

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

CAUSAL FEATURE SELECTION WITH DIMENSION REDUCTION FOR INTERPRETABLE TEXT CLASSIFICATION

A PREPRINT

Guohou Shan¹ James Foulds² Shimei Pan^{2*}

¹Temple University

²University of Maryland, Baltimore County
tul05519@temple.edu, {jfoulds,shimei}@umbc.edu

October 12, 2020

ABSTRACT

Text features that are correlated with class labels, but do not directly cause them, are sometimes useful for prediction, but they may not be insightful. As an alternative to traditional correlation-based feature selection, causal inference could reveal more principled, meaningful relationships between text features and labels. To help researchers gain insight into text data, e.g. for social science applications, in this paper we investigate a class of matching-based causal inference methods for text feature selection. Features used in document classification are often high dimensional, however existing causal feature selection methods use Propensity Score Matching (PSM) which is known to be less effective in high-dimensional spaces. We propose a new causal feature selection framework that combines dimension reduction with causal inference to improve text feature selection. Experiments on both synthetic and real-world data demonstrate the promise of our methods in improving classification and enhancing interpretability.

1 Introduction

In text classification, representing text often requires a large number of features. For example, TF*IDF, and n -grams are frequently used in representing documents. Since even the number of unigram or TF*IDF features is equal to the number of unique words in a text corpus, we can easily extract thousands or even millions of features from text. Moreover, when the ground truth data needed to train a text classifier is limited, the classifier can easily overfit the training data and perform poorly on new data. Dimension reduction methods are thus important for improving both accuracy and interpretability.

Feature selection is a common approach for dimension reduction in which noisy features that may confuse a classifier are removed before or during training. Compared with projection-based dimension reduction methods such as PCA Li et al. [2014] or neural network-based embedding methods such as BERT Devlin et al. [2019] which learn new latent features that may be hard to interpret, feature selection identifies text features that are directly interpretable, since word-based features are meaningful to humans by definition. The most commonly used feature selection techniques are correlation-based Hall [1999], Yu and Liu [2003]. Since an observed correlation could be due to uncontrolled confounding variables, they are not ideal for providing insight. For example, an observed positive correlation between *ice cream sales* and *shark attacks* can be misleading as this may be caused by a confounding variable, *temperature increase*. This can also happen with text, e.g. *ice cream sales*, *shark attacks*, and *temperature increase* are now n -grams. Since causal inference can remove spurious associations and uncover principled relationships between input features and output labels, it has the potential to provide better insight than correlation-based feature selection.

The only prior work known to us on causal text feature selection is that of Paul [2017], who uses Propensity Score Matching (PSM) for causal inference [Rubin, 1974]. However, PSM is not ideal for high-dimensional data such as text King and Nielsen [2019]. In this study, we propose a novel approach to address the main limitations of Paul [2017]. The main contributions of our research are:

- We develop and formalize a causal feature selection framework which generalizes, makes precise, and specifies the assumptions made by an approach previously considered by Paul [2017].
- We develop a new causal text feature selection method, including a principled approach to leverage dimensionality reduction in the process, which addresses one of the main limitations in Paul [2017].
- We evaluate our methods on two real-world and two synthetic datasets. Our results demonstrate that our methods achieve better classification performance as well as improved interpretability compared to widely used correlation and causal-inference based methods.

2 Related Work

Most existing work on text feature selection is correlation-based Le and Ho [2015], Agarwal and Mittal [2012], Hall [1999]. The only work known to us on causal text feature selection employs Propensity Score Matching (PSM) Paul [2017]. Since PSM was shown to not perform well in high dimensional spaces due to issues such as bias, statistical inefficiency, and imbalanced matching Gu and Rosenbaum [1993], Caliendo and Kopeinig [2008], King and Nielsen [2019], it is not ideal for text feature selection.

In NLP, there is also some work on extracting causal relationships expressed in text Dasgupta et al. [2018], Mirza and Tonelli [2016]. However, our task is fundamentally different. Rather than identifying causal relations expressed in text, which is an Information Extraction (IE) task, we are using causal models to identify significant features that have causal relations with outcome variables.

There is also work on applying NLP techniques to aid causal inference Wood-Doughty et al. [2018], and on applying causal inference to aid other NLP tasks Landeiro and Culotta [2016], Roberts et al. [2020]. For example, Roberts et al. [2020] proposed a matching method called Topical Inverse Regression Matching for causal inference controlling for text data. Their work does not study feature selection.

3 Background: Causal Inference

In this section, we describe the causal inference framework our research is based on. We also describe two widely used causal inference methods instantiated from this framework.

3.1 Neyman-Rubin Potential Outcomes Framework

The framework is a popular causal inference formulation Neyman [1923 and 1990], Rubin [1974]. Within the framework, for a sample i , there are two possible outcomes (e.g., class labels): $Y_i(1)$ if i undergoes treatment T and $Y_i(0)$ if not. The *treatment effect* for observation i is defined by $Y_i(1) - Y_i(0)$. Since for each sample i , we observe either $Y_i(1)$ or $Y_i(0)$, one of the two conditions is always missing in observational data.

A naive approach for testing the effect of a treatment T is to directly compare the treated and control groups. This works in randomized experiments where the treatment assignment is random. Since random assignment of samples to treatment groups balances both known and unknown sample characteristics that may affect the outcome, it reduces the likelihood that there will be differences in the sample characteristics between treated and control groups. In an observational study, however, this balance can be distorted by the systematic assignment of the treatment to samples Rosenbaum and Rubin [1983], Pearl [2010], Hill and Su [2013]. Thus, the *average treatment effect (ATE)* for the treated group cannot be directly estimated. To address this issue, typical approaches aim to counteract the relationship between treatment and covariates by matching instances in the treatment and control groups.

3.2 Mahalanobis Distance Matching (MDM) and Propensity Score Matching (PSM)

The most straightforward and nonparametric way to match treatment and control instances is to exactly match on the pre-treatment covariates X . This approach is infeasible if the sample size is limited and the dimensionality of X is large or if X contains continuous covariates. Coarsened exact matching (CEM) performs binning on each variable before matching identical instances Iacus et al. [2012], but does not address the curse of dimensionality. Two common approaches to solve this problem are nearest neighbor matching (NNM), often based on the Mahalanobis distance (MDM) to account for scaling and correlations Cochran and Rubin [1973], Rubin [1979], and propensity score matching (PSM) Rosenbaum and Rubin [1983]. The Mahalanobis distance between any two samples X_i, X_j is defined as

$$MD(X_i, X_j) = \{(X_i - X_j)' S^{-1} (X_i - X_j)\}^{1/2}$$




Figure 1: Assumptions of our Causal Feature Selection (CFS) framework. (A) Causal feature selection with observational features. (B) Causal feature selection with latent variable assumption. (C) Causal feature selection with dimensionality reduction.

where S is the empirical covariance matrix of X . In MDM, to estimate the treatment effect, one matches each treated sample with the closest control sample, as calculated by the Mahalanobis distance.

Another widely used matching approach is Propensity Score Matching (PSM) Rosenbaum and Rubin [1983]. PSM matches on the probability of assignment to treatment, known as the propensity score: $e(X_i) = \Pr(T_i = 1|X_i)$. It is typically estimated using logistic regression. The unidimensional metric acts as a one-number summary of the covariates, such that conditioning on the propensity score $e(X_i)$ leads to conditional independence between the features in the treatment and control groups, thereby approximating a randomized controlled experiment. Although very popular, PSM has some limitations. For example, King et al. [2016] argue that PSM leads to more imbalance between treatment and control groups than matching methods that approximate full blocking such as MDM and CEM. PSM also exhibits poor performance in high-dimensional spaces due to biased estimation, imbalance, and statistical inefficiency Gu and Rosenbaum [1993], Abadie and Imbens [2006], Caliendo and Kopeinig [2008], King and Nielsen [2019].

In the next section, we propose a causal feature selection framework for text classification, which generalizes and formalizes the method proposed by Paul [2017] and aims to address some of its limitations.

4 Problem Formulation

Figure 1 illustrates our problem formulation for causal feature selection in three scenarios. In the general case (Figure 1 (A)), we assume that there are classification features X (e.g., bag-of-words or TF*IDF features for a document), and an outcome Y (i.e., the class label). The features are divided into two main types: causal features X_C , which directly cause the outcome Y , and other non-causal observed features X_O which have no direct causal relationship with Y . We select one of the features as a “treatment” variable X_T , assumed to be binary, e.g. whether a word is present in a document. Our goal is to use causal inference to determine whether the treatment feature X_T directly causes the outcome Y , i.e. is X_T a member of X_C , or of X_O ? In the full causal feature selection process, each feature will be considered as a “treatment” in turn. The features that are found to cause Y will be selected by our method.

Figure 1 (A) shows a graphical model diagram that describes the assumptions we make about the data generating process when applying our causal feature selection method. The diagram combines factor graph notation, where a rectangular factor node indicates a clique of dependent variables, with directed edges that represent causal relationships. We assume that X_O , X_C , and X_T are related to each other in unknown ways, as indicated by their shared factor node in the factor graph. The directed edge from X_C to Y indicates that X_C has a direct causal effect on Y . An edge with a question mark is shown from X_T to Y since X_T may potentially directly cause Y , and our goal is to verify whether this is the case. All the variables in the graph are observed in this setting, so their nodes are shaded.

In a perfect world, we would have “twin” instances that differ only in X_T so that we could observe both potential outcomes $Y_i(1)$ and $Y_i(0)$ for each instance i to directly estimate the average treatment effect (ATE), and hence verify whether X_T is causal for Y . In practice, we can only observe one outcome per instance, and the treatment is also correlated with the other variables, unlike in a randomized controlled trial. We must correct for this using causal inference.

Previously, Paul [2017] proposed a particular instantiation of our general framework in Figure 1 (A) in which PSM is applied to word-level features X .

On the other hand, it is well known that causal inference techniques such as PSM struggle with data that are high-dimensional Gu and Rosenbaum [1993], Abadie and Imbens [2006], Caliendo and Kopeinig [2008], King and Nielsen [2019], as is common with word-level text features. Instead, we propose to leverage a latent variable model (LVM) to reduce the effective dimensionality of the causal inference problem by inferring the latent variables and using them as the features for causal inference algorithms, thereby avoiding the problems with PSM and related techniques in high-dimensional text data.

We aim to simplify the complex high-dimensional causal inference problem in Figure 1 (A). To this end, we can consider simplifying assumptions on the joint distribution of X . A common approach used by latent variable models is to assume that the relationships between the features X are mediated by latent variables L , which provide low-dimensional representations of high-dimensional data such as text. For example, in the LDA topic model the complex dependencies between all words in a document are simplified by assuming that they are generated independently given a latent distribution over topics. Figure 1 (B) shows a latent variable model instantiation of Figure 1 (A), where there are K latent variables L , e.g. the document’s proportions for K topics, which explain the features X . After assuming an LVM model as in Figure 1 (B), we arrive at our formulation for causal feature selection method with dimension reduction by marginalizing out X_O and X_C , then representing the instances via the latent variables L instead of the observable features (Figure 1 (C)). To determine whether X_T is causal of Y , we now only need to control for differences in the lower-dimensional latent features using causal inference techniques, which is expected to be typically easier to achieve compared to using high-dimensional X . To implement this in practice, we will use an algorithm to infer, hence effectively “observe” L based on X .

Note that causal inference controlling for L is different to causal inference controlling for X , and may lead to different conclusions. However, since we assumed that X was generated based on L , in many cases it is reasonable to assume that L captures much of the important information in X . Furthermore, for many NLP models L is semantically meaningful. For example, in LDA L represents the topical content of the document, and conditioning on this representation may just as desirable, or even more desirable, compared to directly controlling for the differences in the frequencies of other word count features X .

5 Proposed Methods

In this section, we introduce our proposed causal feature selection (CFS) method which includes four key steps: dimension reduction, matching, feature selection and classification.

5.1 Dimension Reduction

Given all of the observed features X and the latent dimensionality K , we employ dimension reduction to uncover L which captures the main characteristics of each sample i . Our hypothesis is that matching based on L will help us find better treated and control pairs than directly matching on surface features X , especially when $|X|$ is large. We explore a diverse set of dimension reduction methods including Principal Component Analysis (PCA), Gaussian Random Projection (GRP) and Mini-Batch Dictionary Learning (MBDL). Specifically, **PCA** learns a low-dimensional representation of X under independent identical Gaussian noise. It is one of the most widely used methods for dimension reduction Zhang et al. [2019]. The normal version of PCA (**nPCA**) learns a linear combination of X while sparse PCA (**sPCA**) uses Lasso (or Elastic Net) to derive modified principal components with sparse loading Zou et al. [2006]. Moreover, random projection is a powerful method for dimension reduction, which theoretically preserves distances in the original feature space Bingham and Mannila [2001]. When the random projection is produced by Gaussian distributions, it becomes Gaussian Random Projection (**GRP**). Finally, **MBDL** learns a dictionary (a set of atoms) that can be used to represent data using a sparse code Mairal et al. [2009]. In addition, for text feature selection, we also employ dimension reduction techniques widely used in text analysis, Latent Dirichlet Allocation (**LDA**) Blei et al. [2003] and Document Embedding (e.g., **Doc2Vec** Le and Mikolov [2014]), which both map a document from a high-dimensional sparse vector representation (e.g., 1-hot representations) to a low dimensional latent vector representation. Compared to general dimension reduction methods such as PCA, text-based dimension reduction methods may have advantages as they capture important concepts and semantic relations in a document.

5.2 Matching

Matching plays an important role in reducing the bias due to confounding variables. We investigate different matching techniques: (1) with ground truth matching pairs (e.g., “identical twins”) (2) matching based on surface features using Nearest Neighbor Matching (NNM) and (3) our proposed approach, matching based on latent features using NNM. For a given binary treatment variable X_T (e.g., whether a word appears in a document or not), we assign all the samples

satisfying $X_T = 1$ to the treated group (e.g., those with a particular word) and those with $X_T = 0$ to the control group (those without a particular word). If X_T is continuous (e.g., the TF*IDF score of a word), we first transform it to a binary variable (e.g., $\hat{X}_T = 1$ if the TF*IDF score is non-zero, or above the mean, and $\hat{X}_T = 0$ otherwise).

We then pair each sample in the treated group with a sample from the control group. Assuming ground truth matching pairs are not available, we employ NNM to find the best matching pairs based on a chosen similarity measure (e.g., cosine similarity). For example, in PSM, given the propensity score of a sample from the treated group, we use NNM to identify a sample from the control group whose propensity score is the closest. We may also use NNM to match treated and control samples based on latent or surface features. In both cases we use cosine similarity due to its success in matching text for information retrieval. To improve matching speed when the number of samples or the number of features is large, we employ a KD-tree based search algorithm Bentley [1975] to quickly locate the nearest neighbor.

5.3 Feature Selection

After the samples are paired, we need to decide whether there is any significant difference between the treated and control group w.r.t. the outcome. Following Paul [2017], if the outcome variable is binary, we employ McNemar’s test McNemar [1947], which has its test statistic as

$$\chi^2 = \frac{(TN - CP)^2}{TN + CP}, \quad (1)$$

where TN is the number of samples in the treated group with a negative class label and CP is the number of samples in the control group with a positive class label. Based on the p -value associated with the McNemar’s test, we can select the target feature X_T if the p -value is less than α , where α is a hyperparameter that decides the threshold for the significance test. If the outcome variable is numerical, we can instead compare the outcome means between the treated and control group using a t -test.

5.4 Classification

On the real datasets, we experimented with several commonly used machine learning classifiers: logistic regression (**LR**), random forest (**RF**), support vector machine (**SVM**), and multi-layer perceptron (**MLP**). Since **RF** performed the best on the real datasets, we reported only the results with **RF**. Algorithm 1 summarizes our causal feature selection process.

6 Datasets

To evaluate the effectiveness of our methods, we used two synthetic datasets and two real datasets.

Synthetic Datasets: Our proposed framework is based on two hypotheses: (1) matching quality plays a significant role in causal feature selection, and (2) dimension reduction is effective in matching treated-control pairs. To test these hypotheses directly, we designed two synthetic datasets.

The **Latent** dataset is designed to test the system performance when we are able to observe the “latent” features directly. Since there is no need to employ dimension reduction to uncover latent variables in this setting, we can focus on investigating the impact of matching quality on causal feature selection. Specifically, we create 250 samples with 50 “latent” features L generated from multi-normal distributions with 0 means. To create ground truth treated-control pairs, we generate another 250 samples by adding small noise to the existing samples. We also create a new binary treatment feature X_T to indicate that the first 250 samples belong to the treated group and the second 250 samples belong to the control group. To test the impact of feature selection on classification, we generate a binary outcome variable Y based on L and X_T using a logistic regression function. We also generate 10 random variables called X_O that are not related to Y . We repeat the process 50 times to generate 50 slightly different datasets.

The **Surface** dataset simulates a more realistic setting where latent features are not observable. First we use linear regression functions to generate 100 observable features X_C from the 50 latent features L in the **Latent** dataset. The outcome variable Y is generated from the treatment feature X_T and the 100 observable features X_C using a logistic function. To simulate a real text dataset, we also generated a large number of normally distributed random variables X_O (a total of 2900 of them) that are not relevant to the outcome variable Y . We randomly generated 50 slightly different surface datasets in our study.

Real Datasets: The **movie dataset (Movie)** contains reviews from IMDB Maas et al. [2011]. Movies are rated on a 1-10 scale and the reviews with a rating ≥ 7 are labeled as positive while reviews rated ≤ 4 are labeled as negative. Similarly to Paul [2017], we discard neutral reviews. Since feature selection is most useful in low resource scenarios

Algorithm 1 The Causal Feature Selection Algorithm

Require: A dataset S with labels Y , surface feature vectors X ; a classification method F ; size of latent vectors K ; a dimension reduction method G , a nearest neighbor matching method NMM , a statistical significance test method SST , a significance level threshold α and a similarity threshold β ;

Ensure: Selected feature list f_lst ; Classification performance pf

```

1: Initialize  $f\_lst = []$ ;
2: for each feature  $X_t \in X$  do
3:   Initialize a matching pair list  $p\_lst = []$ ; treated group  $T = []$ ; control group  $C = []$ ;
4:   for each  $s_i \in S$  do
5:     if  $X_{ti} == 1$ , add  $s_i$  to  $T$ 
6:     else add  $s_i$  to  $C$ 
7:   end for
8:    $L_K = G(\{X - X_t\}, K)$  where  $\{X - X_t\}$  includes all the surface features in  $X$  except  $X_t$  and  $L_K$ 
9:   represents a  $K$ -dimensional latent feature space.
10:  for each sample  $t_j \in T$  do
11:     $c_{tj} = NNM(t_j, C, Sim_{LK})$  where  $Sim_{LK}$  is a similarity measure defined over  $K$ -dimensional
12:    latent feature space  $L$ 
13:    if  $Sim_{LK}(t_j, c_{tj}) > \beta$ , add the pair  $(t_j, c_{tj})$  in  $p\_lst$ ;
14:  end for
15:  Assume  $Y_{p\_lst(T)}$  is the aggregated outcomes associated with all the treated samples in  $p\_lst$  and
16:   $Y_{p\_lst(C)}$  is the aggregated outcomes associated with all the control samples in  $p\_lst$ 
17:  Calculate the  $p$ -value using  $SST$  to assess the difference between  $Y_{p\_lst(T)}$  and  $Y_{p\_lst(C)}$ ;
18:  if  $p\text{-value} \leq \alpha$ , add  $X_t$  in  $f\_lst$ ;
19: end for
20:  $\hat{Y} = F(f\_lst)$ ;
21: Compute classification performance  $pf$  based on  $Y$  and  $\hat{Y}$ 

```

	#samples	#features
<i>Latent</i>	500	61
<i>Surface</i>	500	3,001
<i>Movie</i>	400-1,000	3,028
<i>SOTU</i>	54	2,969

Table 1: Summary of the Synthetic and Real Datasets

where the training sample size is smaller than the feature size, we only used a small subset of the IMDB dataset, whose size we varied systematically. We adjusted the infrequent word filtering threshold to keep the feature size similar to that of the **Surface** dataset.

The **State of the Union (SOTU)** dataset contains the annual speech by the presidents of the United States to the congress.² We use this dataset to study the impact of the speech on a president’s post-address approval rate. Each state of the union address is annotated with an “increase” or “decrease” label based on a president’s Gallup Poll approval rate before and right after the SOTU address.³ We also filtered infrequent words and kept the feature size similar to that of the **Surface** dataset. Table 1 summarizes the statistics of these datasets including sample and feature size.

7 Experiments on Synthetic Datasets

We use the **Latent** dataset to test systems when there is no need for dimension reduction. First, we study the performance of Causal Feature Selection (CFS) if it is given “perfect matches (PM)” (**CFS-PM**). We compare it with a system that performs CFS based on a randomly generated matching pairs (**CFS-RM**) or pairs discovered by NNM using latent features (**CFS-LM**). We use the **Surface** dataset to test the impact of various dimension reduction methods such as **nPCA**, **sPCA**, **GRP** and **MBDL**.

Evaluation Metric: An evaluation metric based on feature selection decisions (i.e., whether statistical significance

²<https://www.presidency.ucsb.edu/documents/presidential-documents-archive-guidebook/annual-messages-congress-the-state-the-union>

³<https://www.presidency.ucsb.edu/statistics/data/presidential-job-approval>

Methods	CFS-PM	CFS-LM
RC	0.62	0.62
Methods	PSM	CFS-RM
RC	0.57	0.49
Methods	L1	MDM
RC	0.52	0.54

Table 2: Performance on *Latent*

Methods	CFS-PM	CFS-NM	L1	PSM	CFS-RM
RC	0.78	0.52	0.51	0.55	0.44
Methods	CFS-nPCA10	CFS-nPCA50	CFS-sPCA10	CFS-sPCA50	CFS-GRP10
RC	0.54	0.62	0.54	0.68	0.5
Methods	CFS-GRP50	CFS-MBDL10	CFS-MBDL50	MDM	
RC	0.54	0.52	0.56	0.54	

Table 3: Performance on *Surface*

α is reached) would be heavily dependent on α and the sample size. We circumvent this issue by instead using a rank-based approach, where we rank the features according to their p -values. Since X_T has a causal relationship with Y and the variables in X_O do not, we compare the ranks of X_T and a randomly selected irrelevant feature from X_O . If a system typically ranks X_T higher than X_O , it is good at identifying causal features. Our evaluation measure, called **Rank Correctness (RC)**, is defined as the probability a system correctly ranks X_T higher than X_O on the 50 randomly generated datasets.

Baselines: We compare our methods with **L1** regression, a classic association-based feature selection method, and **PSM** and **MDM**, two of the commonly used causal inference methods based on the potential outcomes framework. **PSM** was also adopted for causal feature selection in Paul [2017].

Results: Table 2 shows the performance of different models on the **Latent** dataset. **CFS-PM**(RC=0.62), **CFS-LM**(RC=0.62) and **PSM**(RC=0.57) all work better than **L1**(RC=0.52, $p < 0.05$ based on t-test). **CFS-PM** and **CFS-LM** also performed significantly better than **PSM** ($p < 0.05$). Unlike the above causal methods, **MDM** did not perform well (RC=0.54). Finally, as expected, random matching **CFS-RM** performed the worst (RC=0.49). On the **Surface** dataset (shown in Table 3), **CFS-PM** continued to be the best (RC= 0.78) and **CFS-RM** the worst (RC=0.44). Among the methods that employ dimension reduction, **CFS-sPCA50** performed the best (RC=0.68). It significantly outperformed **CFS-NM** (RC=0.52, $p < 0.01$), which matches treated and control samples based on surface features without dimension reduction. It also significantly outperformed all the baselines: **PSM** (RC=0.55, $p < 0.01$), **MDM**(RC=0.54, $p < 0.01$) and **L1**(RC=0.51, $p < 0.01$). In addition, it performed much better than **CFS-sPCA10** (RC=0.54, $p < 0.01$), which does not have access to the ground truth K . Models employing nPCA also performed quite well (e.g., **CFS-nPCA50**, RC=0.62).

In summary, based on the experiments on the synthetic datasets, we found that (1) models with access to ground truth “perfect twins” performed the best; (2) when there is an underlying latent space, dimension reduction can be used to improve model performance over those that match directly based on surface features without dimension reduction; (3) our proposed models can significantly outperform all the established baselines including **L1**, **MDM** and **PSM**; and (4) models employing the ground truth dimensionality K performed better than those that do not have access to this information.


8 Experiments on Real-World Datasets

Since treatment/control group assignment is a binary decision, on both the **Movie** and the **SOTU** dataset we binarize the word count for the treatment variable X_T for feature selection (1 if a word appears in a document and 0 if not). During classification, we use the TF*IDF scores of the selected features as the predicting variables. All the classification results reported here are based on the random forest model (**RF**). In addition to the models tested on the **Surface** dataset, we introduced three new models: **CFS-LDA**, **CFS-Doc2Vec** and **Full Model**. Both **LDA** and **Doc2Vec** are widely used dimension reduction methods for text analysis. The **Full Model** employs all the TF*IDF word features without any dimension reduction. Unlike the synthetic datasets, here we do not know the ground truth on whether a feature is causally related to an outcome or not. Thus, instead of **RC**, we adopt extrinsic evaluation measures typically used for text classification: recall, precision, and F1-score.

On the **Movie** dataset, to investigate how system performance may be impacted by sample size, we systematically varied the sample size from 400 to 1,000. For each dataset, we randomly split the samples into training (80%) and development (20%). We used the development data to tune model parameters (e.g., dimensionality K and α). Since the original movie dataset is much larger (with over 50,000 reviews), we randomly selected 100 new reviews from the remaining data as the test dataset. For **SOTU**, since it is a small dataset, we employ nested cross-validation for hyperparameter tuning and testing. In the outer loop, we employ leave-one-out cross validation to split data into training and testing. In the inner loop, we use five-fold cross validation to split data into training and development. The development data is used for hyperparameter tuning.

8.1 Classification Accuracy

Figure 2 shows the classification results on the **Movie** dataset. When the sample size was small (=400), **CFS-LDA** (F1=0.655) performed statistically significantly better than all the baselines, i.e. **L1** (F1=0.652, $p < 0.05$), the **Full model**

Figure 2: System Performance (F1) on the *Movie* dataset.

	CFS-LDA	PSM	L1	MDM
Pos	great	treasure	think	mass
	love	great	hollywood	trust
	pleasure	love	way	doubt
	enjoy	classic	millionaire	attack
	hollywood	wonder	bet	delight
Neg	inaccuracy	thriller	light	complaint
	lose	weak	teacher	tear
	issue	monster	fight	weak
	complaint	spoil	woman	heard
	terrible	technology	pal	realism

Table 5: Top-5 *pos* and *neg* words from the *Movie* data

Metrics Models	Recall	Precision	F1-score
CFS-nPCA	0.575	0.959	0.719
CFS-sPCA	0.568	0.961	0.714
CFS-GRP	0.570	0.941	0.710
CFS-MBDL	0.555	0.951	0.701
CFS-LDA	0.660	0.962	0.783
CFS-Doc2vec	0.608	0.948	0.741
PSM	0.583	0.945	0.721
MDM	0.580	0.939	0.717
L1	0.586	0.947	0.724
Full-model	0.578	0.947	0.718

Table 4: Model Performance on the *SOTU* Dataset.

	CFS-LDA	PSM	L1	MDM
Pos	faith	aspiration	determine	path
	god	courage	wheel	inspire
	motivation	effort	administer	benjamin
	point	fight	undertake	blue
	heart	goal	reinforce	burden
Neg	hitler	air	fate	congressman
	place	cynic	distress	challenge
	commit	place	jurisdiction	act
	history	poverty	terror	battle
	defeat	walk	inequality	bone

Table 6: Top-5 *pos* and *neg* words from *SOTU*

($F1=0.651$, $p<0.01$), **PSM** ($F1=0.651$, $p<0.001$), and **MDM** ($F1=0.653$, $p<0.05$). The pattern remained the same when we increased the sample size to 600, and 800. With sample size 600, the F1 for **CFS-LDA** was 0.678, which is significantly better than **L1** ($F1=0.677$, $p<0.05$), the **Full model** ($F1=0.676$, $p<0.01$), **PSM** ($F1=0.675$, $p<0.001$), and **MDM** ($F1=0.672$, $p<0.001$). With sample size 800, the F1 for **CFS-LDA** was 0.701, statistically significantly better than **L1** ($F1=0.697$, $p<0.001$), the **Full model** ($F1=0.699$, $p<0.05$), **PSM** ($F1=0.699$, $p<0.05$), and **MDM** ($F1=0.691$, $p<0.001$). Finally, when the sample size was increased to 1000, **CFS-LDA** and **CFS-Doc2vec** were the best performing models followed by **CFS-sPCA** and **MDM**, although the differences were not statistically significant. In all cases, **CFS-LDA** consistently performed the best. Table 4 shows the results on the *SOTU* dataset. Again, **CFS-LDA** had the best performance ($F1=0.783$). It performed statistically significantly better than all the baselines: **L1** ($F1=0.724$, $p<0.05$), the **Full model** ($F1=0.718$, $p<0.01$), **PSM** ($F1=0.721$, $p<0.01$), and **MDM** ($F1=0.717$, $p<0.01$).

8.2 Feature Interpretability

In addition to improving classification performance, causal feature selection can potentially improve interpretability by filtering out spurious associations between words and class labels. Tables 5 and 6 show the top-5 most significant positive- (*pos*) and negative-coefficient (*neg*) words per method according to p -value. Due to the page limit, we only report the results for **CFS-LDA**, the best-performing model in our framework, and the three feature selection baselines: **L1**, **PSM** and **MDM**. On the *Movie* dataset, except for “hollywood,” almost all the top words selected by **CFS-LDA** were sentiment words such as *great*, *love*, *pleasure*, *enjoy* and *complaint*. **PSM**, another causal feature selection method, identified sentiment words such as *great*, *love* as well as movie domain words such as *classic*, *thriller*, *monster*, *spoil*. There was no clear pattern in the top words selected by **L1** and **MDM**. Disturbingly, **L1** chose *woman* as a *neg* word. On *SOTU*, the *pos* words identified by **CFS-LDA** contained spiritual and emotional words such as *faith*, *god*, *motivation* and *heart*. The *pos* words identified by **PSM** were related to achieving (e.g., *aspiration*, *courage*, *effort*, *fight*, *goal*). The patterns of the words chosen by **L1** and **MDM** were unclear.

To gain insight into the types of words selected by each method, we used LIWC Pennebaker et al. [2015], a psycholinguistic lexicon, to map the top 100 words per method into psychologically and linguistically meaningful categories including *affective processes*, *cognitive processes*, *drives*, and *personal concerns* (Figure 3). On the *Movie* dataset, the causal methods gave the highest attention to *affective processes* (18.1%–19.6%), which we would apriori expect to be the most relevant category for sentiment analysis. In contrast, the words selected by **L1** were concentrated on *personal concerns* (16.7%), *social processes* (14.3%), and with *affective processes* (13.1%) in third place. On the *SOTU* dataset, except for **PSM** which had the highest concentration on *relativity*, all the other methods share the same top three LIWC categories: *affective processes*, *drives* and *personal concerns*. Finally, since the methods varied in LIWC percentages,

	Affective Proc.	Social Proc.	Cognitive Proc.	Perceptual Proc.	Biological Proc.	Drives	Time Orient.	Relativity	Personal Conc.	Informal Lang.
L1-Movie	13.1	14.3	8.3	8.3	3.6	11.9	9.5	11.9	16.7	2.4
CFS-LDA-Movie	18.1	13.8	11.7	7.4	3.2	11.7	7.4	11.7	12.8	2.1
MDM-Movie	18.1	9.5	12.4	2.9	5.7	16.2	9.5	12.4	13.3	0.0
PSM-Movie	19.6	13.7	10.8	3.9	2.9	17.6	6.9	7.8	16.7	0.0
Average	17.2	12.8	10.8	5.6	3.9	14.4	8.3	11.0	14.9	1.1
L1-SOTU	17.6	8.8	7.4	4.4	1.5	20.6	7.4	10.3	22.1	0.0
CFS-LDA-SOTU	16.5	6.6	9.9	2.2	7.7	19.8	4.4	12.1	20.9	0.0
MDM-SOTU	18.9	6.3	7.4	2.1	6.3	21.1	5.3	14.7	16.8	1.1
PSM-SOTU	15.6	6.2	12.5	2.1	3.1	17.7	5.2	18.8	17.7	1.0
Average	17.2	7.0	9.3	2.7	4.7	19.8	5.6	14.0	19.4	0.5

Figure 3: LIWC category percentages for the top 100 words selected by each method. Bold: highest percentages per method.

we also computed a stability metric, denoted sb , measuring the squared deviation of each method j from the consensus average: $sb_j \triangleq \text{mean}_{d,c}(\text{mean}_m(\text{percent}_{d,c,m} - \text{percent}_{d,c,j})^2)$, where j and m index methods, d indexes datasets, and c indexes LIWC categories. Among these methods, **CFS-LDA** was the most stable, in the sense that its LIWC percentages were closest to the consensus according to the sb metric: **CFS-LDA**: 1.90, **MDM**: 2.76, **PSM**: 4.19, and **L1**: 4.54.

9 Conclusion

We have proposed a novel causal feature selection framework which combines dimension reduction with causal inference to identify predictive and interpretable features for text classification. Our experiments on both real and synthetic datasets demonstrate the importance of dimension reduction in identifying high-quality matching pairs. The results show that employing commonly used dimension reduction techniques for text data within causal feature selection, e.g. via **CFS-LDA**, consistently outperforms baselines in all scenarios tested. We have also shown that the word features identified by our methods are easier to interpret versus correlation-based methods. Causal feature selection for interpretable text analysis is still in its infancy. This work is a small but significant step in advancing the state of the art.

References

- Li Li, Longkai Zhang, and Houfeng Wang. Multi-label text categorization with hidden components. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1816–1821, 2014.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi:10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- Mark Andrew Hall. *Correlation-based feature selection for machine learning*. PhD thesis, University of Waikato Hamilton, 1999.
- Lei Yu and Huan Liu. Feature selection for high-dimensional data: A fast correlation-based filter solution. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, pages 856–863, 2003.
- Michael J Paul. 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, 2017.
- Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Educational Psychology*, 66(5), 1974.
- Gary King and Richard Nielsen. Why propensity scores should not be used for matching. *Political Analysis*, 2019.
- Nguyen Hoai Nam Le and Bao Quoc Ho. A comprehensive filter feature selection for improving document classification. In *Proceedings of the 29th Pacific Asia Conference on Language, Information and Computation*, pages 169–177, Shanghai, China, October 2015. URL <https://www.aclweb.org/anthology/Y15-1020>.
- Basant Agarwal and Namita Mittal. Categorical probability proportion difference (CPPD): A feature selection method for sentiment classification. In *Proceedings of the 2nd Workshop on Sentiment Analysis where AI meets Psychology*, pages 17–26, Mumbai, India, December 2012. The COLING 2012 Organizing Committee. URL <https://www.aclweb.org/anthology/W12-5303>.
- Xing Sam Gu and Paul R. Rosenbaum. Comparison of multivariate matching methods: Structures, distances, and algorithms. *Journal of Computational and Graphical Statistics*, 2(4):405–420, 1993.

- Marco Caliendo and Sabine Kopeinig. Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72, 2008.
- Tirthankar Dasgupta, Rupsa Saha, Lipika Dey, and Abir Naskar. Automatic extraction of causal relations from text using linguistically informed deep neural networks. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 306–316, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi:10.18653/v1/W18-5035. URL <https://www.aclweb.org/anthology/W18-5035>.
- Paramita Mirza and Sara Tonelli. CATENA: CAusal and TEmporal relation extraction from NATural language texts. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 64–75, Osaka, Japan, December 2016. The COLING 2016 Organizing Committee. URL <https://www.aclweb.org/anthology/C16-1007>.
- Zach Wood-Doughty, Ilya Shpitser, and Mark Dredze. Challenges of using text classifiers for causal inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4586–4598, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi:10.18653/v1/D18-1488. URL <https://www.aclweb.org/anthology/D18-1488>.
- Virgile Landeiro and Aron Culotta. Robust text classification in the presence of confounding bias. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, page 186–193. AAAI Press, 2016.
- Margaret E Roberts, Brandon M Stewart, and Richard A Nielsen. Adjusting for confounding with text matching. *American Journal of Political Science (in press)*, 2020.
- Jerzy Neyman. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, 5(4), 1923 and 1990.
- Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- Judea Pearl. The foundations of causal inference. *Sociological Methodology*, 40(1):75–149, 2010.
- Jennifer Hill and Yu-Sung Su. Assessing lack of common support in causal inference using bayesian nonparametrics: Implications for evaluating the effect of breastfeeding on children’s cognitive outcomes. *The Annals of Applied Statistics*, 7(3):1386–1420, 2013.
- Stefano M. Iacus, Gary King, and Giuseppe Porro. Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1):1–24, 2012.
- William G Cochran and Donald B Rubin. Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 417–446, 1973.
- Donald B. Rubin. Using multivariate matched sampling and regression adjustment to control bias in observational studies. *Journal of the American Statistical Association*, 74(366):318–328, 1979.
- Gary King, Richard Nielsen, et al. Why propensity scores should not be used for matching. Copy at <http://j.mp/1sexgVw> Download Citation BibTex Tagged XML Download Paper, 378, 2016.
- Alberto Abadie and Guido W. Imbens. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267, 2006.
- Chihao Zhang, Kuo Gai, and Shihua Zhang. Matrix normal pca for interpretable dimension reduction and graphical noise modeling. *arXiv preprint arXiv:1911.10796*, 2019.
- Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.
- Ella Bingham and Heikki Mannila. Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 245–250. ACM, 2001.
- Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the 26th annual international conference on machine learning*, pages 689–696. ACM, 2009.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196, 2014.
- Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.

- Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947.
- Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*, pages 142–150. Association for Computational Linguistics, 2011.
- James W Pennebaker, Ryan L Boyd, Kayla Jordan, and Kate Blackburn. The development and psychometric properties of liwc2015. Technical report, University of Texas at Austin, 2015.