

BayesSenMC: an R package for Bayesian Sensitivity Analysis of Misclassification

by Jinhui Yang, Lifeng Lin and Haitao Chu

Abstract In case-control studies, the odds ratio is commonly used to summarize the association between a binary exposure and a dichotomous outcome. However, exposure misclassification frequently appears in case-control studies due to inaccurate data reporting, which can produce bias in measures of association. In this article, we implement a Bayesian sensitivity analysis of misclassification to provide a full posterior inference on the corrected odds ratio under both non-differential and differential misclassification. We present an R (R Core Team, 2018) package **BayesSenMC**, which provides user-friendly functions for its implementation. The usage is illustrated by a real data analysis on the association between bipolar disorder and rheumatoid arthritis.

Introduction

Many epidemiological studies are concerned with assessing the risk of an outcome between exposed and non-exposed subjects. For example, in a case-control study, researchers first identify subjects who have the disease of interest (the case group) and subjects who do not (the control group), and then ascertain the exposure status of the subjects in each group. The odds ratio is typically used to assess the association between the exposure and disease in the case-control study; it describes the ratio of the exposure odds in the case group to that in the control group.

However, misclassification of exposure, disease outcome, or covariates appears frequently in observational studies of epidemiological or medical research (Rothman et al., 2008; Brakenhoff et al., 2018). In a case-control study, misclassification is often due to inaccurate reporting of the exposure status (e.g., self-reported data). This can consequently lead to biased estimation of exposure probabilities and odds ratio. To adjust for such biases, we can correct the odds ratio using the observed data from the case-control study and the sensitivity and specificity of correctly classifying exposure status from external data. Here, the sensitivity is the proportion of exposed subjects that are correctly classified as exposed (i.e., true positive), and the specificity is the proportion of non-exposed subjects that are correctly classified as non-exposed (i.e., true negative).

Quantitative assessment of misclassification bias is necessary to estimate uncertainty in study results. There are many statistical methods for misclassification correction; nearly all of them use prior information that maps observed measurements to true values (Greenland, 2005). These methods include regression calibration and multiple imputation (Rosner et al., 1989; Spiegelman et al., 2001; Cole et al., 2006), in which the mapping is based on a validation study. Also, sensitivity analysis can be used to evaluate the effects of uncertainties in measurement on the observed results of the study (Greenland, 1996; Lash and Flink, 2003; Chu et al., 2006), in which the mapping from observed to true measurements may be based on prior information or expert opinion about the accuracy of the measurement. However, when such information or opinion is lacking, researchers may over- or under-adjust for misclassification with an inaccurate guess, which may, in turn, produce a poor estimate (Gustafson et al., 2006).

Moreover, despite the ubiquity of measurement error, these methods remain rarely used due to the complexity of statistical approaches, especially the complexity of prior specifications as well as lack of software package (Lash and Flink, 2003). For example, in a random sample survey of 57 epidemiological studies (Jurek et al., 2006), only one study used quantitative corrections. Sensitivity analysis is simple but limited insofar as it does not provide formal interval estimates that combine uncertainty due to random error with misclassification. Several authors have addressed this deficiency by using probabilistic (Monte Carlo) sensitivity analyses (Greenland, 2005). For example, Fox et al. (2005) proposed a probabilistic sensitivity analysis of misclassified binary variables based on multiple imputation; they provided SAS code and Excel macro for this approach. Such methods can be viewed as means of summarizing the bias over sensitivity analyses using a prior distribution about the bias parameters (Greenland, 2005). Many of them have been implemented in the R package **episensr** (Haine, 2021). Specifically, **episensr** allows for specifications of prior distributions for sensitivity and specificity, such as uniform and logit normal, as well as sequential bias modeling that can be applied to more than one type of bias, such as for both misclassification and selection biases.

Other methods employ a Bayesian implementation of the probabilistic bias analysis or perform an outright Bayesian analysis (Greenland, 2005; Chu et al., 2006; MacLehose and Gustafson, 2012; Gustafson et al., 2006). The Bayesian analysis is substantially different from conventional probabilistic sensitivity analysis and much more flexible for incorporating different types of priors; however, it is

often computationally expensive and more difficult to conduct by general users without a statistical background. [Gustafson et al. \(2006\)](#) accounted for the prior uncertainties of sensitivity and specificity in the evaluation of the results of a case-control study. [MacLehose and Gustafson \(2012\)](#) compared a Bayesian approach with probabilistic bias analysis based on a case-control study of congenital defects, concluding that the two approaches are mostly similar if using similar prior data admissibility as well as uniform priors on exposure probabilities.

This article focuses on an R package correcting for exposure misclassification in a case-control study. Extending from the Bayesian approach introduced by [Gustafson et al. \(2006\)](#), we implement the methods outlined in [Chu et al. \(2006\)](#), which account for the correlation between the sensitivity and specificity in the model specification. The methods can be applied to both non-differential and differential misclassification; that is, the degree of misclassification can be the same across the case and control groups or distinctly different. Furthermore, we use the generalized linear mixed bivariate effects model introduced by [Chu et al. \(2010\)](#) to jointly model the sensitivity and specificity that may be informed by an external meta-analysis on the diagnostic accuracy of the exposure factor.

This article introduces the implementation of the methods for misclassification via our R package **BayesSenMC** (Bayesian sensitivity analysis by Monte Carlo sampling). The package is mainly implemented in Stan, an imperative probabilistic programming language, which uses Hamiltonian Monte Carlo (HMC), a form of efficient Markov Chain Monte Carlo (MCMC) sampling.

An illustrative example

This section presents an illustrative case-control study on the association between bipolar disorder and rheumatoid arthritis, originally investigated by [Farhi et al. \(2016\)](#); this example will also be used to demonstrate the implementation of the methods for misclassification. The exposure is bipolar disorder, and the disease outcome is rheumatoid arthritis, which is a chronic autoimmune disorder that primarily affects joints and occurs in nearly 1% of the population in developed countries ([McInnes and Schett, 2017](#)). Table 1 presents the data.

Table 1: Counts of the case-control study of the association between bipolar disorder and rheumatoid arthritis.

Rheumatoid arthritis	Bipolar Disorder		Total
	Exposed	Unexposed	
Case	66	11,716	11,782
Control	243	57,730	57,973

The unadjusted odds ratio is 1.34 with the 95% confidence interval (CI) (1.02, 1.76), indicating a significant association between rheumatoid arthritis and bipolar disorder. Of note, [Farhi et al. \(2016\)](#) acknowledged the limitation that “lack of validation of the diagnosis of bipolar disorder in the subjects cannot be completely excluded.” For example, bipolar disorder can be classified as type I, type II, etc.; it is especially difficult to diagnose bipolar disorder type II ([Phillips and Kupfer, 2013](#)).

Assuming certain fixed values or prior distributions of sensitivity and specificity, we can use the method by [Chu et al. \(2006\)](#) to correct the odds ratio accounting for the exposure misclassification. The sensitivity and specificity can be either some fixed values or random variables following some prior distributions.

The prior distributions can be estimated from external evidence using a meta-analysis, e.g., using the bivariate generalized linear mixed model approach proposed by [Chu et al. \(2010\)](#). The following section presents the details of these methods. This article uses the meta-analysis performed by [Carvalho et al. \(2015\)](#) to obtain the prior distributions of the sensitivity and specificity of classifying the exposure status of bipolar disorder. The meta-analysis contains three subgroups of screening detection instruments: bipolar spectrum diagnostic scale (8 studies with sensitivity between 0.52 and 0.90 and specificity between 0.51 and 0.97); hypomania checklist (17 studies with sensitivity between 0.69 and 1.00 and specificity between 0.36 and 0.98); and mood disorder questionnaire (30 studies with sensitivity between 0.00 and 0.91 and specificity between 0.47 and 1.00). The dataset from the Clalit Health Services (the largest Health Maintenance Organization in Israel) used in [Farhi et al. \(2016\)](#) does not specify the exact screening detection instrument for identifying bipolar disorder; therefore, we will use all three subgroups’ data (55 studies in total) in [Carvalho et al. \(2015\)](#) for our analysis on the sensitivity and specificity. A subset of the meta-analysis data is shown in Table 2.

According to their definitions, the study-specific sensitivity and specificity can be estimated as (the number of true positives) / (the number of true positives plus the number of false negatives) and

Table 2: The meta-analysis on diagnosis accuracy of bipolar disorder performed by [Carvalho et al. \(2015\)](#).

Study ID	True positive	False negative	True negative	False positive
1	81	9	444	427
2	12	3	44	19
3	74	26	97	3
4	52	16	23	4
5	228	113	18	4
⋮	⋮	⋮	⋮	⋮
55	63	6	32	13

(the number of true negatives) / (the number of true negatives plus the number of false positives). For example, study 1 gives the sensitivity $81/(81 + 9) = 0.90$ and the specificity $444/(444 + 427) \approx 0.51$. The generalized linear mixed-effects model will be used to synthesize all 55 studies to estimate the overall sensitivity and specificity.

Methods

In this section, we introduce the specific models and methods to deal with misclassification.

Bayesian approach to correcting misclassification bias

Consider a case–control study, and we are interested in the odds ratio from this study. Table 3 presents the notation of the observed data. The odds ratio is estimated as

$$\widehat{OR} = \frac{ad}{bc}.$$

When the odds ratio is larger or smaller than 1, the exposure happens more or less likely in the case group, suggesting an association between the disease status and the exposure status. On the other hand, the odds ratio close to 1 suggests that the disease and the exposure are less likely associated.

Table 3: Observed counts of a case–control study.

Group	Exposed	Unexposed	Total
Case	a	b	N_1
Control	c	d	N_0

Assume that the observed exposure probability is P_k , the true exposure probability is π_k , the sensitivity is Se_k , and the specificity is Sp_k for group k ($k = 1$ for the case group and 0 for the control group) in the case–control study. Then, we can represent the observed exposure probability in terms of the true exposure probability, the sensitivity, and the specificity:

$$\begin{aligned} P_k &= P(\text{observed } E \text{ in group } k) \\ &= P(\text{observed } E \mid \text{true } E \text{ in group } k)P(\text{true } E \text{ in group } k) \\ &\quad + P(\text{observed } E \mid \text{true } \bar{E} \text{ in group } k)P(\text{true } \bar{E} \text{ in group } k) \\ &= Se_k \pi_k + (1 - Sp_k)(1 - \pi_k), \end{aligned} \tag{1}$$

where E denotes exposure and \bar{E} denotes non-exposure. This yields

$$\pi_k = (P_k + Sp_k - 1) / (Se_k + Sp_k - 1).$$

Consequently, the misclassification-corrected odds ratio can be calculated as

$$OR_c = \frac{\pi_1 / (1 - \pi_1)}{\pi_0 / (1 - \pi_0)} = \frac{(P_1 + Sp_1 - 1)(Se_0 - P_0)}{(P_0 + Sp_0 - 1)(Se_1 - P_1)}. \tag{2}$$

Based on Equation (1), we can specify the following Bayesian hierarchical model to estimate the corrected odds ratio of the case–control study, with a and c being the observed counts from Table 3:

$$\begin{aligned}
\text{Likelihood: } & a \sim \text{Bin}(N_1, P_1) \text{ and } c \sim \text{Bin}(N_0, P_0); \\
\text{Link: } & P_k = Se_k \pi_k + (1 - Sp_k)(1 - \pi_k), \quad k = 0, 1; \\
& \text{LOR}_c = \text{logit}(\pi_1) - \text{logit}(\pi_0) \text{ and } \text{OR}_c = \exp(\text{LOR}_c); \\
\text{Prior: } & \text{logit}(\pi_0) \sim N(0, 10^2) \text{ and } \text{LOR}_c \sim N(0, 2^2); \\
& Se_k, Sp_k \sim f(\cdot).
\end{aligned} \tag{3}$$

Here, we assume weakly-informative priors for the true exposure probabilities π_0 and π_1 , which give a 95% CI of the true odds ratio between $e^{-2 \times 1.96}$ (≈ 0.02) and $e^{2 \times 1.96}$ (≈ 50.40), and a binomial distribution for the number of observed exposure in case and control studies. Thus, the Bayesian inference can be formulated as a posterior distribution of π_0 , π_1 , and the corrected odds ratio.

Additionally, the function $f(\cdot)$ denotes the joint prior for the sensitivity and specificity (for either non-differential or different misclassification). In practice, the sensitivity and specificity are not available from the case–control study, and they may be estimated as certain fixed values by subjective experts' opinions.

Alternatively, we may consider incorporating evidence-based prior information from existing studies on the diagnostic accuracy of the exposure status (e.g., the data in Table 2). This allows us to account for uncertainties in the sensitivity and specificity and potential correlation between them. The next subsection presents methods to obtain the prior information for the sensitivity and specificity from a meta-analysis.

Estimating prior distributions on the sensitivity and specificity from a meta-analysis

This section briefly discusses the generalized linear mixed-effects model (GLMM) to estimate priors on the sensitivity and specificity. Suppose that a meta-analysis on the diagnostic accuracy of the exposure status is available as external data to inform the priors of sensitivity and specificity that are needed to correct the odds ratio in the case–control study. Denote the number of independent studies in the meta-analysis by m , and let n_{i11} , n_{i00} , n_{i01} and n_{i10} be the number of true positives, true negatives, false positives, and false negatives, respectively, in study i ($i = 1, \dots, m$). Consequently, there are $n_{i11} + n_{i10}$ truly exposed subjects and $n_{i00} + n_{i01}$ truly unexposed subjects.

Assuming that n_{i11} and n_{i00} follow binomial distributions given the number of exposed and unexposed subjects, respectively, the bivariate GLMM can be specified as (Chu et al., 2010; Ma et al., 2016):

$$\begin{aligned}
n_{i11} & \sim \text{Bin}(n_{i11} + n_{i10}, Se_i) \text{ and } n_{i00} \sim \text{Bin}(n_{i00} + n_{i01}, Sp_i), \quad i = 1, \dots, m; \\
g\left(\frac{Se_i - Se^L}{Se^U - Se^L}\right) & = u + \mu_i \text{ and } g\left(\frac{Sp_i - Sp^L}{Sp^U - Sp^L}\right) = v + \nu_i; \\
\begin{bmatrix} \mu_i \\ \nu_i \end{bmatrix} & \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_\mu^2 & \rho\sigma_\mu\sigma_\nu \\ \rho\sigma_\mu\sigma_\nu & \sigma_\nu^2 \end{bmatrix}\right),
\end{aligned} \tag{4}$$

where u and v are the fixed effects implying the overall sensitivity and specificity in all m studies, and μ_i and ν_i are the study-specific random effects. Also, σ_μ^2 and σ_ν^2 describe the heterogeneity of the underlying sensitivity and specificity across studies, and ρ models the correlation between the sensitivity and specificity. We denote the estimated fixed effects by \hat{u} and \hat{v} , the estimated variances as $\hat{\sigma}_\mu^2$ and $\hat{\sigma}_\nu^2$, and the estimated correlation coefficient as $\hat{\rho}$.

The lower and upper bounds Se^L , Se^U , Sp^L , and Sp^U provide constraints on sensitivity and specificity, which are chosen to exclude all improbable values. A smaller difference between Se^L and Se^U (or between Sp^L and Sp^U) indicates higher confidence in the diagnostic accuracy of the exposure status. When there is no confidence for the range, we can set $Se^L = Sp^L = 0$ and $Se^U = Sp^U = 1$. Alternatively, setting $Se^L = Sp^L = 0.5$ indicates that the diagnosis of exposure is better than chance. For simplicity of implementation, we only allow the same lower and upper bounds for Se and for Sp in our package **BayesSenMC**. In addition, $g(\cdot)$ is the link function (e.g., the logit, probit, and complementary log-log). The logit link, $\text{logit}(t) = \log \frac{t}{1-t}$, is commonly used in practice, and our package **BayesSenMC** adopts this link.

Recall that the Bayesian hierarchical model for estimating the corrected odds ratio in the case–control study in Equation (3) specifies a joint prior $f(\cdot)$ for the sensitivity and specificity. We consider six specifications for this prior as follows:

- (i) Crude (uncorrected) odds ratio: no misclassification. The specification of the prior is equivalent

to setting $Se_0 = Se_1 = Sp_0 = Sp_1 = 1$; consequently, $P_1 = \pi_1$ and $P_0 = \pi_0$.

- (ii) Corrected OR: misclassification of the exposure status exists, and the sensitivity and specificity for both cases and controls are assumed to be fixed values. These fixed values can be directly plugged in the Bayesian model in Equation (3).
- (iii) Logit-prior corrected OR: non-differential misclassification of the exposure status exists ($Se_0 = Se_1 = Se$ and $Sp_0 = Sp_1 = Sp$), and the uncertainties of the sensitivity and specificity are considered independently by using normal priors on the logit scale. The evidence of the priors comes from the diagnostic meta-analysis performed in the GLMM in Equation (4). Specifically, we can assign

$$\text{logit} \left(\frac{Se - Se^L}{Se^{UL} - Se^L} \right) \sim N(\hat{u}, \hat{\sigma}_\mu^2) \text{ and } \text{logit} \left(\frac{Sp - Sp^L}{Sp^{UL} - Sp^L} \right) \sim N(\hat{v}, \hat{\sigma}_\nu^2)$$

as the priors in the Bayesian hierarchical model for the case-control study in Equation (3).

- (iv) Fixed-correlation corrected OR: non-differential misclassification of the exposure status exists ($Se_0 = Se_1 = Se$ and $Sp_0 = Sp_1 = Sp$), and the sensitivity and specificity have a joint normal prior on the logit scale to account for their correlation. In practice, the sensitivity is very likely correlated with the specificity when dichotomizing a continuous measurement (Chu and Cole, 2006). Specifically, we use the following bivariate joint prior

$$\begin{bmatrix} \text{logit} \left(\frac{Se - Se^L}{Se^{UL} - Se^L} \right) \\ \text{logit} \left(\frac{Sp - Sp^L}{Sp^{UL} - Sp^L} \right) \end{bmatrix} \sim N \left(\begin{bmatrix} \hat{u} \\ \hat{v} \end{bmatrix}, \begin{bmatrix} \hat{\sigma}_\mu^2 & \hat{\rho}\hat{\sigma}_\mu\hat{\sigma}_\nu \\ \hat{\rho}\hat{\sigma}_\mu\hat{\sigma}_\nu & \hat{\sigma}_\nu^2 \end{bmatrix} \right).$$

Compared with the previous prior specification with independent sensitivity and specificity, the correlation coefficient $\hat{\rho}$ is additionally considered here; it is also estimated from the GLMM in Equation (4).

- (v) Random-correlation corrected OR: in addition to the above bivariate joint prior for the non-differential sensitivity and specificity, we can also consider modeling the uncertainties in the estimated correlation coefficient. We consider applying Fisher's z-transformation to the correlation coefficient in the GLMM. Specifically, instead of directly estimating the correlation coefficient ρ in Equation (4), we reparameterize $\rho = \frac{\exp(2z) - 1}{\exp(2z) + 1}$ and obtain the point estimate of z from the GLMM and its standard error, denoted by \hat{z} and s_z , respectively. These estimates can be subsequently used as the priors for the sensitivity and specificity in the case-control study:

$$\begin{bmatrix} \text{logit} \left(\frac{Se - Se^L}{Se^{UL} - Se^L} \right) \\ \text{logit} \left(\frac{Sp - Sp^L}{Sp^{UL} - Sp^L} \right) \end{bmatrix} \sim N \left(\begin{bmatrix} \hat{u} \\ \hat{v} \end{bmatrix}, \begin{bmatrix} \hat{\sigma}_\mu^2 & \rho\hat{\sigma}_\mu\hat{\sigma}_\nu \\ \rho\hat{\sigma}_\mu\hat{\sigma}_\nu & \hat{\sigma}_\nu^2 \end{bmatrix} \right);$$

$$\rho = \frac{\exp(2z) - 1}{\exp(2z) + 1};$$

$$z \sim N(\hat{z}, s_z^2).$$

- (vi) Differential corrected OR: finally, we consider the differential misclassification of the exposure status, i.e., $Se_0 \neq Se_1$ and $Sp_0 \neq Sp_1$. All above choices of priors can be similarly applied to the four-variate set $\{Se_0, Sp_0, Se_1, Sp_1\}$. For simplicity, we consider a joint prior similar to that in (iv); however, the prior applies to cases and controls separately, and it does not account for the uncertainties in the correlation coefficient as in (v). That is,

$$\begin{bmatrix} \text{logit} \left(\frac{Se_k - Se^L}{Se^{UL} - Se^L} \right) \\ \text{logit} \left(\frac{Sp_k - Sp^L}{Sp^{UL} - Sp^L} \right) \end{bmatrix} \sim N \left(\begin{bmatrix} \hat{u} \\ \hat{v} \end{bmatrix}, \begin{bmatrix} \hat{\sigma}_\mu^2 & \hat{\rho}\hat{\sigma}_\mu\hat{\sigma}_\nu \\ \hat{\rho}\hat{\sigma}_\mu\hat{\sigma}_\nu & \hat{\sigma}_\nu^2 \end{bmatrix} \right), \quad k = 0, 1.$$

Because of the complexity of the Bayesian model in Equation (3) with the above various choices of priors for the sensitivity and specificity, we will use Markov chain Monte Carlo (MCMC) sampling to produce the posterior distribution and thus estimate the misclassification-bias-corrected odds ratio in the case-control study and its credible interval.

Implementation in R

The aforementioned methods can be implemented in the R package **BayesSenMC**. The function `nlmeNDiff` fits a non-differential GLMM and returns a `lme4` (Bates et al., 2021) object, for which commands such as `summary` can be used to extract useful statistics from the model; see `methods(class`

= "merMod") for more details. Users can also call the `paramEst` function to get a list of specific parameter estimates of the fit that can be directly inputted into the model functions of **BayesSenMC** for Bayesian inferences. In addition, the link function used in `nlmeNDiff` can be modified by specifying `lower` and `upper`, which then changes the lower and upper bounds of Se_k and Sp_k ($k = 1$ for cases and 0 for controls).

The package **BayesSenMC** includes six model functions and one graphing function called `plotOR`. The model functions return an S4 object of type `stanfit`, an instance of **rstan** (Stan Development Team, 2020), which is an interface of Stan (Carpenter et al., 2017) in R. Users can call methods such as `print` or `extract` to get detailed information about the posterior samples. The MCMC procedures are implemented with a default of two chains, each with 1000 iterations of burn-in period and 2000 iterations to estimate the posterior parameters. They are fit using `stan`, and the default Monte Carlo algorithm is the No-U-Turn sampler, a variant of Hamiltonian Monte Carlo (Hoffman and Gelman, 2014; Betancourt, 2017). Any additional arguments to the model function call will be passed into `stan`. The returned object can then be inputted into `plotOR` to visualize the posterior distribution of the adjusted odds ratio, as well as the probability density lines of odds ratio in the cases of no misclassification and constant Se/Sp as comparisons to the posterior distribution. It takes optional argument passed into `geom_histogram`, and returns a **ggplot2** (Wickham et al., 2021) object that can be further customized.

The latest version of **BayesSenMC** is available from CRAN. The package can be directly installed via the R prompt:

```
R> install.packages("BayesSenMC")
R> library("BayesSenMC")
```

Example in R

In this section, we use the data in Table 1 as well as Table 2 of meta-analysis data on the diagnosis accuracy of bipolar disorder to demonstrate the capabilities of **BayesSenMC**. The analyses are conducted using R version 4.1.0 (2021-05-18).

We first fit the meta-analysis data using the GLMM procedure implemented in our package, assuming non-differential misclassification. Given the range of Se and Sp of the bipolar disorder meta-analysis data, we must only assume $Se^L = Sp^L = 0$ and $Se^U = Sp^U = 1$ for the GLMM to compute real-value results. However, with more information about the type of diagnoses in Farhi et al. (2016), one can find more informative constraints on Se and Sp to fit a more precise model.

```
R> data(bd_meta)
R> my.mod <- nlmeNDiff(bd_meta, lower = 0)
R> my.mod
```

```
Generalized linear mixed model fit by maximum likelihood (Laplace
Approximation) [glmerMod]
Family: binomial ( logit(0, 1) )
Formula: cbind(Y, N - Y) ~ ((0 + Se + Sp) | sid) + Se
Data: dat_final
      AIC      BIC    logLik deviance df.resid
851.6825 865.1849 -420.8412  841.6825     105
Random effects:
Groups Name Std.Dev. Corr
sid     Se  0.7116
        Sp  0.8935  -0.38
Number of obs: 110, groups: sid, 55
Fixed Effects:
(Intercept)          Se
      1.12626      -0.05746
```

The indicator variable Se has a value of 1 for Se estimates and 0 for Sp estimates. The random effects are grouped within each study, numbered after `sid`.

The fit reports the Akaike information criterion (AIC), which can be used to compare across models. The logit means of Se and Sp are given by the fixed effects, 1.069 and 1.126, which translate to a sensitivity of 0.744 and a specificity of 0.755. The values, larger than 0.5, suggest that overall, the diagnostic accuracy for bipolar disorder given our meta-data is better than random, albeit nowhere

near perfect. The standard deviations of logit Se and Sp are given by the random effects, as well as the correlation. All of the mentioned parameter estimates can be returned in a list by calling `paramEst`.

We then plug the parameter estimates to get the posterior distributions for the corrected odds ratio given different priors of Se and Sp. We run all 6 different models with the case-control study observations in [Farhi et al. \(2016\)](#), shown in Table 1. The model specifications are shown in the previous subsection.

```
R> params <- paramEst(my.mod)
R> m.1 <- crudeOR(a = 66, N1 = 11782, c = 243, N0 = 57973, chains = 3, iter = 10000)
R> m.2 <- correctedOR(a = 66, N1 = 11782, c = 243, N0 = 57973, prior_list = params,
+ chains = 3, iter = 10000)
R> m.3 <- logitOR(a = 66, N1 = 11782, c = 243, N0 = 57973, prior_list = params,
+ chains = 3, iter = 10000)
R> m.4 <- fixedCorrOR(a = 66, N1 = 11782, c = 243, N0 = 57973, prior_list = params,
+ chains = 3, iter = 10000)
R> m.5 <- randCorrOR(a = 66, N1 = 11782, c = 243, N0 = 57973, prior_list = params,
+ chains = 3, iter = 10000)
R> m.6 <- diffOR(a = 66, N1 = 11782, c = 243, N0 = 57973, mu = c(1.069, 1.069, 1.126, 1.126),
+ s.lg.se0 = 0.712, s.lg.se1 = 0.712, s.lg.sp0 = 0.893, s.lg.sp1 = 0.893,
+ corr.sesp0 = -0.377, corr.sesp1 = -0.377, corr.group = 0, chains = 3,
+ iter = 10000, traceplot = TRUE)

# get summary of model output
R> m.1
```




Figure 1: Traceplot of 3 Markov chains with 10,000 iterations for randomly correlated logit bivariate model.

Each model above is implemented with 3 Markov chains, and each chain consists of 5000 burn-in samples and 10,000 iterations to estimate the parameters (Figure 1). The posterior mean, median and 95% confidence limits of adjusted odds ratio are as below: 1.35 (1.34, 1.01, 1.74), 5.63 (0.85, 0.02, 36.10), 8.61 (1.93, 0.03, 62.27), 9.33 (1.95, 0.03, 62.62), 9.05 (2.00, 0.03, 64.33), 5.23 (0.82, 0.02, 32.64). To obtain and analyze the model output, one can simply call the model variable (e.g., `m.1`). The summary displays the parameters for the model as well as the mean and confidence limits of the adjusted odds ratio (i.e., `ORadj`). One can also specify `traceplot = TRUE` to display a plot of sampled corrected log odds ratio values over iterations, such as in the above `diffOR` method call.

The above example demonstrates the significance of sensitivity and specificity in a case-control study. We can examine that by the ratio of upper to lower 95% posterior interval: $1.74/1.01 = 1.72$, $36.10/0.02 = 1805$, $62.27/0.03 = 2075.67$, $62.62/0.03 = 2087.33$, $64.33/0.03 = 2144.33$ and $32.64/0.02 = 1632$. The greatest jump happens when we assume misclassification in the case-control study, and it only differs slightly with more uncertainties in the model. The increase is especially significant in [Farhi et al. \(2016\)](#) because the estimated mean Se and Sp are around only 0.75, as seen from the GLMM.

In the future, we will consider adding other specifications of priors for sensitivity and specificity to our package, such as beta priors.

```
R> library("ggplot2")
R> g1 <- plotOR(m.1, a = 66, N1 = 11782, c = 243, N0 = 57973, se = 0.744,
+             sp = 0.755, x.max = 3, y.max = 5, binwidth = 0.1) + ggtitle("(i)")
#..... please see supplementary R script for rest of code .....
```

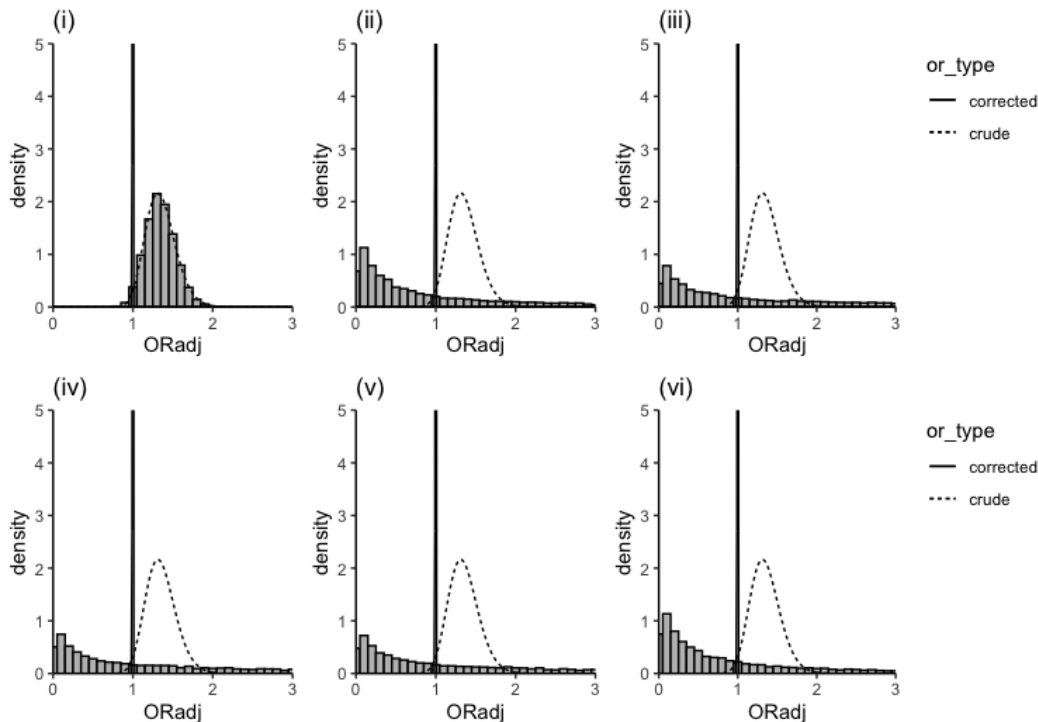


Figure 2: Visualization of posterior distributions of odds ratio for all models. (i) crude (uncorrected) odds ratio with no misclassification; (ii) corrected OR with constant misclassification ($Se = 0.744$ and $Sp = 0.755$); (iii) corrected OR with logit bivariate normal misclassification; (iv) corrected OR that extends from (iii) but with constant correlation between Se and Sp ; (v) corrected OR that extends from (iii) but with Fisher's z -transformed correlation; (vi) corrected OR with differential misclassification. The dotted and solid lines are the probability density lines of crude OR (i) and corrected OR with no misclassification (ii), respectively, assuming log-normality on odds ratio.

We also implement a graphing function, `plotOR`, which takes the input of a model built with one of the above methods, the observations of the same case-control study, and the estimated Se and Sp from the GLMM. The method visualizes the posterior distribution of that model and plots the probability density line of adjusted odds ratio given no misclassification (crude OR) and constant misclassification as specified by Se and Sp (corrected OR). This makes it easy for users to compare the current posterior distribution (especially for models with more uncertainty) with more certain models to visualize the effect of misclassification in a case-control study. In addition, the lines serve as references when comparing across models. The plots and relevant codes are shown in Figure 2. Users can also choose to extract the data from the `rstan` objects by calling functions such as `extract`, `as.data.frame`, etc.

According to the plot, we observe a drastic change to the posterior distribution after taking non-perfect Se and Sp into account. Then, we observe slightly more uniform distributions as there is more uncertainty in the model. What is also worth noting is that in part (ii) of the plot, the posterior density and MCMC sampling do not share the same shape, even though both assume non-perfect constant Se and Sp . This may be a result of low Se and Sp values, which may affect the log-normality assumption in the MCMC posterior samples.

We now show the effects of the number of iterations and chains on the computing speed of our models. All models have been pre-compiled, which reduces the computing time significantly. For example, `randCorrOR`, which is presumably one of the most complex and time-consuming models to compute, takes about 1.32 seconds to run 3 chains with 5000 warm-up periods and 10,000 iterations each. In comparison, it takes about 0.20 seconds to compute 2 chains with 1000 warm-up periods and

2000 iterations each. In practice, a larger number of MCMC chains and iterations leads to more stable and accurate results and thus is recommended. Furthermore, we find that models, such as provided by `randCorrOR`, have smaller target posterior distribution regions in a Markov chain, thus rendering it easy for the algorithm to miss the true distribution and result in “divergent transitions,” which may return biased estimates. Increasing the value of `adapt_delta` parameter up to 1 in the control argument of the methods can effectively make `rstan` take smaller steps to approach the target.

Conclusion

In this article, we introduce and implement the methods for making posterior inferences on the corrected odds ratio by modeling the uncertainty on both differential and non-differential misclassification through appropriate prior distributions. The specific implementation is publicly available using the R package `BayesSenMC`. The process can be divided into two parts. First, one can use the GLMM model with a binomial-logit link to estimate prior information on Se and Sp via a meta-analysis on the misclassification of exposure status. Second, the estimates can be plugged into the modeling functions to provide inferences for the odds ratio. The models can also be visualized side-by-side for better comparisons. The validity of the analyses depends highly on the relevance of meta-analysis, in which irrelevant studies may skew the prior estimates of Se and Sp significantly, and consequentially, the corrected odds ratio. In addition, our models assume normal and independent priors on true exposure probabilities, which may be limited in some cases (Chu et al., 2006).

Bibliography

- D. Bates, M. Maechler, B. Bolker, and S. Walker. **lme4**: *Linear Mixed-Effects Models using ‘Eigen’ and S4*, 2021. URL <https://CRAN.R-project.org/package=lme4>. R package version 1.1-27.1. [p5]
- M. Betancourt. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint arXiv:1701.02434*, 2017. [p6]
- T. B. Brakenhoff, M. Mitroiu, R. H. Keogh, K. G. M. Moons, R. H. H. Groenwold, and M. van Smeden. Measurement error is often neglected in medical literature: a systematic review. *Journal of Clinical Epidemiology*, 98:89–97, 2018. doi: 10.1016/j.jclinepi.2018.02.023. [p1]
- B. Carpenter, A. Gelman, M. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: a probabilistic programming language. *Journal of Statistical Software*, 76(1): 1–32, 2017. doi: 10.18637/jss.v076.i01. [p6]
- A. F. Carvalho, Y. Takwoingi, P. M. G. Sales, J. K. Soczynska, C. A. Köhler, T. H. Freitas, J. Quevedo, T. N. Hyphantis, R. S. McIntyre, and E. Vieta. Screening for bipolar spectrum disorders: a comprehensive meta-analysis of accuracy studies. *Journal of Affective Disorders*, 172:337–346, 2015. doi: 10.1016/j.jad.2014.10.024. [p2, 3]
- H. Chu and S. R. Cole. Bivariate meta-analysis of sensitivity and specificity with sparse data: a generalized linear mixed model approach. *Journal of Clinical Epidemiology*, 59(12):1331–1332, 2006. doi: 10.1016/j.jclinepi.2006.06.011. [p5]
- H. Chu, Z. Wang, S. R. Cole, and S. Greenland. Sensitivity analysis of misclassification: a graphical and a Bayesian approach. *Annals of Epidemiology*, 16(11):834–841, 2006. doi: 10.1016/j.annepidem.2006.04.001. [p1, 2, 9]
- H. Chu, H. Guo, and Y. Zhou. Bivariate random effects meta-analysis of diagnostic studies using generalized linear mixed models. *Medical Decision Making*, 30(4):499–508, 2010. doi: 10.1177/0272989X09353452. [p2, 4]
- S. R. Cole, H. Chu, and S. Greenland. Multiple-imputation for measurement-error correction. *International Journal of Epidemiology*, 35(4):1074–1081, 2006. doi: 10.1093/ije/dyl097. [p1]
- A. Farhi, A. D. Cohen, O. Shovman, D. Comaneshter, H. Amital, and D. Amital. Bipolar disorder associated with rheumatoid arthritis: a case-control study. *Journal of Affective Disorders*, 189:287–289, 2016. doi: 10.1016/j.jad.2015.09.058. [p2, 6, 7]
- M. P. Fox, T. L. Lash, and S. Greenland. A method to automate probabilistic sensitivity analyses of misclassified binary variables. *International Journal of Epidemiology*, 34(6):1370–1376, 2005. doi: 10.1093/ije/dyi184. [p1]

- S. Greenland. Basic methods for sensitivity analysis of biases. *International Journal of Epidemiology*, 25(6):1107–1116, 1996. doi: 10.1093/ije/25.6.1107-a. [p1]
- S. Greenland. Multiple-bias modelling for analysis of observational data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 168(2):267–306, 2005. doi: 10.1111/j.1467-985X.2004.00349.x. [p1]
- P. Gustafson, N. D. Le, and R. Saskin. Case–control analysis with partial knowledge of exposure misclassification probabilities. *Biometrics*, 57(2):598–609, 2006. doi: 10.1111/j.0006-341X.2001.00598.x. [p1, 2]
- D. Haine. **episensr**: *Basic Sensitivity Analysis of Epidemiological Results*, 2021. URL <https://CRAN.R-project.org/package=episensr>. R package version 1.1.0. [p1]
- M. D. Hoffman and A. Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15:1593–1623, 2014. [p6]
- A. M. Jurek, G. Maldonado, S. Greenland, and T. R. Church. Exposure-measurement error is frequently ignored when interpreting epidemiologic study results. *European Journal of Epidemiology*, 21(12):871–876, 2006. doi: 10.1007/s10654-006-9083-0. [p1]
- T. L. Lash and A. K. Flink. Semi-automated sensitivity analysis to assess systematic errors in observational data. *Epidemiology*, 14(4):451–458, 2003. doi: 10.1097/01.EDE.0000071419.41011.cf. [p1]
- X. Ma, L. Nie, S. R. Cole, and H. Chu. Statistical methods for multivariate meta-analysis of diagnostic tests: an overview and tutorial. *Statistical Methods in Medical Research*, 25(4):1596–1619, 2016. doi: 10.1177/0962280213492588. [p4]
- R. F. MacLehose and P. Gustafson. Is probabilistic bias analysis approximately Bayesian? *Epidemiology*, 23(1):151–158, 2012. doi: 10.1097/EDE.0b013e31823b539c. [p1, 2]
- I. B. McInnes and G. Schett. Pathogenetic insights from the treatment of rheumatoid arthritis. *The Lancet*, 389(10086):2328–2337, 2017. doi: 10.1016/S0140-6736(17)31472-1. [p2]
- M. L. Phillips and D. J. Kupfer. Bipolar disorder diagnosis: challenges and future directions. *The Lancet*, 381(9878):1663–1671, 2013. doi: 10.1016/S0140-6736(13)60989-7. [p2]
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2018. URL <http://www.R-project.org/>. [p]
- B. Rosner, W. C. Willett, and D. Spiegelman. Correction of logistic regression relative risk estimates and confidence intervals for systematic within-person measurement error. *Statistics in Medicine*, 8(9):1051–1069, 1989. doi: 10.1002/sim.4780080905. [p1]
- K. J. Rothman, S. Greenland, and T. L. Lash. *Modern Epidemiology*. Lippincott Williams & Wilkins, Philadelphia, PA, 3rd edition, 2008. [p1]
- D. Spiegelman, R. J. Carroll, and V. Kipnis. Efficient regression calibration for logistic regression in main study/internal validation study designs with an imperfect reference instrument. *Statistics in Medicine*, 20(1):139–160, 2001. doi: 10.1002/1097-0258(20010115)20:1<139::AID-SIM644>3.0.CO;2-K. [p1]
- Stan Development Team. **RStan**: *the R Interface to Stan*, 2020. URL <https://CRAN.R-project.org/package=rstan>. R package version 2.21.2. [p6]
- H. Wickham, W. Chang, L. Henry, T. L. Pedersen, K. Takahashi, C. Wilke, K. Woo, H. Yutani, D. Dunnington, and RStudio. **ggplot2**: *Create Elegant Data Visualisations Using the Grammar of Graphics*, 2021. URL <https://CRAN.R-project.org/package=ggplot2>. R package version 3.3.5. [p6]

Jinhui Yang
University of Minnesota Twin Cities
Department of Computer Science and Engineering
200 Union St SE
Minneapolis, MN 55455
yang7004@umn.edu

Lifeng Lin
Florida State University
Department of Statistics
117 N Woodward Ave
Tallahassee, FL 32306
linl@stat.fsu.edu

Haitao Chu
University of Minnesota Twin Cities
School of Public Health
420 Delaware St SE
Minneapolis, MN 55455
chux0051@umn.edu