

Challenges of Using Text Classifiers for Causal Inference

Zach Wood-Doughty*, Ilya Shpitser[†], Mark Dredze*[†]

Department of Computer Science

*Center for Language and Speech Processing

[†]Malone Center for Engineering in Healthcare

Johns Hopkins University, Baltimore, MD 21218

{zach, ilyas, mdredze}@cs.jhu.edu

Abstract

Causal understanding is essential for many kinds of decision-making, but causal inference from observational data has typically only been applied to structured, low-dimensional datasets. While text classifiers produce low-dimensional outputs, their use in causal inference has not previously been studied. To facilitate causal analyses based on language data, we consider the role that text classifiers can play in causal inference through established modeling mechanisms from the causality literature on missing data and measurement error. We demonstrate how to conduct causal analyses using text classifiers on simulated and Yelp data, and discuss the opportunities and challenges of future work that uses text data in causal inference.

1 Introduction

Most scientific analyses, in domains from economics to medicine, focus on low-dimensional structured data. Many such domains also have unstructured text data; advances in natural language processing (NLP) have led to an increased interest in incorporating language data into scientific analyses. While language is inherently unstructured and high dimensional, NLP systems can be used to process raw text to produce structured variables. For example, work on identifying undiagnosed side effects from electronic health records (EHR) uses text classifiers to produce clinical variables from the raw text (Hazlehurst et al., 2009).

NLP tools may also benefit the study of causal inference, which seeks to identify causal relations from observational data. Causal analyses traditionally use low-dimensional structured variables, such as clinical markers and binary health outcomes. Such analyses require assumptions about the data-generating process,

which are often simpler with low-dimensional data. Unlike prediction tasks which are validated by held-out test sets, causal inference involves modeling counterfactual random variables (Neyman, 1923; Rubin, 1976) that represent the outcome of some hypothetical intervention. To rigorously reason about hypotheticals, we use causal models to link our counterfactuals to observed data (Pearl, 2009).

NLP provides a natural way to incorporate text data into causal inference models. We can produce low-dimensional variables using, for example, text classifiers, and then run our causal analysis. However, this straightforward integration belies several potential issues. Text classification is not perfect, and errors in a NLP algorithm may bias subsequent analyses. Causal inference requires understanding how variables influence one another and how correlations are confounded by common causes. Classic methods such as stratification provide a means for handling confounding of categorical or continuous variables, but it is not immediately obvious how such work can be extended to high-dimensional data.

Recent work has approached high-dimensional domains by using random forests (Wager and Athey, 2017) and other methods borrowed from machine learning (Chernozhukov et al., 2016). But even compared to an analysis that requires hundreds of confounders (Belloni et al., 2014), NLP models with millions of variables are very high-dimensional. While physiological symptoms reflect complex biological realities, many symptoms such as blood pressure are one-dimensional variables. While doctors can easily quantify the effect of high blood pressure on some outcome, can we use the “positivity” of a restaurant review to estimate a causal effect? More broadly, is it possible to employ text classification methods in a causal analysis?

We explore methods for the integration of text classifiers into causal inference analyses that consider confounds introduced by imperfect NLP. We show what assumptions are necessary for causal analyses using text, and discuss when those assumptions may or may not be reasonable. We draw on the causal inference literature to consider two modeling aspects: missing data and measurement error. In the missing data formulation, a variable of interest is sometimes unobserved, and text data gives us a means to model the missingness process. In the measurement error formulation, we use a text classifier to generate a noisy proxy of the underlying variable.

We highlight practical considerations of a causal analysis with text data by conducting analyses with simulated and Yelp data. We examine the results of both formulations and show how a causal analysis which properly accounts for possible sources of bias produces better estimates than naïve methods which make unjustified assumptions. We conclude by examining how our approach may enable new research avenues for inferring causality with text data.

2 Causal Inference, Briefly

While randomized control trials (RCT) are the gold standard of determining causal effects of treatments on outcomes, they can be expensive or impossible in many settings. In contrast, the world is filled with observational data collected without randomization. While most studies simply report correlations from observational data, the field of causal inference examines what assumptions and analyses make it possible to identify causal effects.

We formalize a causal statement like “smoking causes cancer” as “if we were to conduct a RCT and assign smoking as a treatment, we would see a higher incidence of cancer among those assigned smoking than among the control group.” In the framework of Pearl (1995), we consider a counterfactual variable of interest: what *would have* been the cancer incidence among smokers if smoking *had been* randomized? Specifically, we consider a causal effect as the counterfactual outcome of a hypothetical intervention on some treatment variable. If we denote smoking as our treatment variable A and cancer as our outcome variable Y , then we are interested in the counterfactual distribution, denoted $p(Y(a))$ or

$p(Y | \text{do}(a))$. We interpret this as “the distribution over Y had A been set, possibly contrary to fact, to value a .” For a binary treatment A , the causal effect of A on Y is denoted $\tau = E[Y(1)] - E[Y(0)]$; the average difference between if you had received the treatment and if you had not. Throughout, we use causal directed acyclic graphs (DAG), which assumes that an intervention on A is well-defined and results in a counterfactual variable $Y(a)$ (Pearl, 1995; Dawid, 2010).

Figure 1a shows an example of simple confounding. This is the simplest DAG in which counterfactual distribution $p(Y(a))$ is not simply $p(Y | A)$, as C influences both the treatment A and the outcome Y . To recover the counterfactual distribution $p(Y(a))$ that would follow an intervention upon A , we must “adjust” for C , applying the so-called “back-door criterion” (Pearl, 1995). We can then derive the counterfactual distribution $p(Y(a))$ and desired causal effect, τ_S as a function of the observed data, (Fig. 4 Eq. 1.) This derivation is shown in Appendix A.

Note that $p(Y(a))$ and τ_S require data on C , and if C is not in fact observed, it is impossible to recover the causal effect. Formally, we say that $p(Y(a))$ is *not identified* in the model, meaning there is no function f such that $p(Y(a))=f(p(A, Y))$. Identifiability is a primary concern of causal inference (Shpitser and Pearl, 2008).

Throughout, we assume for simplicity that A , C , and Y are binary variables. While text classifiers can convert high-dimensional data into binary variables for such analyses, we need to make further assumptions about how classification errors affect causal inferences. We cannot assume that the output of a text classifier can be treated as if it were ground truth. To conceptualize the ways in which a text classifier may be biased, we will consider them as a way to recover from missing data or measurement error.

3 Causal Models

Real-world observational data is messy and often imperfectly collected. Work in causal inference has studied how analyses can be made robust to missing data or data recorded with systematic measurement errors.




Figure 1: DAGs for causal inference without text data. Red variables are unobserved. A is a treatment, Y is an outcome, and C is a confounder.

A	C	Y	R_A	A	C	Y	A^*	C	Y	A^*	A
1	1	0	1	1	1	0	0	1	0	1	1
0	1	1	0	?	1	1	0	1	1	0	1
0	0	1	1	0	0	1	0	0	1	0	0
1	0	1	0	?	0	1	1	0	1	1	1

(a) Simple Confounding (b) Missing Data (c) Measurement Error (d) Mismeasurement

Figure 2: Example data rows for causal inference without text data.

3.1 Missing Data

Our dataset has “missing data” if it contains individuals (instances) for which some variables are unobserved, even though these variables are typically available. This may occur if some survey respondents choose not to answer certain questions, or if certain variables are difficult to collect and thus infrequently recorded. Missing data is closely related to causal inference – both are interested in hypothetical distributions that we cannot directly observe (Robins et al., 2000; Shpitser et al., 2015).

Consider a causal model where A is sometimes missing (Figure 1b). The variable R_A is a binary indicator for whether A is observed ($R_A = 1$) or missing. The variable $A(R_A = 1)$, written as $A(1)$, represents the counterfactual value of A were it never missing. Finally, A is the observed proxy for $A(1)$: it has the same value as $A(1)$ if $R_A = 1$, and the special value “?” if $R_A = 0$.

Solving missingness can be seen as intervening to set R_A to 1. Given $p(A, R_A, C, Y)$, we want to recover $p(A(1), C, Y)$. We may need to make a “Missing at Random” (MAR) assumption, which says that the missingness process is independent of the true missing values, conditional on observed values. Figure 1b reflects the MAR assumption; R_A is independent of $A(1)$ given fully-observed C and Y . If an edge existed from $A(1)$ to R_A , we have “Missing Not at Random” (MNAR) and would not be identified except in special cases (Shpitser et al., 2015).

3.2 Measurement Error

Sometimes a necessary variable is never observed, but is instead proxied by a variable which differs from the truth by some error. Consider the example of body mass index (BMI) as a proxy for obesity in a clinical study. Obesity is a known risk factor for many health outcomes, but has a complex clinical definition and is nontrivial to measure. BMI is a simple deterministic function of height and weight. To conduct a causal analysis of obesity on cancer when only BMI and cancer are measured, we can proceed as if we had measured obesity and then correct our analysis for the known error that comes from using BMI as a proxy for obesity (Hernán and Cole, 2009; Michels et al., 1998).

To generalize this concept, we can replace obesity with our ground truth variable A and replace BMI with a noisy proxy A^* . Figure 1c gives the DAG for this model. Unlike missing data problems, there is no hypothetical intervention which recovers the true data distribution $p(A, C, Y)$. Instead, we manipulate the observed distribution $p(A^*, C, Y)$ with the known relationship $p(A^*, A)$ to recover the desired $p(A, C, Y)$.

Unlike missing data, measurement error conceptualization can be used even when we never observe A (e.g. the table in Figure 2c) as long as we have knowledge about the error mechanism $p(A^*, A)$. Using this knowledge, we can correct for the error using ‘matrix

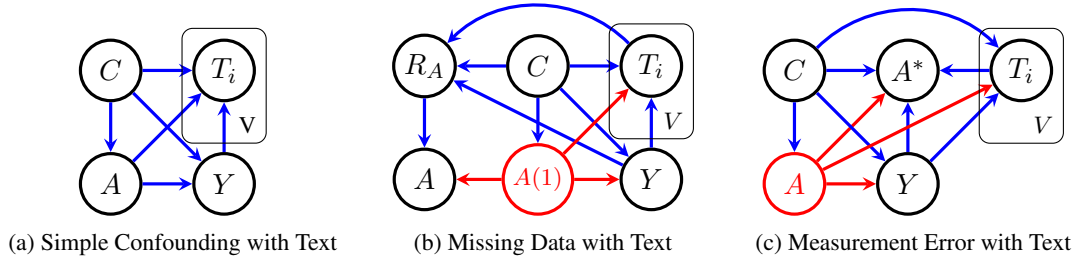


Figure 3: DAGs for causal inference with text data. In the Yelp experiments we discuss, T_i influences Y and not the other way around.

$$\tau_S = \sum_C (p(Y = 1 | A = 1, C) - p(Y = 1 | A = 0, C)) p(C) \quad (1)$$

$$\tau_{MD} = \sum_C \left(\frac{p(A = 1 | \mathbf{T}, C, Y = 1, R_A = 1)}{\sum_{y'} p(A = 1 | \mathbf{T}, C, y', R_A = 1) p(Y = y' | C)} - \frac{p(A = 0 | \mathbf{T}, C, Y = 1, R_A = 1)}{\sum_{y'} p(A = 0 | \mathbf{T}, C, Y = y', R_A = 1) p(Y = y' | C)} \right) p(Y = 1, C) \quad (2)$$

$$\tau_{ME} = \sum_C \left(\frac{-\delta_{c,y=1} q_{c,y=1}(0) + (1 - \delta_{c,y=1}) q_{c,y=1}(1)}{(1 - \epsilon_{c,y=1} - \delta_{c,y=1})} - \frac{(1 - \epsilon_{c,y=1}) q_{c,y=1}(0) - \epsilon_{c,y=1} q_{c,y=1}(1)}{(1 - \epsilon_{c,y=1} - \delta_{c,y=1})} \right) p(C) \quad (3)$$

Define $\epsilon_{c,y} = p(A = 0 | A^* = 1, C = c, Y = y)$, $\delta_{c,y} = p(A = 1 | A^* = 0, C = c, Y = y)$, $q_{c,y}(0) = p(C = c, Y = y, A^* = 0)$, and $q_{c,y}(1) = p(C = c, Y = y, A^* = 1)$.

Figure 4: Functionals for the Causal Effects for Simple Confounding (τ_{SC}), Missing Data (τ_{MD}) and Measurement Error (τ_{ME}). Derivations are in Appendices A, B, and C.

adjustment’ (Pearl, 2010). In practice we might learn $p(A^*, A)$ from data such as that found in Figure 2d. Recent work has also considered how multiple independent proxies of A could allow identification without any data on $p(A^*, A)$ (Kuroki and Pearl, 2014).

4 Causal Models for Text Data

We can use conceptualizations for missing data and measurement error to support causal analyses with text data. The choice of model depends on the assumptions we make about the data-generation process.

We add new variables to our models (Figure 1a) to represent text, which produces the data-generating distribution shown in Figure 3a. This model assumes that the underlying A , C , and Y variables are generated before the text variables; we use text to recover the true relationship between A and Y .

We represent text as an arbitrary set of V variables, which are independent of one another

given the non-text variables. In our implemented analyses we will represent text as a bag-of-words, wherein each T_i is simply the binary indicator of the presence of the i -th word in our vocabulary of V words, and $\mathbf{T} = \cup_i T_i$. The restriction to simple text models allows us to explore connections to causal inference applications, though future work could relax assumptions of the text models to be inclusive of more sophisticated text models (e.g. neural sequence models (Lai et al., 2015; Zhang et al., 2015)), or consider causal relationships between two text variables.

To motivate our explanations, consider the task of predicting an individuals’ smoking status from free-text hospital discharge notes (Uzuner et al., 2008; Wicentowski and Sydes, 2008). Some hospitals do not explicitly record patient smoking status as structured data, making it difficult to use such data in a study on the outcomes of smoking. We will suppose that we are given a dataset with patient data on lung cancer outcome (Y) and age (C), that our data on smoking status (A) is affected

by either missing data or measurement error, but that we have text data (\mathbf{T}) from discharge records that will allow us to infer smoking status with reasonable accuracy.

4.1 Missing Data

To show how we might use text data to recover from missing data, we introduce missingness for A from Figure 3a to get the model in Figure 3b. The missing arrow from $A(1)$ to R_A encodes the MAR assumption, which is sufficient to make it possible to identify the full data distribution from the observed data.

Suppose our motivation is to estimate the causal effect of smoking status (A) on lung cancer (Y) adjusting for age (C). Imagine that missing data arises because hospitals sometimes – but not always – delete explicit data on smoking status from patient records. If we have access to patients’ discharge notes (\mathbf{T}) and know whether a given patient had smoking status recorded (R_A), then the DAG in Figure 3b may be a reasonable model for our setting. Note that we must again assume that A does not directly affect R_A .

The causal effect of A on Y in Figure 3b is identified as τ_{MD} , given in Eq. 2 in Figure 4. The derivation is given in Appendix B.

4.2 Measurement Error

We model text data with measurement error by introducing a proxy A^* to the model in Figure 3c. We assume that the proxied value of A^* can depend upon all other variables, and that we will be able to estimate $p(A^*, A)$ given an external dataset, e.g. text classifier accuracy on held-out data.

Suppose we again want to estimate the causal effect from §4.1, but this time none of our hospital records contain explicit data on smoking status. However, imagine that we have a separate training dataset of medical discharge records annotated by expert pulmonologists for patients’ smoking status. We could then train a classifier to predict smoking status using discharge record text¹.

Working from the derivation for matrix adjustment in binary models given by Pearl (2010), we identify the causal effect of A on Y (Figure 3c) as τ_{ME} (Eq 3 in Figure 4.) The derivation is in Appendix C.

¹This is the precise setting of Uzuner et al. (2008).

5 Experiments

We now empirically evaluate the effectiveness of our two conceptualizations (missing data and measurement error) for including text data in causal analyses. We induce missingness or mismeasurement of the treatment variable and use text data to recover the true causal relationship of that treatment on the outcome. We begin with a simulation study with synthetic text data, and then conduct an analysis using reviews from `yelp.com`.

5.1 Synthetic Data

We select synthetic data so that we can control the entire data-generation process. For each data row, we first sample data on three binary variables (A, C, Y) and then sample V different binary variables T_i representing a V -vocabulary bag-of-words. A graphical model for this distribution appears in Figure 3a. We augment this distribution to introduce either missing data (Figure 3b) or measurement error (Figure 3c.) For measurement error, we sample two datasets. A small training set which gives data on $p(A, C, Y, \mathbf{T})$ and a large test set which gives data on $p(C, Y, \mathbf{T})$.

The full data generating process appears in Appendix D, and the implementation (along with all our code) is provided online².

5.2 Yelp Data

We utilize the 2015 Yelp Dataset Challenge³ which provides 4.7M reviews of local businesses. Each review contains a one- to five-star rating, up to 5,000 characters of text. Yelp users can flag reviews as “Useful” as a mark of quality.

We extract treatment, outcome, and confounder variables from the structured data. The treatment is a binarized user rating that takes value 1 if the review has four or five stars and value 0 if the review has one or two stars. Three-star reviews are discarded from our analysis. The outcome is whether the review received at least one “Useful” flag. The confounder is whether the review’s author has received at least two “Useful” flags across all reviews, according to their user object. In our data, 74.2% of reviews were positive, 42.6% of reviews were flagged as “Useful,” and 56.7% users had received at least

²github.com/zachwooddoughty/emnlp2018-causal

³yelp.com/dataset/challenge

two such flags. We preprocess the text of each review by lowercasing, stemming, and removing stopwords, before converting to a bag-of-words representation with the 4,334 word vocabulary of all words which appeared at least 1000 times in a sample of 1M reviews.

Based on this $p(A, C, Y, \mathbf{T})$ distribution, we assume the data-generating process that matches Figure 3a and introduce missingness and mismeasurement as before, giving us data-generating processes matching Figures 3b and 3c.

Our intention is not to argue about a true real-world causal effect of Yelp reviews on peer behavior: we do not believe that our confounder is the only common cause of the author’s rating and the platform’s response. We leave for future work a case study that jointly addresses questions of identifiability and estimation of a real-world causal effect. In this work, our experiments focus on a simpler task: can a correctly-specified model that uses text data effectively estimate a causal effect in the presence of missing data or measurement error.

5.3 Models

We now introduce several baseline methods which, unlike our correctly specified models τ_{MD} and τ_{ME} , are not consistent estimators of our desired causal effect. We would expect that the theoretical bias in these estimators would result in poor performance in our experiments.

5.3.1 Baseline: Naïve Model

In both the missing data and measurement error settings, our models use some rows that are full observed. In missing data, these are rows where $R_A = 1$; in measurement error, the training set is sampled from the true distribution. The simplest approach to handling imperfect data is to throw away all rows without full data, and calculate Eq 1 from that data. In Figure 5, these are labeled as `*.naive`.

5.3.2 Baseline: Textless Model

In Figure 3b, if we do not condition on T_i to d-separate $A(1)$ from its missingness indicator, that influence may bias our estimate. While we know that ignoring text may introduce asymptotic bias into our estimates of the causal effect, we empirically evaluate how much bias is produced by this “Textless” model compared to a correct

model. This is labeled as `*.no_text` in Figure 5 (a).

In principle, we could conduct a measurement error analysis using a model that does not include text. In practice, we found we could not impute A^* from C and Y alone. The non-textual classifier had such high error that the adjustment matrix was singular and we could not compute the effect. Thus, we have no such baseline in our measurement error results.

5.3.3 Baseline: no_y and unadjusted Models

In Figure 3b, we must also condition on C and Y to d-separate $A(1)$ from its missingness indicator. In our misspecified model for missing data, we do not condition on Y , leaving open a path for $A(1)$ to influence its missingness. In Figure 5 (a), this model is labeled as `*.no_y`.

When correcting for measurement error, a crucial piece of the estimation is the matrix adjustment using the known error between the proxy and the truth. A straightforward misspecified model for measurement error is to impute a proxy for each row in our dataset and then calculate the causal effect assuming no error between the proxy and truth. This approach, while simplistic, can be thought of as using a text classifier as a proxy without regard for the text classifier’s biases. In Figure 5 (b), this approach is labeled as `*.unadjusted`.

5.3.4 Correct Models

Finally, we consider the estimation approaches presented in §4.1 and §4.2. For the missing data causal effect (τ_{MD} from Eq 2) we use a multiple imputation estimator which calculates the average effect across 20 samples from $p(A|\mathbf{T}, C, Y)$ for each row where $R_A = 0$. For the measurement error causal effect (τ_{ME} from Eq 3), we use the training set of $p(A, C, Y, \mathbf{T})$ data to estimate $\epsilon_{c,y}$ and $\delta_{c,y}$ and the larger set of $p(C, Y, \mathbf{T})$ data to estimate $q_{c,y}$ and $p(C)$.

These models are displayed in Figure 5 (a) as `*.full` and in Figure 5 (b) `*.adjusted`.

5.4 Evaluation

Each model takes in a data sample with missingness or mismeasurement, and outputs an estimate of the causal effect of A on Y in the underlying data. Rather than comparing models’ estimates against a population-level estimate,

we compare against an estimate of the effect computed on the same data sample, but without any missing data or measurement error. This ‘perfect data estimator’ may still make errors given the finite data sample. We compare against this estimator to avoid a small-sample case where an estimator gets lucky. In Figure 5, we plot data sample size against the squared distance of each model’s estimate from a perfect data estimator’s estimate, averaged over ten runs. Figure 6 in Appendix E contains a second set of experiments using a larger vocabulary.

6 Results

Given that our correctly-specified models are proven to be asymptotically consistent, we would expect them to outperform misspecified models. However, for any given dataset, asymptotic consistency provides no guarantees.

6.1 Missing Data

The missing data (MD) experiments suggest that the correct `full` model does perform best. The `no_y` model performs approximately as well as the correct model on the synthetic data, but not on the Yelp data. The difference between the `no_y` and `full` missing data models is simply a function of the effect of Y on R_A . We could tweak our synthetic data distribution to increase the influence of Y to make the `no_y` model perform worse.

When we initially considered other data-generating distributions for missing data, we found that when we reduced the influence of the text variables on R_A , the `no_text` and `naive` models approached the performance of the correctly-specified model. While intuitive, this reinforces that the underlying distribution matters a great deal in how modeling choices may introduce biases if incorrectly specified.

6.2 Measurement error

The measurement error results tell a more interesting story. We see enormous fluctuations of the `adjusted` model, and in the synthetic data, the `unadjusted` model appears to be quite superior.

In the synthetic dataset, this is likely because our text classifier had near-perfect accuracy, and so simple approach of assuming its predictions were ground-truth introduced less bias. A broader issue with the `adjusted` model is that the

matrix adjustment approach requires dividing by (potentially very small) probabilities, this sometimes resulted in huge over-corrections. In addition, since those probabilities are estimated from a relatively small training dataset, small changes to the error-estimate can propagate to huge changes in the final casual estimate.

This instability of the matrix adjustment approach may be a bigger problem for text and other high-dimensional data: unlike in our earlier example of BMI and obesity, there are likely no simple relationships between text and clinical variables. However, instead of using matrix adjustment as a way to recover the true effect, we may instead use it to bound the error our proxy may introduce. As mentioned by Pearl (2010), when $p(A | A^*)$ is not known exactly, we can use a Bayesian analysis to bound estimates of a causal effect. In a downstream task, this would let us explore the stability of our `adjusted` results.

7 Related Work

A few recent papers have considered the possibilities for combining text data with approaches from the causal inference literature. Landeiro and Culotta (2016) and Landeiro and Culotta (2017) explored text classification when the relationship between text data and class labels are confounded. Other work has used propensity scores as a way to extract features from text data (Paul, 2017) or to match social media users based on what words they write (De Choudhury et al., 2016). The only work we know of which seeks to estimate causal effects using text data focuses on effects *of* text or effects *on* text (Egami et al., 2018; Roberts et al., 2018). In our work, our causal effects do not include text variables: we use text variables to recover an underlying distribution and then estimate a causal effect within that distribution.

There is a conceptually related line of work in the NLP community on inferring causal relationships expressed in text (Girju, 2003; Kaplan and Berry-Rogghe, 1991). However, our work is fundamentally different. Rather than identify casual relations expressed via language, we are using text data in a causal model to identify the strength of an underlying causal effect.

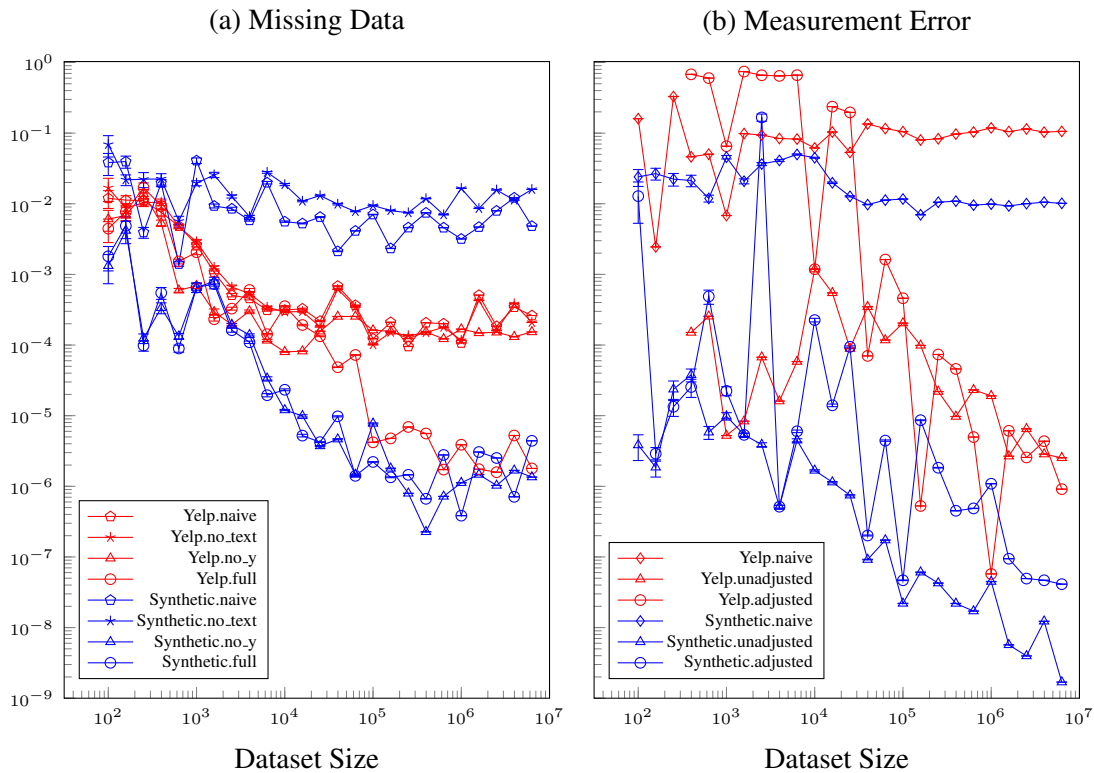


Figure 5: Experimental results. Squared distance (y-axis, lower is better) of the estimated causal effect from τ_{SC} calculated from the full data with no missing data or measurement error. Error bars (negligible for larger datasets) are 1.96 times standard error across 10 experiments. Additional experiments with a larger vocabulary are shown in Appendix E.

8 Future Directions

While this paper addresses some initial issues arising from using text classifiers in causal analyses, many challenges remain. We highlight some of these issues as directions for future research.

We provided several proof-of-concept models for estimating effects, but our approach is flexible to more sophisticated models. For example, a semi-parametric estimator would make no assumptions about the text data distribution by wrapping the text classifier into an infinite-dimensional nuisance model (Tsiatis, 2007). This would enable estimators robust to partial model misspecification (Bang and Robins, 2005).

Choices in the design of statistical models of text consider issues like accuracy and tractability. Yet if these models are to be used in a causal framework, we need to understand how modeling assumptions introduce biases and other issues that can interfere with a downstream causal analysis. To take an example from the medical domain, we

know that doctors write clinical notes throughout the healthcare process, but it is not obvious how to model this data-generating process. We could assume that the doctor’s notes passively record a patient’s progression, but in reality it may be that the content of the notes themselves actively change the patient’s care; causality could work in either direction.

New lines of work in causality may be especially helpful for NLP. In this work, we used simple logistic regression on a bag-of-words representation of text; using state-of-the-art text models will likely require more causal assumptions. Nabi and Shpitser (2017) develops causality-preserving dimensionality reduction, which could help develop text representations that preserve causality.

Finally, we are interested in case studies on incorporating text classifiers into real-world causal analyses. Many health studies have used text classifiers to extract clinical variables from EHR data (Meystre et al., 2008). These works could be extended to study causal effects involving

those extracted variables, but such extensions would require an understanding of the underlying assumptions. In any given study, the necessity and appropriateness of assumptions will hinge on domain expertise. The conceptualizations outlined in this paper, while far from solving all issues of causality and text, will help those using text classifiers to more easily consider research questions of cause and effect.

Acknowledgments

This work was in part supported by the National Institute of General Medical Sciences under grant number 5R01GM114771 and by the National Institute of Allergy and Infectious Diseases under grant number R01 AI127271-01A1. We thank the anonymous reviewers for their helpful comments.

References

- Heejung Bang and James M Robins. 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973.
- Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. 2014. Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, and Whitney K Newey. 2016. Double machine learning for treatment and causal parameters. Technical report, cemmap working paper, Centre for Microdata Methods and Practice.
- A Philip Dawid. 2010. Beware of the dag! In *Causality: Objectives and Assessment*, pages 59–86.
- Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 2098–2110. ACM.
- Naoki Egami, Christian J Fong, Justin Grimmer, Margaret E Roberts, and Brandon M Stewart. 2018. How to make causal inferences using texts. *arXiv preprint arXiv:1802.02163*.
- Roxana Girju. 2003. Automatic detection of causal relations for question answering. In *Proceedings of the ACL 2003 workshop on Multilingual summarization and question answering-Volume 12*, pages 76–83. Association for Computational Linguistics.
- Brian Hazlehurst, Allison Naleway, and John Mullooly. 2009. Detecting possible vaccine adverse events in clinical notes of the electronic medical record. *Vaccine*, 27(14):2077–2083.
- Miguel A Hernán and Stephen R Cole. 2009. Invited commentary: causal diagrams and measurement bias. *American journal of epidemiology*, 170(8):959–962.
- Randy M Kaplan and Genevieve Berry-Rogghe. 1991. Knowledge-based acquisition of causal relationships in text. *Knowledge Acquisition*, 3(3):317–337.
- Manabu Kuroki and Judea Pearl. 2014. Measurement bias and effect restoration in causal inference. *Biometrika*, 101(2):423–437.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Recurrent convolutional neural networks for text classification. In *AAAI*, volume 333, pages 2267–2273.
- Virgile Landeiro and Aron Culotta. 2016. Robust text classification in the presence of confounding bias. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Virgile Landeiro and Aron Culotta. 2017. Controlling for unobserved confounds in classification using correlational constraints. *arXiv preprint arXiv:1703.01671*.
- Stéphane M Meystre, Guergana K Savova, Karin C Kipper-Schuler, and John F Hurdle. 2008. Extracting information from textual documents in the electronic health record: a review of recent research. *Yearbook of medical informatics*, 17(01):128–144.
- Karin B Michels, Sander Greenland, and Bernard A Rosner. 1998. Does body mass index adequately capture the relation of body composition and body size to health outcomes? *American Journal of Epidemiology*, 147(2):167–172.
- Razieh Nabi and Ilya Shpitser. 2017. Semi-parametric causal sufficient dimension reduction of high dimensional treatments. *arXiv preprint arXiv:1710.06727*.
- Jerzy Neyman. 1923. Sur les applications de la thar des probabilités aux expériences agaricales: Essay des principe. excerpts reprinted (1990) in English. *Statistical Science*, 5:463–472.
- Michael J Paul. 2017. Feature selection as causal inference: Experiments with text classification. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 163–172.
- Judea Pearl. 1995. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.
- Judea Pearl. 2009. *Causality*. Cambridge university press.

- Judea Pearl. 2010. On measurement bias in causal inference. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 425–432. AUAI Press.
- Margaret E Roberts, Brandon M Stewart, and Richard A Nielsen. 2018. Adjusting for confounding with text matching.
- James M Robins, Andrea Rotnitzky, and Daniel O Scharfstein. 2000. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical models in epidemiology, the environment, and clinical trials*, pages 1–94. Springer.
- D. B. Rubin. 1976. Causal inference and missing data (with discussion). *Biometrika*, 63:581–592.
- Ilya Shpitser, Karthika Mohan, and Judea Pearl. 2015. Missing data as a causal and probabilistic problem. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 802–811. AUAI Press.
- Ilya Shpitser and Judea Pearl. 2008. Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research*, 9(Sep):1941–1979.
- Anastasios Tsiatis. 2007. *Semiparametric theory and missing data*. Springer Science & Business Media.
- Özlem Uzuner, Ira Goldstein, Yuan Luo, and Isaac Kohane. 2008. Identifying patient smoking status from medical discharge records. *Journal of the American Medical Informatics Association*, 15(1):14–24.
- Stefan Wager and Susan Athey. 2017. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*.
- Richard Wicentowski and Matthew R Sydes. 2008. Using implicit information to identify smoking status in smoke-blind medical discharge summaries. *Journal of the American Medical Informatics Association*, 15(1):29–31.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657.

A Simple Confounding

$$p(Y(a)) = \sum_C p(Y(a) | C)p(C) \quad (4)$$

$$= \sum_C p(Y(a) | A, C)p(C) \quad (5)$$

$$= \sum_C p(Y | A, C)p(C) \quad (6)$$

Eq 5 holds because $Y(a) \perp A | C$, as seen in Figure 1a. Plugging this distribution into $\tau_S = E[Y(1)] - E[Y(0)]$ gives us the causal effect presented in Figure 4, Eq 1.

This assumes that an intervention on A is well-defined; if we *did* conduct a randomized control trial, we *could* assign $A = a$ and break A 's dependence on C .

In general, this step requires that we condition on all “back-door” paths between the treatment and the outcome. In Figure 1(a), if we did not have data on C , we could not block the back-door path between A and Y .

Eq 6 holds due to consistency. We assume that, given we intervened to set $A = a$, if that individual would have been assigned $A = a$ in nature, then the distribution over Y is the same.

B Missing Data

Denote $p(Y(A(1) = a)) = p(Y(a))$.

First, we identify the causal effect in terms of the true $A(1)$.

$$p(Y(a)) = \sum_C p(Y(a) | C)p(C) \quad (7)$$

$$= \sum_C p(Y(a) | A(1), C)p(C) \quad (8)$$

$$= \sum_C p(Y | A(1), C)p(C) \quad (9)$$

Where 7 holds by chain rule, 8 holds by $A(1) \perp Y(a) | C$, and 9 by consistency.

Now, we identify $A(1)$ in terms of observed data.

$$p(A(1), C, Y) = p(A(1) | C, Y)p(C, Y) \quad (10)$$

$$= p(A(1) | C, Y, R_A = 1)p(C, Y) \quad (11)$$

$$= p(A | C, Y, R_A = 1)p(C, Y) \quad (12)$$

Where 10 holds by chain rule, 11 by $A(1) \perp R_A | C, Y$, and 12 by consistency.

Now, use Eq 12 to identify $p(Y | A(1), C)$ from Eq 9 in terms of observed data.

$$p(Y | A(1), C) = \frac{p(Y, A(1), C)}{p(A(1), C)} \quad (13)$$

$$= \frac{p(Y, A(1), C)}{\sum_Y p(Y, A(1), C)} \quad (14)$$

$$= \frac{p(A | C, Y, R_A = 1)p(C, Y)}{\sum_Y p(A | C, Y, R_A = 1)p(C, Y)} \quad (15)$$

$$= \frac{p(A | C, Y, R_A = 1)p(Y | C)}{\sum_Y p(A | C, Y, R_A = 1)p(Y | C)} \quad (16)$$

Where 13 holds by definition, 14 holds by marginalization, 15 holds by an application of 12 twice, and 16 holds by canceling out $p(C)$.

If we include text in this derivation, we simply replace $p(A | C, Y, R_A = 1)$ with $p(A | \mathbf{T}, C, Y, R_A = 1)$, where \mathbf{T} is all our text variables.

Finally, combine Eq 9 and Eq 16 to get:

$$p(Y(A(1) = a)) = \sum_C \frac{p(A | C, Y, R_A = 1)p(Y | C)}{\sum_Y p(A | C, Y, R_A = 1)p(Y | C)}p(C) \quad (17)$$

Plugging this distribution into $\tau_{MD} = E[Y(1)] - E[Y(0)]$ gives us the causal effect presented in Figure 4, Eq 2.

C Measurement Error

Define the following terms for convenience:

$$\epsilon_{c,y} = p(A = 0 | A^* = 1, C = c, Y = y) \quad (18)$$

$$\delta_{c,y} = p(A = 1 | A^* = 0, C = c, Y = y) \quad (19)$$

$$q_{c,y}(0) = p(C = c, Y = y, A^* = 0) \quad (20)$$

$$q_{c,y}(1) = p(C = c, Y = y, A^* = 1) \quad (21)$$

Eq (5) and (7) from Pearl 2010 gives us:

$$p(A = 1, C = c, Y = y) = \frac{-\delta_{c,y}q_{c,y}(0) + (1 - \delta_{c,y})q_{c,y}(1)}{(1 - \epsilon_{c,y} - \delta_{c,y})} \quad (22)$$

$$p(A = 0, C = c, Y = y) = \frac{(1 - \epsilon_{c,y})q_{c,y}(0) - \epsilon_{c,y}q_{c,y}(1)}{(1 - \epsilon_{c,y} - \delta_{c,y})} \quad (23)$$

Now,

$$p(Y | A = 1, C) = \frac{p(Y, A = 1, C)}{p(A = 1, C)} \quad (24)$$

$$= \frac{p(Y, A = 1, C)}{\sum_Y p(Y, A = 1, C)} \quad (25)$$

$$= \frac{-\delta_{c,y}q_{c,y}(0) + (1 - \delta_{c,y})q_{c,y}(1)}{(1 - \epsilon_{c,y} - \delta_{c,y})} \\ = \frac{-\delta_{c,y'}q_{c,y'}(0) + (1 - \delta_{c,y'})q_{c,y'}(1)}{\sum_{y'} \frac{(1 - \epsilon_{c,y'} - \delta_{c,y'})}{(1 - \epsilon_{c,y'} - \delta_{c,y'})}} \quad (26)$$

and then,

$$p(Y | A = 0, C) = \frac{p(Y, A = 0, C)}{p(A = 0, C)} \quad (27)$$

$$= \frac{p(Y, A = 0, C)}{\sum_Y p(Y, A = 0, C)} \quad (28)$$

$$= \frac{(1 - \epsilon_{c,y})q_{c,y}(0) - \epsilon_{c,y}q_{c,y}(1)}{(1 - \epsilon_{c,y} - \delta_{c,y})} \\ = \frac{(1 - \epsilon_{c,y'})q_{c,y'}(0) - \epsilon_{c,y'}q_{c,y'}(1)}{\sum_{y'} \frac{(1 - \epsilon_{c,y'} - \delta_{c,y'})}{(1 - \epsilon_{c,y'} - \delta_{c,y'})}} \quad (29)$$

Plugging this distribution into $\tau_{ME} = E[Y(1)] - E[Y(0)]$ gives us the causal effect presented in Figure 4, Eq 3.

D Synthetic Data Distribution

In the distributions below, $\text{Ber}(p)$ is used as the abbreviation a Bernoulli distribution with probability p .

Below, s_i , u_i and v_i are the effect of C, A, and Y on the probability of word T_i ; each is drawn from $\mathcal{N}(0, \zeta)$, a parameter which controls how correlated words are with the underlying variables. When ζ is close to 0, the words are essentially random. When ζ is large, the words are essentially deterministic functions of the underlying variables. Similarly w_i is the effect of word T_i on R_A , and is drawn from $\mathcal{N}(0, \eta)$.

For both settings, we set vocabulary size to 4,334 (to match Yelp experiments) and $\zeta = 0.5$. For the missing data setting, we set $\eta = 0.1$. We picked these constants by empirically finding a reasonable middle ground between the text data providing only noise and being a deterministic function of their parents. We picked all other constants such that the naïve correlation $p(Y | A)$ was a poor estimate of the counterfactual $p(Y(a))$ in the full-data setting.

D.1 Missing data data-generation

$$C \sim \text{Ber}(0.4)$$

$$A(1) \sim \text{Ber}(-0.3C + 0.4)$$

$$Y \sim \text{Ber}(0.2C + 0.1A + 0.5)$$

$$T_i \sim \text{Ber}(0.5 + u_i A + v_i C)$$

$$R_A \sim \text{Ber}\left(0.7 + 0.2C - 0.4Y + \sum_i w_i T_i\right)$$

D.2 Measurement error data-generation

$$C \sim \text{Ber}(0.4)$$

$$A \sim \text{Ber}(-0.3C + 0.4)$$

$$Y \sim \text{Ber}(0.2C + 0.1A + 0.5)$$

$$T_i \sim \text{Ber}(0.5 + s_i C + u_i A + v_i Y)$$

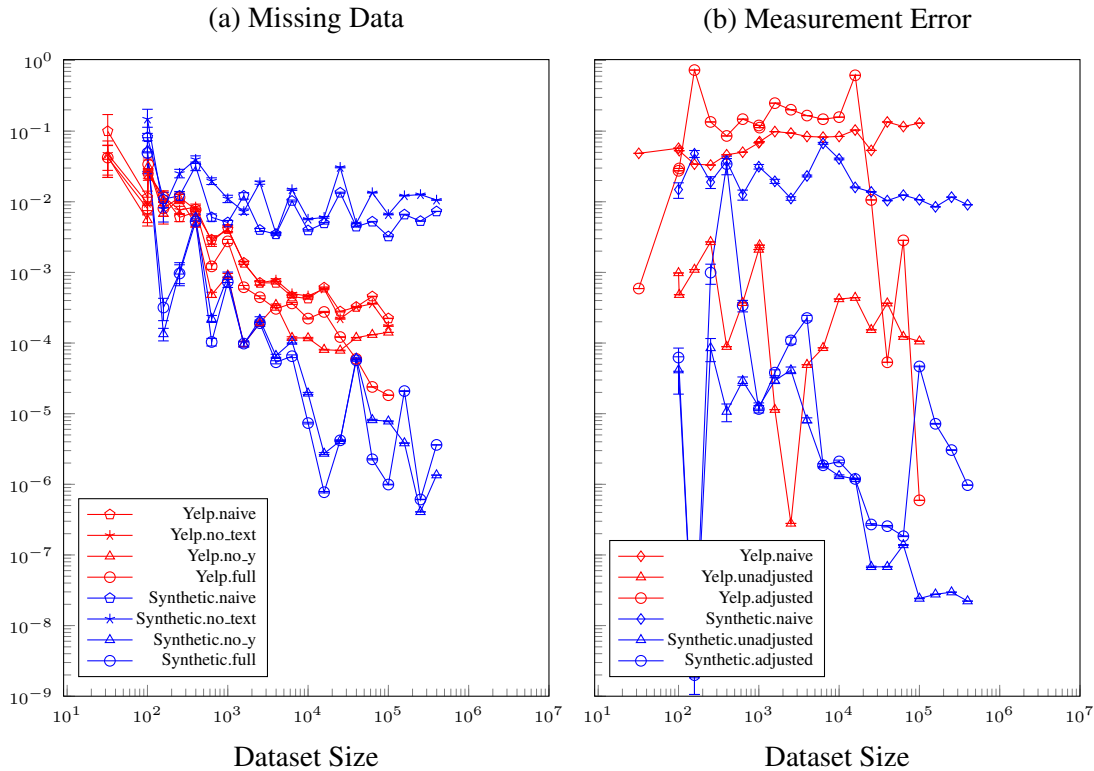


Figure 6: Experimental results with a vocabulary of size 53,197. Squared distance (y-axis, lower is better) of the estimated causal effect from τ_{SC} calculated from the full data with no missing data or measurement error. Error bars (negligible for larger datasets) are 1.96 times standard error across 10 experiments.

E Additional Experiments

Figure 6 shows the results of a second set of experiments, which are identical to those described in §5 except the vocabulary size is now 53,197 instead of 4,334. For the Yelp data, the larger vocabulary consists of all words which appear at least ten times in a sample of 1M reviews. As the larger vocabulary introduced greater memory requirements, we did not run these experiments with as large of datasets.

The results of these experiments show roughly the same patterns as those seen in Figure 5. The adjusted measurement error models again appear erratic, generally performing worse than the unadjusted models though better than the naive models.

The full missing data model appeared to slightly outperform the no_y model on Yelp data but only perform as well on the synthetic data. Both these models appeared better than the naive and no_text models on both datasets.