

SOFTWARE FOCUS

Shiny-MAGEC: A Bayesian R shiny application for meta-analysis of censored adverse events

Zihan Zhou¹, Zizhong Tian¹, Christine Peterson², Le Bao³ and Shouhao Zhou¹

¹Public Health Sciences, Pennsylvania State University, United States

²Biostatistics, The University of Texas MD Anderson Cancer Center, United States

³Statistics, Pennsylvania State University, United States

Corresponding author: Shouhao Zhou; Email: szhou1@pennstatehealth.psu.edu

Received: 12 April 2025; **Revised:** 16 September 2025; **Accepted:** 6 October 2025

Keywords: Bayesian inference; drug safety; incomplete reporting; MAGEC; meta-analysis

Abstract

Accurate assessment of adverse event (AE) incidence is critical in clinical research for drug safety. While meta-analysis serves as an essential tool to comprehensively synthesize the evidence across multiple studies, incomplete AE reporting in clinical trials remains a persistent challenge. In particular, AEs occurring below study-specific reporting thresholds are often omitted from publications, leading to left-censored data. Failure to account for these censored AE counts can result in biased AE incidence estimates. We present an R Shiny application that implements a Bayesian meta-analysis model specifically designed to incorporate censored AE data into the estimation process. This interactive tool provides a user-friendly interface for researchers to conduct AE meta-analyses and estimate the AE incidence probability using an unbiased approach. It also enables direct comparisons between models that either incorporate or ignore censoring, highlighting the biases introduced by conventional approaches. This tutorial demonstrates the Shiny application's functionality through an illustrative example on meta-analysis of PD-1/PD-L1 inhibitor safety and highlights the importance of this tool in improving AE risk assessment. Ultimately, the new Shiny app facilitates more accurate and transparent drug safety evaluations. The Shiny-MAGEC app is available at: <https://zihanzhou98.shinyapps.io/Shiny-MAGEC/>.

Highlights

• What is already known?

- Meta-analysis is a critical tool to synthesize clinical safety data and evaluate the harms of new interventions.
- Adverse event (AE) data in clinical trials are often incompletely reported, particularly for rare events.
- Conventional approaches, such as complete-case analysis, can lead to biased estimates of AE incidence by ignoring or mishandling missing data.

• What is new?

- We present Shiny-MAGEC, an R Shiny application that implements a Bayesian meta-analysis model designed for censored AE data.
- This interactive, user-friendly app enables automated input of study-level data with reporting cutoffs, facilitates bias-corrected safety meta-analyses, and compares results with conventional methods without bias correction.

● Potential impact for RSM readers

- To our knowledge, Shiny-MAGEC is the first software designed for the meta-analysis of drug safety AE data.
- By addressing the limitations of conventional AE reporting, it improves transparency, accuracy, and methodological rigor in safety evidence synthesis, supporting better-informed clinical and regulatory decisions.

1. Introduction

Meta-analysis plays a crucial role in assessing drug safety (or equivalently, quantifying drug-related harms).^{1,2} By synthesizing evidence from multiple studies, it enhances the precision of risk estimates for adverse events (AEs), which is essential for a comprehensive understanding of a drug's safety profile. Since individual clinical studies are often underpowered for detecting rare harm signals, meta-analysis serves as a valuable tool to combine data across studies, increase sample size, and improve power, enabling a more robust evaluation of drug-related risks.^{3,4}

However, a major challenge in safety meta-analysis is the incomplete reporting of AEs in cancer clinical trials, which can lead to biased estimates if not properly addressed.^{5,6} While the NCI Common Terminology Criteria for Adverse Events (CTCAE) v5.0 catalogs over 800 unique AEs,⁷ trial publications typically report only the most common or severe ones, while rare or less frequent events may be omitted due to space constraints or study-specific reporting thresholds. For instance, when extracting AE count data for grade 3 or higher pneumonitis for PD-1 and PD-L1 inhibitors, researchers may not find explicit AE frequency reporting, but instead may see a footnote in the main text or supplementary materials stating that AE counts were only included and reported if they exceeded a certain proportion of the study sample size.^{8,9} In many oncology trials, AEs may only be listed if they occur in more than 5% of patients, resulting in *left-censored* AE data. Indeed, left-censoring is a common incomplete AE reporting phenomenon, as demonstrated by discrepancies between published articles and unpublished trial documents.¹⁰ Moreover, reporting cutoffs are sometimes defined separately for different severity grades (e.g., a non-reporting threshold for all-grade AEs and a separate threshold for grade 3–5 AEs).¹¹ This variability in reporting practices highlights the difficulty of harmonizing AE data in meta-analysis.

Failure to account for left-censored AE data can markedly bias meta-analysis results. One naïve approach is to treat an unreported AE as if it did not occur in that study, which effectively imputes a zero count for missing categories. This approach allows the inclusion of all relevant studies in the meta-analysis and is straightforward to implement. However, the assumption can be problematic since the absence of an AE in a report doesn't always mean it didn't occur. If an AE was unreported only because it fell below the reporting threshold, assuming it to be zero underestimates the true event rate and can paint an overly safe picture of the drug.

Another common naïve approach is to omit data from studies for an AE outcome if that AE wasn't reported, equivalent to performing a complete-case analysis for data with missingness. Essentially, the meta-analysis is run only on the subset of studies that provided data for that outcome. While this avoids making false zero assumptions, it introduces its own bias. If unreported AEs are ignored in a meta-analysis and only the studies with higher observed AE counts are included, the estimated incidence rates will be skewed upward, since only studies with higher event occurrences contribute.^{12–14} Additionally, complete-case analysis wastes the important partial information in those omitted studies where the event count was rare enough to not be reported, and it reduces the overall sample size and power of the meta-analysis.

Subsequently, both of these naïve strategies distort the true safety profile of treatments. If AEs are under-estimated, patients might be unknowingly exposed to risks; conversely, overestimating AEs could deter the use of an effective treatment, emphasizing the need for appropriate statistical approaches tailored to handle censored AE data. Notably, standard methods for missing data in

meta-analyses^{15–17} are not designed or suitable to address this specific form of missing-not-at-random data arising from incomplete AE reporting.

To address this pervasive gap in the meta-analysis of safety data, a Bayesian meta-analytic model, named MAGEC (meta-analysis of adverse drug effects with censored data), was developed.⁶ This model correctly handles the censored AE data and provides exact inferences for a balanced, evidence-based understanding of drug harms. Specifically, the censoring cutoff information is incorporated into the likelihood function to enhance the estimation of AE incidence probabilities. Rather than imputing missing counts externally, their work treats unreported AEs as latent variables; the Bayesian model then elegantly integrates over the uncertainty of unreported AEs and draws inference on those hidden counts as part of the Markov Chain Monte Carlo (MCMC) estimation while simultaneously accounting for model between-study variability. As a result, this approach was shown to effectively produce unbiased estimates of AE incidence.

In this work, we develop an R Shiny application that implements the Bayesian MAGEC model in a user-friendly, interactive platform. This tool provides applied researchers and clinicians with a practical means to conduct meta-analyses of AEs in the presence of incomplete reporting. With the app, users can input aggregated AE data (including the information on reporting cutoffs) and obtain posterior estimates of AE incidence under the Bayesian model. Additionally, the application allows result comparison between models that account for censoring versus those that ignore it, thereby directly illustrating the bias that can arise from naïve complete-case analysis. By packaging advanced methodology into an accessible Shiny interface, our tool lowers the barrier to broader adoption of these statistical best practices in routine safety evidence synthesis.

This article is organized as follows. Section 2 provides a brief overview of the Bayesian MAGEC model.⁶ Section 3 presents the functionalities of the R Shiny application. Section 4 offers an illustrative example using clinical trial data on PD-1/PD-L1 inhibitors to demonstrate the app in practice. Finally, Section 5 concludes with a discussion of key takeaways, the potential impact of this tool on drug safety evaluation, and future directions for research and practice.

2. A brief review of the Bayesian MAGEC model

To address the challenge of left-censored AE data in meta-analysis, a Bayesian MAGEC random-effects meta-analytic model was developed to account for censored AE outcomes.⁶ It explicitly accounts for censoring by incorporating study-specific reporting thresholds into the likelihood function, enabling more accurate estimation of AE incidence rates while appropriately reflecting the uncertainty introduced by missing data.

Let Y_i denote the count of AEs under a target severity interval (e.g., all-grade or grade 3 and above) in study i , and N_i represent the number of patients treated in that study. Y_i is assumed to follow a binomial distribution:

$$Y_i \sim \text{Binomial}(N_i, \theta_i),$$

where θ_i represents the study-specific incidence probability. To account for left-censored AE data, the full likelihood function integrates both completely observed and censored data:

$$\mathcal{L} = \prod_{o=1}^O f_Y(y_o) \prod_{l=1}^L F_Y(c_l) = \prod_{o=1}^O f_Y(y_o) \prod_{l=1}^L \sum_{k_l=0}^{c_l} f_Y(k_l),$$

where O represents the set of studies with fully-observed AE counts, and L represents the set of studies where the AE count is left-censored below a study-specific cutoff c_l . The censored probability component, $F_Y(c_l)$, ensures that the analysis incorporates information about studies where AEs were unreported due to cutoff-based omission rather than true absence.

The incidence probability parameter θ_i can be modeled based on the link function $g(\cdot)$ that can transform values in the probability space into the real-value space. In the Shiny application, a logit link was used, such that

$$g(\theta_i) = \text{logit}(\theta_i) = \mu + \alpha_i + \mathbf{X}_i\boldsymbol{\beta},$$

where μ is the overall log-odds of AE incidence, α_i are the study-specific random effects, $\mathbf{X}_i$ is the design matrix of study-level covariates in the i th study, and $\boldsymbol{\beta}$ are the effect parameters corresponding to the study-level factors. In the current version of the Shiny application, specifications of the study-level covariates are not available. The α_i parameters follow conditionally independent normal distributions with mean 0 and variance τ^2 , where τ quantifies the between-study heterogeneity of the log odds. A noninformative normal prior distribution $\mathcal{N}(0, v_0^2)$ is placed on the overall mean parameter μ ($v_0^2 = 10^4$ by default). A half-Cauchy prior distribution $C^+(0, A)$ is placed for the between-study standard deviation (SD) parameter τ . Such weakly-informative prior with its scale parameter A equal to 2.5, 10, or 25 was generally recommended for the SD parameters in hierarchical models.¹⁸ The estimation for posterior inference is implemented using Just Another Gibbs Sampling (JAGS),¹⁹ and carried out using a data-augmentation strategy for censored data.¹¹

This MAGEC model mitigates bias in incidence estimation while incorporating between-study variability. The hierarchical Bayesian framework avoids relying on asymptotic normality for inference, improves estimation stability through MCMC sampling,²⁰ and enhances the reliability of drug safety assessments. Simulation studies have demonstrated its robustness in small-sample scenarios and for rare events.

3. The R shiny application: Shiny-MAGEC

We have developed an R Shiny application named Shiny-MAGEC that implements the Bayesian meta-analysis model reviewed in Section 2. Users can upload their raw AE data collected from multiple clinical studies to obtain meta-analytic estimates of the overall AE incidence probability and the between-study heterogeneity.

In this section, we provide a walk-through of the application, outlining its main features and functions. Users may refer to Section 4 for an illustrative example using a sample data excerpted from the real data application² and subsequently reanalyzed in a methodological study.⁶ The application consists of an operation panel on the left and a result panel on the right. The right panel is organized into two primary tabs: (1) “User Guide,” providing basic information about the software and showing an overview of the uploaded dataset and (2) “Results,” displaying the analysis outputs. The accepted data format, advanced settings tunable by users, and the result presentations supported by the application are introduced below. The general layout and features are summarized in Figure 1.

3.1. Data preparation

To familiarize themselves with the accepted data format, users can download the sample dataset at the operation panel for reference. The sample data, a subset of an AE meta-analysis dataset,² is displayed in Table 1. The input data should be provided as a CSV file with mandatory column names: “study,” “N,” “Y,” and “cutoff.” The “study” column is a unique identifier for each study. We recommend using the abbreviated names of different studies (i.e., character strings) to label the studies since these labels will be used in presenting and plotting the study-specific results as shown in Section 3.2. The “N” column is the sample size assigned to the treatment in each study, and the “Y” column (when not missing) represents the observed AE count given the prespecified severity grade and category of interest. “Y” should be coded as NA or left blank (preferred) when an AE count is unreported. The “cutoff” column

Shiny-MAGEC: Bayesian Meta-Analysis of Censored Adverse Events




Figure 1. An overview of the operation and result panels in Shiny-MAGEC.

Table 1. An example AE data subset² including the grade 3–5 pneumonitis counts for patients treated with Atezolizumab. “-” indicates unreported (left-censored) in the original publication. The left-censored cutoffs are calculated specific to different studies. Given a cutoff of 0 (e.g., in 2018-Colevas-Ann Oncol), the actual pneumonitis count, though unreported, was exactly 0.

Study source	No. of treated patients (N)	Cutoff	Pneumonitis count (Y)
2014-Herbst-Nature	277	1	0
2016-Fehrenbacher-Lancet	142	14	1
2016-McDermott-J Clin Oncol	70	0	0
2016-Rosenberg-Lancet	310	0	2
2017-Balar-Lancet	119	0	-
2016-Mizugaki-Invest New Drugs	6	1	-
2017-Peters-J Clin Oncol	659	6	11
2017-Rittmeyer-Lancet	609	60	4
2018-Colevas-Ann Oncol	32	0	-
2018-Emens-JAMA Oncol	116	3	1
2018-Horn-Eur J Cancer	89	4	0
2018-Lukas-J Neurooncol	16	0	-
2018-McDermott-Nat Med	103	20	0
2018-Petrylak-JAMA Onc	95	0	-
2018-Powles-Lancet	459	9	-

gives the study-specific left-censoring threshold, which can be extracted from a footnote or methods description in the original publication. Each cutoff is the largest integer not reported; for example, in a study with $N = 459$ and a footnote stating that “AE counts $\geq 2\%$ of the treated size are disclosed,” the study-specific cutoff would be 9, since any count larger than $459 \times 2\% = 9.18$ would be reported. Of note, some studies’ reporting threshold might differ by AE categories (e.g., “2017-Peters-J Clin

Oncol” in Table 1²¹). Users should carefully collect the “cutoff” information in order to match their AE category of synthesis interest.

After uploading the meta-analysis data following the required format, the application will automatically display the data content in the “Data Overview” section in the “User Guide” tab for a routine check. In addition, the app will automatically check the uploaded data and provides error messages to the user in the following situations: (1) if any data entry in the “N,” “cutoff,” or “Y” columns contains non-numeric symbols instead of numeric values, NA, or a blank; (2) if the column names do not match the required format; and (3) if “cutoff” > “N” or “Y” > “N.” This immediate feedback allows users to verify that their data has been read correctly (e.g., all studies are listed and the columns are interpreted as intended) before proceeding.

3.2. Operation and results of Shiny-MAGEC

Once the data is loaded, the analysis can be initiated by clicking the “Run Models” button. Basically, two Bayesian models will be sequentially executed: one that incorporates the censored data and uses the MAGEC approach and another that solely utilizes the fully observed data. The latter serves as an optional supplementary reference highlighting the bias that might be caused by the inappropriate complete-case analysis procedure. The Bayesian model estimations are based on MCMC simulations, and the sampling algorithm is the Metropolis–Hasting algorithm implemented using the JAGS software (Version 4.3.1). A few advanced parameters regarding the prior specifications in the models and the MCMC simulations can be manipulated by checking the box of “Advanced Settings.” By default, a half-Cauchy prior is specified for the between-study SD τ , and the scale parameter λ in the prior distribution of τ described in Section 2 is set to 2.5. The app also provides sensitivity analysis options, allowing the users to specify a uniform prior for the between-study SD τ , or an inverse-Gamma prior for the between-study variance τ^2 . The hyperparameters in these priors can further be customized in “Advanced Settings.” For MCMC, three chains are run in parallel, each with a total of 100,000 iterations, that include a burn-in of 50,000 and a thinning interval of 5 (i.e., keeping 1 out of every 5 samples). These defaults enhance robust estimation and an acceptable convergence performance in general applications. Finally, for reproducibility, a random seed for MCMC sampling can be set under “Advanced Settings.”

When the model fitting is finished, the application generates tabular and figure outputs for the MAGEC model shown in the “Results” tab. Within the “Results” tab, first, summary statistics about the posterior estimates of the overall AE incidence probability, the between-study variation prediction for a new data entry are shown in a table. This includes the posterior median, the SD, boundaries of the 95% credible interval (CrI), and the mean and standard error of the posterior distributions. If any optional sensitivity analysis prior is selected, a separate table presenting the same results for the sensitivity analysis will also be provided. It is followed by a short paragraph briefly summarizing the key meta-analytic results. For the purpose of comparison, the biased overall AE incidence probability from the complete-case analysis is also provided. The summary table also includes the Gelman–Rubin potential scale reduction factors R-hat statistics²² for diagnosing the convergence of the MCMC chains. If the R-hat statistics of any key parameters is larger than 1.01, a warning will be output, suggesting the user increase the lengths of the MCMC chains in the “Advanced settings” to improve the convergence and mixing performance.

In the bottom part of the page, the meta-analytic estimates as well as the study-specific estimates will be comprehensively visualized in a forest plot. Users may navigate to the advanced operation panels on the tab to customize the figure. Finally, an analysis report in Microsoft Word format containing the summary tables and the forest plots of both the MAGEC analysis and the complete-case analysis can be downloaded by clicking the “Download Report” button on the operation panel. To facilitate scientific writing and reporting, users can also find a template for drafting the statistical analysis section in the downloaded analysis report, which details the statistical methods for the MAGEC modeling (e.g., the specification of MCMC simulations) and cites relevant references.

4. An illustrative example

We present an illustrative example to demonstrate the use of Shiny-MAGEC based on the built-in sample dataset. A meta-analysis of 125 clinical studies was conducted to evaluate the incidence probabilities of pneumonitis (a type of AE involving inflammation of lung tissue) for several types of PD-1 and PD-L1 inhibitors.² As shown in Table 1 (Section 3.1), our sample data is a small subset of the full data focused on the incidence of Grade 3–5 pneumonitis for patients who received the PD-L1 drug Atezolizumab. Note that in the practical use of this current version of the Shiny app without accounting for study-level factors, such division of meta-analysis data by specific grade interval and specific AE category, as illustrated in the example, is highly recommended for delivering evidence synthesis within meaningful subgroups.

Once the data are uploaded, a data overview can be found in the “User Guide” tab. By clicking the “Advanced Settings” checkbox on the left, a panel allowing customization of prior hyperparameters and MCMC options will show up. In this illustrative example, we use the default settings. After clicking the “Run Model” button, a progress bar will approximately indicate the models’ execution statuses. In this illustrative example, the running time on an AMD R9-6900HS CPU is approximately 10 seconds. Once the posterior sampling is completed, we can navigate to the “Results” page to review the model outputs.

The main results generated based on the MAGEC methodology are displayed in the “Results” tab, while the parallel results derived from the complete-case analysis are summarized briefly in a separate paragraph only for supplementary comparison purposes. Figure 2 illustrates the summary statistics of the MAGEC model estimates (results associated with the default and alternative prior choices would be paralleled). Furthermore, it will automatically generate paragraphs summarizing the key meta-analytic results as exemplified in Figure 3 for users’ references.

Additionally, a forest plot will be generated to illustrate both the study-specific and overall AE incidence probability estimates (in percentages), as presented in Figure 4. This visualization will follow established scientific reporting structures exemplified in prior literature.^{2,6} As demonstrated in prior research, this presentation approach provides a clear and structured summary of results, enhancing interpretability and comparability of the results.

It is straightforward that by incorporating richer information provided by the censored studies, the MAGEC model provides different point and interval estimates compared to the complete-case analysis that solely uses the non-censored studies. By applying the MAGEC modeling approach, the AE incidence is estimated at 0.38% (95% CrI [0.05%, 0.87%]) in this example, whereas the complete-case

Summary Statistics of Bayesian Model Estimates

Parameter	50%	2.5%	97.5%	mean	sd	Rhat
Overall AE incidence probability (%)	0.3824	0.0546	0.8747	0.403	0.2189	1.0026
Predicted AE incidence probability (%)	0.4129	0.0098	2.9773	0.7409	2.51	1.0013
Between-study standard deviation of log odds	0.8949	0.1468	2.7578	1.0311	0.6751	1.0019

Sensitivity Analysis

Parameter	50%	2.5%	97.5%	mean	sd	Rhat
Uniform prior						
Overall AE incidence probability (%)	0.3403	0.0248	0.8614	0.3673	0.2256	1.0026
Predicted AE incidence probability (%)	0.379	0.0034	3.8087	0.9272	4.2128	1.0016
Between-study standard deviation of log odds	1.0191	0.1601	3.8307	1.2523	0.963	1.0026
Inverse-Gamma prior						
Overall AE incidence probability (%)	0.3536	0.0801	0.8206	0.3778	0.1914	1.0012
Predicted AE incidence probability (%)	0.3588	0.0194	2.8626	0.6664	1.7001	1.0011
Between-study standard deviation of log odds	0.9363	0.5066	2.1107	1.0334	0.4347	1.0011

Figure 2. Shiny app outputs of the summary statistics table.

A Bayesian meta-analysis of AE incidence probability is conducted using the MAGEC method (Qi et al., 2022, Qi et al., 2024, Zhou et al., 2025). A template for describing the statistical methods can be found in the downloaded report.

The posterior median of the overall AE incidence probability is 0.38% and the 95% credible interval (CrI) is [0.05%, 0.87%].

The posterior median of the between-study standard deviation (SD) of log odds is 0.89 and the 95% CrI is [0.15, 2.76].

For a new study entry, the 95% prediction CrI for the AE incidence probability is [0.01%, 2.98%].

As a supplementary analysis, the Bayesian meta-analysis model ignoring study entries with censored AEs suggests that the overall AE incidence probability is 0.51%, with 95% CrI [0.1%, 1.13%]. Note that this complete-case analysis result might be biased and only for the comparison purpose.

Uniform prior

The posterior median of the overall AE incidence probability is 0.34% and the 95% credible interval (CrI) is [0.02%, 0.86%].

The posterior median of the between-study standard deviation (SD) of log odds is 1.02 and the 95% CrI is [0.16, 3.83].

For a new study entry, the 95% prediction CrI for the AE incidence probability is [0%, 3.81%].

As a supplementary analysis, the Bayesian meta-analysis model ignoring study entries with censored AEs suggests that the overall AE incidence probability is 0.48%, with 95% CrI [0.05%, 1.14%]. Note that this complete-case analysis result might be biased and is only for comparison purposes.

Inverse-Gamma prior

The posterior median of the overall AE incidence probability is 0.35% and the 95% credible interval (CrI) is [0.08%, 0.82%].

The posterior median of the between-study standard deviation (SD) of log odds is 0.94 and the 95% CrI is [0.51, 2.11].

For a new study entry, the 95% prediction CrI for the AE incidence probability is [0.02%, 2.86%].

As a supplementary analysis, the Bayesian meta-analysis model ignoring study entries with censored AEs suggests that the overall AE incidence probability is 0.46%, with 95% CrI [0.12%, 1.09%]. Note that this complete-case analysis result might be biased and is only for comparison purposes.

Figure 3. Result descriptions based on the illustrative example.




Figure 4. Comprehensive forest plot of the incidence probabilities (in percentage) of Grade 3–5 pneumonitis related to the PD-L1 drug Atezolizumab based on the MAGEC meta-analysis model.

analysis yields a higher estimate of 0.51% (95% CrI [0.10%, 1.13%]). The discrepancy reflects a 34% over-estimation bias inherent in the complete case approach, reinforcing findings from previous studies that have shown complete case methods tend to inflate incidence estimates due to missing data.⁶

5. Discussion

Meta-analysis has become an indispensable tool for characterizing the safety profiles of medications, but it comes with unique challenges compared to efficacy analysis. Incomplete reporting of AEs, together with possible heterogeneity across studies, can compromise the validity of a drug safety meta-analysis if not properly handled. On the methodological front, researchers are developing more sophisticated models to tackle these challenges. The Bayesian censored-data approach⁶ is one such innovation that addresses partially observed safety data. In this article, we extend these efforts by providing a user-friendly tool that brings advanced methodology to a broader audience of clinicians and researchers.

The Shiny-MAGEC app offers an interactive platform to facilitate the application of the Bayesian meta-analysis method described for censored AE data. By incorporating reporting cutoff information into the analysis, it mitigates the estimation bias of AE incidence rates, accompanied by appropriate uncertainty quantification. This new tool facilitates rigorous systematic review and meta-analysis of safety data by making state-of-the-art bias correction accessible to users without requiring advanced statistical programming skills.

Beyond this specific tool, our work aligns with a broader movement to improve the quality of AE meta-analyses. Regulators (e.g., FDA and EMA) and guideline groups (e.g., Cochrane, CONSORT, and PRISMA) have increasingly emphasized the need for high standards in safety evidence synthesis to support robust conclusions about drug risks.^{23–27} A balanced, evidence-based understanding of drug harms is critical to weighing benefits against risks in clinical decision-making. Particularly, regulatory agencies rely on comprehensive safety evidence when making approval and labeling decisions, as flawed or incomplete safety meta-analyses can have serious consequences.^{1,28} Therefore, safety outcomes should be synthesized with the same rigor and transparency as efficacy data, ensuring scientific completeness and reliability in assessments of drug risk. An understated risk in a meta-analysis might delay regulatory actions on a harmful drug, whereas an overestimated risk could unnecessarily alarm practitioners and patients. In practice, both underestimation and overestimation of risks must be avoided.

While the challenges in meta-analysis of drug safety are non-trivial, they are surmountable with diligent methodology innovation and improved data practices.²⁹ By implementing best practices for AE reporting, leveraging advanced statistical techniques like the Bayesian models employed here, and exploring new tools like natural language processing for data gathering,³⁰ the field can move toward more reliable and informative safety meta-analyses. These improvements will ultimately benefit patients and healthcare providers, as treatment decisions can be made with a clearer understanding of the balance between benefits and harms.

This R Shiny application has several limitations. First, the computation capacity of the app is restricted by the current shinyapps.io hosting plan, which has a maximum instance size of 1 GB. Users with higher computational demands are recommended to refer to the sample JAGS model code of the MAGEC model provided in the Supplementary Material and run the analysis on a local machine. Note that the Shiny-MAGEC focuses on handling the left-censored data, aiming to offer a concise and user-friendly interface that is suitable for most AE meta-analysis studies. Right-censoring and interval-censoring, though also can be accounted for based on the MAGEC model, are rarely encountered in the AE reporting of modern studies. We encouraged those users who have the needs to handle right- or interval-censored AEs to locally run the MAGEC model code detailed in the Supplementary Material. A potential limitation of the MAGEC model is that its performance may be less reliable in settings with very few studies or extremely high between-study heterogeneity. Future research should investigate strategies to improve robustness under these boundary conditions and extend the model's applicability to such challenging scenarios. Finally, one under-addressed topic in the safety meta-analysis is the publication bias, where studies reporting non-significant or unfavorable safety outcomes may be selectively unpublished or underreported, potentially biasing AE incidence estimates.^{31,32} Exploring how the publication bias mechanism might interact with the AE incomplete reporting issue and developing methods to simultaneously adjust for the biases caused by the two sources are insightful future directions.

The future of drug safety meta-analysis is moving toward greater scientific rigor, transparency, and integration of diverse data sources—all aimed at safeguarding public health through better evidence on drug risks. Continued methodological research and consensus-building are needed to refine how we pool and interpret AE data. As tools like our R Shiny app become integrated into researchers' workflows, we anticipate more accurate and trustworthy assessments of drug safety that will support better-informed clinical and regulatory decisions.

Acknowledgements. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We thank Dr. Xinyue Qi for her help on the project.

Author contributions. Z.Z. and Z.T. contributed equally to this work. Conceptualization: S.Z. Methodology: Z.Z., Z.T., and S.Z. Data curation: Z.Z. and Z.T. Software: Z.Z. and Z.T. Writing—original draft: Z.Z., Z.T., and S.Z. Writing—review and editing: Z.Z., Z.T., C.B.P., L.B., and S.Z. All authors approved the final submitted draft.

Competing interest statement. The authors declare that no competing interests exist.

Data availability statement. The illustrative dataset about several studies' AE count reporting of Grade 3–5 pneumonitis related to the PD-L1 drug Atezolizumab is available within the Shiny app that can be accessed at <https://zihanzhou98.shinyapps.io/Shiny-MAGEC/>.

Funding statement. This work is supported in part by NIH/NCI CCSG Grant P30CA016672 and NIH/NLM Grant R21LM014534.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/rsm.2025.10052>.

References

- [1] FDA. *Guidance Document: Meta-Analyses of Randomized Controlled Clinical Trials to Evaluate the Safety of Human Drugs or Biological Products*. 2018: 1–29.
- [2] Wang Y, Zhou S, Yang F, et al. Treatment-related adverse events of PD-1 and PD-L1 inhibitors in clinical trials: A systematic review and meta-analysis. *JAMA Oncol*. 2019;5(7): 1008–1019.
- [3] Stoto MA. Drug safety meta-analysis: Promises and pitfalls. *Drug Saf*. 2015;38(3): 233–243.
- [4] Hong H, Wang C, Rosner GL. Meta-analysis of rare adverse events in randomized clinical trials: Bayesian and frequentist methods. *Clin Trials*. 2021;18(1): 3–16.
- [5] Huang H-Y, Andrews E, Jones J, Skovron ML, Tilson H. Pitfalls in meta-analyses on adverse events reported from clinical trials. *Pharmacoepidemiol Drug Saf*. 2011;20(10): 1014–1020.
- [6] Xinyue QI, Shouhao ZH, Christine BP, et al. Meta-analysis of censored adverse events. *New England J Stat Data Sci*. 2024;2(3): 380.
- [7] National Cancer Institute. Common terminology criteria for adverse events (CTCAE) version 5.0. 2017. https://ctep.cancer.gov/protocolDevelopment/electronic_applications/ctc.htm.
- [8] Powles T, Durán I, Van der Heijden MS, et al. Atezolizumab versus chemotherapy in patients with platinum-treated locally advanced or metastatic urothelial carcinoma (IMvigor211): A multicentre, open-label, phase 3 randomised controlled trial. *Lancet*. 2018;391(10122): 748–757.
- [9] Nanda R, Chow LQ, Dees EC, et al. Pembrolizumab in patients with advanced triple-negative breast cancer: Phase Ib KEYNOTE-012 study. *J Clin Oncol*. 2016;34(21): 2460–2467.
- [10] Golder S, Loke YK, Wright K, Norman G. Reporting of adverse events in published and unpublished studies of health care interventions: A systematic review. *PLoS Med*. 2016;13(9): e1002127.
- [11] Qi X, Zhou S, Plummer M. On Bayesian modeling of censored data in JAGS. *BMC Bioinf*. 2022;23(1): 1–13.
- [12] Grigor EJM, Fergusson D, Kekre N, et al. Risks and benefits of chimeric antigen receptor T-cell (CAR-T) therapy in cancer: A systematic review and meta-analysis. *Transfus Med Rev*. 2019;33(2): 98–110.
- [13] Zhou X, Yao Z, Bai H, et al. Treatment-related adverse events of PD-1 and PD-L1 inhibitor-based combination therapies in clinical trials: A systematic review and meta-analysis. *Lancet Oncol*. 2021;22(9): 1265–1274.
- [14] Fujiwara Y, Horita N, Adib E, et al. Treatment-related adverse events, including fatal toxicities, in patients with solid tumours receiving neoadjuvant and adjuvant immune checkpoint blockade: A systematic review and meta-analysis of randomised controlled trials. *Lancet Oncol*. 2024;25(1): 62–75.
- [15] Higgins JPT, White IR, Wood AM. Imputation methods for missing outcome data in meta-analysis of clinical trials. *Clin Trials* 2008;5(3): 225–239.
- [16] White IR, Welton NJ, Wood AM, Ades AE, Higgins JP. Allowing for uncertainty due to missing data in meta-analysis—Part 2: Hierarchical models. *Stat Med*. 2008;27(5): 728–745.
- [17] Mavridis D, Chaimani A, Efthimiou O, Leucht S, Salanti G. Addressing missing outcome data in meta-analysis. *Evid-Based Mental Health*. 2014;17(3): 85–89.
- [18] Gelman A. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Anal*. 2006;1(3): 515–534.
- [19] Plummer, M. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In: Hornik, K., Leisch, F., Zeileis, A. (editors). *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. Vol. 124. 2003, Vienna, Austria.
- [20] Hamza TH, van Houwelingen HC, Stijnen T. The binomial distribution of meta-analysis was preferred to model within-study variability. *J Clin Epidemiol*. 2008;61(1): 41–51.
- [21] Peters S, Gettinger S, Johnson ML, et al. Phase II trial of atezolizumab as first-line or subsequent therapy for patients with programmed death-ligand 1–selected advanced non–small-cell lung cancer (BIRCH). *J Clin Oncol*. 2017;35(24): 2781–2789.

- [22] Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis*. 3rd ed. Chapman and Hall/CRC; 2013.
- [23] FDA. *Guidance for Industry: Adverse Reactions Section of Labeling for Human Prescription Drug and Biological Products—Content and Format*. 2006: 1–13.
- [24] Ioannidis JPA, Evans SJ, Gøtzsche PC, et al. Better reporting of harms in randomized trials: An extension of the CONSORT statement. *Ann Internal Med*. 2004;141(10): 781–788.
- [25] CIOMS. *Evidence Synthesis and Meta-Analysis: Report of CIOMS Working Group X*. Geneva : World Health Organization; 2016.
- [26] Higgins JPT, Green S. *Cochrane Handbook for Systematic Reviews of Interventions*. John Wiley & Sons; 2019.
- [27] Page MJ, McKenzie JE, Bossuyt PM, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*. 2021;372: n71.
- [28] EMA. *ICH guideline M3(R2) on non-clinical safety studies for the conduct of human clinical trials and marketing authorization for Pharmaceuticals*. 2010.
- [29] FDA. *Guidance for Industry: Hematologic Malignancies: Regulatory Considerations for Use of Minimal Residual Disease in Development of Drug and Biological Products for Treatment*. 2020: 1–21.
- [30] Young IJB, Luz S, Lone N. A systematic review of natural language processing for classification tasks in the field of incident reporting and adverse event analysis. *Int J Med Inform*. 2019;132: 103971.
- [31] Xing X, Zhu J, Shi L, Xu C, Lin L. Assessment of inverse publication bias in safety outcomes: An empirical analysis. *BMC Med*. 2024;22(1): 494.
- [32] Ioannidis JPA. Inverse publication reporting bias favouring null, negative results. *BMJ Evid-Based Med*. 2024;29(1): 6–9.