

# Misclassification in Automated Content Analysis Causes Bias in Regression. Can We Fix It? Yes We Can!

Nathan TeBlunthuis, Valerie Hase & Chung-Hong Chan

**To cite this article:** Nathan TeBlunthuis, Valerie Hase & Chung-Hong Chan (16 Jan 2024): Misclassification in Automated Content Analysis Causes Bias in Regression. Can We Fix It? Yes We Can!, Communication Methods and Measures, DOI: [10.1080/19312458.2023.2293713](https://doi.org/10.1080/19312458.2023.2293713)

**To link to this article:** <https://doi.org/10.1080/19312458.2023.2293713>



View supplementary material [↗](#)



Published online: 16 Jan 2024.



Submit your article to this journal [↗](#)





View related articles [↗](#)



View Crossmark data [↗](#)



# Misclassification in Automated Content Analysis Causes Bias in Regression. Can We Fix It? Yes We Can!

Nathan TeBlunthuis <sup>a,b</sup>, Valerie Hase <sup>c</sup>, and Chung-Hong Chan <sup>d</sup>

<sup>a</sup>School of Information, University of Michigan, Ann Arbor, USA; <sup>b</sup>Department of Communication Studies, Northwestern University, Evanston, USA; <sup>c</sup>Department of Media and Communication, LMU Munich, Munich, Germany; <sup>d</sup>Department of Computational Social Science, Cologne, Germany; GESIS - Leibniz-Institut für Sozialwissenschaften TeBlunthuis, Hase & Chan, Mannheim, Germany

## ABSTRACT

Automated classifiers (ACs), often built via supervised machine learning (SML), can categorize large, statistically powerful samples of data ranging from text to images and video. They have become widely popular measurement devices in communication science and related fields. Despite this popularity, even highly accurate classifiers make errors that cause misclassification bias and misleading results when input to downstream statistical analyses—unless such analyses account for these errors. As we show in a systematic literature review of SML applications, communication scholars largely ignore misclassification bias. In principle, existing statistical methods can use “gold standard” validation data, such as that created by human annotators, to correct misclassification bias. We introduce and test such methods, including a new method we design and implement in the R package `MISCLASSIFICATION_MODELS`, via Monte Carlo simulations designed to reveal each method’s limitations, which we also release. Based on our results, we recommend our new error correction method as it is versatile and efficient. In sum, automated classifiers, even those below common accuracy standards or those making systematic misclassifications, can be useful for measurement with careful study design and appropriate error correction methods.

*Automated classifiers* (ACs) based on supervised machine learning (SML) have rapidly gained popularity as part of the *automated content analysis* toolkit in communication science (Baden et al., 2022). With ACs, researchers can categorize large samples of text, images, video, or other types of data into predefined categories (Scharkow, 2013). Studies, for instance, use SML-based classifiers to study frames (Burscher et al., 2014), tonality (van Atteveldt et al., 2021), or civility (Hede et al., 2021) in news media texts or social media posts.

However, there is an increasing concern that automated classifiers’ imperfections may compromise measurement validity for studying theories and concepts from communication science (Baden et al., 2022; Hase et al., 2022). Research areas where ACs have the greatest potential—e.g., content moderation, social media bots, affective polarization, or radicalization—are haunted by the specter of methodological questions related to misclassification bias (Rauchfleisch et al., 2020): How accurate must an AC be to measure a variable? Can an AC built for one context be used in another (Burscher et al., 2015; Hede et al., 2021)? Is comparing automated classifications to some external ground truth sufficient to claim validity? How do biases in AC-based measurements affect downstream statistical analyses (Millimet & Parmeter, 2022)?

Social scientists using ACs are broadly aware that they have imperfections, and consensus best practices are to acknowledge this fact and report quantifiers of classifier performance such as precision and recall. Such acknowledgment is an important step, but clearly falls short of a satisfying answer to the methodological questions above. We introduce and analyze *misclassification bias*: How misclassifications by ACs distort statistical findings in downstream analysis—unless correctly modeled (Fong & Tyler, 2021).

Our study begins with a demonstration of misclassification bias in a real-world example based on the Perspective toxicity classifier. Next, we provide a systematic literature review of  $N = 48$  studies employing SML-based text classification. Although communication scholars have long scrutinized related questions about manual content analysis for which they have recently proposed statistical corrections (Bachl & Scharkow, 2017; Geiß, 2021), misclassification bias in automated content analysis is largely ignored. Our review demonstrates a troubling lack of attention to the threats ACs introduce and virtually no mitigation of such threats. As a result, in the current state of affairs, researchers are likely to either draw misleading conclusions from inaccurate ACs or avoid ACs in favor of costly methods such as manually coding large samples (van Atteveldt et al., 2021).

Our primary contribution, an effort to rescue ACs from this dismal state, is to *introduce and test methods for correcting misclassification bias* (Buonaccorsi, 2010; Carroll et al., 2006; Yi et al., 2021). We consider three recently proposed methods: Fong and Tyler's (2021) generalized method of moments calibration method, Zhang's (2021) pseudo-likelihood models, and Blackwell et al.'s (2017) application of imputation methods. To overcome these methods' limitations, we draw a general likelihood modeling framework from the statistical literature on measurement error (Carroll et al., 2006) and tailor it to the problem of misclassification bias. Our novel implementation is the experimental R package `MISCLASSIFICATION_MODELS`.

We test these four error correction methods and compare them against ignoring misclassification (the naïve approach) and refraining from automated content analysis using only manual coding (the feasible approach). We use Monte Carlo simulations to model four prototypical situations identified by our review: Using ACs to measure either (1) an independent or (2) a dependent variable where the classifier makes misclassifications that are either (a) easy to correct (when an AC is unbiased and misclassifications are uncorrelated with covariates, i.e., *nonsystematic misclassification*) or (b) more difficult (when an AC is biased and misclassifications are correlated with covariates, i.e., *systematic misclassification*).

According to our simulations, even biased classifiers without high predictive performance can be useful in conjunction with appropriate validation data and error correction methods. As a result, we are optimistic about the potential of ACs and automated content analysis for communication science and related fields—if researchers can correct for misclassification. Current practices of “validating” ACs by making misclassification rates transparent via metrics such as the F1 score, however, provide little safeguard against misclassification bias.

In sum, we make a methodological contribution by introducing the often-ignored problem of misclassification bias in automated content analysis, testing error correction methods to address this problem via Monte Carlo simulations, and introducing a new method for error correction. Profoundly, we conclude that automated content analysis will progress not only—or even primarily—by building more accurate classifiers but through rigorous human annotation and statistical error modeling.

## Why misclassification is a problem: an example based on the Perspective API

There is no perfect AC. All ACs make errors. This inevitable misclassification causes bias in statistical inference (Carroll et al., 2006; Scharkow & Bachl, 2017), leading researchers to make both type-I (false discovery) and type-II (failure to reject the null) errors in hypotheses tests. To illustrate the problematic consequences of this misclassification bias, we focus on real-world data and a specific research area in communication research: detecting and understanding harmful social media content. Communication researchers often employ automated tools such as the Perspective toxicity classifier (cjadams et al., 2019) to detect toxicity in online content (e.g., Hopp & Vargo, 2019; Kim et al., 2021;

Votta et al., 2023). As shown next, however, relying on toxicity scores created by ACs such as the Perspective API as (in-)dependent variables produces different results than using measurements created via manual annotation.

To illustrate this, we use the Civil Comments dataset released in 2019 by Jigsaw, the Alphabet corporation subsidiary behind the Perspective API. Methodological details on the data and our example are available in the Appendix.<sup>1</sup> The dataset has 448,000 English-language comments made on independent news sites. It also includes manual annotations of each comment concerning its toxicity (*toxicity*), whether it discloses aspects of personal identity like race or ethnicity (*identity disclosure*), and the number of likes it received (*number of likes*).

In addition to manual annotations of each comment, we obtained AC-based toxicity classifications from the Perspective API in November 2022. Perspective's toxicity classifier performs very well, with an accuracy of 92% and an F1 score of 0.79. Nevertheless, if we treat human annotations as the ground-truth, the classifier makes systematic misclassifications: It disproportionately misclassifies comments as toxic if they disclose racial or ethnic identity (Pearson's  $\rho = 0.121$ ;  $95\%CI = [0.12, 0.124]$ ). This seemingly small degree of misclassification can have nontrivial consequences for the overall accuracy of results as shown below. We will illustrate this through two separate research questions (RQs) researchers may test based on automated classifications from the Perspective API.

Consider the following exemplary research question that social scientists may ask: How does positive feedback (likes) to more-or-less toxic behavior in online comments depend on whether comments also disclose identity? They may use a logistic regression model to predict whether a comment contains *identity disclosure* using *number of likes*, *toxicity*, and their interaction as independent variables (IVs).


Figure 1a illustrates what happens in this case, i.e., when *ignoring misclassification in an IV*: Researchers may reject a hypothesized direct relationship between likes and identity disclosure because the model indicates that that their correlation is mediated by toxicity, but the direct relationship is not well estimated. In contrast, researchers using human annotations would conclude that the direct relationship between likes and identity disclosure is subtle but statistically significant. Thus, even a very accurate AC can introduce type-II errors, i.e., researchers failing to rejecting a null hypothesis due to misclassification.

Now consider a second exemplary RQ: How does toxic behavior in online comments depend on whether comments disclose identity and on the positive feedback they receive (likes)? Researchers of this question may predict the *toxicity* of a comment with *number of likes*, *identity disclosure* in a comment, and their interaction as IVs. Figure 1b illustrates what happens in this case, i.e., when *ignoring misclassification in a dependent variable (DV)*: Using Perspective's classification of toxicity results in a small negative direct effect of likes. However, there is no detectable relationship when using manual annotations. In this case, misclassification leads to type-I error, i.e., false discovery of a nonzero relationship.


## Why transparency about misclassification is not enough

Although the Perspective API is certainly accurate enough to be useful to content moderators, the example above demonstrates that this does not imply usefulness for social science (Grimmer et al., 2021). Machine learning takes the opposite position on the bias-variance trade-off than conventional statistics does and achieves high predictiveness at the cost of more biased inference (Breiman, 2001). As a growing body of scholarship critical of the hasty adoption of machine learning in areas such as criminal justice, healthcare, or content moderation demonstrates, ACs boasting high performance often have biases related to social categories (Barocas et al., 2019). Such biases often result from non-representative training data and

<sup>1</sup>The appendix is available in an Open Science Framework (OSF) repository <https://osf.io/pyqf8/files/osfstorage/656f977c783ec62762eddabf>.



(a) Example 1: *Misclassification in an independent variable.*



(b) Example 2: *Misclassification in a dependent variable.*

**Figure 1.** Bias through misclassification: a real-world example using the perspective API and the civil comments dataset.

A figure visualizes the results from the real-data example in two panels. The top panel is captioned “Example 1: Misclassification in an independent variable”, titled “Logistic Regression Predicting Racial/Ethnic Identity Disclosure”, and has three facets corresponding to the independent variables in the logistic regression: “Likes”, “Toxicity” and the interaction between “Likes” and “Toxicity”. The bottom panel’s caption is “Example 2: Misclassification in a dependent variable” and it has the title “Logistic regression predicting toxicity”. It also has three facets for the independent variables in the example and they correspond to “Likes”, “Identity Disclosure” and the interaction between “Likes” and “Identity Disclosure”.

The X axes of both panels are title “Coefficients and 95% Confidence Intervals”. On the Y axis of both panels are two values corresponding to models fit using data from “Automated Classifications” and “Manual Annotations”. For each of these values a circle shows the point estimate and a line shows the confidence interval.

spurious correlations that neither reflect causal mechanisms nor generalize to different populations (Bender et al., 2021). Much of this critique targets unjust consequences of these biases to individuals. Our example shows that these biases can also contaminate scientific studies using ACs as measurement devices. Even very accurate ACs can cause both type-I and type-II errors, which become more likely when classifiers are less accurate or more biased or when effect sizes are small.

We argue that current common practices to address such limitations are insufficient. These practices aim to characterize the quality of measurement tools by reporting classifier performance on manually annotated data quantified via metrics like accuracy, precision, recall, or the F1 score (Baden et al., 2022; Hase et al., 2022; Song et al., 2020). While such steps may promote confidence in the reliability of results and are an important first step, they cannot provide conclusive evidence for or against misclassification bias. Despite high predictiveness, bias can still flow downstream into statistical inferences and, ultimately, threaten the validity of results, as our example demonstrates. Instead of relying only on transparency rituals to ward off misclassification bias, researchers can and should use validation data to correct it.

These claims may surprise because of the widespread misconception that misclassification causes only conservative bias (i.e., bias toward null effects). This is believed because it is true for bivariate least squares regression when misclassifications are nonsystematic (Carroll et al., 2006; Loken & Gelman, 2017; van

Smeden et al., 2020).<sup>2</sup> As a result, researchers interested in a hypothesis of a statistically significant relationship may not consider misclassification an important threat to validity (Loken & Gelman, 2017).

As our example shows, however, misclassification bias can be anti-conservative (Carroll et al., 2006; Loken & Gelman, 2017; van Smeden et al., 2020). In regression models with more than one IV or in nonlinear models, such as the logistic regression we used in our example, even nonsystematic misclassification can cause bias away from 0. Moreover, systematic misclassification can bias inference in any direction.

ACs designed in one context and applied in another are likely to commit systematic misclassification. The Perspective API, for example, was developed for social media comments but performs much worse when applied to news data (Hede et al., 2021). Systematic misclassification may also arise when an AC used for measurement shapes behavior in a sociotechnical system under study. As examples, the Perspective API is used for online forum moderation (Hede et al., 2021), as is the ORES API for Wikipedia moderators (TeBlunthuis et al., 2021). Misclassification from such classifiers can be systematic because they have causal effects on outcomes related to moderation.

If ACs become standard measurement devices, for instance Google's Perspective API for measuring toxicity (see critically Hosseini et al., 2017) or Botometer for classifying social media bots (see critically Rauchfleisch et al., 2020), entire research areas may be subject to systematic biases. Even if misclassification bias is usually conservative, it can slow progress in a research area. Consider how Scharnow and Bachl (2017) argue that media's "minimal effects" on political opinions and behavior in linkage studies may be an artifact of measurement errors in both manual content analyses and self-reported media use in surveys. Conversely, if researchers selectively report statistically significant hypothesis tests, misclassification can introduce an upward bias in the magnitude of reported effect sizes and contribute to a replication crisis (Loken & Gelman, 2017).

## Quantifying the problem: error correction methods in SML-based text classification

To understand how social scientists, including communication scholars, engage with the problem of misclassification in automated content analysis, we conducted a systematic literature review of studies using supervised machine learning (SML) for text classification (see the Appendix in our Supplement for details).<sup>3</sup> Our sample consists of studies identified by similar reviews on automated content analysis (Baden et al., 2022; Hase et al., 2022; Jünger et al., 2022; Song et al., 2020). Our goal is not to comprehensively review all SML studies but to provide a picture of common practices, with an eye toward awareness of misclassification and its statistical implications.

We identified a total of 48 empirical studies published between 2013 and 2021, more than half of which were published in communication journals. Studies used SML-based text classification for purposes such as to measure frames (Oppenhuizen et al., 2019) or topics (Vermeer et al., 2020). They often employed SML-based ACs to create dichotomous (50%) or other categorical variables (23%).<sup>4</sup> Of these empirical studies, many used SML-based ACs as IVs

<sup>2</sup>The measurement error literature defines categories of measurement error, of which we will discuss four: (1) Measurement error in an IV is *classical* when the measurement equals the true value plus noise. In other words, it is nonsystematic and the variance of an AC's predictions is greater than the variance of the true value (Carroll et al., 2006). (2) *Berkson* measurement error in an IV is when the truth equals the measurement plus noise. In agreement with prior work, we assume that nonsystematic measurement error from ACs is classical because an AC is an imperfect model of reality (Fong & Tyler, 2021; Zhang, 2021). Similarly, it is hard to imagine how an AC would have Berkson errors as predictions would then have lower variance than the training data. We thus do not consider Berkson error. (3) Measurement error in an IV is called *differential* if it is not conditionally independent of the DV given the other IVs. (4) Measurement error in a DV is called *systematic* when it is correlated with an IV. We use this more general term to simplify our discussions that pertain equally to misclassified IVs and DVs.

<sup>3</sup>Automated content analysis includes a range of methods both for assigning content to predefined categories (e.g., dictionaries) and for assigning content to unknown categories (e.g., topic modeling) (Grimmer & Stewart, 2013). While we focus on SML, our arguments extend to other approaches such as dictionary-based classification and even beyond the specific context of text classification.

<sup>4</sup>Metric variables were created in 35% of studies, mostly via the non-parametric method by Hopkins and King (2010).

**Table 1.** Overview of error correction methods.

Method	Abbrev.	Key Citation	Random Misclass.		Systematic Misclass.	
			IV	DV	IV	DV
GMM calibration	GMM	Fong and Tyler (2021)	Yes	No	No	No
Multiple Imputation	MI	Blackwell et al. (2012)	Yes	Yes	Yes	No
Pseudo-likelihood modeling	PL	Zhang (2021)	No	No	No	No
Maximum likelihood adjustment (ours)	MLA	Carroll et al. (2006)	Yes	Yes	Yes	Yes

(44%) or DVs (40%) in multivariate analyses. Ninety percent reported univariate statistics such as proportions.

Overall, our review reveals that *studies almost never reflected upon nor corrected misclassification bias*. According to our review, 85% of studies reported metrics such as recall or precision, but only 19% of studies explicitly stated that an AC misclassified texts which may introduce measurement error. Only a single article reported using error correction methods. This indicates a clear need to better understand and address misclassification bias. We now turn toward error correction methods for doing so.

### Addressing the problem: approaches for correcting misclassification

Statisticians have extensively studied measurement error (including misclassification), the problems it causes for statistical inference, and methods for correcting these problems (see Carroll et al., 2006; Fuller, 1987). In this paper, we propose *Maximum Likelihood Adjustment* (MLA), a method capable of fixing misclassification bias when an automatic classifier measures either an IV or a DV and makes either random or systematic misclassifications. This section explains our proposed approach and then compares it with three alternatives: Fong and Tyler’s (2021) GMM calibration method (GMM), multiple imputation (MI) (Blackwell et al., 2012), and Zhang’s (2021) pseudo-likelihood models (PL). Table 1 summarizes each method and the conditions where they are effective according to our simulations.<sup>5</sup>

For clarity, we first restate the problem in mathematical notation. In *the IV case*, we want to estimate a regression model  $Y = B_0 + B_1X + B_2Z + \varepsilon$  where  $X$  is only observed in a small sample of manually annotated data  $\tilde{X}$ , but automated classifications  $W$  of  $X$  are observed. To illustrate, in our first real-world example,  $X$  is toxicity,  $\tilde{X}$  are the civil comment annotations,  $W$  are the Perspective API’s toxicity classification,  $Z$  are likes, and  $Y$  is identity disclosure. Similarly, in *the DV case*, we only observe the true value of  $Y$  in a small sample  $\tilde{Y}$ . In our second real-word example,  $Y$  is toxicity.

Across both cases, a sample of annotated data ( $\tilde{X}$  or  $\tilde{Y}$ ) may be too small to provide convincing evidence, but collecting additional annotations is too expensive. In contrast, an AC can make classifications  $W$  for the entire dataset but introduces misclassification bias. How can we correct this bias in an automated content analysis?

### Maximum likelihood adjustment for misclassification

We propose *Maximum Likelihood Adjustment* (MLA) for correcting misclassification bias. MLA tailors Carroll et al.’s (2006) presentation of the general statistical theory of likelihood modeling for measurement error correction in the context of automated content analysis.<sup>6</sup> We give a brief intuitive explanation here and a more rigorous explanation in the Appendix.

<sup>5</sup>Statisticians have studied other methods including simulation extrapolation and score function methods. As we argue in the Appendix, these error correction methods are not advantageous when manually annotated data is available, as is often the case with ACs.

<sup>6</sup>In particular see Chapter 8 (especially example 8.4) and Chapter 15 (especially 15.4.2).



MLA works in slightly different ways depending on if an IV or a DV is automatically classified, but the intuition is similar. In both cases, MLA depends on an *error model* of the automated classifications  $W$  given  $X$ ,  $Y$ , and  $Z$  (Carroll et al., 2006). It may seem surprising that the error model predicts the automated classifications using validation data instead of the other way around, as is the case in the first stages of the MI and GMM approaches (see below). In MLA, the role of the error model is not to make new automated classifications, but to account for the relationship between automated classifications and the annotations and how this relationship varies depending on other observable variables.

In the IV case, MLA constitutes the use of a single *joint model* comprising the error model, the *main model* of scientific interest, and an *exposure model* giving the probability of  $X$ . Unlike the MI and GMM techniques discussed below, MLA estimates the joint model all at once using likelihood maximization. When annotations  $\tilde{X}$  are available, the data are entirely observed and are included directly in the model. Otherwise, when  $X$  is unobserved, the model uses a synthetic copy of the unobserved data for each possible value of  $X$  (e.g., if  $X$  is binary, there are two synthetic datasets: one with  $X = 1$  and another with  $X = 0$ .) The synthetic data are weighted by the probability of  $W$  given  $X$  according to the error model and by the probability of  $X$  according to the exposure model. This way, MLA accounts for all possible values of  $X$  and “integrates out” the unobserved data. Intuitively, when  $X$ ’s true value is unknown, MLA uses information about how the probability of the observed  $W$  depends on the possible values of  $X$  and the probability that  $X$  takes each of these values instead. The difference in the DV case is that there is no exposure model because the main model gives the probability of  $Y$ .

MLA has four advantages in the context of ACs that reflect the benefits of integrating out partially observed discrete variables and maximum likelihood estimation (MLE).<sup>7</sup> First, it is generally applicable to any model with a convex likelihood (i.e., an optimizer can find a unique set of parameters that best fit the data). Such models include generalized linear models (GLMs) and generalized additive models (GAMs). Second, assuming the model is correctly specified, MLA estimators are fully consistent (i.e., as sample size increases, estimates converge to the true value) (Carroll et al., 2006). This is related to MLA’s excellent statistical efficiency; it requires less manually annotated data to make precise estimates. Third, it is applicable when either an IV or the DV is automatically classified. Finally, MLA uses established methods for estimating confidence intervals, such as the Fischer information and profile likelihood methods, that effectively account for uncertainty from all the joint model’s components.

### Other approaches to correcting misclassification

We now compare our proposed MLA approach to others recently discussed by social scientists and discuss how the MLA approach is advantageous.

#### Regression calibration

*Regression calibration* uses observable IVs, including automated classifications  $W$  and other variables measured without error  $Z$ , to approximate the true value of  $X$  (Carroll et al., 2006). Fong and Tyler (2021) propose a regression calibration procedure designed for SML that we refer to as *GMM calibration* or GMM.<sup>8</sup> For their calibration model, Fong and Tyler (2021) use 2-stage least squares (2SLS). They regress observed variables  $Z$  and AC predictions  $W$  onto the manually annotated data and then use the resulting model to approximate  $X$  as  $\hat{X}$ . They then use the generalized method of moments (gmm) to combine estimates based on the approximated IV  $\hat{X}$  and estimates based on the manually annotated data  $\tilde{X}$ . This makes efficient use of manually annotated data and provides an asymptotic theory for deriving confidence intervals. An additional strength compared to the other

<sup>7</sup>It is worth noting that an analogous method can be implemented in a Bayesian framework instead of MLE. A Bayesian approach may provide additional strengths in flexibility and uncertainty quantification, but bears additional complexity.

<sup>8</sup>Fong and Tyler (2021) describe their method within an instrumental variable framework, but it is equivalent to regression calibration, the standard term in measurement error literature.



methods is that GMM does not require the analyst to make assumptions about the distribution of the outcome  $Y$ .

The GMM approach bears limitations that MLA overcomes. Systematic misclassification can invalidate GMM when AC predictions  $W$  are not conditionally independent of  $Y$  given  $X$ . In the MLA framework, including  $Y$  in the error model for  $X$  can account for such systematic misclassification. This technique, however, does not translate to the 2SLS framework. 2SLS assumes that the instrumental variable (i.e.,  $W$ ) is conditionally independent of  $Y$  given  $X$ . In other words, regression calibration assumes that measurement error is non-differential (Carroll et al., 2006).

### Multiple imputation

*Multiple imputation* (MI) treats misclassification as a missing data problem. Communication scholars may be familiar with MI as a method for analyzing datasets with missing values, for instance, in survey studies (Cho, 2011; Myers, 2011). MI can also be applied to misclassification bias when validation data are available. The true value of  $X$  can be treated as observed in manually annotated data  $\tilde{X}$  and missing otherwise (Blackwell et al., 2017).

Like MLA, MI uses a model to infer likely values of possibly misclassified variables. The difference is that multiple imputation samples several (hence *multiple* imputation) entire datasets by filling in the missing data from the predictive probability distribution of  $X$  conditional on other variables  $\{W, Y, Z\}$ . MI then runs a statistical analysis on each of these sampled datasets and pools the results of each of these analyses (Blackwell et al., 2012). Note that  $Y$  can be included among the imputing variables, giving MI the potential to address systematic misclassification. Blackwell et al. (2012) claim that MI is relatively robust when it comes to small violations of the assumption of nondifferential error. Moreover, in theory, MI can be used for correcting misclassifications both in IVs and DVs.

MLA's advantages compared to MI come not from its flexibility (MI also affords flexible specifications), but from its constraints tailored to the problem of misclassification bias. MI is designed for applicability to a wide range of problems involving missing continuous as well as categorical variables. This generality is unnecessary for addressing misclassified variables, which are categorical by definition. Somewhat contradictorily, MI assumes that the covariance matrix of the data  $(Y, X, Z)$  is multivariate normal, an assumption that is never true for categorical data and helps explain why (as shown below) MI struggles particularly in the case where an AC predicts  $Y$ . In sum, when it comes to fixing misclassification bias, we argue that MI is both more complex and requires stronger assumptions than necessary.

### Pseudo-likelihood

"Pseudo-likelihood" methods (PL)—even if not always explicitly labeled this way—are another approach for correcting misclassification bias. Zhang (2021) proposes a method that approximates the error model using quantities from the AC's confusion matrix—the positive and negative predictive values in the case of a mismeasured IV and the AC's false positive and false negative rates in the case of a mismeasured DV. Because quantities from the confusion matrix are neither data nor model parameters, Zhang's (2021) method is technically a "Pseudo-likelihood" method. A clear benefit is that this method does not require validation data, only summary statistics which may be easier to obtain. PL is conceptually similar to MLA. Indeed, for non-systematic misclassification, PL and MLA maximize almost identical likelihoods. The difference is that instead of estimating the error model from data, PL uses parameters derived from quantities that summarize misclassification (e.g., precision and recall). These statistics do not account for systematic misclassification and so PL is unsuited for such. Moreover, PL overstates the precision of its estimates because it does not account for uncertainty in these statistics.

## Evaluating misclassification models: Monte Carlo simulations

We now present four Monte Carlo simulations (*Simulations 1a, 1b, 2a, and 2b*) with which we evaluate existing methods (GMM, MI, PL) and our approach (MLA) for correcting misclassification bias.

Monte Carlo simulations are a tool for evaluating statistical methods, including (automated) content analysis (e.g., Bachl & Scharnow, 2017; Fong & Tyler, 2021; Geiß, 2021; Song et al., 2020; Zhang, 2021). They are defined by a data generating process from which datasets are repeatedly sampled. Repeating an analysis for each of these datasets provides an empirical distribution of results the analysis would obtain over study replications. Monte Carlo simulation affords exploration of finite-sample performance, robustness to assumption violations, comparison across several methods, and ease of interpretability (Mooney, 1997). Such simulations allow exploration of how results depend on assumptions about the data-generating process and analytical choices and are thus an important tool for designing studies that account for misclassification.

### Parameters of the Monte Carlo simulations


The simulations reported in our main text are designed to (1) verify that error correction methods are effective when they should be effective in theory, and (2) to demonstrate the promising potential of error correction methods in automated content analysis. The classifiers in these scenarios have predictive performance substantially lower than is commonly held to be necessary for reliable measurement via ACs. We chose to not study all possible combination of simulation parameters to limit the energy consumption of fitting thousands or even millions of models to generated datasets, but focus on the most illustrative and common cases. Our simulations have relatively small samples and large effect sizes, but simulations with larger samples and smaller effects would inevitably obtain qualitatively similar results. Our substantive findings about the methods' effectiveness are not sensitive to these choices as shown by additional simulations in the Appendix with varying classifier accuracy and systematic error, imbalance in predicted variables, and correlation between  $X$  and  $Y$ .

In our simulations, we tested four error correction methods: *GMM calibration* (GMM) (Fong & Tyler, 2021), *multiple imputation* (MI) (Blackwell et al., 2012), Zhang's *pseudo-likelihood model* (PL) (Zhang, 2021), and our *Maximum Likelihood Adjustment* approach (MLA). We use the `PREDICTIONERROR` R package (Fong & Tyler, 2021) for the GMM method, the `AMELIA` R package with 200 imputed datasets for the MI approach, and our own implementation of Zhang's (2021) PL approach in R. We develop our MLA approach in the R package `MISCLASSIFICATION_MODELS`. For PL and MLA, we quantify uncertainty using the Fisher information quadratic approximation.<sup>9</sup>


In addition, we compare these error correction methods to two common approaches in communication science: the *feasible* estimator (i.e., conventional content analysis that uses only manually annotated data and not ACs) and the *naïve* estimator (i.e., using AC-based classifications  $W$  as stand-ins for  $X$ , thereby ignoring misclassification). According to our systematic review, the *naïve* approach reflects standard practice in studies employing SML for text classification.

We evaluate each of the six analytical approaches in terms of *consistency* (whether the estimates of parameters  $\hat{B}_X$  and  $\hat{B}_Z$  have expected values nearly equal to the true values  $B_X$  and  $B_Z$ ), *efficiency* (how precisely the parameters are estimated and how precision improves with additional data), and *uncertainty quantification* (in terms of the empirical coverage probability of the 95% confidence intervals estimated in our simulations. Ideally, 95% confidence intervals contain the true parameter value in 95% of simulations). To evaluate efficiency, we repeat each simulation with different amounts of total observations, i.e., unlabeled data to be classified by an AC (ranging from 1000 to 10,000 observations), and manually annotated observations (ranging from 100 to 400 observations). Since our review indicated that ACs are most often used to create binary variables, we restrict our simulations to misclassifications related to a binary (in-)dependent variable.


<sup>9</sup>The code for reproducing our simulations and our experimental R package is available here: [https://osf.io/pyqf8/?view\\_only=c80e7b76d94645bd9543f04c2a95a87e](https://osf.io/pyqf8/?view_only=c80e7b76d94645bd9543f04c2a95a87e).




(a) In Simulation 1a, classifications  $W$  are conditionally independent of  $Y$  so a model using  $W$  as a proxy for  $X$  has non-differential error.




(b) In Simulation 1b, the edge from  $W$  to  $Y$  signifies that the automatic classifications  $W$  are not conditionally independent of  $Y$  given  $X$ , indicating differential error.



(c) In Simulation 2a, an unbiased classifier measures the outcome.



(d) In Simulation 2b, the edge connecting  $W$  and  $Z$  signifies that the predictions  $W$  are not conditionally independent of  $Z$  given  $Y$ , indicating systematic misclassification.



**Figure 2.** Bayesian networks representing the conditional independence structure of our simulations.

Four bayesian networks ('a', 'b', 'c', and 'd') illustrate the four simulations. All the networks have 4 variables:  $W, X, Y$ , and  $Z$ . Gray shading indicates that  $X$  is the automatically classified variable in the first two networks while  $Y$  is the automatically classified variable in the last two networks. All of the networks have connections between  $X$  and  $Z$ ,  $X$  and  $Y$ , and  $Z$  and  $Y$ . Network 'a' corresponds to simulation 1a and illustrates non-differential error because  $W$ , the proxy variable, is connected to  $X$  but not to  $Y$ . Network 'b', corresponding to simulation 1b, is like network 'a' but with a connection between  $W$  and  $Y$  illustrating differential error. In network 'c', which represents simulation 2a,  $W$  is disconnected from  $X$  and  $Z$  to signify that the classifiers errors are not correlated with the independent variables. Simulation 2b corresponds to network 'd' in which  $W$  is connected to  $Z$ .

### Four prototypical scenarios for our Monte Carlo simulations

We simulate regression models with two IVs ( $X$  and  $Z$ ). This sufficiently constrains our study's scope, but the scenario is general enough to be applied in a wide range of research studies.

Whether the methods we evaluate below are effective or not depend on the conditional dependence structure among IVs  $X$  and  $Z$ , the DV  $Y$ , and automated classifications  $W$ . This structure determines if adjustment for systematic misclassification is required (Carroll et al., 2006). Bayesian networks provide a useful graphical representation of conditional independence structure (Pearl, 1986) and we use such networks to illustrate our scenarios in Figure 2.<sup>10</sup> An edge between two variables in these networks indicates that they have a direct relationship. Two nodes are *conditionally independent* given a set of other variables that *separate* them (i.e., the two nodes would be disconnected if the set of nodes

<sup>10</sup>Bayesian networks are similar to causal directed acyclic graphs (DAGs), an increasingly popular representation in causal inference methodology. DAGs are Bayesian networks with directed edges that indicate the direction of cause and effect. We use Bayesian networks for generality and because causal direction is not important in the measurement error theory we use.

is removed from the network). For example, in [Figure 2a](#), the automatic classifications  $W$  are conditionally independent of  $Y$  given  $X$  because all paths between  $W$  and  $Y$  contain  $X$ . This indicates that the model  $Y = B_0 + B_1 W + B_2 Z$  (the *naïve estimator*) has non-differential error because the automated classifications  $W$  are conditionally independent of  $Y$  given  $X$ . However, in [Figure 2b](#), there is an edge between  $W$  and  $Y$  to indicate that  $W$  is not conditionally independent of  $Y$  given other variables. Therefore, the naïve estimator has differential error. To illustrate, this figure corresponds to our first real-data example in which Perspective’s toxicity classifications  $W$  are correlated with identity disclosure  $Y$  even after accounting for  $X$  (as follows from the correlation of misclassification ( $W - X$ ) with  $Y$ ).

We first simulate two cases where an AC measures an IV without (*Simulation 1a*) and with differential error (*Simulation 1b*). Then, we simulate using an AC to measure a DV, either one with misclassifications that are uncorrelated (*Simulation 2a*) or correlated with an IV (*Simulation 2b*). GMM is not designed to correct misclassifications in DVs, so we omit this method in *Simulations 2a* and *2b*.

Our main simulations below are designed to demonstrate how error correction methods can use validation data together with low-quality automated classifications to obtain precise and consistent estimates and to clarify the conditions where our proposed method is effective and others are not. Additional simulations presented in the Appendix evaluate the methods with varying degrees of classifier accuracy, systematic misclassification, skewness in the classified variable, error model misspecification, violation of parametric assumptions, and correlation between the IVs. These additional simulations show that our main qualitative results are not sensitive to our choices of simulation parameters and help support our recommendations.

### Misclassification in an independent variable (*Simulations 1a and 1b*)

We first consider studies with the goal of testing hypotheses about the coefficients for  $X$  and  $Z$  in a least squares regression. In this simulated example,  $Y$  is a continuous DV,  $X$  is a balanced ( $P(X) = 0.5$ ) binary IV measured with an AC,  $Z$  is a normally distributed IV with mean 0 and standard deviation 0.5 measured without error, and  $X$  and  $Z$  are correlated (Pearson’s  $\rho = 0.24$ ). To represent a study design where an AC is needed to obtain sufficient statistical power,  $Z$  and  $X$  can explain only 10% of the variance in  $Y$ .


In *Simulation 1a* ([Figure 2a](#)), we simulate an AC with 72% accuracy.<sup>11</sup> This reflects a situation where  $X$  may be difficult to predict, but the AC, represented as a logistic regression model having linear predictor  $W^*$ , provides a useful signal. We simulate nondifferential misclassification because  $W = X + \xi$ ,  $\xi$  is normally distributed with mean 0, and  $\xi$  and  $W$  are conditionally independent of  $Y$  given  $X$  and  $Z$ .

In our real-data example, we included an example where the Perspective API disproportionately misclassified comments as toxic if they disclosed aspects of identities which resulted in differential misclassification. In *Simulation 1b* ([Figure 2b](#)), we test how error correction methods can handle such differential error by making AC predictions similarly depend on the DV  $Y$ . This simulated AC has 74% accuracy and makes predictions  $W$  that are negatively correlated with the residuals of the linear regression of  $X$  and  $Z$  on  $Y$  (Pearson’s  $\rho = -0.17$ ). As a result, this AC makes fewer false-positives and more false-negatives at greater levels of  $Y$ .

### Measurement error in a dependent variable (*Simulations 2a and 2b*)

We then simulate using an AC to measure the DV  $Y$  which we aim to explain given a binary IV  $X$  and a continuous IV  $Z$ . The goal is to estimate coefficients for  $X$  and  $Z$  in a logistic regression model. In

<sup>11</sup>Classifier accuracy varies between our simulations because it is difficult to jointly specify classifier accuracy and the required correlations among variables and due to random variation between simulation runs. We report the median accuracy over simulation runs.



**Figure 3.** Simulation 1a: Nonsystematic misclassification of an independent variable. Error correction methods, except for PL, obtain precise and accurate estimates given sufficient manually annotated data.


A figure presents results from Simulation 1a in a 3x2 grid of subplots. Each subplot visualizes how each of the 6 methods estimate the coefficient of  $X$ . A black point shows the average estimate and a black line shows the region where 95% of the estimates fall. A quantile dotplot overlaid on the black point and lines visualizes the distribution of estimates over repetitions of the simulation. Moving left to right are simulations of datasets having 1000, 5000, and 10,000 total data points. The top row of subplots shows simulations with 400 gold-standard annotations and the bottom row has simulations with 100 annotations.

*Simulation 2a* (see Figure 2c) and *Simulation 2b* (see Figure 2d)  $X$  and  $Z$  are, again, balanced ( $P(X) = P(Z) = 0.5$ ) and correlated (Pearson's  $\rho = 0.24$ ).

In *Simulation 1*, we chose the variance of the normally distributed outcome given our chosen coefficients  $B_X$  and  $B_Z$ , but this is not appropriate for *Simulation 2*'s logistic regression. We therefore choose, somewhat arbitrarily,  $B_X = 0.7$  and  $B_Z = -0.7$ . We again simulate ACs with moderate predictive performance. The AC in *Simulation 2a* is 72% accurate, and the AC in *Simulation 2b* is 73% accurate. In *Simulation 2a*, the misclassifications are nonsystematic as  $\xi$  has mean 0 and is independent of  $X$  and  $Z$ . However, in *Simulation 2b* the misclassifications  $\xi$  are systematic and correlated with  $X$  (Pearson's  $\rho = 0.1$ ).

## Simulation results

For each method, we visualize the consistency and efficiency of estimates across prototypical scenarios by using quantile dot plots with 95% coverage intervals (Fernandes et al., 2018). We report our evaluation of uncertainty quantification in the form of empirical coverage probability in the Appendix to save space. For example, Figure 3 visualizes the results for *Simulation 1a*. Each subplot shows a simulation with a given total sample size (No. observations) and a given sample of manually annotated observations (No. annotations). To assess a method's consistency, we locate the expected value of the point estimate across simulations with the center of the black circle and observe if the estimates become more precise as sample sizes increase. As an example, see the leftmost column in the bottom-left subplot of Figure 3. For the naïve estimator, the circle is far to the left of the dashed line indicating the true value of  $B_X$ . Here, ignoring misclassification causes bias toward 0 and the estimator



**Figure 4.** Simulation 1b: systematic misclassification of an independent variable. Only the MLA and MI approaches obtain consistent estimates of  $B_X$ .

A figure presents results from Simulation 1b in a 3x2 grid of subplots. Each subplot visualizes how each of the 6 methods estimate the coefficient of  $X$ . A black point shows the average estimate and a black line shows the region where 95% of the estimates fall. A quantile dotplot overlaid on the black point and lines visualizes the distribution of estimates over repetitions of the simulation. Moving left to right are simulations of datasets having 1000, 5000, and 10,000 total data points. The top row of subplots shows simulations with 400 gold-standard annotations and the bottom row has simulations with 100 annotations.


is inconsistent. To assess a method's efficiency, we mark the region in which point estimate falls in 95% of the simulations with black lines. The black lines in the bottom-left subplot of Figure 3, for example, show that the feasible estimator, which uses only manually annotated data, is consistent but less precise than estimates from error correction methods.

### Simulation 1a: nonsystematic misclassification of an independent variable

Figure 3 presents results from Simulation 1a. Here, the naïve estimator is severely biased in its estimation of  $B_X$ . Fortunately, error correction methods (GMM, MI, MLA) produce consistent estimates and acceptably accurate confidence intervals. Notably, the PL method is inconsistent, and considerable bias remains when the sample of annotations is much smaller than the entire dataset. This is likely due to  $P(X = x)$  missing from the PL estimation.<sup>12</sup> Figure 3 also shows that MLA and GMM estimates become more precise in larger datasets. As Fong and Tyler (2021) also observed, this precision improvement is less pronounced for MI estimates, indicating that GMM and MLA use automated classifications more efficiently than MI.

In brief, when misclassifications cause nondifferential error, MLA and GMM are effective, efficient, and provide accurate uncertainty quantification. They complement each other due to different assumptions: MLA depends on correctly specifying the likelihood but its robustness to incorrect specifications is difficult to analyze (Carroll et al., 2006). The GMM approach depends on the exclusion restriction instead of distributional assumptions (Fong & Tyler, 2021). In cases similar to

<sup>12</sup>Compare Equation D4 in the Appendix to Equations 24–28 from Zhang (2021).



**Figure 5.** Simulation 2a: Nonsystematic misclassification of a dependent variable. Only the MLA approach is consistent.


A figure presents results from Simulation 2a in a 3x2 grid of subplots. Each subplot visualizes how each of the 6 methods estimate the coefficient of  $X$ . A black point shows the average estimate and a black line shows the region where 95% of the estimates fall. A quantile dotplot overlaid on the black point and lines visualizes the distribution of estimates over repetitions of the simulation. Moving left to right are simulations of datasets having 1000, 5000, and 10,000 total data points. The top row of subplots shows simulations with 400 gold-standard annotations and the bottom row has simulations with 100 annotations.

*Simulation 1a*, we therefore recommend both GMM and MLA to correct for misclassification. The Appendix further explores the trade-off between the two approaches.

### Simulation 1b: systematic misclassification of an independent variable

Figure 4 visualizes *Simulation 1b*. Here, systematic misclassification gives rise to differential error and creates extreme misclassification bias that is more difficult to correct. As Figure 4 shows, the naïve estimator is opposite in sign to the true parameter. Of the four methods we test, only the MLA and the MI approach provide consistent estimates. This is expected because they use  $Y$  to adjust for misclassifications. The bottom row of Figure 4 shows how the precision of the MI and MLA estimates increases with additional observations. As in *Simulation 1a*, MLA uses this data more efficiently than MI does. However, due to the low accuracy and bias of the AC, additional unlabeled data improves precision less than one might expect. Both methods provide acceptably accurate confidence intervals as shown in Figure 4. Figure F3 in the Appendix shows that, as in *Simulation 1a*, effective correction for misclassifications of  $X$  is required to consistently estimate  $B_Z$ , the coefficient of  $Z$  on  $Y$ . Inspecting results from methods that do not correct for differential error is useful for understanding their limitations. When few annotations of  $X$  are observed, GMM is nearly as bad as the naïve estimator. PL is also visibly biased. Both improve when a greater proportion of the data is labeled since they combine AC-based estimates with the feasible estimator.





**Figure 6.** Simulation 2b: systematic misclassification of a dependent variable. Only the MLA approach is consistent.

A figure presents results from Simulation 2b in a 3x2 grid of subplots. Each subplot visualizes how each of the 6 methods estimate the coefficient of X. A black point shows the average estimate and a black line shows the region where 95% of the estimates fall. A quantile dotplot overlaid on the black point and lines visualizes the distribution of estimates over repetitions of the simulation. Moving left to right are simulations of datasets having 1000, 5000, and 10,000 total data points. The top row of subplots shows simulations with 400 gold-standard annotations and the bottom row has simulations with 100 annotations.

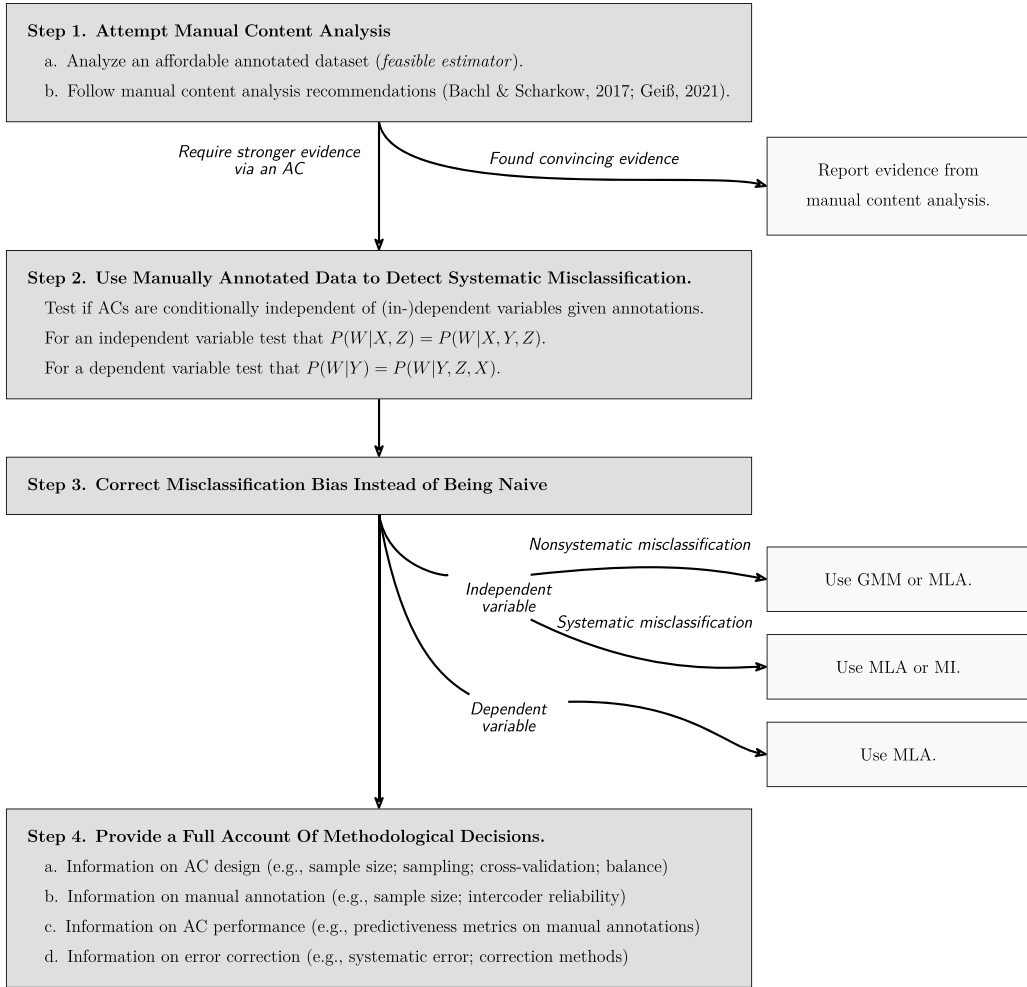
In sum, our simulations suggest that the MLA approach is superior in conditions of differential error. Although estimates by the MI approach appear consistent, the method's practicality is limited by its inefficiency.

### Simulation 2a: nonsystematic misclassification of a dependent variable

Figure 5 illustrates *Simulation 2a*: nonsystematic misclassification of a dependent variable. This also introduces bias as evidenced by the naïve estimator's inaccuracy. Our MLA method is able to correct this error and provide consistent estimates.

Surprisingly, the MI estimator is inconsistent and does not improve with more human-labeled data. The PL approach is also inconsistent, especially when only few of all observations are annotated manually. It is closer to recovering the true parameter than the MI or the naïve estimator, but provides only modest improvements in precision compared to the feasible estimator. It is clear that the precision of the MLA estimator improves with more observations data to a greater extent than the PL estimator. Figure F4 in the Appendix shows that when the amount of human-labeled data is low, the 95% confidence intervals of both the MLA and PL become inaccurate due to the poor finite-sample properties of the quadratic approximation for standard errors.

In brief, our simulations suggest that MLA is the best error correction method when random misclassifications affect the DV. It is the only consistent option and more efficient than the PL method, which is almost consistent.



**Figure 7.** Recommendations for automated content analysis study design.

A flow chart summarizes our recommendations from this section.

Step 1 is to attempt a manual content analysis on an affordable annotated dataset following recommendations from prior work;). If convincing evidence is found this should be reported, but if not strong evidence using an automatic classifier should be sought.

Step 2 is to use the data from step 1 to test for systematic misclassification via conditional independence tests. If an independent variable is to be classified test that the conditional probability of the classifications given the independent variables equals the probability of the automatic classifications given the dependent and independent variables. If a dependent variable is to be classified test that the conditional probability of the classifications given the outcome equals the conditional probability of the classifications given the outcome and the independent variables.

Step 3 is to correct for misclassification bias following our recommendations based on our simulations. In the independent variable case with nonsystematic misclassification we recommend the GMM or MLA strategies. If such a variable is systematically misclassified then we recommend the MLA or MI strategies. If a dependent variable is misclassified then we recommend MLA. Step 4 is to provide a full account of methodological decisions to include information on automatic classifier design (e.g., sample size; sampling; cross-validation; balance), information on manual annotation (e.g., sample size; intercoder reliability), information on AC performance (e.g., predictiveness metrics on manual annotations), an information on error correction (e.g., systematic error; correction methods).

### **Simulation 2b: systematic misclassification of a dependent variable**

In *Simulation 2b*, misclassifications of the dependent variable  $Y$  are correlated with an independent variable  $Z$ . As shown in [Figure 6](#), this causes dramatic bias in the naïve estimator. Similar to *Simulation 2a*, MI is inconsistent. PL is also inconsistent because it does not account for  $Y$  when correcting for misclassifications. As in *Simulation 1b*, our MLA method is

consistent, but only does much better than the feasible estimator when the dataset is large. Therefore, our simulations suggest that MLA is the best method when misclassifications in the DV are correlated with an IV. Figure F9 in the Appendix shows that the precision of estimates for the coefficient for *Z* improves with additional data to a greater extent.

## **Transparency about misclassification is not enough—we have to fix it!**

### **Recommendations for automated content analysis**

“Validate, Validate, Validate” (Grimmer & Stewart, 2013, p. 269) is one of the guiding mantras for automated content analysis. It reminds us that ACs can produce misleading results and of the importance of steps to ascertain validity, for instance, by making misclassification transparent.

Like Grimmer and Stewart (2013), we are deeply concerned that computational methods may produce invalid evidence. In this sense, their validation mantra animates this paper. Critically evaluating the performance of measurement tools by providing transparency about misclassification rates via metrics such as precision or recall is the meaningful first step to build confidence on the tools, but it leaves unanswered an important question: Is comparing automated classifications to some external ground truth sufficient to claim that results are valid? Or is there something else we can do and should do?

We think there is: Using statistical methods to not only quantify but also correct for misclassification. Our study provides several recommendations for using automated classifiers for statistical inference in confirmatory data analysis as summarized in Figure 7.

#### **Step 1: Attempt manual content analysis**

Manual content annotation is often done *post facto*, for instance, to calculate predictiveness of an already existing AC such as Google’s Perspective classifier. We propose to instead use manually annotated data *ante facto*, i.e., before building or validating an AC. Practically speaking, the main reason to use an AC is feasibility: to avoid the costs of manually coding a large dataset. One may, for example, need a large dataset to study an effect one assumes to be small. Manually labeling such a dataset is expensive. Often, ACs are seen as a cost-saving procedure, but scholars often fail to consider the threats to validity posed by misclassification. Moreover, validating an existing AC or building a new AC is also expensive, for instance, due to costs of computational resources or manual annotation of (perhaps smaller) test and training datasets.

We therefore caution researchers against preferring automated over manual content analysis unless doing so is necessary to obtain useful evidence. We agree with Baden et al. (2022) who argue that “social science researchers may be well-advised to eschew the promises of computational tools and invest instead into carefully researcher-controlled, limited-scale manual studies” (p. 11). In particular, we recommend using manually annotated data *ante facto*: Researchers should begin by examining human-annotated data so to discern if an AC is necessary. In our simulations, the feasible estimator is less precise but consistent in all cases. So if fortune shines and this estimate sufficiently answers one’s research question, manual coding is sufficient. Here, scholars should rely on existing recommendations for descriptive and inferential statistics when using manual content analysis (Bachl & Scharkow, 2017; Geiß, 2021). If the feasible estimator, however, fails to provide convincing evidence, for example, by not rejecting the null, manually annotated data is not wasted. It can be reused to build an AC or to correct misclassification bias.

#### **Step 2: Use manually annotated data to detect systematic misclassification**

As demonstrated in our simulations, knowing whether an AC makes systematic misclassifications is important: It determines which correction methods can work. Fortunately, manually

annotated data can be used to detect systematic misclassification. For example, Fong and Tyler (2021) suggest using Sargan's J-test of the null hypothesis that the product of the AC's predictions and regression residuals has an expected value of 0. More generally, one can test if the data's conditional independence structures can be represented by Figures 2a or 2c. This can be done, for example, via likelihood ratio tests of  $P(W|X, Z) = P(W|X, Y, Z)$  (if an AC measures an IV  $X$ ) or of  $P(W|Y) = P(W|Y, Z, X)$  (if an AC measures a DV  $Y$ ) or by visual inspection of plots of relating misclassifications to other variables (Carroll et al., 2006). We strongly recommend using such methods to test for systematic misclassification and to design an appropriate correction.

### **Step 3: Correct for misclassification bias instead of being naïve**

Across our simulations, we showed that the naïve estimator is biased. Testing different error correction methods, we found that these generate different levels of consistency, efficiency, and accuracy in uncertainty quantification. That said, our proposed MLA method should be considered as a versatile method because it is the only method capable of producing consistent estimates in prototypical situations studied here. We recommend the MLA method as the first “go-to” method. As shown in the Appendix, this method requires specifying a valid error model to obtain consistent estimates. One should take care that the model does not have omitted variables including nonlinearities and interactions. Our misclassificationmodels R package provides reasonable default error models and a user-friendly interface to facilitate adoption of our MLA method (see the Appendix).

When feasible, we recommend comparing the MLA approach to another error correction method. Consistency between two correction methods shows that results are robust independent of the correction method. If the AC is used to predict an IV, GMM is a good choice if error is nondifferential. Otherwise, MI can be considered. Unfortunately, if the AC is used to predict a DV, our simulations do not support a strong suggestion for a second method. PL might be useful reasonable choice with enough manually annotated data and non-differential error. This range of viable choices in error correction methods also motivates our next recommendation.

### **Step 4: Provide a full account of methodological decisions**

Finally, we add our voices to those recommending that researchers report methodological decisions so others can understand and replicate their design (Pipal et al., 2022; Reiss et al., 2022), especially in the context of machine learning (Mitchell et al., 2019). These decisions include but are not limited to choices concerning test and training data (e.g., size, sampling, split in cross-validation procedures, balance), manual annotations (size, number of annotators, intercoder values, size of data annotated for intercoder testing), and the classifier itself (choice of algorithm or ensemble, different accuracy metrics). They extend to reporting error correction methods and the choices involved as proposed by our third recommendation. In our review, we found that reporting such decisions is not yet common, at least in the context of SML-based text classification. When correcting for misclassification, uncorrected results will often provide a lower-bound on effect sizes; corrected analyses will provide more accurate but less conservative results. Therefore, both corrected and uncorrected estimates should be presented as part of making potential multiverses of findings transparent.

## **Conclusion, limitations, and future directions**

Misclassification bias is an important threat to validity in studies that use automated content analysis to measure variables. As we showed in an example with data from the Perspective API, widely used and very

accurate automated classifiers can cause type-I and type-II errors. As evidenced by our literature review, this problem has attracted enough attention neither within communication science (but see Bachl & Scharkow, 2017) nor in the broader computational social science community. Although reporting metrics of classifier performance on manually annotated validation data such as precision or recall, is important, such practices provide little protection from misclassification bias. They use annotations to enact a transparency ritual to ward against misclassification bias, but annotations can do much more. With the right statistical model, they can correct misclassification bias.

We introduce *Maximum Likelihood Adjustment* (MLA), a new method we designed to correct misclassification bias. We then use Monte Carlo simulations to evaluate the MLA approach in comparison to other error correction methods. Our MLA method is the only one that is effective across the full range of scenarios we tested. It is also straightforward to use. Our implementation in the R package `MISCLASSIFICATION_MODELS` provides a familiar formula interface for regression models. Remarkably, our simulations show that our method can even use an automated classifier below common accuracy standards to obtain consistent estimates. Therefore, low accuracy is not necessarily a barrier to using an AC.

Based on these results, we provide four recommendations for the future of automated content analysis: Researchers should (1) attempt manual content analysis before building or validating ACs to see whether human-labeled data is sufficient, (2) use manually annotated data to test for systematic misclassification and choose appropriate error correction methods, (3) correct for misclassification via error correction methods, and (4) be transparent about the methodological decisions involved in AC-based classifications and error correction.

Our study has several limitations. First, the simulations and methods we introduce focus on misclassification by automated tools. They provisionally assume that human annotators do not make errors, especially not systematic ones. This assumption can be reasonable if intercoder reliability is very high, but, as with ACs, this may not always be the case. Thus, it may be important to account for measurement error by human coders (Bachl & Scharkow, 2017) and by automated classifiers simultaneously. In theory, it is possible to extend our MLA approach in order to do so (Carroll et al., 2006). However, because the true values of content categories are never observed, accounting for automated and human misclassification at once requires latent variable methods that bear considerable additional complexity and assumptions (Pepe & Janes, 2007). We leave the integration of such methods into our MLA framework for future work. In addition, our method requires an additional assumption that the error model is correct. As we argue in the Appendix, this assumption is often acceptable.

Second, the simulations we present do not consider all possible factors that may influence the performance and robustness of error correction methods. We extend our simulations to investigate factors including classifier accuracy, heteroskedasticity, and violations of distributional assumptions in the Appendix.

Third, the methods we test and the simulations we present can provide little guidance about how much validation data is required for a given level of statistical precision. The question of how a researcher can optimally allocate a budget between costly validation data and cheap but error-prone data has been studied in the statistical literature on measurement error (e.g., Amorim et al., 2021). Future work should develop and apply such methods for automated content analysis.

More generally, we have shown that misclassification bias can result in misleading real-world findings, that it is too often unconsidered and unaddressed, and that statistical methods can be effective at correcting it. Yet, we wish to know how consequential misclassification bias is to real-world empirical studies so as to give more precise recommendations about when error correction methods are necessary. Future work that replicates past studies employing automated content analysis and extends them with error correction methods can show how important error correction methods are to a range of real-world study designs.

## Author note

We have uploaded a preprint of this article (TeBlunthuis et al., 2023). An online appendix with additional results and discussion complements the main text and is available in an Open Science Framework (OSF) repository <https://osf.io/pyqf8/files/osfstorage/656f977c783ec62762eddabf>.

## Acknowledgement

Dr. TeBlunthuis lead the project conception, wrote the initial draft, and developed the simulations, real-data example, and error correction methods. Dr. Hase designed and lead the literature review. Dr. Chan contributed to the project conception, to the literature review, by leading the R package development, and by advising. All authors contributed to writing.

We thank the Computational Methods Hackathon pre-conference at the 2022 annual meeting of the International Communication Association (ICA). Dr. TeBlunthuis presented a workshop on misclassification bias at this hackathon and connected with Dr. Chan, who was already working on this problem, and with Dr. Hase to begin our collaboration. We also thank our colleagues who provided feedback on this manuscript at various stages including: Benjamin Mako Hill, Nick Vincent, Aaron Shaw, Ruijia Cheng, Kaylea Champion, and Jeremy Foote with the Community Data Science Collective as well as Abigail Jacobs, Manuel Horta Ribero, Nicolai Berk, and Marko Bachl. We also thank the anonymous reviewers with the computational methods division at ICA 2023 whose comments helped improve this manuscript and to the division's awards committee for the recognition. We similarly thank the editorial team and anonymous reviewers at Communication Methods & Measures for the invaluable comments. Any remaining errors are our own. Additional thanks to those who shared data with Dr. TeBlunthuis during the project's conception: Jin Woo Kim, Sandra Gonzalez-Bailon, and Manlio De Domenico. This work was facilitated through the use of advanced computational, storage, and networking infrastructure provided by the Hyak supercomputer system at the University of Washington.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## Notes on contributors

**Dr. Nathan TeBlunthuis** (PhD University of Washington, 2021) is a Postdoctoral Research Fellow at the University of Michigan School of Information. His scholarship focuses on online communities, digitally mediated organization, collaborative knowledge work, and computational methods for communication research.

**Dr. Valerie Hase** is a postdoctoral researcher at the Department of Media and Communication at LMU Munich, Germany. Her research focuses on automated content analysis, digital trace data, bias in computational methods, and digital journalism.

**Dr. Chung-hong Chan** (PhD University of Hong Kong, 2018) is Senior Researcher in the Department of Computational Social Science, GESIS – Leibniz Institute for the Social Sciences, Cologne, Germany, and External Fellow at the Mannheim Center for European Social Research, University of Mannheim (Germany). An epidemiologist by training, he is interested in developing new quantitative methods for communication research.

## ORCID

Nathan TeBlunthuis  <http://orcid.org/0000-0002-3333-5013>

Valerie Hase  <http://orcid.org/0000-0001-6656-4894>

Chung-Hong Chan  <http://orcid.org/0000-0002-6232-7530>

## References

- Amorim, G., Tao, R., Lotspeich, S., Shaw, P. A., Lumley, T., & Shepherd, B. E. (2021). Two-phase sampling designs for data validation in settings with covariate measurement error and continuous outcome. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 184(4), 1368–1389. <https://doi.org/10.1111/rssa.12689>
- Bachl, M., & Scharkow, M. (2017). Correcting measurement error in content analysis. *Communication Methods and Measures*, 11(2), 87–104. <https://doi.org/10.1080/19312458.2017.1305103>



- Baden, C., Pipal, C., Schoonvelde, M., & van der Velden, M. A. C. G. (2022). Three gaps in computational text analysis methods for social sciences: A research agenda. *Communication Methods and Measures*, 16(1), 1–18. <https://doi.org/10.1080/19312458.2021.2015574>
- Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness in machine learning. fairmlbook.org.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Conference on Fairness, Accountability, and Transparency (FACt '21)*, New York, NY, USA (pp. 610–623). <https://doi.org/10.1145/3442188.3445922>.
- Blackwell, M., Honaker, J., & King, G. (2012). Multiple Overimputation: A Unified Approach to Measurement Error and Missing Data. *Sociological Methods & Research*, 46(3), 303–341. <https://doi.org/10.1177/0049124115585360>
- Blackwell, M., Honaker, J., & King, G. (2017). A unified approach to measurement error and missing data: Details and extensions. *Sociological Methods & Research*, 46(3), 342–369. <https://doi.org/10.1177/0049124115589052>
- Breiman, L. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science*, 16(3), 199–231. <https://doi.org/10.1214/ss/1009213726>
- Buonaccorsi, J. P. (2010). *Measurement error: Models, methods, and applications*. Chapman and Hall/CRC.
- Burscher, B., Odijk, D., Vliegthart, R., de Rijke, M., & de Vreese, C. H. (2014). Teaching the computer to code frames in news: Comparing two supervised machine learning approaches to frame analysis. *Communication Methods and Measures*, 8(3), 190–206. <https://doi.org/10.1080/19312458.2014.937527>
- Burscher, B., Vliegthart, R., & De Vreese, C. H. (2015). Using supervised machine learning to code policy issues: Can classifiers generalize across contexts? *The ANNALS of the American Academy of Political and Social Science*, 659(1), 122–131. <https://doi.org/10.1177/0002716215569441>
- Carroll, R. J., Ruppert, D., Stefanski, L. A., & Crainiceanu, C. M. (2006). *Measurement error in nonlinear models* (2nd ed.). Chapman & Hall/CRC.
- Cho, J. (2011). The geography of political communication: Effects of regional variations in campaign advertising on Citizen communication. *Human Communication Research*, 37(3), 434–462. <https://doi.org/10.1111/j.1468-2958.2011.01406.x>
- cjadams.,Borkan, D., inversion., Sorensen, J., Dixon, L., Vasserman, L., & nithum. (2019). Jigsaw Unintended Bias in Toxicity Classification. Kaggle. <https://kaggle.com/competitions/jigsaw-unintended-bias-in-toxicity-classification>
- Fernandes, M., Walls, L., Munson, S., Hullman, J., & Kay, M. (2018). Uncertainty displays using quantile dotplots or cdfs improve transit decision-making. *The 2018 CHI Conference on Human Factors in Computing Systems*, April 21–16, 2018, Montreal, QC, Canada (pp. 1–12). Association for Computing Machinery. <https://doi.org/10.1145/3173574.3173718>
- Fong, C., & Tyler, M. (2021). Machine learning predictions as regression covariates. *Political Analysis*, 29(4), 467–484. <https://doi.org/10.1017/pan.2020.38>
- Fuller, W. A. (1987). *Measurement error models*. Wiley.
- Geiß, S. (2021). Statistical Power in content analysis designs: How effect size, sample size and coding accuracy jointly affect hypothesis testing – a monte carlo simulation approach. *Computational Communication Research*, 3(1), 61–89. <https://doi.org/10.5117/CCR2021.1.003.GEIS>
- Grimmer, J., Roberts, M. E., & Stewart, B. M. (2021). Machine Learning for Social Science: An Agnostic Approach. *Annual Review of Political Science*, 24(1), 395–419. <https://doi.org/10.1146/annurev-polisci-053119-015921>
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297. <https://doi.org/10.1093/pan/mps028>
- Hase, V., Mahl, D., & Schäfer, M. S. (2022). Der „Computational Turn“: Ein „interdisziplinärer Turn“? Ein systematischer Überblick zur Nutzung der automatisierten Inhaltsanalyse in der Journalismusforschung. *Medien & Kommunikationswissenschaft*, 70(1–2), 60–78. <https://doi.org/10.5771/1615-634X-2022-1-2-60>
- Hede, A., Agarwal, O., Lu, L., Mutz, D. C., & Nenkova, A. (2021). From toxicity in online comments to incivility in American News: Proceed with caution. *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume* (pp. 2620–2630). <https://doi.org/10.18653/v1/2021.eacl-main.225>
- Hopkins, D. J., & King, G. (2010). A method of automated nonparametric content analysis for social science. *American Journal of Political Science*, 54(1), 229–247. <https://doi.org/10.1111/j.1540-5907.2009.00428.x>
- Hopp, T., & Vargo, C. J. (2019). Social Capital as an inhibitor of online political incivility: An analysis of behavioral patterns among politically active facebook users. *International Journal of Communication*, 13, 21.
- Hosseini, H., Kannan, S., Zhang, B., & Poovendran, R. (2017). Deceiving Google’s Perspective API Built for Detecting Toxic Comments. *arXiv preprint*. <https://arxiv.org/abs/1702.08138>
- Jünger, J., Geise, S., & Hännelt, M. (2022). Unboxing computational social media research from a datahermeneutical perspective: How do scholars address the tension between automation and interpretation? *International Journal of Communication*, 16, 1482–1505. <https://ijoc.org/index.php/ijoc/article/view/17477/3714>
- Kim, J. W., Guess, A., Nyhan, B., & Reifler, J. (2021). The distorting prism of social media: How self-selection and exposure to incivility fuel online comment toxicity. *Journal of Communication*, 71(6), 922–946. <https://doi.org/10.1093/joc/jqab034>
- Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science*, 355(6325), 584–585. <https://doi.org/10.1126/science.aal3618>



- Millimet, D. L., & Parmeter, C. F. (2022). Accounting for skewed or one-sided measurement error in the dependent variable. *Political Analysis*, 30(1), 66–88. <https://doi.org/10.1017/pan.2020.45>
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model Cards for Model Reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, New York, NY, USA (pp. 220–229). Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287596>
- Mooney, C. Z. (1997). *Monte Carlo simulation*. Sage Publications, Inc.
- Myers, T. A. (2011). Goodbye, listwise deletion: Presenting hot deck imputation as an easy and effective tool for handling missing data. *Communication Methods and Measures*, 5(4), 297–310. <https://doi.org/10.1080/19312458.2011.624490>
- Oppenhuizen, A. E., Schouten, K., & Klijn, E. H. (2019). Framing a Conflict! How media report on earthquake risks caused by gas drilling: A longitudinal analysis using machine learning techniques of media reporting on gas drilling from 1990 to 2015. *Journalism Studies*, 20(5), 714–734. <https://doi.org/10.1080/1461670X.2017.1418672>
- Pearl, J. (1986). Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29(3), 241–288. [https://doi.org/10.1016/0004-3702\(86\)90072-X](https://doi.org/10.1016/0004-3702(86)90072-X)
- Pepe, M. S., & Janes, H. (2007). Insights into latent class analysis of diagnostic test performance. *Biostatistics*, 8(2), 474–484. <https://doi.org/10.1093/biostatistics/kxl038>
- Pipal, C., Song, H., & Boomgaarden, H. G. (2022). If you have choices, why not choose (and share) all of them? A multiverse approach to Understanding news engagement on Social Media. *Digital Journalism*, 11(2), 255–275. <https://doi.org/10.1080/21670811.2022.2036623>
- Rauchfleisch, A., Kaiser, J., & Zollo, F. (2020). The false positive problem of automatic bot detection in social science research. *PLOS ONE*, 15(10), e0241045. <https://doi.org/10.1371/journal.pone.0241045>
- Reiss, M., Kobilke, L., & Stoll, A. (2022). Reporting supervised text analysis for communication science.
- Scharkow, M. (2013). Thematic content analysis using supervised machine learning: An empirical evaluation using German online news. *Quality & Quantity*, 47(2), 761–773. <https://doi.org/10.1007/s11135-011-9545-7>
- Scharkow, M., & Bachl, M. (2017). How measurement error in content analysis and self-reported media use leads to minimal media effect findings in linkage analyses: A simulation study. *Political Communication*, 34(3), 323–343. <https://doi.org/10.1080/10584609.2016.1235640>
- Song, H., Tolochko, P., Eberl, J.-M., Eisele, O., Greussing, E., Heidenreich, T., Lind, F., Galyga, S., & Boomgaarden, H. G. (2020). In validations we trust? The impact of imperfect human annotations as a gold standard on the quality of validation of automated content analysis. *Political Communication*, 37(4), 550–572. <https://doi.org/10.1080/10584609.2020.1723752>
- TeBlunthuis, N., Hase, V., & Chan, C.-H. (2023). *Misclassification in automated content analysis causes bias in regression. Can We Fix It? Yes We Can!*.
- TeBlunthuis, N., Hill, B. M., & Halfaker, A. (2021). Effects of algorithmic flagging on fairness: Quasi-experimental evidence from Wikipedia. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), :56:1–:56:27. <https://doi.org/10.1145/3449130>
- van Atteveldt, W., van der Velden, M. A. C. G., & Boukes, M. (2021). The validity of sentiment analysis: Comparing manual annotation, crowd-coding, dictionary approaches, and machine learning algorithms. *Communication Methods and Measures*, 15(2), 121–140. <https://doi.org/10.1080/19312458.2020.1869198>
- van Smeden, M., Lash, T. L., & Groenwold, R. H. H. (2020). Reflection on modern methods: Five myths about measurement error in epidemiological research. *International Journal of Epidemiology*, 49(1), 338–347. <https://doi.org/10.1093/ije/dyz251>
- Vermeer, S., Trilling, D., Kruikemeier, S., & de Vreese, C. (2020). Online news user journeys: The role of social media, news websites, and topics. *Digital Journalism*, 8(9), 1114–1141. <https://doi.org/10.1080/21670811.2020.1767509>
- Votta, F., Noroozian, A., Dobber, T., Helberger, N., & de Vreese, C. (2023). Going micro to go negative?: Targeting toxicity using Facebook and Instagram Ads. *Computational Communication Research*, 5(1), 1–50. <https://doi.org/10.5117/CCR2023.1.001.VOTT>
- Yi, G. Y., Delaigle, A., & Gustafson, P. (Eds.). (2021). *Handbook of measurement error models*. Chapman and Hall/CRC.
- Zhang, H. (2021). *How using machine learning classification as a variable in Regression Leads to attenuation bias and what to do about it* (Preprint ed.). SocArXiv.