

Adjusting for Confounders with Text: Challenges and an Empirical Evaluation Framework for Causal Inference

Galen Weld¹, Peter West¹, Maria Glenski², David Arbour³, Ryan A. Rossi³ and Tim Althoff¹

¹University of Washington

²Pacific Northwest National Laboratory

³Adobe Research

Abstract

Causal inference studies using textual social media data can provide actionable insights on human behavior. Making accurate causal inferences with text requires controlling for confounding which could otherwise impart bias. Recently, many different methods for adjusting for confounders have been proposed, and we show that these existing methods disagree with one another on two datasets inspired by previous social media studies. Evaluating causal methods is challenging, as ground truth counterfactuals are almost never available. Presently, no empirical evaluation framework for causal methods using text exists, and as such, practitioners must select their methods without guidance. We contribute the first such framework, which consists of five tasks drawn from real world studies. Our framework enables the evaluation of *any* causal inference method using text. Across 648 experiments and two datasets, we evaluate every commonly used causal inference method and identify their strengths and weaknesses to inform social media researchers seeking to use such methods, and guide future improvements. We make all tasks, data, and models public to inform applications and encourage additional research.

1 Introduction

The massive volume of social media data offers significant potential to help researchers better understand human behavior by making causal inferences. Researchers often formalize causal inference as the estimation of the *average treatment effect (ATE)* of a specific treatment variable (*e.g.* therapy) on a specific outcome (*e.g.* suicide) (Rubin 2005; Rosenbaum 2010; Keith, Jensen, and O’Connor 2020). A major challenge is adjusting for *confounders* (*e.g.* comments mentioning depression) that affect both the treatment and outcome (depression affects both an individual’s propensity to receive therapy and their risk of suicide) (Keith, Jensen, and O’Connor 2020). Without adjusting for depression as a confounder, we might look at suicide rates among therapy patients and those not receiving therapy, and wrongly conclude that therapy causes suicide.

The gold standard for avoiding confounders is to assign treatment via a *randomized controlled trial (RCT)*. Unfortun-

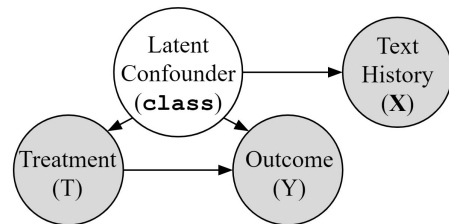


Figure 1: Causal graph representing the the context of our evaluation framework. All edges have known probabilities. While our framework naturally generalizes to more complex scenarios, we chose binary treatments (T) and outcomes (Y), and a binary latent confounder (`class`), as even in this simple scenario, current methods struggle.

nately, in many domains, assigning treatments in this manner is not feasible (*e.g.* due to ethical or practical concerns). Instead, researchers conduct *observational studies* (Rosenbaum 2010), using alternate methods to adjust for confounders.

Text (*e.g.* users’ social media histories) can be used to adjust for confounding by training an NLP model to recognize confounders (or proxies for confounders) in the text, so that similar treated and untreated observations can be compared. However, a recent review (Keith, Jensen, and O’Connor 2020) finds that evaluating the performance of such methods is “a difficult and open research question” as *true ATEs* are almost never known, and so, unlike in other NLP tasks, we cannot know the correct answer. We find that this challenge is amplified, as methods disagree with one another on real world tasks (§3) – how do we know which is correct?

Theoretical bounds on the performance of methods are almost never tight enough to be informative. We derive such bounds for the methods included here (Appendix A¹) and find that our empirical evaluation framework produces tighter bounds more than 99% of the time. As ground truth is almost never available, the only² practical method to evaluate causal inference methods is with semi-synthetic data, where synthetic treatments and/or outcomes are assigned to

¹Link to Extended Paper with Appendix

²With the extremely rare exception of *constructed* observational studies, conducted with a parallel RCT.

real observations, as in Fig. 1 (Dorie et al. 2019; Jensen 2019; Gentzel, Garant, and Jensen 2019). While widely-used semi-synthetic benchmarks have produced positive results in the medical domain (Dorie et al. 2019), no such benchmark exists for causal inference methods using text (Keith, Jensen, and O’Connor 2020).

In this work, we contribute the first evaluation framework for causal inference with text (§5). Our framework is simple, principled, and can be applied to any method that produces an *ATE* estimate given observed treatment assignments, outcomes, and text covariates. It includes a broad range of tasks challenging enough to identify where current methods fail and the most promising avenues for improvement. However, no single benchmark can adequately evaluate methods for every application. As such, our framework can be easily extended to include additional tasks relevant to any application and potential confounder (§5.3).

Inspired by challenges from a wide range of studies (Johansson, Shalit, and Sontag 2016; De Choudhury et al. 2016; Choudhury and Kiciman 2017; Falavarjani et al. 2017; Olteanu, Varol, and Kiciman 2017; Kiciman, Counts, and Gasser 2018; Sridhar et al. 2018; Saha et al. 2019; Veitch, Sridhar, and Blei 2019; Roberts, Stewart, and Nielsen 2020), our framework consists of five tasks (§5.2): Linguistic Complexity, Signal Intensity, Strength of Selection Effect, Sample Size, and Placebo Test. Each semi-synthetic task is generated from public social media users’ profiles, perturbed with synthetic posts to create increasing levels of difficulty. To increase the robustness of our approach, we evaluate methods on these tasks using real data from both Reddit and Twitter. Not all of these tasks are exclusive to causal inference with text, yet all are important to a great deal of textual causal inference studies. As such, their evaluation in this context is important. Using these tasks, we evaluate the specific strengths and weakness of 9 widely-used methods and 3 common estimators, conducting 648 experiments.

Concerningly, we find that almost every method predicts a false significant treatment effect when none is present, which could be greatly misleading to unwary practitioners (§7). While we find that each method struggles with at least one challenge, methods leveraging recent, hierarchical, transformer-based architectures perform best, although such methods are not yet widely used (Keith, Jensen, and O’Connor 2020). These limitations and findings highlight the importance of continued research on the evaluation of causal inference methods for text.

The ICWSM community consists of researchers who work both on the development of causal inference methods, and practitioners who solve real world problems using causal inference.

For methods developers: We make our framework publicly available³ to enable the evaluation of any causal inference method which uses text, and encourage the development of more robust methods. Our framework can be easily extended to include additional tasks.

For practitioners: We identify strengths and weaknesses of commonly used methods, identifying those best suited for

specific applications, and make these publicly available³.

2 Background and Related Work

Causal Inference with Social Media Data We formalize causal inference using notation from Pearl (1995). Given a series of n observations (in our context, a social media user), each observation is a tuple $O_i = (Y_i, T_i, \mathbf{X}_i)$, where Y_i is the outcome (e.g. did user i develop a suicidal ideation?), T_i is the treatment (e.g. did user i receive therapy?), and \mathbf{X}_i is the vector of observed covariates (e.g. user i ’s textual social media history).

The *Fundamental Problem of Causal Inference* is that each user is either treated or untreated, and so we can never observe *both* outcomes. Thus, we cannot compute the $ATE = \frac{1}{n} \sum_{i=1}^n Y_i [T_i = 1] - Y_i [T_i = 0]$ directly, and must estimate it by finding comparable treated and untreated observations. To do so, it is common practice to use a model to estimate the **propensity score**, $\hat{p}(\mathbf{X}_i) \approx p(T_i = 1 | \mathbf{X}_i)$, for each observation i . As treatments are typically known, propensity score models are effectively supervised classifiers, predicting T_i , given \mathbf{X}_i . Matching, stratifying, or weighting using these propensity scores will produce an unbiased *ATE* estimate if four assumptions hold: all confounders must be observed, outcomes for each user must not be affected by treatment assignments to other users (SUTVA), propensity scores must be accurate, and there must be overlap in the distribution of covariates in the treated and untreated groups (*common support assumption*) (Rosenbaum 2010; Hill and Su 2013). In practice, verifying these assumptions is difficult. In particular, ensuring that all confounding factors are observed is practically impossible in real world applications, hence the need for empirical evaluation.

Causal Inference and NLP Until recently, there has been little engagement between causal inference researchers and the NLP research community (Keith, Jensen, and O’Connor 2020). There are many ways to consider text in a causal context, such as text as a mediator (Veitch, Sridhar, and Blei 2019; Landeiro and Culotta 2016), text as treatment (Wood-Doughty, Shpitser, and Dredze 2018; Egami et al. 2018; Fong and Grimmer 2016; Tan, Lee, and Pang 2014; Zhang, Mullainathan, and Danescu-Niculescu-Mizil 2020), text as outcome (Egami et al. 2018; Zhang et al. 2018), and causal discovery from text (Mani and Cooper 2000; Mirza and Tonelli 2016). However, we narrow our focus to text as a *confounder*. This is an important area of research because the challenge of adjusting for confounding underlies most causal contexts, such as text as treatment or outcome (Keith, Jensen, and O’Connor 2020). Effective adjusting for confounding with text enables causal inference in any situation where observations can be represented with text – e.g. social media, news articles, and dialogue.

Adjusting for Confounding with Text A recent review (Keith, Jensen, and O’Connor 2020, Table 1) summarizes common practices across a diverse range of studies. Almost every method used in practice consists of two parts: a **propensity score model**, which uses some text represen-

³<https://behavioral-data.github.io/CausalInferenceChallenges/>

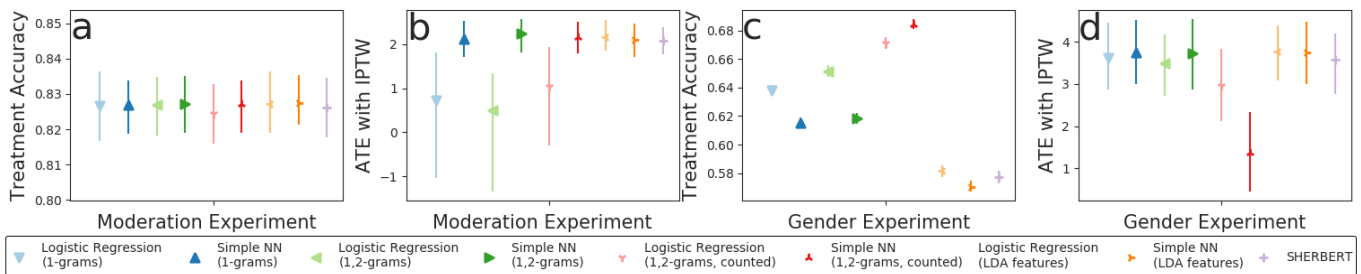


Figure 2: Treatment accuracy and ATE for both real world experiments, with bootstrapped 95% confidence intervals. Note that for the Gender Experiment, the models with the highest accuracy have the lowest ATE .

tation to estimate propensity scores, and an **ATE estimator**. Since such propensity-score based methods are by far the most widely used, in this work, we focus on these methods. While we do not evaluate non-propensity score methods such as doubly-robust methods (Kang and Schafer 2007), TMLE (Schuler and Rose 2017), and matching on values other than propensity scores (Roberts, Stewart, and Nielsen 2020; Mozer et al. 2020), our framework’s structure enables evaluation of *any* ATE estimation method that produces an ATE estimate given observed treatment assignments, outcomes, and text covariates. We believe disentangling the challenges of widely-used methods is a key contribution before moving to more complex and less common methods. Evaluating these other methods is an important area of future work.

Text representations used in propensity score models generally do not yet leverage recent breakthroughs in NLP, and roughly fall into three groups: those using uni- and bi-gram representations (De Choudhury et al. 2016; Johansson, Shalit, and Sontag 2016; Olteanu, Varol, and Kiciman 2017), those using LDA or topic modeling (Falavarjani et al. 2017; Roberts, Stewart, and Nielsen 2020; Sridhar et al. 2018), and those using neural word embeddings such as GLoVe (Pham and Shen 2017), fastText (Joulin et al. 2017; Chen, Montano-Campos, and Zadrozny 2020), or BERT (Veitch, Sridhar, and Blei 2019), (Pryzant et al. 2018). Three classes of estimators are commonly used to compute the ATE : **inverse probability of treatment weighting** (IPTW), propensity score **stratification**, and **matching**, either using propensity scores or, less frequently, some other distance metric. In our evaluation, we separate the propensity score models from the ATE estimators to better understand each component’s individual impact.

Evaluation of Causal Inference In rare specialized cases, researchers can use the unbiased outcomes of a parallel RCT to evaluate those of an observational study, as in Eckles and Bakshy (2017). This practice is known as a *constructed* observational study, and, while useful, is only possible where parallel RCTs can be conducted. Outside these limited cases, proposed models are typically evaluated on synthetic data generated by their authors. These synthetic datasets often favor the proposed model, and do not reflect the challenges faced by real applications (Keith, Jensen, and O’Connor 2020).

Theoretical evaluation of causal inference methods is

generally unhelpful, as theoretically derived performance bounds are mostly much less informative than those derived empirically (Arbour and Dimmery 2019). In this work, we compute the theoretical bounds and find that they are so loose as to not effectively guide practitioners in selecting methods. Our empirical evaluation framework based on realistic tasks produces tighter bounds in more than 99% of cases (Appendix A).

Outside of the text domain, *widely used* empirical evaluation datasets have been successful, most notably the 2016 Atlantic Causal Inference Competition (Dorie et al. 2019), and a strong case has been made for the empirical evaluation of causal inference models (Gentzel, Garant, and Jensen 2019; Jensen 2019; Lin et al. 2019). In the text domain, matching approaches have been evaluated empirically (Mozer et al. 2020), but this approach evaluates only the quality of *matches*, not the causal effect estimates. In contrast, our work applies to all estimators, not just matching, and evaluates the entire causal inference pipeline.

3 Current Models Disagree

Recent causal inference papers (Veitch, Sridhar, and Blei 2019; Roberts, Stewart, and Nielsen 2020; De Choudhury et al. 2016; Chandrasekharan et al. 2017; Bhattacharya and Mehrotra 2016) have used social media histories to adjust for confounding. Each of these papers uses a different propensity score model: BERT in Veitch, Sridhar, and Blei (2019), topic modeling in Roberts, Stewart, and Nielsen (2020), logistic regression in De Choudhury et al. (2016), Mahalanobis distance matching in Chandrasekharan et al. (2017), and Granger Causality in Bhattacharya and Mehrotra (2016). For all of these studies, ground truth causal effects are unavailable, and so we cannot tell if the chosen model was correct. However, we *can* compute the accuracy of their propensity scores (accuracy of a binary classifier predicting treatment assignment), and see if their ATE estimates agree—if they don’t, then at most one disagreeing model can be correct.

Methods We conducted two experiments using real world data from Reddit, inspired by these recent papers. In the **Moderation Experiment**, we test if having a post removed by a moderator impacts the amount a user later posts to the same community again. In the **Gender Experiment**, we use data from Veitch, Sridhar, and Blei (2019) to study the impact of the author’s gender on the score of their posts. For

details on data collection, see Appendix B.

Results Comparing the performance of nine different methods (Fig. 2), we find that all models have similar treatment accuracy in the Moderation Experiment. However, the models using 1,2-gram features perform better in the Gender Experiment than the LDA and SHERBERT models. Most importantly, we see that while many confidence intervals⁴ overlap, there are notable differences between ATE estimates for different models, even when treatment accuracy is nearly identical (Fig. 2a,b). That some ATE estimates' confidence intervals overlap 0 while others do not (Fig. 2b) indicates that some models find nonzero treatment effects at the common $p = 0.05$ threshold while others do not. Lastly, we note that models with the highest treatment accuracy tend to have the lowest ATE estimates (Fig. 2c,d).

Implications This should come as a great concern to the computational social science research community. We do not know which model may be correct, and we do not know whether there may be a more accurate model that would even further decrease the estimated treatment effect. We derive theoretical bounds and compute them (Appendix A), finding that in 99+% of cases, these bounds are looser than those computed empirically using our framework, making them less useful for model selection. This concern underlines the importance and urgency of empirical evaluation for causal inference with text, and motivates our contribution. Next, we describe key challenges in adjusting for confounding with text and present a principled evaluation framework that highlights these challenges and generates actionable insights for future research.

4 Challenges for Causal Inference with Text

Using the common setting of real social media histories (De Choudhury et al. 2016; Olteanu, Varol, and Kiciman 2017; Veitch, Sridhar, and Blei 2019; Choudhury and Kiciman 2017; Falavarjani et al. 2017; Kiciman, Counts, and Gasser 2018; Saha et al. 2019; Roberts, Stewart, and Nielsen 2020), we identify five challenges consistently present when representing natural language for causal inference:

1. **Linguistic Complexity:** Natural language uses a diverse set of tokens to express related underlying meaning. Someone who struggles with mental health might write “I feel depressed” or “I am isolated from my peers,” which have distinct tokens but both may be indicative of depression. *Can models recognize a range of expressions which are correlated with treatment?*
2. **Signal Intensity:** Some users only have a few posts that contain a specific signal (such as poor mental health) whereas others may have many posts with this signal. Signals are especially weak when posts containing the signal constitute only a small fraction of a user's posts. *Can models detect weak signals?*

⁴In these experiments, as well as all following experiments, results are reported with bootstrapped 95% confidence intervals computed by resampling from the population and recomputing the estimators (§6.2) and evaluation metrics.

3. **Strength of Selection Effect:** Many studies have few comparable treated and untreated users (Li, Thomas, and Li 2018; Crump et al. 2009). *Can models adjust for strong selection effects?*
4. **Sample Size:** Observational studies often face data collection limitations.⁵ *Can models perform well with limited data samples?*
5. **Placebo Test:** Oftentimes, no causal effect is present between a given treatment and an outcome. *Do models falsely predict causality when none is present?*

While natural language is far more complex than any finite set of challenges can capture, the five we have chosen to highlight are challenges that regularly need to be addressed in many causal inference applications that use natural language. This set of challenges was developed by reviewing a broad set of existing studies (Johansson, Shalit, and Sontag 2016; De Choudhury et al. 2016; Choudhury and Kiciman 2017; Falavarjani et al. 2017; Olteanu, Varol, and Kiciman 2017; Kiciman, Counts, and Gasser 2018; Sridhar et al. 2018; Saha et al. 2019; Veitch, Sridhar, and Blei 2019; Roberts, Stewart, and Nielsen 2020) and identifying commonalities. While the strength of selection effect, sample size, and placebo test challenges are not *exclusive* to causal inference with text, these challenges are present in most real world studies, and as such, a holistic evaluation framework must consider them. These five challenges also cover three key concepts of model performance: generalizability (*linguistic complexity*), sensitivity (*signal intensity, strength of selection effect*), and feasibility (*sample size, placebo test*) that are critical for comprehensive evaluation. To produce our evaluation framework, we derive a concrete task from each challenge.

5 Framework for Evaluation

We generate five tasks, each with discrete levels of difficulty, and corresponding semi-synthetic task datasets based on real social media histories. Without the semi-synthetic component, it would not be possible to empirically evaluate a model, as we would not know the true *ATE* or propensity scores. By basing our user histories on real data, we are able to include much of the realism of unstructured text found ‘in the wild.’ This semi-synthetic approach to evaluation preserves the best of both worlds: the empiricism of synthetic data with the realism of natural data (Jensen 2019; Gentzel, Garant, and Jensen 2019; Jensen 2019).

5.1 Semi-Synthetic Dataset Generation

While the method for generating a semi-synthetic dataset can be arbitrarily complex, we seek the simplest approach which is able to identify and explain where existing methods fail. We generate our datasets according to a simplified model of the universe; where all confounding is present in the text, and where there are only two types of people, `class 1` and `class 2` (Fig. 1). In the context of mental health, for example, these two classes could simply be people who struggle with depression (`class 1`), and those

⁵In Keith, Jensen, and O'Connor (2020, Table 1), 8/12 studies had fewer than 5,000 observations, and 4/12 had fewer than 1,000.

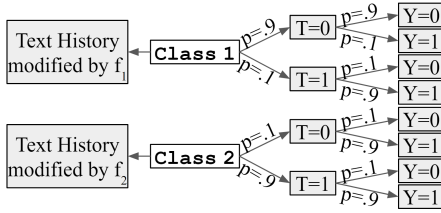


Figure 3: Users are first randomly divided into two latent (unobserved) classes with a 50/50 split, and their text histories have synthetic posts inserted specific to each task. Observed binary treatments and outcomes are assigned with conditional probabilities such that `Class 1` has a true *ATE* of .8, and `Class 2` has a true *ATE* of 0. Since the classes are balanced, the overall true *ATE* is .4.

who don't (`class 2`). If models struggle on even this simple two-class universe, as we find, then it is highly unlikely they will perform better in the more complex real world. In this universe, the user's (latent) class determines the probability of treatment and outcome conditioned on treatment. Dependent on class, but independent of treatment and outcome is the user's comment history, which contains both synthetic and real posts that are input to the model to produce propensity scores. As such, the comment history is an observed proxy for the `class` confounder.

We produce each dataset using a generative process, as shown in Fig. 3. For each task, we start with the same collection of real world user histories from public Reddit or Twitter profiles. We randomly assign an equal number of users to `class 1` and `class 2`. Into each profile, we insert synthetic posts using a function f_n for `class n` specific to each task, described in §5.2. We assign binary treatments (conditioned on class) and binary outcomes (conditioned on class and treatment) according to a known probability distribution (Fig. 3). These outcomes and treatments could represent anything of interest, and they need not be binary.

To estimate the *ATE*, there must be overlap between the treated and untreated groups (*common support*), so we cannot make all users in `class 1` treated and all users in `class 2` untreated. Instead, users in `class 1` are predominantly but not always assigned to treatment (with a .9/.1 split), and users in `class 2` are predominantly but not always assigned to control (also with a .9/.1 split), in order to ensure overlap of covariates between the treated and control groups. This overlap provides common support, and thus our observations do not necessitate trimming (Crump et al. 2009; Lee, Lessler, and Stuart 2011; Yang and Ding 2018).

Once a treatment has been assigned according to the class' probabilities, a positive outcome is assigned with probability .9 (treated) and .1 (untreated) for `class 1`, and .9 regardless of treatment for `class 2`. Thus, `class 1` has a true *ATE* of .8, and `class 2` has a true *ATE* of 0. Since the two classes are balanced, the overall true *ATE* is .4. The objective for propensity score models is to recover the treatment probabilities for each class, which are then used to estimate the true *ATE*.

Real World User Histories In order to maximize generalizability, we experiment with both Reddit and Twitter user histories as the real world component of our semi-synthetic datasets. Reddit and Twitter are natural data sources as they are both publicly accessible, and widely used for relevant research (Medvedev, Lambiotte, and Delvenne 2019; Yu and Muñoz-Justicia 2020).

We downloaded all Reddit comments for the 2014 and 2015 calendar years from the Pushshift archives (Baumgartner et al. 2020) and grouped comments by user. After filtering out users with fewer than 10 comments, we randomly sampled 8,000 users and truncated users' histories to a maximum length of 60 posts for computational practicality.⁶ These users were randomly partitioned into three sets: a 3,200 user training set, an 800 user validation set, and a 4,000 user test set used to compute Treatment Accuracy and *ATE* Bias.

To gather Twitter data, we used a similar method as that used for Reddit. We used the Streaming API to produce a random sample of all public tweets posted in December 2020, then randomly sampled users from this sample and used the Twitter API to gather their complete Tweet histories. As with Reddit histories, we filtered out users with fewer than 10 Tweets, and truncated the histories of users with more than 60 Tweets, then randomly partitioned the resulting users into a training set, a validation set, and a test set, each of the same size as their Reddit data counterparts.⁷

Synthetic Posts When generating semi-synthetic tasks, we insert three types of synthetic posts, representative of major life events that could impact mental health, into real users' Reddit and Twitter histories. Examples are given here, and are listed completely in Appendix C:

- *Sickness Posts* describe being ill (e.g. 'The doctor told me I have AIDS', 'How do I tell my parents I have leukemia?'). We vary both the illness, as well as way the it is expressed.
- *Social Isolation Posts* indicate a sense of isolation or exclusion. ('I feel so alone, my last friend said they needed to stop seeing me.', 'My wife just left me.')
- *Death Posts* describe the death of companion (e.g. 'I just found out my Mom died', 'I am in shock. My son is gone.'). We vary the phrasing as well as the companion.

While these synthetic posts are drawn from the mental health domain, which is commonly represented among previous studies (De Choudhury et al. 2016; Choudhury and Kiciman 2017; Saha et al. 2019), the applicability of our framework is not limited to this specific context. These synthetic posts test a model's ability to recognize different tokens, and the specific tokens used are less critical. Furthermore, the concrete choice of mental health language has the benefit of making the user histories human-readable. Our framework can be

⁶The resulting set of Reddit users had a mean of 41.07 posts per user, mean of 37.37 tokens per post, and a mean of 1523.28 tokens per user.

⁷The resulting set of Twitter users had a mean of 57.76 posts per user, mean of 19.90 tokens per post, and a mean of 1149.59 tokens per user.

easily extended by modifying the synthetic posts to include tokens from different domains (§5.3).

5.2 Tasks

We consider five tasks focused around the previously described common challenges for text-based causal inference methods.

Linguistic Complexity This task tests a model’s ability to recognize a diverse set of tokens as being correlated with treatment. We increase the difficulty in four steps by increasing the diversity of synthetic sentences inserted into user histories assigned to `class 1` (*i.e.* the linguistic complexity of the dataset): f_1 initially appends the same Sickness Post to the end of each `class 1` user’s history; At the second level of difficulty, f_1 selects a Sickness Post uniformly at random; At the third level, f_1 selects either a Sickness or Social Isolation Post; and at the fourth level, f_1 selects a Sickness, Social Isolation, or Death Post. For each level of difficulty, f_2 is the identity function, *i.e.* user histories assigned to `class 2` are unchanged.

Signal Intensity This task tests a model’s ability to distinguish between the number of similar posts in a history. There are two levels of difficulty. At the easier level, f_1 appends 10 randomly sampled (with replacement) Sickness Posts, while f_2 is the identity function. At the harder level, f_1 appends only three Sickness Posts, while f_2 appends one.

Strength of Selection Effect In this and the following tasks, we do not vary f_1 or f_2 . For *Strength of Selection Effect*, we make causal inference more challenging by increasing the strength of the selection effect, decreasing the overlap between treated and untreated users. We test two levels of difficulty: a weaker selection effect (easier) with the same .9/1 split to assign the majority of `class 1` to the treated group and `class 2` to the control group. For the stronger selection effect (harder), we modify the generation framework to increase this split for `class 1` to .95/.05. For both the weak and strong selection effects, we use f_1 to append a single random Sickness Post and f_2 as the identity function. Outcome probabilities, conditioned on treatment, are identical to previous tasks.

Sample Size In this task, we test how the models’ performance drops off as the amount of available training data is reduced.⁸ As before, we use f_1 to append a single random Sickness Post and f_2 as the identity function. For the easiest case, we train on all 3,200 users’ histories in the training set. We then create smaller training sets by randomly sampling subsets with 1,600 and 800 users.

Placebo Test The final task assesses a model’s tendency to predict a treatment effect when none is present. To do so, we must have asymmetric treatment probabilities between `class 1` and `class 2`. Without this asymmetry, the unadjusted estimate would be equal to the true *ATE* of zero. We use the same asymmetric `class 1` treatment split as in the Strength of Selection Effect task.

⁸In Keith, Jensen, and O’Connor (2020, Table 1), 8/12 studies had fewer than 5,000 observations, and 4/12 had fewer than 1,000.

We set $P(Y = 1|T = 0, \text{class}=1) = .05$, $P(Y = 1|T = 1, \text{class}=2) = .95$, and the opposite for $Y = 0$. This gives a treatment effect of +.9 to `class 1` and a treatment effect of -.9 to `class 2`, making the true *ATE* for the entire task equal 0. As in previous tasks, f_1 appends one random Sickness Post and f_2 is the identity function.

A potential limitation of these tasks may be the placement of synthetic posts at the end of histories, or differences in the length of histories. However in §7.1 we show that our results suggest that these potential limitations are very unlikely to affect the validity of evaluation. Furthermore, the framework may be extended (§5.3) to further evaluate methods on these aspects.

5.3 Generalizability of Evaluation Framework

A key tenet of our evaluation framework is its extensibility. Here, we provide a summary of how the framework can be extended, with step by step instructions for addition of continuous outcome values and new tasks. Additional details and resources are available on our website.⁹

Non-binary Treatments and Outcomes The evaluation framework presented here tests methods’ ability to handle five distinct challenges which are relevant to many real world studies. It can be applied to *any* causal inference method which uses text to adjust for confounding. As our framework is the first such framework, we focus on the simplest and most broadly applicable cases: binary treatments and outcomes, as these are the most commonly used in practice. Out of 14 recent casual inference studies, 13 used binary treatments (Keith, Jensen, and O’Connor 2020). While the methods evaluated here have been generalized to handle observations with continuous treatment values (*i.e.*, dose-response observations) (Hirano and Imbens 2004), these methods are rarely used in practice—we are not aware of a single such application which uses text data. However, our framework trivially generalizes to cases with continuous outcome values, and can be easily modified to include continuous treatments and multiple confounders, including unobserved confounders.

To modify the framework to use continuous outcomes:

1. Select a random distribution to use to assign outcome values, conditioned on treatment. A straightforward option would be to use normal distributions with $\sigma^2 = 0.3$ and mean = .1 for (`class 1`, $T = 0$) and mean = .9 for (`class 1`, $T = 1$), (`class 2`, $T = 0$), and (`class 2`, $T = 1$). Any distribution may be used that provides reasonable common support as well as sufficient confusion and selection effects.
2. Train propensity score models as with binary outcomes.
3. Compute ATE estimates using IPTW, matching or stratification estimators, all of which can be applied to continuous outcomes without modification.
4. Compute bias by taking the difference from the true ATE, which can be derived from the known treatment and outcome probability distributions.

⁹<https://behavioral-data.github.io/CausalInferenceChallenges/>

Steps 2-4 are identical to the existing framework, as previously described.

Types of Confounding It is difficult or impossible to know exactly what types of confounders threaten the accuracy of real world causal inferences, and even simply making a reasonable guess requires substantial domain expertise (Rosenbaum 2010). As such, our evaluation framework simulates confounders which are common to many topics of interest to the ICWSM community, where specific ‘indicator’ passages or posts in a longer text document are correlated with a treatment and/or outcome of interest. These ‘indicator’ confounders have been shown to be common in the mental health (Choudhury and De 2014), personal identity (Haimson 2018), and fake news (Talwar et al. 2019) domains. While our five tasks focus on these common types of confounder, tasks can be easily added to cover other types of confounders, for example by adding a task which applies some rewriting transformation (*e.g.* desirability (Wang and Culotta 2019; Pryzant et al. 2018) or gender (Wang and Culotta 2019)) to the histories of `class 1` and/or `class 2`.

To add an additional task which applies a rewriting transformation:

1. Write a new f_1 function which, for example, replaces all occurrences of the token ‘happy’ with ‘sad.’
2. Write a new f_2 function which, for example, replaces all occurrences of the token ‘sad’ with ‘happy.’
3. Use the same treatment and outcome probabilities as used for other tasks, or, optionally, modify these probabilities to decrease or increase the overlap between the treated and untreated groups, as in the Strength of Selection Effect task.
4. Train and evaluate methods’ treatment accuracy and bias of the ATE using the known true ATE.

6 Causal Inference Pipeline

So many methods for adjusting for confounding with text have been proposed recently that is not possible to evaluate every one in a single paper. Instead, we focus on the most commonly used methods, those which are based upon propensity scores. A recent review found that propensity score-based methods are used in 12/13 recent studies (Keith, Jensen, and O’Connor 2020, Table 1). Evaluating less commonly used methods (such as doubly robust methods (Mayer et al. 2019)) is an important area of future work, and our evaluation framework can be applied to *any* method for adjusting for confounding with text. The methods we evaluate here consist of three parts: a text representation, a propensity score model, and an *ATE* estimator.

6.1 Text Representations & Propensity Score Models

The **Oracle** uses the true propensity scores, which are known in our semi-synthetic evaluation framework (Fig. 3). The Oracle provides an upper-bound on model performance, only differing from the theoretical optimum due to finite sample effects.

We include an **Unadjusted Estimator**, which uses the naive method of not adjusting for selection effects, producing an estimated treatment effect of $\bar{Y}_{T=1} - \bar{Y}_{T=0}$, and as such is a lower-bound for models that attempt to correct for selection effects.

We train a **Simple Neural Net** (with one fully connected hidden layer) in four variants with different text representations: 1-grams with a binary encoding, 1,2-grams with a binary encoding, 1,2-grams with counts, and Latent Dirichlet Allocation (LDA) features (Blei, Ng, and Jordan 2003) based on 1,2-grams, counted. We also train **Logistic Regression** models on the same four text representations. Vocabulary sizes for n-gram methods are included in Appendix D.

Finally, we propose and evaluate a novel causal HiERarchical variant of BERT, which we call **SHERBERT**. SHERBERT expands upon Causal BERT proposed by Veitch, Sridhar, and Blei (2019), which is too computationally intensive to scale to user histories containing more than 250 tokens, let alone ones orders of magnitude longer, such as in our tasks. In SHERBERT, we use one pretrained BERT model per post to produce a post-embedding (Appendix E.1 Fig. 5), followed by two hierarchical attention layers to produce a single embedding for the entire history, with a final linear layer to estimate the propensity score. This architecture is similar to HIBERT (Zhang, Wei, and Zhou 2019), but is faster to train on long textual histories, as SHERBERT fixes the pretrained BERT components. More details on SHERBERT are given in Appendix E.

6.2 Average Treatment Effect Estimators

We consider three commonly used *ATE* estimators – IPTW, stratification, and matching. All three estimators use propensity scores but differ in how they weight or group relevant samples.

Inverse Propensity of Treatment Weighting estimates the *ATE* by weighting each user by their relevance to selection effects:

$$\widehat{ATE}_{IPTW} = \sum_{i=1}^n \frac{(2 * T_i - 1) * Y_i}{\hat{p}_{T_i}(\mathbf{X}_i) * \left[\sum_{j=1}^n \frac{1}{\hat{p}_{T_j}(\mathbf{X}_j)} \right]}$$

where T_i , Y_i , and X_i are treatment, outcome, and features for sample i , and $\hat{p}_T(\mathbf{X})$ is the estimated propensity for treatment T on features \mathbf{X} . Use of the Hajek estimator (1970) adjustment improves stability compared to simple inverse propensity.

Stratification divides users into strata based on their propensity score, and the *ATE* for each is averaged:

$$\widehat{ATE}_{strat} = \frac{1}{n} \sum_k n_k * \widehat{ATE}_k$$

where n is the total number of users, n_k is the number of users in the k -th stratum, and \widehat{ATE}_k is the unadjusted *ATE* within the k -th stratum. We report results on 10 strata divided evenly by percentile, but results are qualitatively similar for other k .

Matching can be considered as a special case of stratification, where each strata contains only one treated user. We find that matching produces extremely similar results to stratification, and therefore we include details of our matching approach and results in Appendix F.1.

6.3 Metrics for Evaluation

Our semi-synthetic tasks are generated such that we know the true ATE and thus can compute the **Bias of \widehat{ATE}** . A bias of zero is optimal, indicating a correct estimated ATE . The greater the bias, positive or negative, the worse the model performance. This is the primary metric we use in evaluation, and we compute it for both $\widehat{ATE}_{\text{strat}}$ and $\widehat{ATE}_{\text{IPTW}}$. We also consider **Treatment Accuracy**, the accuracy of the propensity score model’s predictions of binary treatment assignment. While higher accuracy is often better, high accuracy does not guarantee low bias and often is instead indicative of strong selection effects. Furthermore, we include two additional metrics in Appendix F.2: First, the **Mean Squared Error of IPTW weights**, which captures the calibration of propensity scores probabilities and resulting weights, and second, the **Spearman’s Rank Correlation**. In some cases, even if *absolute* propensity scores are incorrect, their *relative* rank may still contain useful information that could be exploited in stratification-based methods. The Spearman Correlation measures the correlation between the true and estimated propensity scores for each model, with a value of 1 indicating a perfect ranking.

7 Results of Evaluation

We apply all five tasks of our evaluation framework to both Reddit and Twitter social media data. We found virtually identical results between the two datasets (*e.g.* Treatment Accuracies within 1.9% of one another), and so for brevity we report primarily on Reddit results here. Complete Twitter results and discussion are included in Appendix G.

Transformers better model relevant linguistic variation

The Linguistic Complexity task shows many trends in the results manifest in other tasks, including treatment accuracy clustering by text representation (Fig. 4a). SHERBERT performs well, with uni- and bi-gram methods somewhere in between. Accuracy correlates fairly well with bias (Fig. 4b,c). As in nearly all tasks, LDA methods perform worst, not even outperforming the unadjusted estimator. This is likely because LDA uses an unsupervised (agnostic to treatment) method to generate a compressed feature set that is likely to miss key features when they comprise only a small part of the overall user history.

Transformer models struggle with counting and ordering

The Signal Intensity task requires methods to effectively ‘count’ the number of posts to distinguish between classes.’ Here, n-gram methods outperform SHERBERT (Fig. 4e,f) This suggests that order embeddings are an important inclusion for future transformer-based methods. LDA methods perform slightly better than unadjusted, due to the stronger presence of tokens correlated with treatment.

High accuracy often reflects strong selection effects, not low ATE bias

In the Strength of Selection Effect task, we decrease the overlap in propensity scores between treated and untreated users, which makes it *easier* to distinguish between the two groups. We see corresponding *increases*

in Treatment Accuracy (Fig. 4g), however, bias worsens (Fig. 4h,i). In Appendix F.2, we also consider the Spearman Correlation, which evaluates propensity scores not on their *absolute* accuracy, but on their *relative* ranking, as in theory it is possible for inaccurate propensity scores to carry useful information. We expect the Stratified and Matching Estimators to perform better than IPTW in these cases, *e.g.* in the Signal Intensity Task, where the stratified estimate has lower bias than IPTW (Fig. 4e,f) and the Spearman Correlation is higher than Treatment Accuracy at the harder difficulty level (Appendix F.2 Fig. 7d,f).

In context of observational studies, methods with high treatment accuracy should be used with extreme caution — high accuracy likely reflects treated and control groups that are too disjoint for any meaningful comparison to be drawn. In this case, the common support assumption is violated, preventing causal inference. This highlights the importance of empirical evaluation of the *complete* causal inference pipeline.

Transformer models fail with limited data

The Sample Size task explores methods’ performance on small datasets, a common occurrence in real world applications. Generally, SHERBERT performs quite well. In this task, SHERBERT outperforms other methods when trained on the full 3,200 observation training set, but its bias and accuracy quickly deteriorate to worse than n-gram features when training data is reduced (Fig. 4j,k,l). When data is especially scarce, practitioners should carefully consider the data-hungry nature of modern transformer architectures even when they are pre-trained. A more sophisticated model is not always the best choice. Furthermore, transfer learning and other means to reduce training data requirements are an important area of future work for causal inference method developers.

Methods estimate non-zero ATEs when the true ATE is zero

Alarmingly, in the Placebo Test, every method except SHERBERT’s stratified estimate failed to include the (correct) null hypothesis ($ATE = 0$) in their 95% confidence intervals across both datasets (Fig. 4n,o), including high accuracy methods using bigram features (Fig. 4m). This corresponds to incorrectly rejecting the null hypothesis of no causal treatment effect at the common p-value threshold of 0.05. This result is of greatest concern, as 8/9 methods falsely claim a non-zero effect.

Text representations and propensity score models have greater impact than estimators

Each estimator evaluated produced overall similar results (Fig. 6), with the quality of the propensity scores being far more impactful. Methods often cluster based on their text representations, with bigram representations generally performing better than n-grams, which generally perform better than LDA representations. There was generally little difference between Logistic Regression and Simple NN methods when trained on the same text representation. However, the choice of an ATE estimator is still important. IPTW is more sensitive to extreme or miscalibrated propensity scores. This is visible in the Strength of Selection Effect task, where the confidence intervals for IPTW are much larger than for the stratified es-

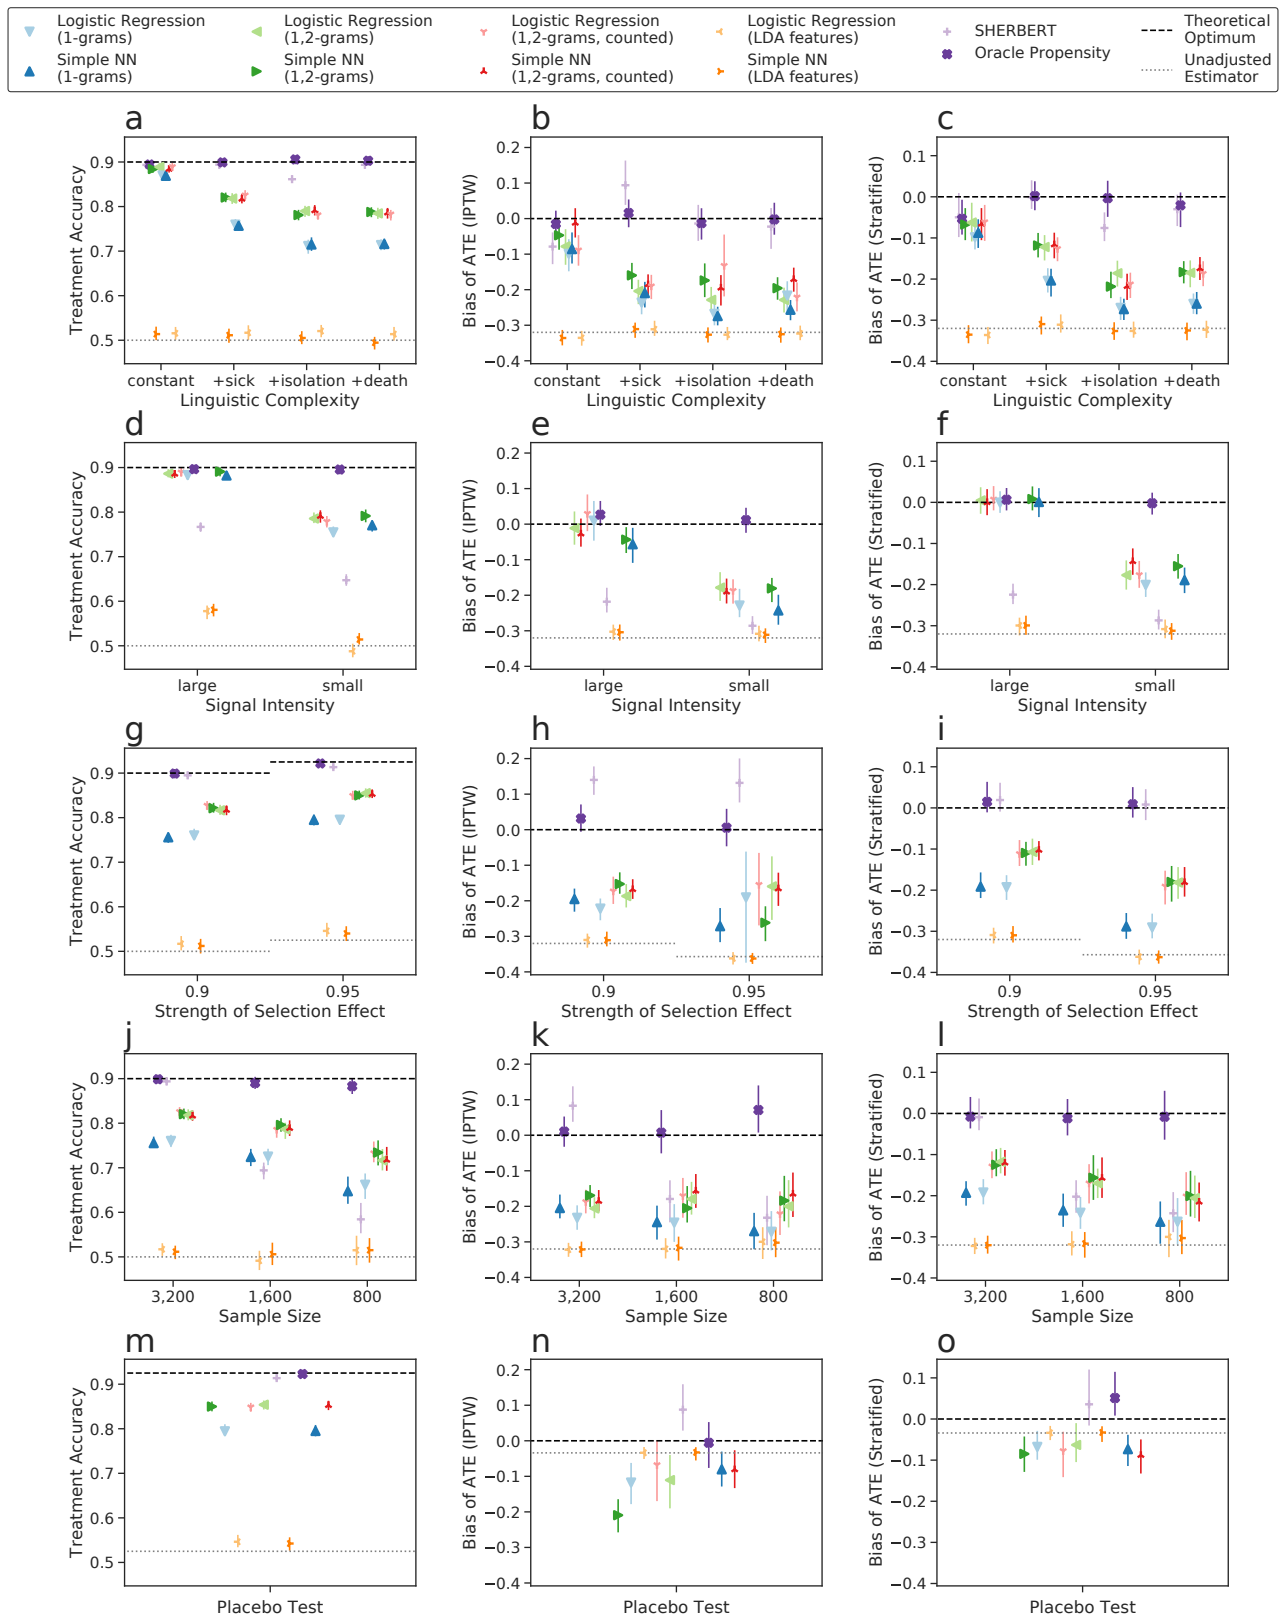


Figure 4: Results for tasks computed with Reddit data, with bootstrapped 95% confidence intervals, perturbed along the x-axis for readability. Columns represent metrics, and rows correspond to tasks. Within each plot, difficulty increases from left to right. SHERBERT generally does well, especially on *Strength of Selection Effect* and *Placebo Test*, but struggles on *Signal Intensity*.

timator (Fig. 4h,i).

7.1 Potential Threats to Evaluation Framework Validity

Many of our evaluation tasks (§5.2) append more synthetic posts to `class 1` histories than `class 2` histories. As a result, `class 1` histories will be 1-3 posts longer, in expectation, than those from `class 1`. Relative to the lengths of these histories (mean 41.1, standard deviation 19.3), this is a small difference, but in principle, a causal inference method may be able to pick up on the length of the history, an artifact of our framework, rather than the textual clues provided by the synthetic posts themselves. In practice, however, evidence from the Signal Intensity Task suggests that this is very unlikely, as every method tested, including SHERBERT, struggles with the to differentiate between histories by counting the number of synthetic posts (Fig. 4e,f). Thus, if we replaced a random post with a synthetic post, instead of appending, we would find very similar results. As methods improve and are better able to differentiate sequence lengths, such a random replacement strategy would be a reasonable and straightforward extension to our framework.

Additionally, our evaluation tasks (§5.2) always *append* synthetic posts to users' histories. We conducted additional experiments using an 'Order of Text' task (Appendix H) to evaluate widely used methods' ability to represent the order of posts in a user's history, and find that every method evaluated, including SHERBERT, completely fails to represent order. This is mostly a result of currently used text representations, primarily n-grams, which aggregate across histories by counting occurrences of tokens. For brevity, details of this task are included in Appendix H. We make this task public, along with all other tasks, to assist in the evaluation of future methods. We invite extensions and adaptations of our framework.

8 Implications & Conclusions

Causal inferences are difficult to evaluate in the absence of ground truth causal effects – a limitation of virtually all real world observational studies. Despite this absence, we *can* compare different methods' estimates and demonstrate that different methods regularly disagree with one another.

Empirical evaluation requires knowledge of the true treatment effects. Our proposed evaluation framework is reflective of five key challenges for causal inference in natural language, and is easily extensible to include different forms of confounders, different synthetic text content, and non-binary treatments and outcomes (§5.3).

Our goal with this work is not to unilaterally pronounce one method as superior to another. Instead, we hope that methods developers and practitioners will comprehensively consider the challenges of making valid causal inferences we have described here, as well as assumptions they may be relying upon, and will use our framework to evaluate their methods empirically. To this end, we evaluate every commonly used propensity score method to produce key insights:

For **methods developers**, we find that continued development of transformer-based models offers a promising path towards rectifying deficiencies of existing models. Models are needed that can effectively represent the order of text, variability in expression, and the counts of key tokens. Given the limited availability of training data in many causal inference applications, more research is needed in adapting pre-trained transformers to small data settings (Gururangan et al. 2020). We hope our public framework¹⁰ will provide a principled method for evaluating future NLP models for causal inference.

For **practitioners**, we find that transformer-based methods such as SHERBERT, which we make publicly available,⁹ perform the best in most cases except those with very limited data. Propensity score models with high accuracy should be applied with great care, as this is likely indicative of a strong and unadjustable selection effect. Many methods failed our placebo test by making false causal discoveries, a major problem (Aarts et al. 2015; Freedman, Cockburn, and Simcoe 2015).

Acknowledgments

This research was supported in part by the Laboratory Directed Research and Development Program at Pacific Northwest National Laboratory, a multiprogram national laboratory operated by Battelle for the U.S. Department of Energy. This research was supported in part by Adobe Data Science Research Award, the Office for Naval Research (#N00014-21-1-2154), NSF grant IIS-1901386, NSF grant CNS-2025022, the Bill & Melinda Gates Foundation (INV-004841), a Microsoft AI for Accessibility grant, and a Garvey Institute Innovation grant.

References

- Aarts, A.; Anderson, J.; Anderson, C.; Attridge, P.; Attwood, A.; Axt, J.; Babel, M.; Bahník, S.; Baranski, E.; Barnett-Cowan, M.; Bartmess, E.; Beer, J.; Bell, R.; Bentley, H.; Beyan, L.; Binion, G.; Borsboom, D.; Bosch, A.; Bosco, F.; and Penuliar, M. 2015. Estimating the Reproducibility of Psychological Science. *Science* 349.
- Abadie, A.; and Imbens, G. W. 2016. Matching on the Estimated Propensity Score. *Econometrica* 84(2): 781–807.
- Arbour, D.; and Dimmery, D. 2019. Permutation Weighting. *arXiv:1901.01230*.
- Baumgartner, J.; Zannettou, S.; Keegan, B.; Squire, M.; and Blackburn, J. 2020. The Pushshift Reddit Dataset. *arXiv:2001.08435*.
- Bhattacharya, P.; and Mehrotra, R. 2016. The Information Network: Exploiting Causal Dependencies in Online Information Seeking. CHIIR '16.
- Blei, D.; Ng, A.; and Jordan, M. I. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3: 993–1022.
- Brier, G. W. 1950. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review* 78(1): 1–3.
- Chandrasekharan, E.; Pavalanathan, U.; Srinivasan, A.; Glynn, A.; Eisenstein, J.; and Gilbert, E. 2017. You Can't Stay Here: The

¹⁰<https://behavioral-data.github.io/CausalInferenceChallenges/>

- Efficacy of Reddit's 2015 Ban Examined Through Hate Speech. *CSCW* 1.
- Chandrasekharan, E.; Samory, M.; Jhaver, S.; Charvat, H.; Bruckman, A.; Lampe, C.; Eisenstein, J.; and Gilbert, E. 2018. The Internet's Hidden Rules: An Empirical Study of Reddit Norm Violations at Micro, Meso, and Macro Scales. *CSCW* 2.
- Chen, V. Z.; Montano-Campos, F.; and Zadrozny, W. 2020. Causal Knowledge Extraction from Scholarly Papers in Social Sciences. *arXiv:2006.08904*.
- Choudhury, M. D.; and De, S. 2014. Mental Health Discourse on reddit: Self-Disclosure, Social Support, and Anonymity. In *ICWSM*.
- Choudhury, M. D.; and Kiciman, E. 2017. The Language of Social Support in Social Media and Its Effect on Suicidal Ideation Risk. *ICWSM 2017*: 32–41.
- Crump, R. K.; Hotz, V. J.; Imbens, G. W.; and Mitnik, O. A. 2009. Dealing with limited overlap in estimation of average treatment effects. *Biometrika* 96(1): 187–199.
- De Choudhury, M.; Kiciman, E.; Dredze, M.; Coppersmith, G.; and Kumar, M. 2016. Discovering Shifts to Suicidal Ideation from Mental Health Content in Social Media. In *CHI '16*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- Dorie, V.; Hill, J.; Shalit, U.; Scott, M.; and Cervone, D. 2019. Automated versus Do-It-Yourself Methods for Causal Inference: Lessons Learned from a Data Analysis Competition. *Statist. Sci.* 34(1): 43–68.
- Eckles, D.; and Bakshy, E. 2017. Bias and high-dimensional adjustment in observational studies of peer effects. *arXiv:1706.04692*.
- Egami, N.; Fong, C. J.; Grimmer, J.; Roberts, M. E.; and Stewart, B. M. 2018. How to Make Causal Inferences Using Texts. *arXiv:1802.02163*.
- Falavarjani, S. A. M.; Hosseini, H.; Noorian, Z.; and Bagheri, E. 2017. Estimating the Effect Of Exercising On Users' Online Behavior. In *AAAI 2017*.
- Fong, C.; and Grimmer, J. 2016. Discovery of Treatments from Text Corpora. In *ACL '16*. Berlin, Germany: Association for Computational Linguistics.
- Freedman, L. P.; Cockburn, I. M.; and Simcoe, T. S. 2015. The Economics of Reproducibility in Preclinical Research. *PLOS Biology* 13(6): 1–9.
- Gentzel, A.; Garant, D.; and Jensen, D. 2019. The Case for Evaluating Causal Models Using Interventional Measures and Empirical Data. In *NeurIPS '19*.
- Gururangan, S.; Marasović, A.; Swayamdipta, S.; Lo, K.; Beltagy, I.; Downey, D.; and Smith, N. A. 2020. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks.
- Haimson, O. L. 2018. *The Social Complexities of Transgender Identity Disclosure on Social Media*. Ph.D. thesis, UC Irvine.
- Hájek, J. 1970. A characterization of limiting distributions of regular estimates. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*.
- Hill, J.; and Su, Y.-S. 2013. Assessing lack of common support in causal inference using Bayesian nonparametrics: Implications for evaluating the effect of breastfeeding on children's cognitive outcomes. *Ann. Appl. Stat.* 7(3).
- Hirano, K.; and Imbens, G. W. 2004. *The Propensity Score with Continuous Treatments*, chapter 7, 73–84. John Wiley & Sons, Ltd.
- Jensen, D. 2019. Comment: Strengthening Empirical Evaluation of Causal Inference Methods. *Statist. Sci.* 34(1).
- Johansson, F.; Shalit, U.; and Sontag, D. 2016. Learning representations for counterfactual inference. In *ICML*.
- Joulin, A.; Grave, E.; Bojanowski, P.; and Mikolov, T. 2017. Bag of Tricks for Efficient Text Classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, 427–431. Valencia, Spain: Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/E17-2068>.
- Kang, J. D. Y.; and Schafer, J. L. 2007. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statist. Sci.* 22(4): 523–539.
- Keith, K. A.; Jensen, D.; and O'Connor, B. 2020. Text and Causal Inference: A Review of Using Text to Remove Confounding from Causal Estimates. In *ACL*.
- Kiciman, E.; Counts, S.; and Gasser, M. 2018. Using Longitudinal Social Media Analysis to Understand the Effects of Early College Alcohol Use. In *ICWSM*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980*.
- Landeiro, V.; and Culotta, A. 2016. Robust Text Classification in the Presence of Confounding Bias. *AAAI*.
- Lee, B. K.; Lessler, J.; and Stuart, E. A. 2011. Weight Trimming and Propensity Score Weighting. *PLOS ONE* 6(3).
- Li, F.; Thomas, L. E.; and Li, F. 2018. Addressing Extreme Propensity Scores via the Overlap Weights. *Am. J. Epidemiol.* 188(1): 250–257.
- Lin, A.; Merchant, A.; Sarkar, S. K.; and D'Amour, A. 2019. Universal Causal Evaluation Engine: An API for empirically evaluating causal inference models. In Le, T. D.; Li, J.; Zhang, K.; Cui, E. K. P.; and Hyvärinen, A., eds., *Proceedings of Machine Learning Research*, volume 104 of *Proceedings of Machine Learning Research*, 50–58. Anchorage, Alaska, USA: PMLR. URL <http://proceedings.mlr.press/v104/lin19a.html>.
- Mani, S.; and Cooper, G. F. 2000. Causal discovery from medical textual data. *Proceedings. AMIA Symposium*.
- Mayer, I.; Wager, S.; Gauss, T.; Moyer, J.-D.; and Josse, J. 2019. Doubly robust treatment effect estimation with missing attributes. *arXiv:1910.10624*.
- Medvedev, A. N.; Lambiotte, R.; and Delvenne, J.-C. 2019. The Anatomy of Reddit: An Overview of Academic Research. *Springer Proceedings in Complexity* 183–204.
- Mirza, P.; and Tonelli, S. 2016. CATENA: CAusal and TEmporal relation extraction from NATural language texts. In *COLING '16*. Osaka, Japan.
- Mozer, R.; Miratrix, L.; Kaufman, A. R.; and Jason Anastasopoulos, L. 2020. Matching with Text Data: An Experimental Evaluation of Methods for Matching Documents and of Measuring Match Quality. *Political Analysis* 1–24.
- Olteanu, A.; Varol, O.; and Kiciman, E. 2017. Distilling the Outcomes of Personal Experiences: A Propensity-Scored Analysis of Social Media. In *CSCW*.
- Pearl, J. 1995. Causal Diagrams for Empirical Research. *Biometrika* 82(4). ISSN 00063444.

Pham, T. T.; and Shen, Y. 2017. A Deep Causal Inference Approach to Measuring the Effects of Forming Group Loans in Online Non-profit Microfinance Platform.

Pryzant, R.; Shen, K.; Jurafsky, D.; and Wagner, S. 2018. Deconfounded Lexicon Induction for Interpretable Social Science. In *NAACL*.

Roberts, M. E.; Stewart, B. M.; and Nielsen, R. A. 2020. Adjusting for confounding with text matching. *American Journal of Political Science* .

Rosenbaum, P. R. 2010. *Design of Observational Studies*. Springer.

Rubin, D. B. 2005. Causal Inference Using Potential Outcomes. *Journal of the American Statistical Association* 100(469): 322–331. doi:10.1198/016214504000001880. URL <https://doi.org/10.1198/016214504000001880>.

Saha, K.; Sugar, B.; Torous, J. B.; Abrahao, B. D.; Kiciman, E.; and Choudhury, M. D. 2019. A Social Media Study on the Effects of Psychiatric Medication Use. *ICWSM* 13.

Schuler, M.; and Rose, S. 2017. Targeted Maximum Likelihood Estimation for Causal Inference in Observational Studies. *American Journal of Epidemiology* 185: 65–73.

Sridhar, D.; Springer, A.; Hollis, V.; Whittaker, S.; and Getoor, L. 2018. Estimating Causal Effects of Exercise from Mood Logging Data. *FAIM'18 Workshop on CausalML* .

Talwar, S.; Dhir, A.; Kaur, P.; Zafar, N.; and Alrasheedy, M. 2019. Why do people share fake news? Associations between the dark side of social media use and fake news sharing behavior. *Journal of Retailing and Consumer Services* 51: 72–82.

Tan, C.; Lee, L.; and Pang, B. 2014. The effect of wording on message propagation: Topic- and author-controlled natural experiments on Twitter. In *ACL*.

Veitch, V.; Sridhar, D.; and Blei, D. M. 2019. Using Text Embeddings for Causal Inference. *arXiv:1905.12741* .

Wang, Y.; Cai, H.; Li, C.; Jiang, Z.; Wang, L.; Song, J.; and Xia, J. 2013. Optimal Caliper Width for Propensity Score Matching of Three Treatment Groups: A Monte Carlo Study. *PLOS ONE* 8: 1–7.

Wang, Z.; and Culotta, A. 2019. When Do Words Matter? Understanding the Impact of Lexical Choice on Audience Perception Using Individual Treatment Effect Estimation. In *AAAI*.

Wood-Doughty, Z.; Shpitser, I.; and Dredze, M. 2018. Challenges of Using Text Classifiers for Causal Inference. In *EMNLP*.

Yang, S.; and Ding, P. 2018. Asymptotic inference of causal effects with observational studies trimmed by the estimated propensity scores. *Biometrika* .

Yu, J.; and Muñoz-Justicia, J. 2020. A Bibliometric Overview of Twitter-Related Studies Indexed in Web of Science. *Future Internet* 12: 91.

Zhang, J.; Chang, J.; Danescu-Niculescu-Mizil, C.; Dixon, L.; Hua, Y.; Taraborelli, D.; and Thain, N. 2018. Conversations Gone Awry: Detecting Early Signs of Conversational Failure. In *ACL*.

Zhang, J.; Mullainathan, S.; and Danescu-Niculescu-Mizil, C. 2020. Quantifying the Causal Effects of Conversational Tendencies. *Proc. ACM Hum.-Comput. Interact.* 4(CSCW2). doi:10.1145/3415202. URL <https://doi.org/10.1145/3415202>.

Zhang, X.; Wei, F.; and Zhou, M. 2019. HIBERT: Document Level Pre-training of Hierarchical Bidirectional Transformers for Document Summarization. In *ACL*.