

RESEARCH

Open Access



Analytical method for detecting outlier evaluators

Yujie Wu¹, Sharon Curhan^{2,3}, Bernard Rosner^{1,2}, Gary Curhan^{2,3,4,5} and Molin Wang^{1,2,4*}

Abstract

Background Epidemiologic and medical studies often rely on evaluators to obtain measurements of exposures or outcomes for study participants, and valid estimates of associations depends on the quality of data. Even though statistical methods have been proposed to adjust for measurement errors, they often rely on unverifiable assumptions and could lead to biased estimates if those assumptions are violated. Therefore, methods for detecting potential ‘outlier’ evaluators are needed to improve data quality during data collection stage.

Methods In this paper, we propose a two-stage algorithm to detect ‘outlier’ evaluators whose evaluation results tend to be higher or lower than their counterparts. In the first stage, evaluators’ effects are obtained by fitting a regression model. In the second stage, hypothesis tests are performed to detect ‘outlier’ evaluators, where we consider both the power of each hypothesis test and the false discovery rate (FDR) among all tests. We conduct an extensive simulation study to evaluate the proposed method, and illustrate the method by detecting potential ‘outlier’ audiologists in the data collection stage for the Audiology Assessment Arm of the Conservation of Hearing Study, an epidemiologic study for examining risk factors of hearing loss in the Nurses’ Health Study II.

Results Our simulation study shows that our method not only can detect true ‘outlier’ evaluators, but also is less likely to falsely reject true ‘normal’ evaluators.

Conclusions Our two-stage ‘outlier’ detection algorithm is a flexible approach that can effectively detect ‘outlier’ evaluators, and thus data quality can be improved during data collection stage.

Keywords Evaluator, False discovery rate, Outlier detection, Quality control, Reviewer

Introduction

Many medical and epidemiological studies that investigate relationships between risk factors and disease outcomes rely on multiple evaluators (e.g. clinicians, technicians) to measure the exposures or outcomes of

interest among study participants. For example, in large epidemiologic studies of hearing loss, pure-tone audiometry measurements are typically obtained by multiple audiologists or trained technicians in sound-treated booths [1–3]. Similarly, in large studies of vision, vision tests are often conducted by multiple evaluators in a clinic setting [4, 5]. Further, potential issues related to the collection of data by multiple evaluators may also extend to studies that rely on data collected by non-human testing methods, such as automated audiometers [6], to obtain test measurements. Obtaining precise estimates of the association between risk factors and disease outcomes not only depends on the statistical methods used, but also the quality of data itself. Although many analytical methods have been proposed

*Correspondence:

Molin Wang

stmow@channing.harvard.edu

¹ Department of Biostatistics, Harvard University, Boston, USA

² Channing Division of Network Medicine, Brigham and Women’s Hospital, Boston, USA

³ Harvard Medical School, Boston, USA

⁴ Department of Epidemiology, Harvard University, Boston, USA

⁵ Renal Division, Department of Medicine, Brigham and Women’s Hospital, Boston, USA



© The Author(s) 2023. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

to adjust for measurement errors arose from data collected with poor quality, those methods typically rely on unverifiable assumptions [7], and pays a cost of the precision of estimates. Therefore, collecting data with better quality is preferred over using statistical methods to adjust for the biases induced by data of worse quality during statistical analysis stage. In this paper, we propose methods for quality control during data collection stage so that problems with the measurements of exposures or outcomes can be discovered and addressed promptly.

Our work is motivated by the Conservation of Hearing Study (CHEARS), an investigation of risk factors for hearing loss among participants in the Nurses' Health Studies II (NHS II), an ongoing cohort study consisting of 116,430 registered female nurses in the US, aged 25-42 years at enrollment in 1989 [8]. The CHEARS Audiology Assessment Arm (AAA) assessed the longitudinal change in the pure-tone air and bone conduction audiometric hearing thresholds (the sound intensity of a pure tone at which it is first perceived) measured in decibels in hearing level, or dB HL, across the full range of conventional frequencies (0.5-8 kHz) [9]. Baseline testing was conducted on 3,749 women whose self-reported hearing status was either 'excellent', 'very good' or had 'a little hearing trouble', and resided within proximity of one of 19 CHEARS testing sites across the US [9]. The 3-year follow-up testing was completed on 3,136 participants (84%). In order to obtain reliable hearing measurements, detecting potential 'outlier' audiologists who tend to have higher or lower hearing test measurements than other audiologists is critical. Once an 'outlier' audiologist is identified, devices used by this audiologist can be examined and an early intervention can be carried out during the data collection stage if necessary. Moreover, this outlier information may have important implications for the approach of data analysis.

To the best of our knowledge, there are no existing statistical methods for detecting 'outlier' evaluators. In this paper, we develop an innovative two-stage algorithm for detecting 'outlier' evaluators. In the first stage, rather than directly evaluating the observed measurements, we extract evaluators' effects on the measurements through regression analysis where the influences of other variables can be accounted for. In the second stage, we perform hypothesis tests to detect 'outlier' evaluators based on the estimated coefficients and variances from the first-stage regression analysis.

The paper is organized as follows. In Section 'Methods', we present the two-stage algorithm to detect 'outlier' evaluators for scenarios when each study participant has either single or multiple measurements. In Section

'Simulation', we perform a simulation study to investigate the performance of our two-stage algorithm. Section 'Application' presents a real data analysis to detect 'outlier' audiologists in the CHEARS AAA. The section 'Discussion' concludes the paper.

Methods

First stage regression

We first consider the scenario when each study participant only has one measurement to be obtained by an evaluator. Throughout the paper, we assumed that the exposure or test outcome of each study participant will be measured by only one evaluator, but one evaluator can measure multiple study participants. Let $i \in \{1, 2, \dots, N\}$ index the study participants; $j \in \{1, 2, \dots, M\}$ index the evaluators who measure the exposure or test outcome. Let n_j denote the number of study participants who are evaluated by the j -th evaluator, such that $\sum_{j=1}^M n_j = N$.

To estimate the effects of evaluators on the measurements, in the first stage, we fit the following linear regression:

$$E(Y_i|X_i, T_i^{(1)}, \dots, T_i^{(M)}) = \sum_{j=1}^M \beta_j T_i^{(j)} + \boldsymbol{\gamma}^T X_i, \quad (1)$$

where Y_i is the measurement for the i -th study participant, $T_i^{(j)}$ is an evaluator indicator which is 1 if the i -th study participant's exposure or outcome is evaluated by the j -th evaluator, and 0 otherwise, X_i is a p -dimensional vector containing potential confounders for the evaluator- Y_i relationship and predictors of Y_i , and $\boldsymbol{\gamma}^T$ is the transpose of the p -dimensional coefficient vector $\boldsymbol{\gamma}$. We use T to denote the transpose of a vector or matrix throughout the paper. Without further specification, all vectors are column vectors throughout this paper. Note that the first stage regression can go beyond linearity, where some nonlinear forms of X_i can be included for more accurate account of the effects of the covariates on the measurement. The regression coefficient β_j represents the mean effect of evaluator j on the measurement after adjusting for X , and in the absence of 'outlier' evaluators, $\beta_j, j = 1, \dots, M$, should be similar across different evaluators.

In practice, there may be multiple measurements for all or part of study participants. Let $k \in \{1, 2, \dots, t_i\}$ index the measurements for the i -th study participant. For example, in the CHEARS AAA, study participants have both ears tested by audiologists, and therefore we have $t_i = 2$ for each participant at each frequency.

In the CHEARS AAA, the Pearson correlation coefficients between the hearing test outcomes of the left and right ear are over 0.7 regardless of frequencies. To take into account the correlation between multiple

measurements while in the meantime being able to estimate the mean effect of evaluators on the measurements after controlling for potential confounders, we propose to apply the Generalized Estimating Equations (GEE) method in the first-stage regression analysis to estimate the effects of evaluators [10, 11]. The model for the multiple correlated measurements can be written as:

$$E[Y_{i,k}|X_i, Z_{i,k}, T_i^{(1)}, \dots, T_i^{(M)}] = \sum_{j=1}^M \beta_j T_i^{(j)} + \gamma^T X_i + \eta^T Z_{i,k}, \tag{2}$$

where $Y_i = [Y_{i,1}, Y_{i,2}, \dots, Y_{i,t_i}]^T$, $\text{Cov}(Y_i) = \Sigma_i$, with Σ_i being the unknown $t_i \times t_i$ variance-covariance matrix of the measurements of the i -th study participant, and $Z_{i,k}$ contains information that is specific to the k -th measurement of the i -th study participant.

The parameters $\theta = [\gamma^T, \beta^T, \eta^T]^T$, with $\beta = [\beta_1, \dots, \beta_M]^T$, can be estimated by solving the following estimating equation [10, 11]:

$$\sum_{i=1}^M D_i^T(\theta) V_i^{-1}(\theta, \alpha)(Y_i - \mu_i(\theta)) = \mathbf{0}, \tag{3}$$

where $\mu_i = E[Y_i|X_i, Z_i, T_i^{(1)}, \dots, T_i^{(M)}]$, $D_i = \frac{\partial}{\partial \theta} \mu_i(\theta)$, $V_i(\theta, \alpha)$ is the working variance-covariance matrix, and α contains parameters characterizing the correlation structure between multiple measurements. Some common working correlation structures for $k_1 \neq k_2 \in \{1, \dots, t_i\}$ are independent, defined as $\text{Corr}(Y_{i,k_1}, Y_{i,k_2}) = 0$; exchangeable, defined as $\text{Corr}(Y_{i,k_1}, Y_{i,k_2}) = \alpha$, and unstructured, defined as $\text{Corr}(Y_{i,k_1}, Y_{i,k_2}) = \alpha_{k_1, k_2}$. The variance of $\hat{\theta}$, $\text{Var}(\hat{\theta})$, can be estimated based on the sandwich variance estimator [10, 11].

The coefficients β_1, \dots, β_M reflect evaluators' effects on the measurements. An 'outlier' evaluator will have a different coefficient than the remaining 'normal' ones. Thus, in the second stage, we perform hypothesis tests to detect 'outlier' evaluators based on estimated $\hat{\beta}$ and $\widehat{\text{Var}}(\hat{\beta})$.

Hypothesis testing

In the second stage, we detect 'outlier' evaluators who give different measurements than their counterparts after adjusting for true predictors and confounders of the outcome. We now formally define 'outlier' evaluators as those evaluators whose effects on the measurements are different from the averaged effect among all the evaluators in the study. Recall that $\beta_j, j = 1, \dots, M$ represents the effect of the j -th evaluator on the measurements after controlling for study participants' characteristics. 'Outlier' evaluators can be detected through testing whether evaluator effects on the measurements are statistically different from the mean effect averaged across all

evaluators. Therefore, for a given evaluator j , the hypothesis can be formulated as:

$$H_{0j} : \beta_j - \frac{1}{M} \sum_{q=1}^M \beta_q = 0, \quad j = 1, 2, \dots, M. \quad \text{v.s.} \quad H_{1j} : \beta_j - \frac{1}{M} \sum_{q=1}^M \beta_q \neq 0 \tag{4}$$

which can be written as $H_{0j} : L_j^T \beta = 0$ v.s. $H_{1j} : L_j^T \beta \neq 0$, with

$$L_j = \begin{bmatrix} -\frac{1}{M} & -\frac{1}{M} & \dots & \underbrace{\frac{M-1}{M}}_{j\text{-th location}} & -\frac{1}{M} & \dots & -\frac{1}{M} \end{bmatrix}^T. \tag{5}$$

Note that, $\beta_j - \frac{1}{M} \sum_{q=1}^M \beta_q$ can be interpreted as the difference between the mean measurement of the j -th evaluator and the average mean measurements over all evaluators adjusting for the characteristics of the study participants being evaluated. The test statistic of the Wald χ^2 test under the null hypothesis H_{0j} is [12]:

$$\left(L_j^T \hat{\beta}\right)^T \left[L_j^T \widehat{\Sigma} L_j\right]^{-1} \left(L_j^T \hat{\beta}\right) \xrightarrow{D} \chi_1^2, \tag{6}$$

where $\widehat{\Sigma}$ is the estimated variance-covariance matrix of $\text{Var}(\hat{\beta})$.

A more robust approach is to compute a truncated mean of the coefficients where potential 'outliers' can be prevented from contaminating the average effect. Let $\beta_{(1)}, \beta_{(2)}, \dots, \beta_{(M)}$ be the ordered values of the regression coefficients. A $\delta \times 100\%$ truncated mean can be calculated as follows [13]:

$$\bar{\beta}_{\text{truncated}} = \frac{1}{M - 2[M \cdot \delta]} \sum_{q=[M \cdot \delta]+1}^{M-[M \cdot \delta]} \beta_{(q)}, \tag{7}$$

where $[x]$ denotes the integer part of x .

The null hypothesis that the j -th evaluator is not an 'outlier' is now to compare the regression coefficient of the j -th evaluator to the $\delta \times 100\%$ truncated mean:

$$H_{0j} : \beta_j - \bar{\beta}_{\text{truncated}} = 0, \quad j = 1, 2, \dots, M. \tag{8}$$

We refer the readers to Supplementary Material Section 1 for technical details on constructing the design matrix $L_{\delta \times 100\%, j}^T$ to perform hypothesis testing in (8).

Since our goal is to detect as many potential 'outlier' evaluators as possible, we would like to achieve sufficient power when the evaluators are true 'outliers'. Therefore, to complete the hypothesis testing procedure, different from the traditional approach where emphasis is placed upon controlling the type-I error α at an acceptable level, we also attach importance to ensuring an appropriate level of type-II error.

Type-I error determination

Ideally, when performing hypothesis tests to detect potential ‘outlier’ evaluators, there is sufficient power to reject the null hypotheses $H_{0,j}$ when a pre-specified alternative hypothesis $H_{1,j}$ is true. Denote the pre-specified alternative hypothesis as $H_{1,j} : \left| \mathbf{L}_j^T \boldsymbol{\beta} \right| = c$, where c can be determined based on subject matter knowledge. For instance, in the CHEARS AAA, the ‘hearing threshold’ for each individual ear is measured by the lowest sound intensity of a pure-tone signal presented individually to each ear, to which the listener reliably responds, and the pure-tone signal was measured in 5-dB steps [9]. As a result, hearing loss was defined as a greater than 5-dB HL increase in the pure-tone averages of testing frequencies at low-frequency (0.5, 1, 2 kHz), mid-frequency (3, 4 kHz), and high-frequency (6, 8 kHz) [9]. Therefore, it is important to identify audiologists who consistently gave 5-dB larger or smaller hearing test results than their counterparts after controlling for study participants’ characteristics. Thus, a reasonable value for the alternative hypothesis for which we hope to have sufficient power to detect is $c = 5$ for the CHEARS AAA. For presentational simplicity, we do not distinguish between \mathbf{L}_j and $\mathbf{L}_{\delta \times 100\%,j}$ in this section, and we use \mathbf{L}_j to denote the contrast matrix of both tests.

In general, the power formula for the hypothesis test: $H_{0,j} : \mathbf{L}_j^T \boldsymbol{\beta} = 0$ v.s. $H_{1,j} : \left| \mathbf{L}_j^T \boldsymbol{\beta} \right| = c$ is:

$$P\left(\left(\mathbf{L}_j^T \hat{\boldsymbol{\beta}}\right)^T \left[\mathbf{L}_j^T \hat{\Sigma} \mathbf{L}_j\right]^{-1} \left(\mathbf{L}_j^T \hat{\boldsymbol{\beta}}\right) > \chi_{1,1-\alpha}^2 \left| \mathbf{L}_j^T \boldsymbol{\beta} \right| = c\right) = \phi, \tag{9}$$

where α is a two-sided type-I error rate, and ϕ is the power of the test.

Under alternative hypothesis, test statistic $\left(\mathbf{L}_j^T \hat{\boldsymbol{\beta}}\right)^T \left[\mathbf{L}_j^T \hat{\Sigma} \mathbf{L}_j\right]^{-1} \left(\mathbf{L}_j^T \hat{\boldsymbol{\beta}}\right)$ follows a noncentral χ^2 distribution with one degree of freedom and noncentral parameter $\lambda_j = \frac{c^2}{\mathbf{L}_j^T \hat{\Sigma} \mathbf{L}_j}$ [14]; we denote this distribution as $\chi_1^2(\lambda_j)$. Let $F_{\chi_1^2(\lambda_j)}$ be the cumulative distribution function of $\chi_1^2(\lambda_j)$. It follows that the power of the test under the significance level α and alternative hypothesis $H_{1,j} : \left| \mathbf{L}_j^T \boldsymbol{\beta} \right| = c$ is

$$\phi = 1 - F_{\chi_1^2(\lambda_j)}(\chi_{1,1-\alpha}^2). \tag{10}$$

To ensure sufficient power for each evaluator at a pre-specified alternative hypothesis, we can first fix the power ϕ of the tests, and solve Eq. (10) to obtain the corresponding significance levels $\alpha_j(\phi)$ for rejecting the null hypothesis $H_{0,j} : \mathbf{L}_j^T \boldsymbol{\beta} = 0$. Under the same power and alternative hypothesis, each evaluator has an evaluator-specific significance level instead of a unified one due to the differences in the estimated variances of the coefficient estimates.

False discovery rate estimation

The null hypotheses that we are testing are $H_{0,1}, H_{0,2}, \dots, H_{0,M}$. Due to multiple testing, using a traditional significance level such as 0.05 in each test may lead to a high rate of finding ‘outlier’ evaluators even if they are ‘normal’ ones (i.e. making false discoveries) [15, 16]. In our setting, since the evaluator-specific significance levels are determined by ensuring a pre-specified power of the tests, we are more likely to make false discoveries than the traditional α -level hypothesis tests when the pre-specified power is large. To protect us from falsely classifying too many ‘normal’ evaluators as ‘outliers’, we propose to adopt the concept of the false discovery rate (FDR) [15] to control the rate of making false positive decisions.

We provide an approximation of FDR by:

$$\hat{E}(\mathbf{Q}; \phi) = \frac{\sum_{j=1}^M \alpha_j(\phi)}{\sum_{j=1}^M I(p_j < \alpha_j(\phi))}, \tag{11}$$

where \mathbf{Q} is defined as the proportion of true null hypotheses being falsely rejected among the total rejected null hypotheses and we refer the readers to Supplementary Material Section 2 for technical details.

Note that, in our approach, instead of using a unified significance level for all tests, such as $\alpha = 0.05$, each null hypothesis has its own evaluator-specific significance level such that a pre-specified power for detecting a pre-specified alternative hypothesis is achieved for all the hypothesis tests. The estimated FDR, $\hat{E}(\mathbf{Q}; \phi)$, on the other hand, can inform us of the number of false discoveries that may be made. Therefore, when choosing an appropriate set of significance levels, apart from ensuring sufficient power for the tests, the estimated FDR can be used as another criterion reflecting our tolerance towards making false discoveries.

FDR vs. Power decision plot

As described in previous sections, for a given power, we could solve Eq. (10) to get the corresponding evaluator-specific significance levels for rejecting the null hypotheses $H_{0,j}, j = 1, \dots, M$, and based on these significance levels, the corresponding FDR can be estimated using Eq. (11). Therefore, the relationship between power and FDR can be reflected by a decision plot where the power (ϕ) is on the x-axis, and the corresponding estimated FDR ($\hat{E}(\mathbf{Q}, \phi)$) is on the y-axis. Based on the decision plot, we can pick up the significance levels at which an acceptable trade-off between power and the FDR is achieved.

We could also first select a relatively low FDR and find the corresponding power along with the evaluator-specific significance levels from the decision plot; we can

then reject the null hypotheses with p -values of the tests less than the thresholds. Alternatively, if we are less concerned about making false discoveries but would like to be able to detect as many potential ‘outlier’ evaluators as possible, then we could first specify a relatively large power, and reject the null hypotheses by comparing the p -values with the corresponding evaluator-specific significance levels; the estimated FDR from the decision plot can inform us of the number of false discoveries we might have made.

FDR-based adjustment

We may further adjust the set of rejected null hypotheses based on the estimated FDR, especially when $\hat{E}(\mathbf{Q}; \tilde{\phi})$ is large under the chosen power $\tilde{\phi}$.

Let \mathcal{R} be the set of the rejected null hypotheses, and k be the number of hypotheses in \mathcal{R} . Denote the rejected hypotheses as $H_{0,(1)}, H_{0,(2)}, \dots, H_{0,(k)}$, where they are ordered by their p -values in an ascending order. Since $\hat{E}(\mathbf{Q}; \tilde{\phi}) \times k$ approximates the expected number of true null hypotheses that are falsely rejected among $H_{0,(1)}, H_{0,(2)}, \dots, H_{0,(k)}$, an ad hoc approach to further adjust the rejected null hypotheses based on the estimated FDR is to move the latter $\lceil \hat{E}(\mathbf{Q}; \tilde{\beta}^p) \times k \rceil$ null hypotheses $H_{0,(k - \lceil \hat{E}(\mathbf{Q}; \tilde{\beta}^p) \times k \rceil + 1)}, \dots, H_{0,(k)}$ out of set \mathcal{R} , where $\lceil x \rceil$ rounds x to the nearest integer. Finally we would only reject $H_{0,(1)}, H_{0,(2)}, \dots, H_{0,(k - \lceil \hat{E}(\mathbf{Q}; \tilde{\beta}^p) \times k \rceil)}$, and the corresponding ‘outliers’ are evaluators (1), (2), ..., and $(k - \lceil \hat{E}(\mathbf{Q}; \tilde{\beta}^p) \times k \rceil)$. An algorithm statement that summaries the complete quality control procedure is provided in Supplementary Material Section 3.

Simulation

We perform a simulation study to assess the proposed quality control procedure for detecting ‘outlier’ evaluators. As a demonstration, we base our simulations on the audiometrically-assessed hearing threshold measurements at 8 kHz that were obtained in the CHEARS AAA in 2014, where 3,568 participants had assessments in both ears that were measured by 68 different licensed audiologists. Note that, the AAA was still in data collection stage in 2014, and detecting the ‘outlier’ audiologists would help investigators make prompt adjustment to obtain accurate measurements for tests conducted afterwards. We evaluate the performance of the proposed FDR estimator in Eq. (11), as well as true positives (successfully detecting true ‘outlier’ evaluators) and false positives (falsely classifying ‘normal’ evaluators as ‘outliers’) yielded by our quality control method compared with using a traditional and unified significance level such as $\alpha = 0.05$ to reject the null hypotheses.

Data generation

We first consider the scenario when evaluators measure a single outcome for each study participant. We generate data based on the model below, mimicking the right ear data obtained from the CHEARS AAA:

$$\begin{aligned}
 Y_i = & \gamma_1 \text{age}_i + \gamma_2 \text{age}_i^2 + \gamma_3 I(\text{very good}_i) \\
 & + \gamma_4 I(\text{a little hearing trouble}_i) + \beta_1 \text{Audio}_i^{(1)} + \beta_2 \text{Audio}_i^{(2)} \\
 & + \dots + \beta_M \text{Audio}_i^{(M)} + \epsilon_i,
 \end{aligned}
 \tag{12}$$

where age is generated from a normal distribution with mean 56.6 years and standard deviation (SD) 4.4; we set the ‘excellent’ self-reported hearing status as the reference group and the prevalences of the other two categories ‘very good’ and ‘a little hearing trouble’ were 0.44 and 0.25, respectively. These values are the same as those in the CHEARS AAA. $\text{Audio}_i^{(j)}, j = 1, \dots, M$, is 1 if the hearing test outcome of the i -th study participant is measured by the j -th audiologist, and 0 otherwise.

The coefficients corresponding to age, age^2 , $I(\text{very good})$, and $I(\text{a little hearing trouble})$ are set to be $\gamma_1 = -2.7, \gamma_2 = 0.03, \gamma_3 = 3.3$ and $\gamma_4 = 10.3$, same as the point estimates from the regression analysis on the CHEARS data. The number of audiologists M are set to be 100, and each measures the hearing outcomes on 40 study participants. We set the coefficients as $\beta_1 = \beta_2 = \dots = \beta_5 = 75, \beta_6 = \beta_7 = \beta_8 = 70$ and $\beta_9 = \beta_{10} = \dots = \beta_{100} = 67$. Since the averaged audiologist effect is approximately 67, the 92 audiologists with true effect 67 are considered as ‘normal’ audiologists, and the 3 audiologists with effect 70 and the 5 with effect 75 are considered as true outliers. Note that, here, five ‘outlier’ audiologists have very different effects on the hearing test outcomes from ‘normal’ audiologists and three ‘outlier’ audiologists are slightly different from ‘normal’ audiologists. The values 75 and 67 are determined by the averages of the estimated regression coefficients in the regression analysis on the CHEARS data for the audiologists in the upper 10th percentile and those between the lower and upper 10th percentiles, respectively. The residual ϵ_i is assumed to be normally distributed with mean 0 and standard deviation (SD) $\sigma = 8, 10, 12$, respectively.

Simulation results

The simulation is performed for 300 replicates. Shown in Fig. 1 are the FDR vs. Power decision plots under different standard deviation (SD) of the residuals. We set the alternative hypothesis as $H_{1,j} : \left| \mathbf{L}_{10\%,j}^T \boldsymbol{\beta} \right| = 5$. The solid curve is the estimated FDR based on Eq. (11) averaged over the 300 simulation replicates under powers (ϕ) ranging from 0.1 to 0.95 with step size of 0.01; a loess curve with the default smoothing span 0.75 is fitted to connect

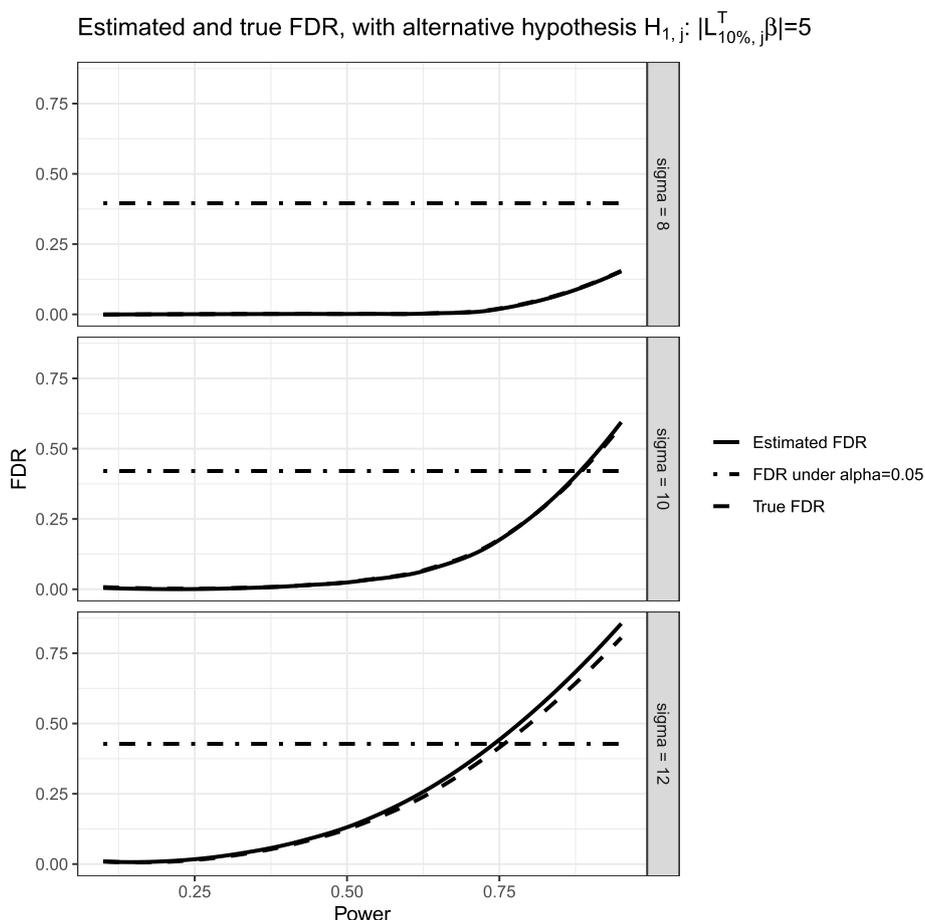


Fig. 1 FDR vs. Power decision plot for single measurement simulation. The alternative hypothesis is $H_{1,j} : |L_{10\%,j}^T \beta| = 5$. The solid curve is the estimated FDR based on Eq. (11) averaged over 300 simulation replicates, and the dashed curve is the empirical true FDR calculated by averaging the proportions of false discoveries $\frac{V(\phi)}{R(\phi)}$ over 300 simulation replicates. The black horizontal dot-dash line represents the empirical true FDR calculated by averaging the proportions of false discoveries over 300 simulation replicates when using $\alpha = 0.05$ as the significance level. The solid and dashed curves are overlapped on the top panel

the points. The dashed curve is an empirical version of the true FDR, which for each ϕ , is the ratio of the number of ‘normal’ audiologists (Audiologists 9 - 100) being falsely detected as ‘outlier’ audiologists to the total number of detected ‘outlier’ audiologists, averaged over the 300 simulation replicates. The horizontal dot-dash line is the empirical version of the true FDR if we use $\alpha = 0.05$ as the significance level for rejecting the null hypotheses averaged over the 300 simulation replicates.

As shown in the decision plot, the estimated FDR is very close to the true FDR when $\sigma = 8$ and 10; while it slightly overestimate the true value when $\sigma = 12$. Moreover, as the SD of the residual increases, the FDR also increases. For example, when $\sigma = 8$, the FDR is less than 0.165 under power 0.95, while if σ increases to 12, the FDR is greater than 0.8 under the same power. Define the noise ratio as $\frac{\sigma^2}{\text{Var}(Y)}$, which is the proportion of the variance of the residual among the total

variance of the outcome measurement. The corresponding noise ratios are approximately 0.52, 0.64, and 0.72 for $\sigma = 8, 10$ and 12. When the noise ratio increases, we are more likely to make false discoveries. Therefore, when performing quality control, including all the possible predictors and confounders in the first stage regression is crucial; this way, we can minimize the residual of the first stage regression and, as a result, minimize the FDR.

Compared with an approach that uses a fixed significance level $\alpha = 0.05$, our method enjoys more flexibility since we can choose the evaluator-specific significance levels by considering both the power and FDR. When $\sigma = 8$, under any power, our approach has a much lower FDR than using $\alpha = 0.05$ as the threshold; and when $\sigma = 10$ and 12, even though the FDR increases, it is still smaller than the FDR if using $\alpha = 0.05$ as the threshold, when the power is chosen to be less than 0.8 and 0.75, respectively.

Since the goal of the method is to detecting as many potential ‘outlier’ evaluators as possible while making the type-I error rate under an acceptable level, we define the true positive proportion for each true ‘outlier’ audiologist (i.e., Audiologists 1 to 8) as the proportion of simulation replicates that correctly detect the audiologist as an ‘outlier’ over the 300 simulation replicates, and the false positive proportion for each true ‘normal’ audiologist (i.e., Audiologists 9 to 100) as the proportion of simulation replicates that falsely identify the audiologist as an ‘outlier’ over the 300 simulation replicates. Figure 2a and b show the true positive proportions for Audiologists 1 to 8, and false positive proportions for the ‘normal’ audiologists (For illustration, we select Audiologists 9 to 16.), where $\sigma = 8$ when generating the data, and the alternative hypothesis is set as $H_{1,j} : \left| \mathbf{L}_{10\%,j}^T \boldsymbol{\beta} \right| = 5$. The black points are the proportions based on our quality control procedure under different powers of the tests; while the horizontal dotted lines are the proportions calculated using $\alpha = 0.05$ as the threshold for rejecting the null hypotheses. We consider both the unadjusted procedure and the FDR-based adjusted procedure.

For the unadjusted procedure, as the power increases, the true positive proportions for Audiologists 1 to 5 reach to 1 quickly, which is expected since the difference between their coefficients and those of the ‘normal’ audiologists are set to be 8, greater than the difference used in the alternative hypothesis $H_{1,j} : \left| \mathbf{L}_{10\%,j}^T \boldsymbol{\beta} \right| = 5$. However, for Audiologists 6 to 8, since their coefficients are only 3 larger than the ‘normal’ audiologists, the true positive proportions are far less than 1 even when the power is large. Compared to the approach that uses $\alpha = 0.05$ as the threshold, our quality control procedure has smaller true positive proportions when the power of test is smaller than 0.3, 0.6, 0.7 for $\sigma = 8, 10, 12$, but gradually they will increase to approximately the same or even higher level. For the ‘normal’ audiologists (Audiologists 9 to 16), the false positive proportions are approximately 0.05 if using $\alpha = 0.05$ as the threshold. Our quality control procedure has even smaller false positive proportions when $\sigma = 8$ and 10 under nearly every power considered. When $\sigma = 12$, the false positive proportions are still smaller than those from using $\alpha = 0.05$ as the threshold, if the power is no larger than 0.9.

Compared with the unadjusted procedure, the FDR-based adjusted true positive proportions for the true ‘outlier’ audiologists and false positive proportions for ‘normal’ audiologists do not change much in the case of $\sigma = 8$ since the FDR is small, and the adjustment is minor. As σ increases, for example, when $\sigma = 10$, the FDR is large enough to yield sufficient number of adjustments for power larger than 0.75. Apart from a decrease

in the false positive proportions for the true ‘normal’ audiologists (Audiologists 9 to 16), we also observe a decrease in the true positive proportions for the true ‘outlier’ audiologists (Audiologists 1 to 8). Therefore, the ad hoc FDR-based adjustment helps to reduce the chances of making false discoveries, with a price of a reduction in the probability of making true positive decisions.

Moreover, we also conducted a simulation study for the scenarios when outcomes are correlated. The data generation process and simulation results are presented in Supplementary Material Section 1. The simulation results are similar with the single measurement scenarios; our outlier detection procedure typically has lower false positive proportions for the true ‘normal’ audiologists and higher true positive proportions for the true ‘outlier’ audiologists compared with the approach that fix the significance level at $\alpha = 0.05$.

Application

To illustrate our method, we apply our method to detect ‘outlier’ audiologists for the audiometrically-assessed hearing threshold measurements in the CHEARS AAA collected in 2014, when the baseline testing was completed on 3,749 participants. We focus on the test results at 8 kHz. We use the GEE approach in the first stage regression analysis and we include age, age², self-reported hearing status (‘excellent’, ‘very good’ and ‘a little hearing trouble’), and dummy variables for the 68 audiologists in the regression model. This regression is fitted using SAS proc genmod, assuming an exchangeable working variance-covariance structure.

We display the scatter plots of $\hat{\beta}_i - \frac{1}{M} \sum_{q=1}^M \hat{\beta}_q$ and $\hat{\beta}_i - \frac{1}{M-2[M-\delta]} \sum_{q=[M-\delta]+1}^{M-[M-\delta]} \hat{\beta}_q$, with $M = 68, \delta = 0.1$, in Fig. 3. Regardless of whether we are comparing with the untruncated mean or the 10% truncated mean, the plots are similar. As shown in Fig. 3a and b, Audiologist 13 has a much larger (> 10 dB) coefficient estimate than their counterparts, and Audiologist 4 has a much smaller (< 10 dB) coefficient estimate than the rest of the audiologists. Moreover, Audiologists 14, 15, 22, 47, 48, 54, 55 and 59 have a mildly different (5-10dB) coefficient estimates from the average effect.

Figure 4a to d show the FDR vs. Power decision plots, where the hypothesis tests are performed to compare each audiologist’s regression coefficient with both the untruncated mean and the 10% truncated mean. We fix the alternative hypothesis as $H_{1,j} : \left| \mathbf{L}_j^T \boldsymbol{\beta} \right| = 5$ and 10, and $H_{1,j} : \left| \mathbf{L}_{10\%,j}^T \boldsymbol{\beta} \right| = 5$ and 10, respectively, for $j = 1, \dots, 68$. Based on the decision plots, ‘outlier’ audiologists can be detected by choosing an appropriate set of significance levels that correspond to reasonable power and FDR. The results are similar between the untruncated mean and

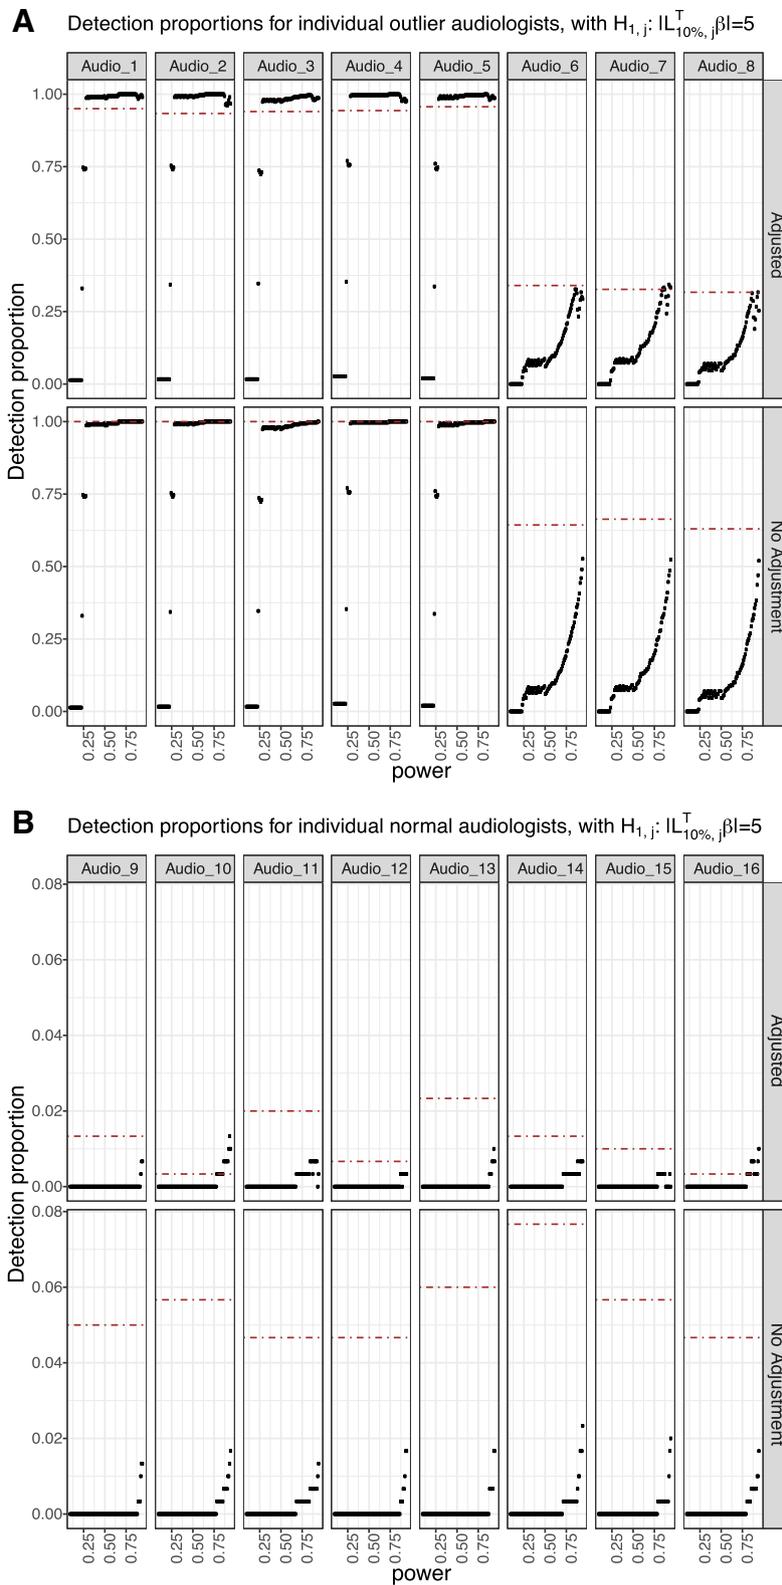


Fig. 2 This figure shows the true positive proportions for the true ‘outlier’ audiologists and false positive proportions for the true ‘normal’ audiologists for single measurement simulation with $\sigma = 8$. The top panel in each subfigure is the result by performing the FDR-based adjustment, while the bottom panel in each subfigure is the result without FDR-based adjustment. The horizontal dot-dash line represents the corresponding true or false positive proportion for each audiologist if we use $\alpha = 0.05$ as the significance level for rejecting the null hypotheses

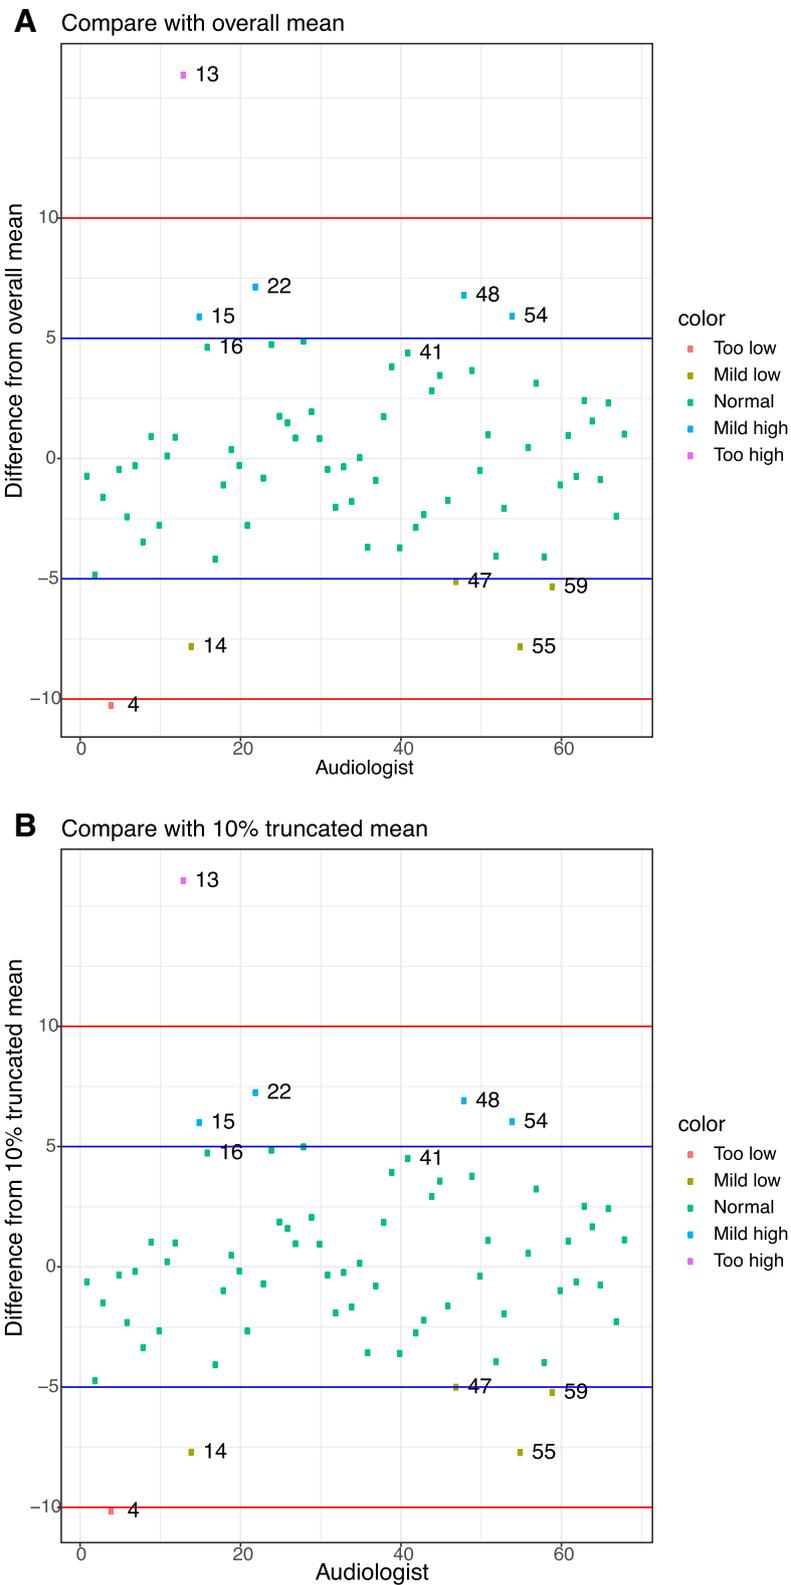


Fig. 3 **a** Subtracting each audiologist’s coefficient estimate by the untruncated mean of all audiologists’ coefficient estimates; **b** Subtracting each audiologist’s coefficient estimate by the 10% truncated mean of all audiologists’ coefficient estimates

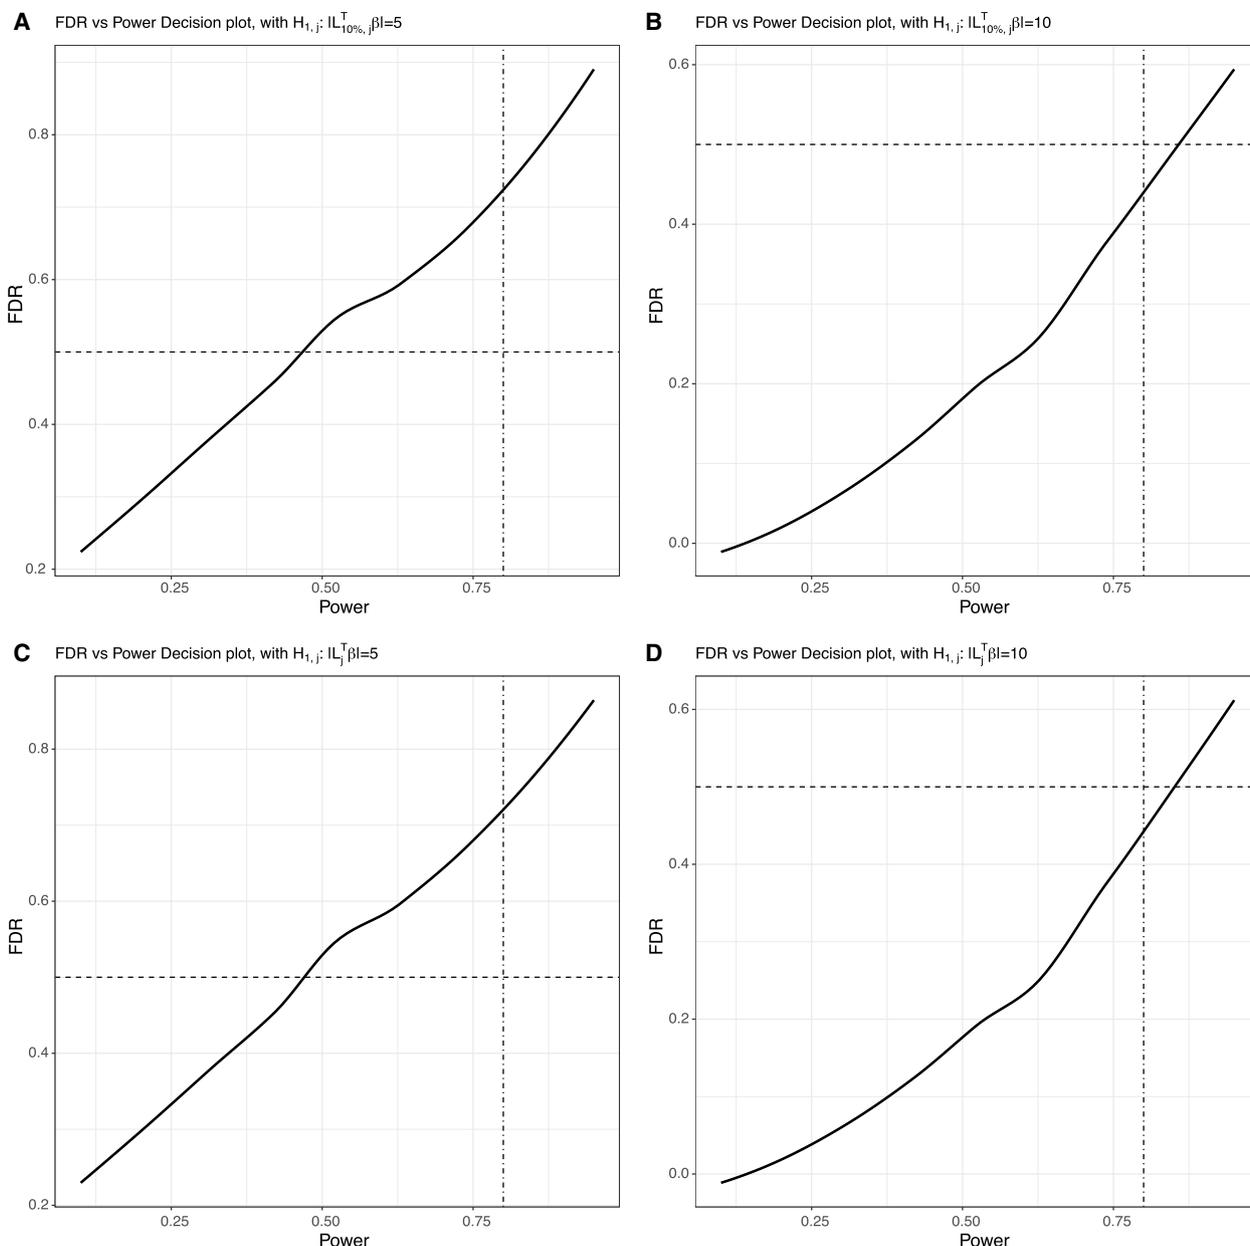


Fig. 4 FDR vs. Power decision plot for detecting ‘outlier’ audiologists, where **a**: $H_{1,j} : |L_{10\%,j}^T \beta| = 5$; **b**: $H_{1,j} : |L_{10\%,j}^T \beta| = 10$; **c**: $H_{1,j} : |L_j^T \beta| = 5$; and **d**: $H_{1,j} : |L_j^T \beta| = 10$. The dot-dash and dashed lines are produced by fixing power at 0.8 or the FDR at 0.5, respectively

truncated mean approach. Table 1 summarize the results when setting the power at 0.8 or the estimated FDR at 0.5. As shown in the table, Audiologists 4 and 13 are detected as ‘outliers’ by all of the approaches regardless of the power, FDR or the alternative hypothesis considered, and Audiologist 48 is detected by all of the approaches under the alternative hypothesis $H_{1,j} : |L_{10\%,j}^T \beta| = 5$ and $H_{1,j} : |L_j^T \beta| = 5$. Therefore, Audiologists 4, 13 and 48 are

likely to be ‘outlier’ audiologists, suggesting close scrutiny may be merited. However, for the approach of using $\alpha = 0.05$ to reject the null hypotheses as shown in the last two rows of the tables, apart from being not flexible as compared with our method, it also suffers from the problem that the power of tests for different audiologists varies significantly with a minimum of 0.55 and a maximum of 1.00.

Table 1 Detected ‘outlier’ audiologists from AAA of CHEARS. Each audiologist’s coefficient estimate is compared with the 10% truncated mean of all audiologists’ coefficient estimates

Alternative	Power	\widehat{FDR}	Outlier Audiologists	Outlier Audiologists Corrected by FDR
$H_{1,j} : L_{10\%j}\beta = 5$	0.80	0.72	2, 4, 8, 13, 14, 15, 16, 17, 22, 24, 28, 36, 39, 41, 42, 47, 48, 49, 52, 54, 55, 57, 58, 59	4, 13, 16, 41, 48
$H_{1,j} : L_{10\%j}\beta = 5$	0.47	0.50	4, 13, 14, 15, 22, 48, 54, 55 ^a	4, 13, 48 ^b
$H_{1,j} : L_{10\%j}\beta = 10$	0.80	0.44	4, 13, 14, 55	4, 13
$H_{1,j} : L_{10\%j}\beta = 10$	0.86	0.50	4, 13, 14, 22, 55	4, 13
-	(0.55, 1.00)	0.28	4, 13, 15, 16, 17, 22, 24, 40, 41, 45, 48, 63	4, 13, 16, 22, 24, 40, 41, 48, 63

To compare the results to standard practice where 0.05 is used as the significance level, last row reports the results from using $\alpha = 0.05$ as the threshold for rejecting tests $H_{0,1}, \dots, H_{0,68}$, where 68 is the number of audiologists in AAA of CHEARS; The range of the power of these 68 hypothesis tests under $\alpha = 0.05$ was reported in the Power column.

^a When we compare the coefficient estimates to the untruncated mean, an additional Audiologist 59 is detected;

^b An additional Audiologist 22 is detected after the FDR correction when we compare the coefficient estimates to the untruncated mean. For other situations, using the untruncated mean for comparison yields the same results as using the truncated mean

Discussion

In this paper, we propose a novel method to address a common issue in large epidemiologic studies that rely on multiple evaluators to obtain exposure or outcome measurements to optimize data quality during data collection stage. Specifically, we developed a two-stage algorithm to detect ‘outlier’ evaluators, who may tend to have higher or lower measurements than those of their counterparts. In the first stage, we fit a regression model for the measurements against evaluators and study participants’ characteristics that could predict the measurements. In the second stage, based on the regression coefficients in the first stage, we perform hypothesis tests to compare the mean measurement of each evaluator with the average mean measurements over all evaluators adjusting for the characteristics of the individuals evaluated. Different from the traditional hypothesis testing procedure where controlling type-I error is the primary focus, we also attach equal importance to ensuring an appropriate level of type-II error since our goal is to detect as many potential ‘outlier’ evaluators as possible for quality control purpose. We derive the evaluator-specific significance levels for rejecting the null hypotheses under selected powers of the tests. These significance levels are not necessarily 0.05 and are different across audiologists due to the differences in the variances of the coefficient estimates. To account for the issue of multiple comparisons, we also derive an FDR-estimator. An FDR vs. Power decision

plot can be created, and based on this plot, the evaluator-specific significance levels for rejecting the null hypotheses can be determined such that both FDR and Power are acceptable.

When performing hypothesis tests to detect ‘outlier’ evaluators, we proposed to compare the coefficient estimates to the truncated mean to prevent those ‘outlier’ evaluators from contaminating the estimated normal effect. Alternatively, we can consider an interval null, that is $H_0 : |\beta_i - \frac{1}{M} \sum_{j=1}^M \beta_j| \leq a$ for some constants $a > 0$. A challenge of this method might be how to select a . We will consider this method in our future research and compare it with the current method. Moreover, when calculating the evaluator-specific significance level, the knowledge of the alternative hypothesis is needed. However, if the prior knowledge is not available, we recommend performing sensitivity analysis for a series of reasonable values of the alternative hypothesis. In addition, the FDR approximation in Eq. (2) holds when the number of hypotheses (M) being conducted is large. However, when M is small, alternatively, we can use the Benjamini-Hochberg (BH) procedure to control the FDR [15]. The BH procedure proceeds by first specifying an FDR level α , and sort the null hypothesis based on p-values in ascending order $(P_{(1)}, P_{(2)}, \dots, P_{(M)})$. Then the largest k such that $P_{(k)} \leq \frac{k}{M}\alpha$ is obtained, and the first k null hypotheses will be rejected. The BH procedure can ensure that the FDR is controlled at level α . However, different from

our approach, the BH procedure does not consider the power of tests and to be conservative, we might use a relatively larger α level such as 0.1 when conducting the BH procedure.

There are several important points for consideration based on our work. First, an increase in the noise ratio $\frac{\sigma^2}{\text{Var}(Y)}$ will increase FDR, especially when the power of the test is large. Therefore, in the first stage regression, it is crucial to include all potential predictors of the measurements as regressors. Second, the proposed method assumes that the evaluator effect on the measurements is not modified by the participants' characteristics. In the case when this assumption is violated, we can estimate the evaluator effect in each category of the potential effect modifier by including the evaluator indicator-effect modifier interactions in the first stage regression model, and then we can regard the same evaluator for testing study participants in different categories of the effect modifier as if they were different evaluators. This way, an evaluator could be detected as an 'outlier' only when testing study participants in a specific category of the effect modifier. Third, to accommodate situations where the measurements are not continuous, a link function can be used in the first stage regression, such as the logit link for binary measurements, and log link for count measurements.

Our quality control procedure is used to detect potential 'outlier' evaluators and once they are detected, quality check on those evaluators should be performed to ensure future measurements can be measured accurately. However, a correction of measurement errors in existing measurements obtained by 'outlier' evaluators is beyond the scope of this paper. We will develop measurement error correction methods in future research; one idea could be to calibrate the measurements from 'outlier' evaluators to 'normal' measurements using information from the first-stage regression models, taking into account participants' characteristics.

The regular regression and GEE approach may not lead to reliable β -estimator if the numbers of study participants tested by some evaluators are small. In this case, an alternative method is to treat the measurements from the same evaluator as a cluster and to use the mixed effects model in the first stage regression analysis. In the scenario where each participant has a single measurement, this mixed effects model may include an evaluator-specific random intercept in addition to the fixed effect participants' characteristics; the estimated value of the j -th evaluator-specific intercept is $\hat{\beta}_j$. Similarly, in the scenario where the participants have multiple measurements, the mixed effects model may include both evaluators and participants (nested within evaluator) as random effects. Once the mixed effects model obtains $\hat{\beta}$ and $\widehat{\text{Var}}(\hat{\beta})$, the

rest of the methods are the same as those stated in Subsection 'Hypothesis testing' to Subsection 'FDR-based adjustment' of this paper.

In addition to the contribution to quality control during the data collection stage of epidemiologic studies, our outlier detection method can also be valuable in clinical settings for the detection of 'outlier' evaluators (e.g. health providers or technicians); for example, clinical diagnoses often rely on measurements from evaluators, and inaccurate measurements may lead to wrong diagnoses. Furthermore, our method can be used in statistical analysis procedures. For example, for studies based on laboratory measurements of biomarkers such as plasma or urine metabolites that are measured in different batches, our method can help to identify potential 'outlier' batches, and a sensitivity analysis can be conducted by excluding those 'outlier' batches and re-estimating the parameters of interests.

R code for implementing the proposed method is available at <https://github.com/molingwang/Analytical-Methods-for-Hearing-Studies/branches>.

Conclusions

Our two-stage algorithm is a useful method for detecting 'outlier' evaluators who tend to give higher or lower measurements than their counterparts after adjusting for study participants' characteristics. Compared with traditional hypothesis tests that focus on type-I error, we also attach importance to the type-II error so that as many potential 'outliers' can be identified, and an estimated FDR is used to control for the false positive rate. We recommend applying our method for 'outlier' detection during data collection stage to improve data quality.

Abbreviations

FDR	False discovery rate
CHEARS	Conservation of Hearing Study
AAA	Audiology Assessment Arm
NHS	Nurses' Health Study
GEE	Generalized Estimating Equations

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12874-023-01988-4>.

Additional file 1.

Acknowledgements

We are thankful to the study participants in CHEARS.

Authors' contributions

Y.W., B.R. and M.W. developed the methods; Y.W. designed and conducted the simulation study, wrote the first draft of the manuscript. S.C., B.R., G.C., and M.W. reviewed the manuscript critically. All authors read and approved the final manuscript.

Funding

This work is supported by NIH grant R01DC017717.

Availability of data and materials

The data that support the findings of this study are available from Nurses' Health Study (NHS) II but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Data are however available from the authors upon reasonable request and with permission of Nurses' Health Study (NHS) II.

Declarations**Ethics approval and consent to participate**

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Received: 30 November 2021 Accepted: 11 July 2023

Published online: 01 August 2023

References

- Cruickshanks KJ, Wiley TL, Tweed TS, Klein BE, Klein R, Mares-Perlman JA, et al. Prevalence of hearing loss in older adults in Beaver Dam, Wisconsin: The epidemiology of hearing loss study. *Am J Epidemiol*. 1998;148(9):879–86.
- Shargorodsky J, Curhan SG, Curhan GC, Eavey R. Change in prevalence of hearing loss in US adolescents. *JAMA*. 2010;304(7):772–8.
- Gopinath B, McMahon CM, Rochtchina E, Karpa MJ, Mitchell P. Incidence, persistence, and progression of tinnitus symptoms in older adults: the Blue Mountains Hearing Study. *Ear Hear*. 2010;31(3):407–12.
- Zhang X, Bullard KM, Cotch MF, Wilson MR, Rovner BW, McGwin G, et al. Association between depression and functional vision loss in persons 20 years of age or older in the United States, NHANES 2005–2008. *JAMA Ophthalmol*. 2013;131(5):573–81.
- Klein R, Lee KE, Gangnon RE, Klein BE. Relation of smoking, drinking, and physical activity to changes in vision over a 20-year period: the Beaver Dam Eye Study. *Ophthalmology*. 2014;121(6):1220–8.
- McCullough ML, Zoltick ES, Weinstein SJ, Fedirko V, Wang M, Cook NR, et al. Circulating vitamin D and colorectal cancer risk: an international pooling project of 17 cohorts. *JNCI: J Natl Cancer Inst*. 2019;111(2):158–69.
- Carroll RJ, Ruppert D, Stefanski LA, Crainiceanu CM. Measurement error in nonlinear models: a modern perspective. Chapman and Hall/CRC; 2006.
- Curhan SG, Wang M, Eavey RD, Stampfer MJ, Curhan GC. Adherence to healthful dietary patterns is associated with lower risk of hearing loss in women. *J Nutr*. 2018;148(6):944–51.
- Curhan SG, Halpin C, Wang M, Eavey RD, Curhan GC. Prospective Study of Dietary Patterns and Hearing Threshold Elevation. *Am J Epidemiol*. 2020;189(3):204–14.
- Liang KY, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika*. 1986;73(1):13–22.
- Zeger SL, Liang KY. Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*. 1986;42:121–30.
- Harrell Jr FE. Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis. Springer; 2015.
- Wilcox RR. Introduction to robust estimation and hypothesis testing. Academic Press; 2011.
- Lehmann EL, Romano JP. Testing statistical hypotheses. Springer Science & Business Media; 2006.
- Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Ser B Methodol*. 1995;57(1):289–300.
- Benjamini Y, Drai D, Elmer G, Kafkafi N, Golani I. Controlling the false discovery rate in behavior genetics research. *Behav Brain Res*. 2001;125(1–2):279–84.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

