$, the value obtained by taking the exponential of the mean difference would explain the ratio of the probability of scoring a “Positive” between two methods, ie, $\Phi(\hat{\mu}_{i1})/\Phi(\hat{\mu}_{i2})$. The mean difference in Figure 3B is -1.15 , which indicates the probability of “Positive” measured by the 3D-CAM is on average $1/\exp\{-1.15\} = 3.16$ times the probability by the CAM. Although this point estimate on the log-transformed scale may show that it is more likely to score “Positive” by the 3D-CAM than by the CAM, the agreement of methods is evaluated by judging whether the bounds of the 95% confidence interval are unacceptably large on the log-transformed scale. The model-based Cohen's kappa calculated by (9) is exactly 1, which indicates perfect agreement between the CAM and 3D-CAM based on this dataset. In other

words, the CAM and 3D-CAM are predicted to give the same values over these 20 patients, which supports the agreement between CAM and 3D-CAM. Our analysis result is consistent with earlier studies in the literature.^{39,40}

5 CONCLUSION AND DISCUSSION

In this paper, we propose an approach for comparing two measuring methods with paired repeated binary data over time. Our GLMM-based framework incorporates both assessing method agreement and evaluating interrater reliability. By treating methods as fixed effects, assessing method agreement is equivalent to testing the equality of fixed effects of methods. Both simulation studies and applications to real data demonstrate the ability of our approach to make correct decisions on method agreement.

Provided that users are not satisfied with simply a decision on whether the two methods agree or not, we further illustrate a novel way to implement the Bland-Altman diagram and Cohen's kappa on the latent variables based upon the GLMMs. Traditional scaled or unscaled agreement indices may provide misleading conclusions with repeated data over time because of correlations among repeated measurements on the same subject, but our approach correctly measures the method agreement by accommodating the dependency in the GLMM.

Our methodology is versatile based on the GLMM framework in the sense that it can be extended to various data structures in more complicated cases. In general, the measurements could be continuous, binary, or ordinal, while this paper focuses on repeated binary measurements. The theory of our approach requires the assumption $J \rightarrow \infty$, which is a common condition for approaches based upon GLMMs to guarantee asymptotically consistent estimates of fixed effects and variance components for random effects. Although our approach theoretically intends to deal with large-scale studies in which there are large numbers of raters and subjects, the additional simulation study in the Supplementary Materials shows that our approach still performs well with only a few raters in the dataset as long as there is reasonably good agreement among raters. Throughout this paper, we consider the rater's effect as random by the intention of generalizing the method agreement results to a large population of raters. On the other hand, if the rater's effect is assumed fixed, there is a straightforward extension of the hypothesis testing procedure based on the GLMM. Meanwhile, the newly developed Bland-Altman diagram and Cohen's kappa could be operated for each rater from a fixed set of raters.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGEMENTS

The ENGAGES study was funded by a National Institutes of Health grant to support pragmatic trials (1 UH2 HL125141 and 5 UH3 AG050312). This study was also funded by the National Institutes of Health NIDUS (NIA R24AG054259) and the Dr. Seymour and Rose T. Brown Endowed Chair at Washington University in St. Louis.

Funding information

APPENDIX A

THE POSTERIOR DISTRIBUTION OF INDIVIDUAL-LEVEL SUMMARY MEASURE

Let τ_{ijm} be the average of μ_{ijtm} over time adjusted for the fixed time effect $g(x_t)$,

$$\tau_{ijm} = \frac{1}{T_i} \sum_{t=1}^{T_i} \mu_{ijtm} - \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t).$$

Then, τ_{ijm} is modeled by the following mixed-effect model:

$$\tau_{ijm} = \beta_m + \gamma_i + \alpha_{jm}.$$

Notice that, for any rater $j \in \{1, \dots, J\}$, we have

$$\tau_{ijm} | \gamma_i \stackrel{iid}{\sim} N(\beta_m + \gamma_i, \sigma_{am}^2),$$

and

$$\beta_m + \gamma_i \sim N(\beta_m, \sigma_\gamma^2).$$

Let $\bar{\tau}_{im} = \sum_{j=1}^J \tau_{ijm} / J$. Then,

$$\beta_m + \gamma_i | \bar{\tau}_{im} \sim N\left(\frac{\sigma_{am}^2}{J\sigma_\gamma^2 + \sigma_{am}^2} \beta_m + \frac{J\sigma_\gamma^2}{J\sigma_\gamma^2 + \sigma_{am}^2} \bar{\tau}_{im}, \left(\frac{1}{\sigma_\gamma^2} + \frac{J}{\sigma_{am}^2}\right)^{-1}\right),$$

so

$$\beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \gamma_i | \bar{\tau}_{im} \sim N\left(\bar{\mu}_{im}, \left(\frac{1}{\sigma_\gamma^2} + \frac{J}{\sigma_{am}^2}\right)^{-1}\right),$$

where

$$\bar{\mu}_{im} = \beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \frac{J\sigma_\gamma^2}{J\sigma_\gamma^2 + \sigma_{am}^2} \left(\gamma_i + \frac{1}{J} \sum_{j=1}^J \alpha_{jm} \right), \quad m = 1, 2.$$

APPENDIX B

PROOF OF THEOREM 1

Proof. From (7), for $m \in \{1, 2\}$, the variance of $\bar{\mu}_{im}$ is

$$\text{Var}(\bar{\mu}_{im}) = \frac{J\sigma_\gamma^4}{J\sigma_\gamma^2 + \sigma_{am}^2} \rightarrow \sigma_\gamma^2, \text{ as } J \rightarrow \infty.$$

Then, by checking

$$\text{Cov}(\bar{\mu}_{i1} - \bar{\mu}_{i2}, \bar{\mu}_{i1} + \bar{\mu}_{i2}) = \text{Var}(\bar{\mu}_{i1}) - \text{Var}(\bar{\mu}_{i2}) \rightarrow 0,$$

the difference $\bar{\mu}_{i1} - \bar{\mu}_{i2}$ is uncorrelated with the average $(\bar{\mu}_{i1} + \bar{\mu}_{i2})/2$ as the number of raters $J \rightarrow \infty$.

For $m \in \{1, 2\}$

$$\bar{\mu}_{im} \xrightarrow{a.s.} \beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t) + \gamma_i, \text{ as } J \rightarrow \infty.$$

Therefore, $(\bar{\mu}_{i1} - \bar{\mu}_{i2}) - (\beta_1 - \beta_2) \xrightarrow{a.s.} 0$. If $\beta_1 = \beta_2$, then $\bar{\mu}_{i1} - \bar{\mu}_{i2} \xrightarrow{a.s.} 0$. Similarly, $(\hat{\mu}_{i1} - \hat{\mu}_{i2}) - (\hat{\beta}_1 - \hat{\beta}_2) \xrightarrow{a.s.} 0$, where $\hat{\beta}_1, \hat{\beta}_2$ are EBLUEs of β_1, β_2 . Since $\hat{\beta}_m - \beta_m \xrightarrow{P} 0$, we have $\hat{\mu}_{i1} - \hat{\mu}_{i2} \xrightarrow{P} 0$ if $\beta_1 = \beta_2$. By the normality assumption and assuming that $\beta_1 = \beta_2$,

$$E(\hat{\mu}_{i1} - \bar{\mu}_{i1}) - E(\hat{\mu}_{i2} - \bar{\mu}_{i2}) = E(\hat{\mu}_{i1} - \hat{\mu}_{i2}) - E(\bar{\mu}_{i1} - \bar{\mu}_{i2}) \rightarrow 0, \text{ as } J \rightarrow \infty.$$

By theorem 13.2 in the work of Jiang about the mean squared prediction error (MSPE) of the EBLUP, the difference between the MSPEs

$$E(\hat{\mu}_{i1} - \bar{\mu}_{i1})^2 - E(\hat{\mu}_{i2} - \bar{\mu}_{i2})^2 \rightarrow 0, \text{ as } J \rightarrow \infty.$$

Therefore,

$$\text{Var}(\hat{\mu}_{i1} - \bar{\mu}_{i1}) - \text{Var}(\hat{\mu}_{i2} - \bar{\mu}_{i2}) \rightarrow 0, \text{ as } J \rightarrow \infty.$$

Then, the difference $(\hat{\mu}_{i1} - \bar{\mu}_{i1}) - (\hat{\mu}_{i2} - \bar{\mu}_{i2})$ is also uncorrelated with $(\hat{\mu}_{i1} - \bar{\mu}_{i1})/2 + (\hat{\mu}_{i2} - \bar{\mu}_{i2})/2$ as the number of raters $J \rightarrow \infty$.

Since $\bar{\mu}_{i1} - \bar{\mu}_{i2}$ is uncorrelated with $\bar{\mu}_{i1} + \bar{\mu}_{i2}$, and $\bar{\mu}_{i1} - \bar{\mu}_{i2} \xrightarrow{a.s.} 0$, $\hat{\mu}_{i1} - \hat{\mu}_{i2} \xrightarrow{P} 0$ as $J \rightarrow \infty$, by the fact that

$$\begin{aligned}
& \text{Cov}((\hat{\mu}_{i1} - \bar{\mu}_{i1}) - (\hat{\mu}_{i2} - \bar{\mu}_{i2}), (\hat{\mu}_{i1} - \bar{\mu}_{i1}) + (\hat{\mu}_{i2} - \bar{\mu}_{i2})) - \text{Cov}(\hat{\mu}_{i1} - \hat{\mu}_{i2}, \hat{\mu}_{i1} + \hat{\mu}_{i2}) \\
&= \text{Cov}((\hat{\mu}_{i1} - \hat{\mu}_{i2}) - (\bar{\mu}_{i1} - \bar{\mu}_{i2}), (\hat{\mu}_{i1} + \hat{\mu}_{i2}) - (\bar{\mu}_{i1} + \bar{\mu}_{i2})) - \text{Cov}(\hat{\mu}_{i1} - \hat{\mu}_{i2}, \hat{\mu}_{i1} + \hat{\mu}_{i2}) \\
&= \text{Cov}(\hat{\mu}_{i1} - \hat{\mu}_{i2}, \hat{\mu}_{i1} + \hat{\mu}_{i2}) - 2\text{Cov}(\hat{\mu}_{i1}, \bar{\mu}_{i1}) + 2\text{Cov}(\hat{\mu}_{i2}, \bar{\mu}_{i2}) - \text{Cov}(\hat{\mu}_{i1} - \hat{\mu}_{i2}, \hat{\mu}_{i1} + \hat{\mu}_{i2}) \\
&\rightarrow 0, \text{ as } J \rightarrow \infty,
\end{aligned}$$

the difference between the EBLUPs, ie, $\hat{\mu}_{i1} - \hat{\mu}_{i2}$, is also uncorrelated with $(\hat{\mu}_{i1} + \hat{\mu}_{i2})/2$, as $J \rightarrow \infty$.

Notice that, for $m = 1, 2$, the mean of $\bar{\mu}_{im}$ is

$$E(\bar{\mu}_{im}) = \beta_m + \frac{1}{T_i} \sum_{t=1}^{T_i} g(x_t).$$

If $\beta_1 = \beta_2$, then $E(\bar{\mu}_{i1}) = E(\bar{\mu}_{i2})$. As $E(\hat{\mu}_{i1} - \hat{\mu}_{i2}) - E(\bar{\mu}_{i1} - \bar{\mu}_{i2}) \rightarrow 0$, $E(\hat{\mu}_{i1}) - E(\hat{\mu}_{i2}) \rightarrow 0$ if $\beta_1 = \beta_2$. $\square$

APPENDIX C

SAS CODE

```

/*For confidentiality, we cannot show the real data.*/
/*Instead, we provide the following pseudodata for user convenience.*/
data a;
  input id $ time $ cam $ dcam $ rater_cam $ rater_dcam $; datalines;
  1 1 Negative Negative A B
  1 2 Negative Postive C D
  ... more data lines...
;

data b;set a;
y=cam; method='CAM';rater=rater_cam;id=id;time=time;timel=time;output;
y=dcam;method='3DCAM';rater=rater_dcam;id=id;time=time;timel=time;output;
keep y method rater id time timel;

proc glimmix data=b;
  class id time method rater;
  model y(event='Positive')= method timel/dist=binary link=probit ddfm=kr;
  random id;
  random rater/group=method;
  /*The correlation matrix of within-subject errors is AR(1) structured.*/
  random time/subject=id(method) type=ar(1) rside;
  /*Output linear predictor (p) and marginal linear predictor (np).*/
  output out=agreeout pred=p pred(noblup)=np;

```

```
/*Test if the coefficients of method's effects are the same.*/
lsmeans method/diff ilink cl;
run;
```

REFERENCES

1. Choudhary PK, Nagaraja HN. Measuring Agreement: Models, Methods, and Applications. New York, NY: John Wiley & Sons; 2017.
2. Inouye S, van Dyck CH, Alessi CA, Balkin S, Siegel AP, Horwitz RI. Clarifying confusion: the confusion assessment method: a new method for detection of delirium. *Ann Intern Med*. 1990;113(12):941–948. [PubMed: 2240918]
3. Marcantonio ER, Ngo LH, O'Connor M, et al. 3D-CAM: derivation and validation of a 3-minute diagnostic interview for CAM-defined delirium: a cross-sectional diagnostic test study. *Ann Intern Med*. 2014;161(8):554–561. [PubMed: 25329203]
4. Balakrishnan N Methods and Applications of Statistics in the Life and Health Sciences. New York, NY: John Wiley & Sons; 2010.
5. Barnhart HX, Haber MJ, Lin LI. An overview on assessing agreement with continuous measurements. *JBioPharmStat*. 2007;17(4):529–569.
6. Carstensen B Comparing Clinical Measurement Methods: A Practical Guide. New York, NY: John Wiley & Sons; 2011.
7. Lin L, Hedayat A, Wu W. Statistical Tools for Measuring Agreement. New York, NY: Springer Science & Business Media; 2012.
8. Altman DG, Bland JM. Measurement in medicine: the analysis of method comparison studies. *J R Stat Soc D Stat*. 1983;32(3):307–317.
9. Bland JM, Altman D. Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*. 1986;327(8476):307–310.
10. Bland JM, Altman DG. Measuring agreement in method comparison studies. *Stat Methods Med Res*. 1999;8(2):135–160. [PubMed: 10501650]
11. Bland JM, Altman DG. Agreement between methods of measurement with multiple observations per individual. *J Biopharm Stat*. 2007;17(4):571–582. [PubMed: 17613642]
12. Bartko JJ. On various intraclass correlation reliability coefficients. *Psychological Bulletin*. 1976;83(5):762–765.
13. Eliasziw M, Young SL, Woodbury MG, Fryday-Field K. Statistical methodology for the concurrent assessment of interrater and intrarater reliability: using goniometric measurements as an example. *Physical Therapy*. 1994;74(8):777–788. [PubMed: 8047565]
14. McGraw KO, Wong SP. Forming inferences about some intraclass correlation coefficients. *Psychological Methods*. 1996;1(1):30–46.
15. Müller R, Büttner P. A critical discussion of intraclass correlation coefficients. *Statist Med*. 1994;13(23–24):2465–2476.
16. Shrout PE, Fleiss JL. Intraclass correlations: uses in assessing rater reliability. *Psychological Bulletin*. 1979;86(2):420–428. [PubMed: 18839484]
17. Lawrence I, Lin K. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*. 1989;45(1):255–268. [PubMed: 2720055]
18. Barnhart HX, Song J, Haber MJ. Assessing intra, inter and total agreement with replicated readings. *Statist Med*. 2005;24(9):1371–1384.
19. Lin L, Hedayat A, Wu W. A unified approach for assessing agreement for continuous and categorical data. *J Biopharm Stat*. 2007;17(4):629–652. [PubMed: 17613645]
20. Cohen J A coefficient of agreement for nominal scales. *Educ Psychol Meas*. 1960;20(1):37–46.
21. Cohen J Weighted kappa: nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin*. 1968;70(4):213–220. [PubMed: 19673146]
22. Fleiss JL. Measuring nominal scale agreement among many raters. *Psychological Bulletin*. 1971;76(5):378–382.

23. Carrasco JL, Jover L. Concordance correlation coefficient applied to discrete data. *Statist Med.* 2005;24(24):4021–4034.
24. Gao J, Pan Y, Haber M. Assessment of observer agreement for matched repeated binary measurements. *Comput Stat Data Anal.* 2012;56(5):1052–1060.
25. Nelson KP, Edwards D. Measures of agreement between many raters for ordinal classifications. *Statist Med.* 2015;34(23):3116–3132.
26. Nelson KP, Mitani AA, Edwards D. Assessing the influence of rater and subject characteristics on measures of agreement for ordinal ratings. *Statist Med.* 2017;36(20):3181–3199.
27. Kiernan RJ, Mueller J, Langston JW, Van Dyke C. The neurobehavioral cognitive status examination: a brief but differentiated approach to cognitive assessment. *Ann Intern Med.* 1987;107(4):481–485. [PubMed: 3631786]
28. Tombaugh TN, McIntyre NJ. The mini-mental state examination: a comprehensive review. *J Am Geriatr Soc.* 1992;40(9):922–935. [PubMed: 1512391]
29. Shea T, Kane C, Mickens M. A review of the use and psychometric properties of the cognistat/neurobehavioral cognitive status examination in adults post-cerebrovascular accident. *Rehabilitation Psychology.* 2017;62(2):221–222. [PubMed: 28569535]
30. Wildes T, Winter A, Maybrier H, et al. Protocol for the electroencephalography guidance of anesthesia to alleviate geriatric syndromes (ENGAGES) study: a pragmatic, randomised clinical trial. *BMJ Open.* 2016;6(6):e011505.
31. Roy A An application of linear mixed effects model to assess the agreement between two methods with replicated observations. *JBioPharm Stat.* 2009;19(1):150–173. [PubMed: 19127473]
32. Roy A, Fuller CD, Rosenthal DI, Thomas CR Jr. Comparison of measurement methods with a mixed effects procedure accounting for replicated evaluations (COM 3 PARE): method comparison algorithm implementation for head and neck IGRT positional verification. *BMC Med Imaging.* 2015;15(1):35–45. [PubMed: 26310853]
33. Barnhart HX, Yow E, Crowley AL, et al. Choice of agreement indices for assessing and improving measurement reproducibility in a core laboratory setting. *Stat Methods Med Res.* 2016;25(6):2939–2958. [PubMed: 24831133]
34. Jiang J *Asymptotic Analysis of Mixed Effects Models: Theory, Applications, and Open Problems.* Boca Raton, FL: CRC Press; 2017.
35. SAS Institute Inc *SAS/STAT 9.4 user's guide.* Cary, NC: SAS Institute Inc; 2017.
36. Fleiss JL, Cohen J, Everitt BS. Large sample standard errors of kappa and weighted kappa. *Psychological Bulletin.* 1969;72(5):323–327.
37. Revelle W *psych: procedures for psychological, psychometric, and personality research R package version 1.8.10.* Evanston, IL: Northwestern University; 2018 <https://cran.r-project.org/web/packages/psych/citation.html>
38. Oleckno WA. *Epidemiology: Concepts and Methods.* Long Grove, IL: Waveland Press; 2008.
39. Kuczmarska A, Ngo LH, Guess J, et al. Detection of delirium in hospitalized older general medicine patients: a comparison of the 3D-CAM and CAM-ICU. *J Gen Intern Med.* 2016;31(3):297–303. [PubMed: 26443577]
40. Vasunilashorn SM, Guess J, Ngo L, et al. Derivation and validation of a severity scoring method for the 3-minute diagnostic interview for confusion assessment method-defined delirium. *J Am Geriatr Soc.* 2016;64(8):1684–1689. [PubMed: 27374833]
41. Jiang J *Large Sample Techniques for Statistics.* New York, NY: Springer Science & Business Media; 2010.

FIGURE 1.

The Bland-Altman diagrams. A, Model 1 on the latent scale; B, Model 1 on the probability scale; C, Model 2 on the latent scale; D, Model 2 on the probability scale

FIGURE 2. Power comparison. The light gray dotted line indicates the significance level at 0.05

FIGURE 3.
The Bland-Altman diagrams for real data. A, The latent scale; B, The log-transformed probability scale

TABLE 1

Contingency table of outcomes

	$Y_1=0$	$Y_1=1$
$Y_2=0$	a	b
$Y_2=1$	c	d

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 2Estimation for the difference between methods, $\beta_1 - \beta_2$

Model	Estimate	SE	<i>p</i> -value	95% CI
Model 1	0.0276	0.1731	0.8739	(−0.3211, 0.3764)
Model 2	0.5706	0.1627	0.0010	(0.2437, 0.8975)

Abbreviations: CI, confidence interval; SE, standard error.

TABLE 3

Estimation for variance components

Model	Parameter	Estimate	SE
Model 1	σ_{γ}^2	0.7578	0.1511
	$\sigma_{\alpha 1}^2$	0.1804	0.0770
	$\sigma_{\alpha 2}^2$	0.4886	0.1707
	ρ	0.0626	0.0410
	σ_{γ}^2	0.7179	0.1424
Model 2	$\sigma_{\alpha 1}^2$	0.1842	0.0835
	$\sigma_{\alpha 2}^2$	0.3756	0.1316
	ρ	0.0686	0.0435

Abbreviation: SE, standard error.

TABLE 4

The intraclass correlation coefficients (ICCs) for interrater reliability

Model	Method 1	Method 2
Model 1	0.9069	0.7825
Model 2	0.9032	0.8206

TABLE 5

Estimation for β_1 , β_2 , and intraclass correlation coefficients (ICCs). The results are based on 1000 replicates

Scenario	Parameter	Estimate	SE
Model 1	β_1	1.5659	0.1968
	β_2	1.5697	0.2087
	ICC ₁	0.8888	0.0427
	ICC ₂	0.8857	0.0440
Model 2	β_1	2.1657	0.2189
	β_2	1.5943	0.2187
	ICC ₁	0.8161	0.0565
	ICC ₂	0.8102	0.0568

Abbreviation: SE, standard error.

TABLE 6

Comparison of the confusion assessment method (CAM) and 3-minute diagnostic interview for confusion assessment method (3D-CAM)

Difference	Estimate	<i>p</i> -value	95% CI
$\beta_{\text{CAM}} - \beta_{\text{3D-CAM}}$	-0.6446	0.5230	(-3.7335, 2.4444)

Abbreviation: CI, confidence interval.