Published in final edited form as:

Stat Med. 2015 August 30; 34(19): 2695–2707. doi:10.1002/sim.6509.

Detecting outlying trials in network meta-analysis

Jing Zhang^{*}, Haoda Fu[†], and Bradley P. Carlin[‡]

^{*}Department of Epidemiology and Biostatistics, School of Public Health, University of Maryland, College Park, MD 20740 USA.

[†]Eli Lilly and Company, Lilly Corporate Center, Indianapolis, IN 46285 USA.

[‡]Division of Biostatistics, University of Minnesota School of Public Health, Minneapolis, MN 55455 USA.

Summary

Network meta-analysis (NMA) expands the scope of a conventional pairwise meta-analysis to simultaneously handle multiple treatment comparisons. However, some trials may appear to deviate markedly from the others, and thus be inappropriate to be synthesized in the NMA. In addition, the inclusion of these trials in evidence synthesis may lead to bias in estimation. We call such trials *trial-level outliers*. To the best of our knowledge, while heterogeneity and inconsistency in NMA have been extensively discussed and well addressed, few previous papers have considered the proper detection and handling of trial-level outliers. In this paper, we propose several Bayesian outlier detection measures, which are then applied to a diabetes data set. Simulation studies comparing our approaches in both arm- and contrast-based model settings are provided in two supporting appendices.

Keywords

Network meta-analysis; Trial-level outliers; Detection measures

1 Introduction

In clinical practice, and at a wider societal level, treatment decisions need to consider all relevant alternative health care technologies. However, traditional meta-analysis limits informed decision making, because it allows combination of evidence on only two treatments. This limitation can be overcome if *all* randomized clinical trials (RCTs) evaluating interventions relevant to the treatment decision are considered collectively. *Network meta-analysis* (NMA) allows this, expanding the scope of a conventional pairwise meta-analysis to simultaneously handle multiple treatment comparisons. NMA synthesizes both direct and indirect information, leading to more accurate treatment effect estimates. There have been many recent papers exploring hierarchical Bayesian methods for NMA; see for example [1][2][3][4][5][6][7]. A large frequentist literature on NMA is also available; see for instance [8][9][10]. However, some trials may appear to deviate markedly from the

others, and thus be inappropriate to be synthesized. The key NMA assumptions of heterogeneity and inconsistency [11][12][13], as well as approaches for outlier detection for fixed and random/mixed effect models, appear to have been well studied for traditional meta-analysis [14][15][16]. However, few previous authors have considered the issue of trial-level outliers, their detection, and guidance on whether or not to discard them from an NMA.

Our primary objective in this article then is to propose and evaluate Bayesian approaches to detect trial-level outliers in NMA evidence structures. First, vaguely reminiscent of “Cook’s distance”, we propose a measure called *relative distance* (RD), which calculates the relative difference in estimates obtained when leaving the potential outlying trial in, versus taking it out. Another possible measure we explore is the Bayesian *standardized trial residual* (STR), which implements the idea of a Bayesian standardized cross-validatory (“leave one out”) residual [17]. Yet another alternative is the *Bayesian p-value*, which calculates the posterior predictive probability of some discrepancy measure being as extreme as that actually observed for each potentially outlying trial. The last measure we consider models the data using *scale mixtures of normals* (SMN), with a trial identified as an outlier if the posterior estimate of its scale parameter is significantly larger than 1 or if the probability that this estimate is larger than 1 is bigger than some threshold value.

The remainder of our article is organized as follows. We begin by introducing our motivating diabetes drug data set in Section 2, followed by two NMA model parameterizations (arm-based and contrast-based) in Section 3. Then in Section 4 we describe the various outlying trial detection measures, and apply them to the diabetes data in Section 5. We close with a discussion of our findings and suggested avenues for future research in Section 6. Two appendices provide support for use of our approaches: our primary simulations in Appendix A, and supplementary supporting simulations for a broader set of models and NMA network types in Appendix B.

2 Illustrative Diabetes Data

The diabetes network meta-analysis shown in Table 1 comprises efficacy responses over 12 internal industry-sponsored trials of 5 potential diabetes treatments (1: PIO (pioglitazone), 2: Placebo, 3: MET (metformin), 4: SU (sulfonylurea), and 5: ROSI (rosiglitazone)). The major efficacy outcome is the mean change in HbA1c (denoted by “mean” in Table 1), which is a lab measurement indicating the average level of blood sugar (glucose) over the previous 3 months, and is thought of as a measure of how well a patient is controlling his or her diabetes. The columns “*n*” and “*sd*” in Table 1 represent sample size and sample standard deviation, respectively.

Figure 1 is an undirected graph plotted with the *netgraph* function in the *netmeta* R package [18] which elucidates the network of comparative relations for the 5 drugs in our NMA. Every edge has different thickness, proportional to the frequency of each comparison (i.e., the number of studies including this pair of treatments). Six trials (Trials 1–6) compare PIO versus Placebo, 4 trials (Trials 9–12) compare PIO versus SU, 1 trial (Trial 8) compares

MET and PIO, and 1 trial (Trial 7) compares PIO, ROSI, and Placebo. Trial 7 is the only multi-arm trial, and is highlighted with an opaque triangle in the figure.

3 Statistical Models for NMA of Continuous Data

In this section, we present two potential NMA model parameterizations: *arm-based* and *contrast-based*. Arm-based methods model a response for each treatment arm, while contrast-based methods focus on merely the relative change from some baseline treatment.

Arm-based method

Among other authors, Zhang et al. [6] have proposed an arm-based NMA model for binary data. Here we modify their model to adapt to the Gaussian diabetes data as follows:

$$y_{ik} \sim N\left(\Delta_{ik}, \frac{\sigma^2}{n_{ik}}\right)$$

$$Sd_{ik}^2 \sim \text{Gamma}\left(\frac{n_{ik}-1}{2}, \frac{n_{ik}-1}{2\sigma^2}\right) \quad (1)$$

$$\text{and } \Delta_{ik} = \mu_k + \gamma\nu_{ik},$$

where the observations y_{ik} (most often thought of as group means $\bar{y}_{ik} = \frac{1}{n_{ik}} \sum_{j=1}^{n_{ik}} y_{ikj}$) represent the mean response change of HbA1c over n_{ik} patients assigned to the k^{th} treatment in the i^{th}

trial. This response is assumed to have a normal distribution with mean Δ_{ik} and variance $\frac{\sigma^2}{n_{ik}}$, where Δ_{ik} and σ^2 are the true population mean and variance. Sd_{ik}^2 is the group-level sample

variance, which then has a *Gamma* distribution with scale parameter $\frac{n_{ik}-1}{2}$ and rate

parameter $\frac{n_{ik}-1}{2\sigma^2}$. Finally, μ_k is the treatment-specific fixed effect and γ is the standard deviation of ν_{ik} , implemented via random effects ν_{ik} independently. Again, this model is a homogeneous-variance. Since γ is the same across k , (1) is a homogeneous-variance model. If we instead use γ_k , we obtain a heterogeneous-variance model.

Note that in (1), $Sd_{ik}^2 = \frac{1}{n_{ik}-1} \sum_{j=1}^{n_{ik}} (y_{ikj} - y_{ik})^2$ is the sample variance of n_{ik} observations, where j represents the subject. Thus by Basu's Theorem, y_{ik} and Sd_{ik}^2 are statistically independent, and $\frac{(n_{ik}-1)Sd_{ik}^2}{\sigma^2} \sim \chi^2(n_{ik}-1) \equiv \text{Gamma}\left(\frac{n_{ik}-1}{2}, \frac{1}{2}\right)$. This in turn implies that $Sd_{ik}^2 \sim \text{Gamma}\left(\frac{n_{ik}-1}{2}, \frac{n_{ik}-1}{2\sigma^2}\right)$. In some previous work [19][20], y_{ik} is

assumed to have a known and constant variance, leading to a somewhat more standard inverse gamma formulation.

Contrast-based method

Following Spiegelhalter et al. [19] and Ding and Fu [20], a contrast-based network meta-analysis model for Gaussian data can be written as $y_{ik} = \mu_i + X_{ik}\delta_{ibk}$ with $\delta_{ibk} \stackrel{\text{ind}}{\sim} N(d_{bk}, \varepsilon^2)$. Here, X_{ik} is an indicator taking value 0 when $k = b$ and value 1 when $k \neq b$, μ_i is the baseline mean response for the i^{th} trial with a weakly informative $N(0, 100)$ prior, and δ_{ibk} measures the effect of treatment k relative to the baseline treatment b , which is permitted to change across studies. Note that when $X_{ik} = 0$, $y_{ik} = \mu_i$ represents the response in the baseline group in the i^{th} trial. Thus d_{bk} is the mean contrast effect of treatment k versus b , and ε^2 is the

variance. Here y_{ik} and Sd_{ik}^2 have the same distribution as in (1), i.e., $y_{ik} \sim N(\Delta_{ik}, \frac{\sigma^2}{n_{ik}})$ and $Sd_{ik}^2 \sim \text{Gamma}(\frac{n_{ik} - 1}{2}, \frac{n_{ik} - 1}{2\sigma^2})$. Again, this model is a homogeneous-variance model, while $\delta_{ibk} \sim N(d_{bk}, \varepsilon_{bk}^2)$ corresponds to a heterogeneous-variance model.

Arm-based methods are not new to network meta-analysis, and have been previously used for conventional pairwise meta-analysis (e.g. Chu et al. [21]) as well. However, many researchers dislike arm-based methods, arguing that they “break randomization” since they assume that not just the treatment effects (differences or odds ratios) are exchangeable across studies, but the actual event rates are as well, potentially allowing the control rate in one study to influence estimation of the treatment effect in another. While we agree that arm-based models do assume more, when appropriate they do offer advantages in ease of interpretation, prior specification, and model fitting. Moreover, contrast-based models’ assumption of exchangeability of effects relative to an arbitrary baseline that is usually changing across trials remains a lot to assume. Suppose for example an NMA of three treatments, the first trial compares treatments A and B, and the second trial compares treatments A, B and C. Contrast-based methods assume that the contrasts between A and B across trials are exchangeable, but this may influence the estimation of the A–C contrast in the second trial and thus have undesirable effects as well.

Since comparison between arm-based and contrast-based methods have been discussed elsewhere [5][6][7], in the main part of this paper, we will focus only on the arm-based method. However, we emphasize our methods could apply equally well to the contrast-based setting (though they might well identify different outlying studies in that case). Section 6 includes further discussion of arm- versus contrast-based approaches in our context, while Appendix B offers a simulation study confirming that several of our methods can actually perform similarly under both model frameworks.

4 Outlier Detection Measures

4.1 Relative Distance

We define cross-validators (or “leave one out”) *relative distance* statistics, RD_i^k , to measure the effect of deleting trial i on our NMA estimate for a particular treatment k as

$$RD_i^k = \left| \frac{\hat{\eta}_k - \hat{\eta}_{k(i)}}{\hat{\eta}_k} \right|, \quad (2)$$

where $\hat{\eta}_k$ is some estimate of interest (e.g., the posterior mean treatment effect) for treatment k from the full data, and $\hat{\eta}_{k(i)}$ is the estimate for treatment k from the data where trial i has been omitted. The bigger RD_i^k is, the higher the relative effect of deleting trial i is, and thus the higher the likelihood that trial i is influential and may be a “trial-level outlier” in this sense. We can also define an *average relative distance* (ARD) to measure the *average* effect of deleting trial i as:

$$ARD_i = \frac{1}{K} \sum_{k=1}^K \left| \frac{\hat{\eta}_k - \hat{\eta}_{k(i)}}{\hat{\eta}_k} \right|, \quad (3)$$

where $\hat{\eta}_k$ and $\hat{\eta}_{k(i)}$ have the same representations as in (2). The bigger the ARD_i is, the greater the average effect of deleting trial i is. We may define trial i as an outlier if RD_i^k or ARD_i is large relative to the full collection of RD or ARD values, respectively. Formal “probabilities of being an outlier” could also be computed; say, $P(RD_i^k > T | \text{data})$ for some preselected threshold T , say $T = 0.1$. Note that the selection of T depends on the particular data and context.

4.2 Standardized Trial Residuals

Again, based on cross-validatory thinking, we might calculate the fitted value for y_{ik} by conditioning on all data except $\mathbf{y}_i$, namely, $\mathbf{y}_{(i)} = (\mathbf{y}_1, \dots, \mathbf{y}_{i-1}, \mathbf{y}_{i+1}, \dots, \mathbf{y}_I)'$, where $\mathbf{y}_i = \{y_{i,k} : k \in S_i\}$ and S_i represents the set of treatments compared in trial i . We then compute the difference between the observed and fitted values for y_{ik} and standardize it as follows:

$$STR_i^k = \frac{y_{ik} - E(y_{ik} | \mathbf{y}_{(i)})}{\sqrt{Var(y_{ik} | \mathbf{y}_{(i)})}}, \quad (4)$$

where STR_i^k stands for the Bayesian *standardized trial residual* for the k^{th} treatment in the i^{th} trial. The *average absolute standardized trial residual* (ASTR) can be defined correspondingly as:

$$ASTR_i = \frac{1}{n_{S_i}} \sum_{k \in S_i} \left| \frac{y_{ik} - E(y_{ik} | \mathbf{y}_{(i)})}{\sqrt{Var(y_{ik} | \mathbf{y}_{(i)})}} \right|, \quad (5)$$

where S_i has the same representation as before, and n_{S_i} represents its cardinality. Note that here we index by $k \in S_i$ instead of $k = 1, \dots, K$ because each trial compares a subset of treatments of interest, and only arms $k \in S_i$ have observed y_{ik} values. Large STR_i^k and $ASTR_i$, say larger than 1.5, suggest observation y_{ik} and trial i may be outliers respectively. Note that in formulas (4) and (5), we compute the posterior mean and variance with respect to the *conditional predictive distribution*,

$$f(y_{ik}|\mathbf{y}_{(i)}) = \frac{f(\mathbf{y})}{f(\mathbf{y}_{(i)})} = \int f(y_{ik}|\boldsymbol{\theta}, \mathbf{y}_{(i)})p(\boldsymbol{\theta}|\mathbf{y}_{(i)})d\boldsymbol{\theta},$$

where $\boldsymbol{\theta}$ is the entire parameter collection. In other words, $f(y_{ik}|\mathbf{y}_{(i)})$ is the posterior predictive density of y_{ik} given the remainder of the data except that concerning trial i .

4.3 Bayesian p-value

An alternative to cross-validatory approaches as described in Sections 4.1 and 4.2 is the use of *posterior predictive* model checks, an approach initially promoted by Rubin [22] and popularized by Gelman et al. [23]. The key idea is to construct some “discrepancy measures” that capture departures of the observed data from the assumed model (likelihood and prior distribution). Note that though such measures must be functions of observed data alone in the classical frequentist framework, Bayesian model checking based on posterior predictive distributions allows more general measures that depend on both data and parameters. Gelman et al. [24] suggest an omnibus goodness of fit discrepancy measure

$D_i^k(y_{ik}, \boldsymbol{\theta})$ that depends on the parameters $\boldsymbol{\theta}$ and the data y_{ik} ,

$$D_i^k = \frac{[y_{ik} - E(Y_{ik}|\boldsymbol{\theta})]^2}{\text{Var}(Y_{ik}|\boldsymbol{\theta})}.$$

We subsequently define an *average discrepancy* measure $AD_i(\mathbf{y}_i, \boldsymbol{\theta})$ as

$$AD_i = \sum_{k \in S_i} \frac{[y_{ik} - E(Y_{ik}|\boldsymbol{\theta})]^2}{\text{Var}(Y_{ik}|\boldsymbol{\theta})},$$

where S_i is the set of treatments that are compared in trial i , and $E(Y_{ik}|\boldsymbol{\theta})$ is the expectation of y_{ik} under the posited model. We now can compare the distribution of $D_i^k(y_{ik}, \boldsymbol{\theta})$ and $AD_i(\mathbf{y}_i, \boldsymbol{\theta})$ for the observed data y_{ik} and $\mathbf{y}_i$ with that of $D_i^k(y_{ik}^*, \boldsymbol{\theta})$ and $AD_i(\mathbf{y}_i^*, \boldsymbol{\theta})$ for hypothetical future values y_{ik}^* and $\mathbf{y}_i^*$. Note that y_{ik}^* and $\mathbf{y}_i^*$ are defined as another “copy” of the observed data point y_{ik} and vector $\mathbf{y}_i$, which are not observed but instead generated from their posterior predictive distributions as part of the MCMC sampling order [17]. D_i^k and AD_i^k computed using the observed data that are extreme relative to this reference distribution indicate poor model fit and merit closer examination in the analysis.

A convenient summary measure of the extremeness of the $D_i^k(y_{ik}^*, \boldsymbol{\theta})$ with respect to the $D_i^k(y_{ik}, \boldsymbol{\theta})$ is the posterior predictive tail area, defined as the *Bayesian p-value* for discrepancy,

$$p_{D_i^k} \equiv P[D_i^k(y_{ik}^*, \boldsymbol{\theta}) > D_i^k(y_{ik}, \boldsymbol{\theta}) | \mathbf{y}] = \int P[D_i^k(y_{ik}^*, \boldsymbol{\theta}) > D_i^k(y_{ik}, \boldsymbol{\theta}) | \boldsymbol{\theta}] p(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}. \quad (6)$$

Similarly, the Bayesian p-value for average discrepancy is defined as

$$p_{AD_i} \equiv P[AD_i(\mathbf{y}_i^*, \boldsymbol{\theta}) > AD_i(\mathbf{y}_i, \boldsymbol{\theta}) | \mathbf{y}] = \int P[AD_i(\mathbf{y}_i^*, \boldsymbol{\theta}) > AD_i(\mathbf{y}_i, \boldsymbol{\theta}) | \boldsymbol{\theta}] p(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}. \quad (7)$$

Note that $p_{D_i^k}$ and p_{AD_i} should not be used to compare models, since while they do offer an indirect model comparison, they fail to obey the Likelihood Principle. Thus, strictly speaking, they serve only as measures of discrepancy between the proposed model and the observed data, and therefore provide information concerning overall model adequacy and outlier detection. Other summaries focused on other measures of poor fit (say, in the tail of the distribution) can also be defined; see Gelman, Meng, and Stern [24].

4.4 Scale Mixtures of Normals

The conditioning feature of MCMC computational methods enables another approach related to models employing *scale mixtures of normals* (SMN) (see [17], p.184) to investigate outlyingness. Here we expand model (1) to

$$y_{ik} \sim N(\Delta_{ik}, \lambda_i \frac{\sigma^2}{n_{ik}}), \quad (8)$$

where the λ_i are unknown scale parameters. We then specify prior distributions for λ_i , for example,

$$\text{SMN}_1: \lambda_i \equiv 1 \Rightarrow \text{Normal errors}$$

$$\text{SMN}_2: \lambda_i \sim IG\left(\frac{v}{2}, \frac{2}{v}\right) \Rightarrow \text{students' } t_v \text{ errors}$$

$$\text{and SMN}_3: \lambda_i \sim \text{Expo}(2) \Rightarrow \text{Double exponential errors,}$$

where the distributions following the arrow symbols identify the possible departures from normality for the error terms. Since extreme observations will correspond to extreme fitted values of these scale parameters λ_i , potential outliers can be identified by examining the λ_i posterior distributions. Doubt is cast on the commensurability of trial i with the rest if the posterior mean (or median) of λ_i is much bigger than 1 (i.e., the error distribution is further from normality), or if $P(\lambda_i > 1 | \mathbf{y})$ is larger than some threshold value, say 0.95.

5 Application to Diabetes Data

5.1 Relative Distance

We first fit model (1) to our diabetes NMA data and record the posterior estimates. Then we fit model (1) 12 more times with the 1st through 12th trials omitted, respectively, and record all necessary posterior estimates. Finally, we calculated RD_i^k according to (2) with $\hat{\mu}_k$ obtained from the full data and $\hat{\mu}_{k(i)}$ obtained from the data with the i^{th} trial deleted. Note

that here we let $\hat{\eta}_k = \hat{\mu}_k$, the mean treatment effect, when we calculate RD_i^k . ARD_i is calculated similarly.

Figure 2 shows the relative distances separately for the 5 treatments (PIO, Placebo, MET, SU, and ROSI). The vertical axes show relative distances ranging from 0.0 to 1.0, and the horizontal axes index trials that are deleted in the calculation of the relative distances. Using 0.2 as a significance threshold since most relative distances are no bigger than this, Trial 6 is mildly influential for PIO; Trials 2, 4, 6 are influential for Placebo; Trial 8 is influential for MET; Trials 6 and 9 are mildly influential for SU; and Trials 2, 4, 6, and 7 are influential for ROSI. Note that Trials 2, 4, and 6 are influential for ROSI even though they do not directly compare this treatment. In short, with 0.2 as the cutoff for RD_i^k , Trials 2, 4, 6, 7, and 8 seem to be potential outliers.

Figure 3 shows the average relative distances. ARD_i for $i = 6, 7$, and 8 are above threshold $T = 0.1$, while those for the others are below this level (though some only narrowly). Thus Trials 6, 7, and 8 are more influential with ARD_i as the evaluation criteria. This is roughly consistent with the results in Figure 2.

We further investigate the possibility that Trials 2, 4, 6, 7, and 8 are outliers. First, it seems fair to call Trials 2 and 4 outliers since the mean change values of HbA1c for patients treated with Placebo are all negative except in these two trials. Second, for Trials 6 and 7, the mean change values of HbA1c for patients treated with PIO from both trials are much smaller than those from the other 10 trials. Thus Trials 6 and 7 are extreme in this sense. Third, the mean HbA1c change responses for patients in Trial 8 do not seem abnormal; however, we observe that MET is only contained in Trial 8, so that deleting Trial 8 will of course have a big impact on the estimates. We thus can infer that RD_8^3 and ARD_8 for Trial 8 are big largely due to lack of information, rather than true “outlyingness”.

5.2 Standardized Trial Residuals

Table 2 shows that the Bayesian standardized trial residuals for Trials 2, 4, and 6 are larger than 1.5 in absolute value (more specifically, $STR_2^2 = -2.24$, $STR_4^2 = -2.29$, and $STR_6^1 = -1.58$). In Trials 2 and 4, the mean changes of HbA1c for patients taking placebo are positive, while those from the other trials that contain placebo ($k = 2$) are negative. Thus it seems reasonable to call Trial 2 and Trial 4 outliers. For STR_6^1 , the Trial 6 Bayesian standardized residual for PIO ($k = 1$), we see that the mean change of HbA1c for patients taking PIO in this trial is -0.17 , much smaller than that from the other trials. Thus Trial 6 would appear to highly underestimate the efficacy of PIO, i.e., it may also be a legitimate outlier. However, ASTR in Table 2 does not identify any significant trial-level outliers, with no values larger than 1.5 and previously unidentified Trial 9 emerging along with Trials 2, 4, and 6 with $ASTR > 1$.

5.3 Bayesian p -values

Table 3 shows that Trials 2, 4, 5, 6, and 7 have at least one Bayesian p -value smaller than 0.05. In the case of Trials 2 and 4, this is likely due to the positive responses for placebo in

these trials. By contrast, Trials 6 and 7 are likely flagged because the PIO mean changes in HbA1c values in these two trials are much smaller than that from the other trials, i.e., the presence of Trials 6 and 7 in the NMA would underestimate the efficacy of PIO. Oddly, the Bayesian p -values for Trial 5 are also smaller than 0.05, an apparent significance that may be inflated by small variances and merits further investigation. At any rate, it suggests the omnibus goodness of fit measure adopted in Section 4.3 may not be optimal in this particular setting.

5.4 Scale Mixtures of Normals

Figure 4 shows that Trials 2, 4, 6, and 7 once again emerge as outliers under both models SMN2 and SMN3, since the posterior estimates for the scale parameters λ_i are significantly larger than 1 (with their log-scale 95% CIs significantly higher than 0). The specific estimated values of λ_i are listed in Table 4. In addition, Table 4 shows that the probabilities that the scale parameters λ_i are larger than 1 in Trials 2, 4, 6, and 7, are all 0.99 or greater, which further suggests the outlyingness of these trials, in broad agreement with the results of the previous detection approaches. Note that we used an Inverse Gamma (1, 1) prior for SMN₂ and an Exponential (2) prior for SMN₃.

5.5 Results With and Without Outliers

We compare the estimated values for parameters of interest from full data with those computed without our identified trial-level outliers (Trials 2, 4, 6, and 7). A value for μ_5 is not computable without outliers since Trial 7, an outlier, is the only source of information on Treatment 5 (ROSI). As shown in Table 5, μ_k estimates are quite different before and after deleting the outliers. For example, in the case of μ_2 , the relative difference is

$\frac{-0.67 - (-0.47)}{-0.47} = 42.6\%$. Relative changes for the other μ 's are similarly meaningful, though are less impressive for the variance γ parameters and σ^2 . In a nutshell, when trial-level outliers exist in an NMA, they can wield significant influence on estimates of the parameters of interest.

6 Discussion and Future Work

Though methods for network meta-analysis have been extensively discussed and explored in the current literature, few previous papers appear to have mentioned trial-level outliers. In this paper, we proposed four detection measures, including RD (ARD), STR (ASTR), Bayesian p -values and SMN, for trial-level outliers in network meta-analysis, and applied them to a diabetes data set (with performance extensively explored via simulation studies in Appendices A and B). Our results suggest RD (ARD), STR (ASTR), and SMN perform well and are promising tools for detecting trial-level outliers. Though we focused on network meta-analysis in this paper, our detection measures apply equally well to pairwise evidence synthesis, and can thus complement existing detection measures for traditional meta-analysis [14][15][16].

Our detection measures can also be extended to binary data. For example, instead of (1) for Gaussian data, Zhang et al. [6] proposed an arm-based method for binary data wherein

$\Phi^{-1}(p_{ik}) = \mu_k + \gamma v_{ik}$, where p_{ik} is the event rate for the k^{th} treatment in the i^{th} trial, and μ_k , γ , and v_{ik} have the same representations as (1). In this binary data setting, RD (ARD) and STR (ASTR) can be defined similarly as in the continuous data setting. The implementation of SMN for binary data relies on rewriting (1) as

$$Y_{ik} \sim \text{Bin}(n_{ik}, p_{ik}), k \in S_i, i=1, \dots, I$$

$$\text{where } p_{ik} = P(Y_{ik}^* > 0)$$

$$\text{and } Y_{ik}^* = \mu_k + \gamma v_{ik} + \varepsilon_{ik},$$

where the Y_{ik}^* are latent variables and $\varepsilon_{ik} \stackrel{iid}{\sim} N(0, 1)$; see Albert and Chib [25]. Using this formulation, we can adapt the SMN method in Section 4.4 accordingly.

We acknowledge that cutoff values for detecting the practical significance of a potential outlier appear hard to select, as they can be context-, model-, and data-specific. We have done simulations (not shown in this paper) to investigate different cutoff values for the SMN method and found that the $P(\lambda_i > c | y)$ varies with different predetermined values for the cutoff c . For example, in the unbalanced design setting of Appendix A, with $c = 1$ as the cutoff criterion, $P(\lambda_4 > 1 | y)$ is 1.00 for Trial 4 but around 0.45 for the other trials, as shown in Table A.4. When the cutoff value is set to be 5, however, $P(\lambda_4 > 5 | y)$ keeps being 1 for Trial 4, but $P(\lambda_i > 5 | y)$ is around 0.06 for the other trials. If we let the cutoff value be 10, the $P(\lambda_4 > 10 | y)$ is again 1 for Trial 4 but around 0.02 for the other trials. The cutoff selection issue also plagues the other measures. However, simulations like those in our appendices can still help guide this selection, based on the design of interest and information content of each trial (as measured by sample sizes and variance estimates).

There are four important issues worthy of discussion. The first involves the selection of detection measures. This article has only concentrated on four commonly used standard Bayesian outlier detection methods in an attempt to illustrate their adaptation to the network meta-analysis setting. Several previous papers have investigated outlier detection in evidence synthesis, but primarily focused on traditional pairwise meta-analysis [16], e.g., using cross-validators p -values. The basic idea is similar to our cross-validators measures, including relative distance and standardized trial residuals. Mixed predictive p -values, as approximations to the cross-validation p -values, have also been proposed for outlier detection [16][26][27]. The idea is to fit the model with *all* the observations, then make predictions for each observation that are compared with the observed values to form a p -value. This is computationally less expensive than the cross-validators p -value, and is similar conceptually to our Bayesian p -value. However, this mixed predictive p -value emerges as very conservative, while our Bayesian p -value is often too eager to conclude lack of fit (c.f. Appendix A). A small improvement, proposed by Welton et al. [16], is to predict

the true treatment effect in each study using the predictive distribution for a “new” study. This approach is still conservative, though less so.

Second, the detection of outlying trials is intimately related to inconsistency checking and trial inclusion criteria. Inconsistency refers to disagreement between direct and indirect evidence, and has been investigated by several authors, almost exclusively in the contrast-based model setting. For example, Jackson et al. [11] developed a random-effects implementation of the design-by-treatment interaction model to handle inconsistency. Higgins et al. [13] introduced the distinction between “loop inconsistency” and “design inconsistency”, and used design-by-treatment interaction terms to investigate inconsistency. White et al. [12] proposed two frequentist approaches for estimating consistency and inconsistency models by expressing them as multivariate random-effects meta-regressions. In our context, outlying trials can be the primary source of inconsistency, in which case it is hard to differentiate outlyingness and inconsistency. On the other hand, deleting outlying trials may affect inconsistency checking due to lack of information, especially in sparse evidence structures with just one or two trials per comparison [26][28]. Outlier detection is in turn related to the inclusion criteria for trials in the network. To be good criteria, study design, the type of interventions, outcome measures and instruments, and language should be explicitly described; participants should also be defined explicitly in terms of age, gender, race, duration of symptoms, localization of symptoms, and type of symptoms [29]. More stringent and proper inclusion criteria may lower the number of outlying trials in an NMA, while blurry or incomplete criteria may suggest more outlyingness.

The third issue is related to multiple hypothesis tests for the Bayesian p -value. We are in danger of incorrectly interpreting the significance of these multiple tests, because we calculate several p -values on the same data. Under the null (perfect fit), we would expect these to follow a Uniform (0, 1) distribution [16][26]. In a plot of ordered p -values, say those in Table 3, against the relevant Uniform order statistics, falling far from the line of equality would indicate outlyingness.

Last but not the least, detection of outlyingness becomes difficult when there is not enough information. Under the CB framework, it is hard to implement our detection measures when there is only one trial making a particular comparison and no further indirect evidence to inform that treatment effect. This issue also plagues the AB modeling framework when a particular treatment is compared in only one trial, which is why we do not think Trial 8 in the diabetes data analysis is an outlier for Treatment 3 (MET) even though RD_8^3 is large.

Turning to future work, we are interested in developing methods for automatic downweighting of outlying trials, building on previous work in bias adjustment in pairwise and network meta-analysis. For instance, Turner et al. [30] use elicited priors to correct for biases and downweight less rigorous or relevant studies. Subsequent hierarchical refinements of this approach [31] uses empirically based prior distributions to adjust and downweight studies that are deemed to be at high risk of bias. A third approach [32] estimates the probability that a particular study is biased, and produces bias-adjusted estimates of treatment effects, under the NMA framework. Alternatively, borrowing the idea of Ibrahim and Chen [33], power priors offer a simple and intuitive approach, by raising the

outlying likelihood to a power $\alpha_0 \in [0, 1]$, and re-standardizing the result to a proper distribution. Hobbs et al. [34] proposed an extension called hierarchical commensurate and power priors for adaptive incorporation of information, which could also be applied here. Future work also looks toward extension to outlier detection for models incorporating baseline covariates and individual-level patient data. Note that when baseline covariates are present, the definition of outliers ought to be modified, since a trial could then be outlying simply by having an unusual population (e.g., more older enrollees).

Of course, criticisms of our methods can be made. In this paper, we have only considered the trials that were already collected in an NMA. However, trials that were candidates for inclusion but were omitted is another issue worthy of attention. This can be broadly related to the issue of publication bias, the concern that studies with significant results are more likely to be published, and published studies (especially those in the meta-analyst's own language) are more likely to be included in an NMA [35]. Another limitation is that our focus here has been the detection of outlyingness under the Bayesian framework, but there is a large frequentist literature on the subject [8][9][10]. For instance, Senn et al. [8] proposed frequentist statistical algorithms using the technique of Aitken estimator [36], and argued that it was inappropriate to treat the main effects of trial as random, and that the generalization from a classic random-effect meta-analysis to a network meta-analysis involved strong assumptions about the variance components. Our detection measures RD (ARD) will still work in a frequentist setting, but STR (ASTR), Bayesian p -values and SMN are inapplicable. Finally, we have not considered *arm-level* outliers, even though it may be that some treatment arms may not be suitable to be synthesized with the others in an NMA. Approaches for these issues and their evaluation await further exploration.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgment

This third author was supported in part by NCI grant 1R01-CA157458-01A1, while all three authors were supported in part by a grant from the Lilly Research Award Program (LRAP).

References

1. Lu G, Ades AE. Combination of direct and indirect evidence in mixed treatment comparisons. *Statistics in Medicine*. 2004; 23(20):3105–3124. [PubMed: 15449338]
2. Lu G, Ades AE. Assessing evidence inconsistency in mixed treatment comparisons. *JASA*. 2006; 101(474):447–459.
3. Lu G, Ades AE, Sutton AJ, Cooper NJ, Briggs AH, Caldwell DM. Meta-analysis of mixed treatment comparisons at multiple follow-up times. *Statistics in Medicine*. 2007; 26(20):3681–3699. [PubMed: 17285571]
4. Lu G, Ades AE. Modeling between-trial variance structure in mixed treatment comparisons. *Biostatistics*. 2009; 10(4):792–805. [PubMed: 19687150]
5. Hong H, Chu H, Zhang J, Carlin BP. A Bayesian missing data framework for generalized multiple outcome mixed treatment comparisons. Research-Report 2012-018, Division of Biostatistics, University of Minnesota, Submitted to Research Synthesis Methods. 2015

6. Zhang J, Carlin BP, Neaton JD, Soon GG, Nie L, Kane R, Virnig BA, Chu H. Network meta-analysis of randomized clinical trials: Reporting the proper summaries. *Clinical Trials*. 2014; 11(2): 246–262. [PubMed: 24096635]
7. Zhang J, Chu H, Hong H, Neaton JD, Virnig BA, Carlin BP. Bayesian hierarchical models for network meta-analysis incorporating nonignorable missingness. Research Report 2013–018, Division of Biostatistics, University of Minnesota. 2015
8. Senn S, Gavini F, Magrez D, Scheen A. Issues in performing a network meta-analysis. *Statistical Methods in Medical Research*. 2013; 22(2):169–189. [PubMed: 22218368]
9. Lu G, Welton NJ, Higgins J, White IR, Ades AE. Linear inference for mixed treatment comparison meta-analysis: A two-stage approach. *Research Synthesis Methods*. 2011; 2(1):43–60. [PubMed: 26061599]
10. Piepho HP, Williams ER, Madden LV. The use of two-way linear mixed models in multitreatment meta-analysis. *Biometrics*. 2012; 68(4):1269–1277. [PubMed: 22845838]
11. Jackson D, Barrett JK, Rice S, White IR, Higgins JPT. A design-by-treatment interaction model for network meta-analysis with random inconsistency effects. *Statistics in medicine*. 2014; 33(21): 3639–3654. [PubMed: 24777711]
12. White IR, Barrett JK, Jackson D, Higgins J. Consistency and inconsistency in network meta-analysis: model estimation using multivariate meta-regression. *Research Synthesis Methods*. 2012; 3(2):111–125. [PubMed: 26062085]
13. Higgins JPT, Jackson D, Barrett JK, Lu G, Ades AE, White IR. Consistency and inconsistency in network meta-analysis: concepts and models for multi-arm studies. *Research Synthesis Methods*. 2012; 3(2):98–110. [PubMed: 26062084]
14. Hedges, LV.; Olkin, L. *Statistical Methods for Meta-Analysis*. New York: Academic Press; 1985.
15. Viechtbauer W, Cheung MWL. Outlier and influence diagnostics for meta-analysis. *Research Synthesis Methods*. 2010; 1(2):112–125. [PubMed: 26061377]
16. Welton, NJ.; Sutton, AJ.; Cooper, NJ.; Abrams, KR.; Ades, AE. *Evidence synthesis for decision making in healthcare*. West Sussex: John Wiley and Sons; 2012.
17. Carlin, BP.; Louis, TA. *Bayesian Methods for Data Analysis*. 3rd edn.. Boca Raton: Chapman and Hall/CRC; 2009.
18. Rücker G, Schwarzer G, Krahn U. Package ‘netmeta’: network meta-analysis with R. R package version 0.4-2. 2014
19. Spiegelhalter, DJ.; Abrams, KR.; Myles, JP. *Bayesian approaches to clinical trials and health-care evaluation*. Chichester: John Wiley & Sons; 2004.
20. Ding Y, Fu H. Bayesian indirect and mixed treatment comparisons across longitudinal time points. *Statistics in Medicine*. 2012; 32(15):2613–2628. [PubMed: 23229717]
21. Chu H, Nie L, Chen Y, Huang Y, Sun W. Bivariate random effects models for meta-analysis of comparative studies with binary outcomes: Methods for the absolute risk difference and relative risk. *Statistical Methods in Medical Research*. 2012; 21(6):621–633. [PubMed: 21177306]
22. Rubin DB. Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics*. 1984; 12(4):1151–1172.
23. Gelman, A.; Carlin, JB.; Stern, HS.; Dunson, DB.; Vehtari, A.; Rubin, DB. *Bayesian Data Analysis*. 3rd edn.. Boca Raton: Chapman and Hall/CRC; 2013.
24. Gelman A, Meng XL, Stern HS. Posterior predictive assessment of model fitness via realized discrepancies (with discussion). *Statistica Sinica*. 1996; 6(4):733–807.
25. Albert JH, Chib S. Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association*. 1993; 88(422):669–679.
26. Madan J, Stevenson MD, Cooper KL, Ades AE, Whyte S, Akehurst R. Consistency between direct and indirect trial evidence: is direct evidence always more reliable? *Value in Health*. 2011; 14(6): 953–960. [PubMed: 21914518]
27. Marshall EC, Spiegelhalter DJ. Approximate cross-validators predictive checks in disease mapping models. *Statistics in medicine*. 2003; 22(10):1649–1660. [PubMed: 12720302]

28. Ohlssen D, Price K, Xia H, Hong H, Kerman J, Fu H, Quartey G, Heilmean C, Ma H, Carlin B. Guidance on the implementation and reporting of a drug safety bayesian network meta-analysis. *Pharmaceu. Statist.* 2014; 13:55–70.
29. van Tulder MW, Assendelft WJJ, Koes BW, Bouter LM, et al. Method guidelines for systematic reviews in the Cochrane Collaboration Back Review Group for spinal disorders. *Spine.* 1997; 22(20):2323–2330. [PubMed: 9355211]
30. Turner RM, Spiegelhalter DJ, Smith G, Thompson SG. Bias modelling in evidence synthesis. *Journal of the Royal Statistical Society: Series A (Statistics in Society).* 2009; 172(1):21–47.
31. Welton NJ, Ades AE, Carlin JB, Altman DG, Sterne JAC. Models for potentially biased evidence in meta-analysis using empirically based priors. *Journal of the Royal Statistical Society: Series A (Statistics in Society).* 2009; 172(1):119–136.
32. Dias S, Welton NJ, Caldwell DM, Ades AE. Checking consistency in mixed treatment comparison meta-analysis. *Statistics in medicine.* 2010; 29(7–8):932–944. [PubMed: 20213715]
33. Ibrahim JG, Chen MH. Power prior distributions for regression models. *Statistical Science.* 2000; 15(1):46–60.
34. Hobbs BP, Carlin BP, Mandrekar SJ, Sargent DJ. Hierarchical commensurate and power prior models for adaptive incorporation of historical information in clinical trials. *Biometrics.* 2011; 67(3):1047–1056. [PubMed: 21361892]
35. Borenstein, M.; Hedges, LV.; Higgins, JPT.; Rothstein, HR. *Introduction to Meta-Analysis.* West Sussex: John Wiley and Sons; 2011.
36. Aitken AC. On least squares and linear combination of observations. *Proceedings of the Royal Society of Edinburgh.* 1936; 55:42–48.

Figure 1.

Graphical representation for the network of the diabetes dataset. The thickness of each link is proportional to the number of trials investigating the relation. The opaque triangle highlights the multi-arm trial (Trial 7).

Figure 2.
Relative distances versus deleted trials for each treatment.

Figure 3.
Average relative distances versus deleted trials.

Figure 4.
Posterior λ_i in log scale for SMN₂ and SMN₃.

Table 1
Diabetes dataset. *n* denotes sample size; mean denotes sample mean; sd denotes standard sample deviation.

Trials	1 (PIO)			2 (Placebo)			3 (MET)			4 (SU)			5 (ROSI)		
	<i>n</i>	mean	sd	<i>n</i>	mean	sd	<i>n</i>	mean	sd	<i>n</i>	mean	sd	<i>n</i>	mean	sd
1	103	-0.76	0.97	115	-0.16	0.92									
2	248	-0.91	1.53	56	0.65	1.27									
3	73	-1.02	1.46	13	-1.10	1.63									
4	131	-1.03	1.57	65	0.66	1.38									
5	285	-1.25	1.15	138	-0.38	1.05									
6	379	-0.17	1.31	193	-1.06	1.67									
7	124	-0.26	0.58	441	-0.41	0.54							145	-0.35	0.49
8	533	-1.59	1.15				539	-1.79	1.13						
9	551	-1.66	1.01							541	-1.84	1.12			
10	283	-1.28	0.97							275	-1.4	0.98			
11	51	-1.07	1.36							56	-1	1.03			
12	41	-1.11	1.32							38	-0.78	1.21			

Table 2

Results for Bayesian standard trial residuals. Bold cells have $|STR| > 1.5$.

Trial	Treatment	STR	ASTR	Trial	Treatment	STR	ASTR
1	1	-0.35	0.64	7	1	-1.20	0.51
	2	-0.92			2	-0.33	
	—	—			5	0.00	
2	1	-0.06	1.15	8	1	1.13	0.60
	2	-2.24			3	0.06	
	—	—			—	—	
3	1	0.11	0.40	9	1	1.28	1.38
	2	0.69			4	1.48	
	—	—			—	—	
4	1	0.09	1.19	10	1	0.55	0.54
	2	-2.29			4	0.53	
	—	—			—	—	
5	1	0.52	0.56	11	1	0.17	0.16
	2	-0.60			4	-0.16	
	—	—			—	—	
6	1	-1.58	1.38	12	1	0.23	0.38
	2	1.18			4	-0.53	
	—	—			—	—	

Table 3

Bayesian p -values for discrepancy. Bold cells have p -value < 0.05 .

Bayesian p -values									
$p_{D_1^1}$	0.25	$p_{D_3^2}$	0.37	$p_{D_6^1}$	0.00	$p_{D_8^1}$	0.50	$p_{D_{10}^4}$	0.57
$p_{D_1^2}$	0.23	$p_{D_4^1}$	0.00	$p_{D_6^2}$	0.00	$p_{D_8^3}$	0.50	$p_{D_{11}^1}$	0.50
$p_{D_2^1}$	0.02	$p_{D_4^2}$	0.00	$p_{D_7^1}$	0.01	$p_{D_9^1}$	0.50	$p_{D_{11}^4}$	0.50
$p_{D_2^2}$	0.00	$p_{D_5^1}$	0.03	$p_{D_7^2}$	0.20	$p_{D_9^4}$	0.48	$p_{D_{12}^1}$	0.31
$p_{D_3^1}$	0.51	$p_{D_5^2}$	0.00	$p_{D_7^5}$	0.50	$p_{D_{10}^1}$	0.56	$p_{D_{12}^4}$	0.27

Table 4

Results for scale mixtures of normals. λ_i denotes the scale parameter; and $P(\lambda_i > 1|y)$ denotes the probability that the scale parameter λ_i is larger than 1 given the data. Bold cells represent the outliers.

Trial	λ_i		$P(\lambda_i > 1 y)$	
	SMN ₂	SMN ₃	SMN ₂	SMN ₃
1	1.57(0.30,20.06)	1.03(0.16,15.68)	0.67	0.51
2	10.22(1.89,97.33)	9.90(1.92,95.03)	0.99	0.99
3	3.11(0.67,26.84)	2.71(0.56,24.46)	0.92	0.89
4	12.71(2.52,114.50)	12.39(2.60,110.30)	1.00	1.00
5	1.69(0.29,28.37)	1.08(0.15,22.77)	0.68	0.53
6	100(24.66,788.60)	99.38(24.89,760.60)	1.00	1.00
7	24.82(4.27,222.20)	24.24(4.36,212.60)	1.00	0.99
8	1.44(0.27,36.64)	0.73(0.14,19.00)	0.63	0.40
9	1.35(0.27,19.96)	0.78(0.15,12.65)	0.62	0.41
10	1.03(0.24,11.63)	0.54(0.12,6.37)	0.51	0.28
11	1.16(0.28,10.98)	0.76(0.18,7.54)	0.57	0.39
12	2.08(0.51,18.51)	1.70(0.41,15.74)	0.82	0.74

Posterior summaries for parameters of interest with and without outliers. Note that μ_5 is not available in Row 2 because the only trial (Trial 7) containing Treatment 5 is an outlier.

Table 5

	μ_1	μ_2	μ_3	μ_4	μ_5	γ	σ^2
With Outliers	-0.95 (-1.28, -0.62)	-0.67 (-1.01, -0.33)	-1.16 (-1.51, -0.80)	-1.08 (-1.42, -0.74)	-0.70 (-1.09, -0.31)	0.57 (0.39, 0.86)	1.34 (1.29, 1.39)
Without Outliers	-1.19 (-1.65, -0.71)	-0.47 (-0.95, 0.02)	-1.39 (-1.87, -0.90)	-1.32 (-1.78, -0.84)	—	0.60 (0.38, 1.01)	1.20 (1.15, 1.26)