

Fighting Spurious Correlations in Text Classification via a Causal Learning Perspective

Yuqing Zhou and Ziwei Zhu

George Mason University
Fairfax, VA, USA

{yzhou31, zzhu20}@gmu.edu

Abstract

In text classification tasks, models often rely on spurious correlations for predictions, incorrectly associating irrelevant features with the target labels. This issue limits the robustness and generalization of models, especially when faced with out-of-distribution data where such spurious correlations no longer hold. To address this challenge, we propose the Causally Calibrated Robust Classifier (CCR), which aims to reduce models' reliance on spurious correlations and improve model robustness. Our approach integrates a causal feature selection method based on counterfactual reasoning, along with an unbiased inverse propensity weighting (IPW) loss function. By focusing on selecting causal features, we ensure that the model relies less on spurious features during prediction. We theoretically justify our approach and empirically show that CCR achieves state-of-the-art performance among methods without group labels, and in some cases, it can compete with the models that utilize group labels. Our code can be found at: <https://github.com/yuqing-zhou/Causal-Learning-For-Robust-Classifier>.

1 Introduction

Despite their success on standard benchmarks, neural networks often struggle to generalize to out-of-distribution (OOD) data. A primary cause of this problem is their tendency to rely on features not causally related to tasks but holding spurious correlations to labels, which can reduce model robustness when the data distribution shifts (Hovy and Søgaard, 2015; Ribeiro et al., 2016; Tatman, 2017; Buolamwini and Gebu, 2018; Hashimoto et al., 2018). For example, in natural language inference (NLI) tasks, if contradictory sentences in a dataset frequently contain negation words, a model trained on this dataset might predict contradiction simply based on the presence of negation words rather than relying on the true underlying features.

When encountering data where such spurious correlations do not hold, the model is likely to make incorrect predictions.

Prior works divide the data into different groups based on combinations of class labels and spurious features. The distribution of these groups is often unbalanced, and the correlation between labels and spurious features in the majority groups leads to spurious correlations. As a result, a model trained on such a dataset inevitably performs worse on the minority groups compared to the majority groups. Existing methods have focused on improving the worst-group accuracy. Some methods tackle this issue by utilizing group labels, which provide information about both the true labels and the spurious features (Sagawa et al., 2019; Goel et al., 2020; Zhang et al., 2020; Idrissi et al., 2022). However, obtaining annotations for group labels is costly and not always feasible, limiting the utility and applicability of these methods.

Other approaches employ a two-stage process, where group labels are first inferred, followed by an adjustment of the loss function based on this inferred information. Examples include reweighting samples (Liu et al., 2021; Qiu et al., 2023) or incorporating supervised contrastive loss (Zhang et al., 2022). However, these methods attempt to calibrate the model learning process by only adjusting the loss function and hope the model's over-reliance on spurious features can be alleviated. The effects are indirect and limited. A more direct and effective approach would involve modifying the model architecture itself to explicitly identify causal features and filter out spurious ones.

Existing works show that models trained using standard Empirical Risk Minimization (ERM) can learn high-quality representations of both spurious and causal features (Kirichenko et al., 2022; Izmailov et al., 2022). Therefore, our goal is to encourage models to rely more on causal features than spurious ones. We propose a method, Causally Cal-

ibrated Robust Classifier (CCR), to identify causal feature representations through counterfactual reasoning. Specifically, we first disentangle the representation vector from the model’s last layer and then construct counterfactual representation vectors. We adapt causal inference theory (Pearl, 2009) to calculate the probabilities of necessity and sufficiency for each feature within the representation vector and the entire representation. By retraining the last layer to maximize the probability of necessity and sufficiency of the overall feature representation, the model is guided to assign higher weights to causal features, ultimately relying more on these features. Additionally, we derive a theoretically unbiased loss function based on inverse propensity weighting (IPW), further enhancing the model’s robustness to spurious correlations.

Our contributions are summarized as follows:

1. We integrate causal inference theory into feature representation selection by calculating the probabilities of necessity and sufficiency (PNS) for each feature, and propose a method that enhances the weights of causal features in model predictions through counterfactual reasoning by maximizing the PNS for the entire representation vector.
2. We derive a general unbiased loss function with the inverse propensity weighting (IPW) and find that existing methods, which reweight samples to mitigate the effects of spurious correlations, are in fact variants of this unbiased loss function.
3. We propose a Causally Calibrated Robust Classifier (CCR) for text classification tasks that combines causal feature selection with unbiased IPW loss, without the need for group label annotations of training datasets.
4. Comparing our methods CCR with standard ERM, JTT (Liu et al., 2021), and AFR (Qiu et al., 2023) on 4 text classification tasks with varying spurious correlations, the proposed method achieves superior performance on most tasks and even outperforms methods that rely on full group labels (Sagawa et al., 2019; Kirichenko et al., 2022).

2 Method

In this paper, we focus on the impact of spurious correlations on transformer-based text classifiers and aim to enhance their robustness. While implemented on BERT (Devlin et al., 2018), our method can be generalized to other models as well.

Our method, CCR, consists of two stages. In stage 1, we apply standard ERM with a covariance

matrix regularization to disentangle the representations in the model’s last layer, ensuring each feature representation is independent. In stage 2, we retrain the last layer to encourage the model to rely on causal features by introducing a sample-wise causal constraint based on counterfactual representation vectors. Additionally, we reweight each sample’s loss in stage 2 to achieve an unbiased loss, ensuring accurate causal feature selection.

2.1 Problem Formulation

In text classification tasks, models are trained to assign a label $y \in \mathcal{Y}$ (such as sentiment, topic, and intent) to a given input text. Transformer-based models, such as BERT, encode the original text input into a fixed-dimensional embedding, regarded as hidden features of the text input. Among these features, some are causally related to the true label, referred to as causal features, while others may be coincidentally correlated with the label, referred to as spurious features. These spurious features do not have a causal connection to the label but can mislead the model. Although models trained by standard ERM can rely on spurious correlation, existing works show that they learn high-quality representations of both causal and spurious features (Kirichenko et al., 2022).

We use capital letters \mathbf{X} , \mathbf{S} , and \mathbf{Y} to denote the variable of causal features, spurious features, and labels, respectively, and \mathbf{x} , \mathbf{s} , and y represent their values. Then a text dataset of N samples can be denoted as $D = \{(\mathbf{x}_0, \mathbf{s}_0, y_0), \dots, (\mathbf{x}_{N-1}, \mathbf{s}_{N-1}, y_{N-1})\}$. For simplicity, we assume $\mathbf{s}_i \in \{\mathbf{s}^0, \mathbf{s}^1\}$ where \mathbf{s}^0 and \mathbf{s}^1 represent two distinct spurious features. Then, the dataset D can be divided into groups based on the different combinations of \mathbf{S} , and \mathbf{Y} . For example, in a binary classification task, the dataset would be divided into four groups: $D_{0,0} = \{(\mathbf{x}_i, \mathbf{s}^0, y^0)\}$, $D_{0,1} = \{(\mathbf{x}_i, \mathbf{s}^1, y^0)\}$, $D_{1,0} = \{(\mathbf{x}_i, \mathbf{s}^0, y^1)\}$, and $D_{1,1} = \{(\mathbf{x}_i, \mathbf{s}^1, y^1)\}$, where y^0 and y^1 represent two classes. To avoid ambiguity, in the following text, we will refer to the input text features only as \mathbf{x} and \mathbf{s} , and use the term "features" exclusively to denote their representation vectors from the model’s last layer.

2.2 Automatic Causal Feature Selection

In this section, we introduce how to encourage the model to focus on causal features. This is achieved through structural modifications to the model, specifically by incorporating a counterfac-

tual feature selection module, which helps the model prioritize features with true causal significance over those that are merely correlated.

Existing works show that the final layer of a classifier contains both high-quality representations of spurious and causal features (Kirichenko et al., 2022; Izmailov et al., 2022). To disentangle these representations, we apply covariance matrix regularization (Cogswell et al., 2015), aiming to maximize the independence of each feature in the final layer. This encourages the separation of spurious and causal features. Next, we conduct automatic causal feature selection on this disentangled feature layer using a counterfactual method. In the following section, we will provide a detailed explanation of this counterfactual approach.

2.2.1 Probabilistic Causality: Necessity and Sufficiency of Features

To identify causal features, we will estimate the causality of each feature. To quantify this causality, we introduce the probability of necessity and sufficiency as a metric to estimate the causal effect of features on labels, based on the causation theory of Pearl (2009). First, we introduce the definitions of the probability of necessity and sufficiency using the counterfactuals. Let \mathbf{E} denote the variable of the feature embedding (the last layer of the classifier) and \mathbf{e} denote its value. Assume the feature embedding vector contains h elements, each representing a feature. Additionally, we define a mask \mathbf{M} for the embedding vector, with the same dimensions as \mathbf{E} , where each element \mathbf{M}_i takes a value of 1 or 0, representing whether a feature is selected (1 means the corresponding feature is selected). Thus, the final picked feature representation vector for prediction can be expressed as $\tilde{\mathbf{E}} = \mathbf{E} \odot \mathbf{M}$, where \odot denotes the element-wise multiplication.

Definition 1. Probability of Necessity (PN) for a single feature \mathbf{E}_j .

$$PN := P(Y_{\mathbf{M}_j \neq 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} \neq y | \mathbf{M}_j = 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, Y = y, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})) \quad (1)$$

where $\mathbf{M}_j = 1$ indicates that \mathbf{E}_j is selected and will contribute to the final prediction, while $\mathbf{M}_j \neq 1$ indicates that it is not selected. Equation 1 represents the probability of $Y \neq y$ for the same input in the absence of the feature \mathbf{E}_j , given that feature \mathbf{E}_j exists and $Y = y$ in reality.

Definition 2. Probability of Sufficiency (PS) for

a single feature \mathbf{E}_j .

$$PS := P(Y_{\mathbf{M}_j = 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} = y | \mathbf{M}_j \neq 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, Y \neq y, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})) \quad (2)$$

Equation 2 measures the contribution of the feature \mathbf{E}_j to the prediction of y for the input (\mathbf{x}, \mathbf{s}) .

Definition 3. Probability of Necessity and Sufficiency (PNS) for a single feature \mathbf{E}_j .

$$PNS := P(Y_{\mathbf{M}_j \neq 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} \neq y, Y_{\mathbf{M}_j = 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} = y) \quad (3)$$

Equation 3 evaluates both the necessity and sufficiency of the feature \mathbf{E}_j for the label $Y = y$, given the input (\mathbf{x}, \mathbf{s}) . If the feature \mathbf{E}_j is a causal feature in this task, its PNS should have a high value, and We aim to encourage the classifier’s final layer to assign a high weight to this feature.

Directly calculating Equation 3 can be difficult. However, by applying and extending Theorem 9.2.10 of Pearl (2009) to a more general case, we can get a lower bound of PNS .

Theorem 1. The lower bound of PNS for causal features is given as follows:

$$PNS \geq \max[0, P(Y_{\mathbf{M}_j = 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} = y) - P(Y_{\mathbf{M}_j \neq 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} \neq y)] \quad (4)$$

We can achieve a higher PNS for the feature \mathbf{E}_j by maximizing this lower bound PNS . In practice, $P(Y_{\mathbf{M}_j = 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} = y)$ is estimated from the model’s output using the original feature vector, while $P(Y_{\mathbf{M}_j \neq 1, \mathbf{M}_{i \neq j} = \mathbf{m}_{i \neq j}, (\mathbf{X}, \mathbf{S}) = (\mathbf{x}, \mathbf{s})} \neq y)$ is estimated using the model’s output based on the counterfactual feature vector.

2.2.2 Causal Constraints with PNS

We introduce a module to generate the counterfactual feature vectors. Let \mathbf{M} denote a mask with all values set to 1, except for one value set to 0. By $\tilde{\mathbf{E}} = \mathbf{E} \odot \mathbf{M}$ with $\mathbf{M}_j = 0$ for each $j \in \{1, 2, \dots, h\}$, we create a counterfactual feature vector, a copy of \mathbf{E} with the j -th element discarded. This process produces h counterfactual feature vectors.

Definition 4. Let PNS_j denote the PNS of the j -th counterfactual feature vector. We define the total PNS of the feature embedding vector \mathbf{E} as follows:

$$PNS = \prod_{j=1}^h PNS_j \quad (5)$$

We take the mean of the negative log of Equation 5 as the causal constraint and incorporate it into the loss function, where PNS_j is replaced by its lower bound. In this way, the final layer can be trained to rely more on causal features, i.e., the feature with higher PNS values.

2.3 Unbiased Loss through IPW

Accurately selecting causal features requires an unbiased loss. To achieve this goal, we derive a reweighting method using inverse propensity weighting. Previous methods (Liu et al., 2021; Qiu et al., 2023) that use heuristics for weight computation can be regarded as special implementations of our framework.

Assume we have a text classification dataset with C classes $\{y^1, y^2, \dots, y^C\}$ and K spurious features $\{s^1, \dots, s^K\}$. Then the dataset can be divided into $C * K$ groups, denoted as $D_{j,k} = \{(\mathbf{x}_i, \mathbf{s}_i, y_i) | y_i = y^j, \mathbf{s}_i = \mathbf{s}^k\}$. In an ideal scenario, the size of group $D_{j,k}$ should be same for all $k = 1, \dots, K$, within the class y^j , i.e., $|D_{j,k}| = |D_{j,t}|, \forall k, t \in \{1, \dots, K\}$. In this case, the spurious features \mathbf{s} are not correlated with the labels y . An ideal loss function is as follows:

$$\begin{aligned} L_{ideal} &= \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K L_{D_{j,k}} \\ &= \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^{|D|} L(y_i, f(\mathbf{x}_i, \mathbf{s}_i)) \mathbb{I}_{\{y_i=y^j, \mathbf{s}_i=\mathbf{s}^k\}}, \end{aligned} \quad (6)$$

where $L_{D_{j,k}}$ is the sum of the losses over all samples in the subset $D_{j,k}$, $|D|$ is the size of the whole dataset D , $L(y_i, f(\mathbf{x}_i, \mathbf{s}_i))$ is the loss on the i -th sample, which will be simplified as L_i in the following text. Additionally, \mathbb{I} is an indicator function. $\mathbb{I}_{\{A\}} = 1$ when A is true. Ideally, all subgroups $D_{j,*}$ within the class y^j contribute equally to the total loss. A classifier f learned on this dataset would be robust and make predictions of y solely based on \mathbf{x} . However, in reality, the group distribution of a dataset D' is unbalanced, i.e., $\exists l, |D'_{j,l}| \gg |D'_{j,k}|, \forall k \neq l$. In this case, we call $|D'_{j,k}|$ as the majority group when $k = l$, and as the minority group otherwise. A classifier trained on such a dataset would falsely rely on the spurious correlation between s^l and label y^j leading to biased predictions.

We can assume that not all samples in the ideal datasets can be observed and the ratio of samples observed for different groups varies. We use the \mathcal{O} to denote the observation of the whole ideal dataset

D and o_i to denote the observation result for i -th sample in the ideal dataset, where $o_i = 1$ means the sample is in the reality dataset D' , and 0 means the sample is not in D' . The probability of $p(o_i = 1)$ depends on its group membership, thus leading to an unbalanced dataset D' . We use $p_{j,k}(o = 1)$ to represent the probability of a sample in $D_{j,k}$ being observed. The majority groups in D' have a much larger $p_{j,k}$ than the minority group. With these notations, the actual loss is given by

$$L_{real} = \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^{|D|} L_i \mathbb{I}_{\{y_i=y^j, \mathbf{s}_i=\mathbf{s}^k\}} \mathbb{I}_{\{o_i=1\}}, \quad (7)$$

In this case,

$$\begin{aligned} \mathbb{E}[L_{real}] &= \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^{|D|} L_i \mathbb{I}_{\{y_i=y^j, \mathbf{s}_i=\mathbf{s}^k\}} p_{j,k}(o_i = 1) \\ &\neq L_{ideal} \end{aligned} \quad (8)$$

i.e., L_{real} is biased.

To achieve an unbiased loss on an unbalanced dataset, the key is to balance the weights of the majority and minority groups in the final loss function. This can be achieved by normalizing the losses for each subset and averaging their contributions across the groups. The loss function is defined as

$$\begin{aligned} L &= \frac{1}{|D'|} \sum_{j=1}^C \frac{|D'_j|}{K} \sum_{k=1}^K \frac{L_{D'_{j,k}}}{|D'_{j,k}|} \\ &= \frac{1}{|D'|} \sum_{j=1}^C \frac{1}{K} \sum_{k=1}^K \frac{L_{D'_{j,k}}}{\frac{|D'_{j,k}|}{|D'_j|}} \end{aligned} \quad (9)$$

where $D'_j = \cup_{k=1}^K D'_{j,k}$ and $D' = \cup_{j=1}^C D'_j$. Here, the normalization ensures that within each group D'_j , the contributions from subsets $D'_{j,k}$ for $k = 1, \dots, K$ are equal, leading to a more equitable contribution to the overall loss. When the dataset is large enough, we can get

$$p(\mathbf{S} = \mathbf{s}^k | Y = y^j) = \frac{|D'_{j,k}|}{|D'_j|}, \quad (10)$$

which is a propensity for a sample with label j to contain the spurious feature \mathbf{s}^k . This suggests that the Equation 9 can be expressed in the following

form:

$$L = \frac{1}{K * |D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^{|D|} \frac{L_i \mathbb{I}_{\{y_i=y^j, s_i=s^k\}} \mathbb{I}_{\{o_i=1\}}}{p(s^k|Y=y^j)} \quad (11)$$

As we have no knowledge of group information, so the true $p(s^k|Y=y^j)$ is unknown. However, if we can find a way to get its estimation \hat{p} , then we can get the following loss function:

$$L_{IPW} = \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^{|D|} \frac{L_i \mathbb{I}_{\{y_i=y^j, s_i=s^k\}} \mathbb{I}_{\{o_i=1\}}}{K * \hat{p}(s^k|Y=y^j)} \quad (12)$$

The expectation of L_{IPW} is given as follows:

$$\begin{aligned} \mathbb{E}[L_{IPW}] &= \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^N \frac{L_i \mathbb{I}_{\{y_i=y^j, s_i=s^k\}} \mathbb{E}[\mathbb{I}_{\{o_i=1\}}]}{K * \hat{p}(s^k|Y=y^j)} \\ &= \frac{1}{|D|} \sum_{j=1}^C \sum_{k=1}^K \sum_{i=1}^N \frac{L_i \mathbb{I}_{\{y_i=y^j, s_i=s^k\}} p(o_i=1)}{K * \hat{p}(s^k|Y=y^j)} \end{aligned} \quad (13)$$

When $\hat{p}(s^k|y^j) = \frac{p_{j,k}(o=1)}{K}$, $\mathbb{E}[L_{IPW}] = L_{ideal}$, which is unbiased. It demonstrates that we can get the unbiased loss function through inverse propensity weighting (IPW). The state-of-the-art robust methods JTT (Liu et al., 2021) and AFR (Qiu et al., 2023) can be considered as special implementations of L_{IPW} , which rely on different heuristics for estimating $\hat{p}(s^k|y^j)$.

Our estimation of $\hat{p}(s^k|y^j)$ is based on the performance of the first-stage model. For each class, we treat the samples correctly classified by the first-stage model as the major groups (e.g., (s^1, y^1) and (s^2, y^2)), while the misclassified samples are taken as minor groups (e.g., (s^2, y^1) and (s^1, y^2)). This process divides the dataset into $2C$ groups, and we use each group's proportion to the whole dataset as the estimated $\hat{p}(s^k|y^j)$.

2.4 Training Framework

The overall training process is as follows.

Stage 1. A linear layer is introduced to the pre-trained model, with input dimensions h (a hyperparameter, representing the number of features) and output dimensions corresponding to the number of classes. The model is then finetuned using ERM, combining cross-entropy loss with covariance matrix regularization on the feature vectors. This process encourages the model to learn an effective feature extractor, which will be frozen in stage 2,

and ensures that the final feature vector consists of h disentangled and independent features.

Stage 2. We freeze all layers of the model except the last one, taking the frozen layers as a fixed feature extractor, while retraining only the last layer in this stage. Before retraining, we evaluate the model's performance on the training set, identifying the correctly and incorrectly classified samples within each class. Based on this, we estimate the data groups and calculate their ratios as the estimated propensities, which are used to assign weights to each group. Then, we finetune the last layer according to the following loss function:

$$L = \frac{1}{|D'|} \sum_{i=1}^{|D'|} \frac{1}{\hat{p}} (CE(y_i, f(\mathbf{x}_i, \mathbf{s}_i)) - \lambda * \frac{1}{h} \sum_{j=1}^h \log P N S_j) \quad (14)$$

where the λ is the coefficient of causality constraints.

3 Experiments

In this section, we investigate the effect of CCR through three research questions. **RQ1:** What is the performance of CCR compared with other state-of-the-art methods on the text classification benchmarks? **RQ2:** How is the contribution of each component of CCR to the model's robustness? **RQ3:** What is the effect of causal feature selection under different hyperparameters?

3.1 Experiments Setup

3.1.1 Datasets

We evaluate our method, CCR, on two real-world datasets, CivilComments (Borkan et al., 2019) and MultiNLI (Williams et al., 2017), following previous works (Liu et al., 2021; Qiu et al., 2023). Additionally, we use two semi-synthetic datasets, which are real-world datasets manually introduced with spurious correlations: the Yelp dataset (Zhang et al., 2015) with author style shortcuts (Zhou et al., 2024) and the Beer dataset (Bao et al., 2018) with concept occurrence shortcuts (Zhou et al., 2024). These four datasets cover four types of spurious correlations (Zhou et al., 2024).

CivilComments (Borkan et al., 2019) is a text classification dataset, where the target is to classify a comment as "toxic" or "not toxic". The WILDS benchmark (Koh et al., 2021) provides a version of CivilComments that contains spurious correlations

between the label "toxic" and some sensitive characteristics such as gender and race, which are used in our experiments. These words describing sensitive characteristics such as gender (male, female) construct a category-word shortcut.

MultiNLI (Williams et al., 2018) is designed for the task of natural language inference, where each sample contains sentence pairs labeled with one of three categories that describe the relationship between the sentences: entailment, contradiction, or neutral. In this dataset, spurious correlations exist between the label "contradiction" and negation words, which lead to a synonym shortcut.

Yelp review dataset (Zhang et al., 2015) consists of review-rating pairs from Yelp, with ratings ranging from 1 to 5. Zhou et al. (2024) constructed a **Yelp-Author-Style** dataset based on the Yelp reviews, where each review is rewritten in the distinct styles of different authors (Shakespeare and Hemingway) according to the ratings. We adapt this multiclass classification task into a binary classification task by selecting reviews with ratings of 2 and 4, where lower ratings are correlated with Hemingway’s style and higher ratings with Shakespeare’s style.

Beer-Concept-Occurrence dataset (Zhou et al., 2024) consists of review-rating pairs about beer, where each sample contains two parts—one for the palate review (the target) and another representing a spurious feature (comments on either "aroma" or "appearance"). Comments on the appearance of the beer are correlated with lower palate ratings, while comments on the aroma are correlated with higher ratings.

3.1.2 Model and Baselines

Model. Following the settings in AFR (Qiu et al., 2023), we use the pretrained model "bert-base-uncased" (Devlin et al., 2018) for the text classification tasks. The hyperparameter settings are provided in Table 5 of Appendix A.1.

Baselines. In our experiments, we primarily focus on the following three baselines, which do not require group labels: standard ERM (Liu et al., 2021), JTT (Liu et al., 2021), and AFR (Qiu et al., 2023). JTT trains two models: the first model is trained via standard ERM, while the second model is trained by upweighting the samples misclassified by the first model. AFR also trains twice but on the same model: it retrains the last layer of a standard ERM-trained model by upweighting the samples

on which the model previously performed poorly.

Additionally, we compare the performance against two baselines that rely on group labels: Group-DRO (Sagawa et al., 2019), which minimizes the loss of the worst group while applying an additional penalty based on group size to avoid overfitting, and DFR (Kirichenko et al., 2022), which is similar to AFR but retrains the last layer on a small dataset with group labels.

3.2 Strong Group Robustness of CCR (RQ1)

We evaluate CCR on four benchmarks, measuring both overall accuracy and worst-group accuracy (WGA), with the results presented in Table 1. Compared to other state-of-the-art robust methods that do not require group labels, CCR achieves the highest WGA across all four datasets, demonstrating strong robustness to spurious correlations. Specifically, CCR achieves the best performance in terms of both mean accuracy over the entire test dataset and WGA on Yelp-Author-Style, improving WGA by 3% and mean accuracy by 0.5%, while on Beer-Concept-Occurrence, its mean accuracy is only 0.2% lower than the highest recorded performance. This performance shows that CCR can improve the WGA without sacrificing the overall performance, suggesting its effect on helping the model mitigate the reliance on spurious correlations and utilize causal features.

Furthermore, we compare CCR with two state-of-the-art (SOTA) methods, Group-DRO and DFR, which utilize group labels. As shown in Table 2, CCR achieves comparable performance without the need for group labels in the training data, and even slightly outperforms Group-DRO and DFR on CivilComments, achieving SOTA WGA and a 1% improvement in mean accuracy. It also surpasses DFR in terms of WGA on MultiNLI.

3.3 Ablation Study (RQ2)

After showing the effect of CCR on improving model robustness, we aim to understand how each component of CCR contributes to the overall performance. Using the Yelp-Author-Style dataset, we investigate the impact of the following operations: feature representation disentanglement, causal feature constraints (CFC) for causal feature selection (CFS), and loss debiasing with inverse propensity weighting (IPW). The results are shown in Table 3 and we have the following observations from it:

First, representation disentanglement, causal feature selection (CFS), and loss debiasing with IPW

Datasets	CivilComments		MultiNLI		Yelp-Author-Style		Beer-Concept-Occur	
Methods	Mean	WGA	Mean	WGA	Mean	WGA	Mean	WGA
ERM	0.9260	0.5740	0.8240	0.6790	0.9200	0.8500	0.9538	0.9000
JTT	0.9110	0.6930	0.7860	0.7260	0.9210	0.8583	0.9554	0.9000
AFR	0.8980	0.6870	0.8140	0.7340	0.9205	0.8542	0.9506	0.9200
CCR (ours)	0.9000	0.7067	0.8072	0.7517	0.9260	0.8874	0.9530	0.9304

Table 1: Comparisons of different methods that do not need group labels. We report both the overall mean accuracy and the worst-group accuracy (WGA). The results for ERM, JTT, and AFR on the CivilComments and MultiNLI datasets are sourced from Qiu et al. (2023). We implement ERM, JTT, and AFR on the Yelp-Author-Style and Beer-Concept-Occurrence datasets.

Datasets	CivilComments		MultiNLI	
Methods	Mean	WGA	Mean	WGA
Group-DRO	0.889	0.699	0.814	0.777
DFR	0.872	0.701	0.821	0.747
CCR(ours)	0.900	0.707	0.807	0.752

Table 2: Comparing CCR with the methods need group labels. The results for Group-DRO and DFR are sourced from Qiu et al. (2023).

all contribute to improving the model’s robustness compared to standard ERM, although the improvements are subtle. CFS alone increases WGA by only 0.8%, but when combined with feature disentanglement, it achieves a 2% improvement in WGA. This suggests that disentangling spurious and causal features as much as possible is essential. If these features are entangled, it may diminish the evaluation of each feature’s necessity and sufficiency to labels, limiting the effectiveness of CFS.

Second, CFS and IPW all get more than 1% improvement in WGA by combining with disentanglement. Meanwhile, disentanglement alone can improve the WGA of ERM by 1.6%, outperforming CFS and IPW. These results suggest that disentanglement may transform the input features into a more effective representation space, which CFS and IPW can leverage to enhance performance.

Third, while both CFS and IPW individually enhance the model’s robustness, their combination is the most effective. After disentanglement, the joint effect of CFS and IPW improves both the mean accuracy and WGA of the model, even increasing WGA by approximately 4%.

Additionally, we explore three different approaches for propensity estimation in IPW: one proposed by us and two others derived from AFR and JTT’s methods for calculating weights. Our method involves counting the samples correctly and incorrectly classified by the first-stage model within each class, then using the inverse of their

ratio as the weights, normalized across the entire training dataset. All three methods show similar performance.

Methods	Mean	WGA
ERM	0.9200	0.8500
disentangle	0.9235	0.8658
CFS	0.9210	0.8583
IPW	0.9210	0.8571
disentangle + CFS	0.9235	0.8701
disentangle + IPW	0.9245	0.8701
disentangle + CFS + IPW (ours)	0.9260	0.8874
disentangle + CFS + IPW (AFR)	0.9265	0.8916
disentangle + CFS + IPW (JTT)	0.9250	0.8831

Table 3: Comparisons of the effect of each component in CCR on the Yelp-Author-Style dataset. For simplicity, we omit ERM from each row. CFS refers to the module for causal feature selection. IPW (ours) represents our proposed method for obtaining weights, while IPW (AFR) and IPW (JTT) use the weight calculation algorithms from AFR (Qiu et al., 2023) and JTT (Liu et al., 2021), respectively.

3.4 Hyperparameter Study (RQ3)

In this section, we investigate the effect of the causality constraint coefficient λ in Equation 14 on model performance. Specifically, we aim to investigate how much emphasis should be placed on the PNS of the feature representation in the final loss function for optimal performance. The experiment results on the Yelp-Author-Style dataset are shown in Figure 1. From Figure 1, we can see that as λ increases from 0.001 to 0.5, WGA increases gradually. However, adding more PNS constraints does not necessarily lead to better performance. As λ increases from 0.5 to 3, WGA fluctuates slightly, reaching its peak at $\lambda = 3$, but gradually decreases as λ increases further.




Figure 1: Study on the effect of different coefficients for causality constraints in CFS. (* = 0.001)

3.5 Analyzing Model Behavior Through Explainability

In Section 3.2, we demonstrated that CCR improves model robustness against spurious correlations. In this section, we further explore model behavior through explainability techniques. We use the SHAP (Lundberg and Lee, 2017) analysis tool to examine how classifiers rely on specific tokens for their predictions.

We use the Beer-Concept-Occurrence dataset as an example, where descriptions of beer appearance are correlated with low palate ratings, while descriptions of beer aroma are associated with high palate ratings. We sampled 200 test instances and computed the average SHAP absolute values for the spurious features (i.e., comments on beer appearance and aroma). Table 4 shows the SHAP analysis results for each method. The SHAP absolute value quantifies a token’s influence on the model’s prediction for the corresponding label. For the spurious features, SHAP values closer to zero indicate less reliance on these features by the model. Although the task is relatively straightforward and all methods yield low SHAP values, we observe that CCR’s SHAP values for spurious features are significantly lower than those of other methods—approximately an order of magnitude lower—demonstrating strong robustness and effectively reducing reliance on spurious features in predictions. Besides, we provide an example of SHAP analysis in Appendix A.2.

Methods	Label	
	0	1
ERM	1.63E-3	1.54E-3
JTT	3.10E-3	2.54E-3
AFR	1.47E-3	1.47E-3
CCR (Ours)	1.85E-4	1.85E-4

Table 4: SHAP values different methods

4 Related Work

Spurious correlations challenge the robustness of machine learning models trained via ERM, especially when there are shifts between the training and test data distribution (Hovy and Søgaard, 2015; Ribeiro et al., 2016; Tatman, 2017; Buolamwini and Gebru, 2018; Hashimoto et al., 2018).

Sagawa et al. (2019) addresses this problem by training a classifier via a regularized group distributionally robust optimization (Group-DRO) method, which minimizes the loss of the worst-case group. This approach is to train the entire model. In contrast, some methods propose that simply retraining the last layer of an ERM-trained model can effectively improve robustness. For example, DFR (Kirichenko et al., 2022) retrain the last layer of the model trained via ERM on a small dataset where the spurious correlations break. However, both Sagawa et al. (2019) and Kirichenko et al. (2022) require prior knowledge of spurious correlations to define groups in the training data, making it costly and limiting its ability to generalize to unseen spurious attributes. Some methods upweight samples during retraining where the initially trained model performs poorly (Liu et al., 2021; Qiu et al., 2023). Another approach uses contrastive learning to improve robustness by aligning same-class representations and separating misclassified negatives, allowing the model to learn similar representations of same-class samples and ignore spurious attributes (Zhang et al., 2022). These methods do not require prior knowledge of group labels. All these approaches focus on reducing spurious correlations by modifying the loss function, while our method tackles the underlying issue by directly identifying causal features using a counterfactual feature generation module.

Our method for identifying causal features is based on the theorem of probability of necessity and sufficiency from causality theory (Pearl, 2009). A related work also applies this theory (Zhang et al., 2023), but their goal is to identify minimal tokens

in input text that lead to correct predictions. In contrast, our approach operates on feature representations within the model to identify causal features, aiming to build a robust classifier that performs well across varying data distributions.

5 Conclusion

In this paper, we address the problem of models relying on spurious correlations in their predictions, particularly in text classification tasks. To mitigate this issue, we propose the Causally Calibrated Robust Classifier (CCR). Specifically, we introduce a causal feature selection model based on counterfactual reasoning, combined with an unbiased inverse propensity weighting (IPW) loss function. We provide theoretical justification for our method and empirically demonstrate that it achieves state-of-the-art performance among approaches that do not require group labels, and in some cases, it can even compete with models that utilize group labels.

6 Limitations

Apart from the contributions mentioned in the paper, there are still some limitations of this work.

First, the loss debiasing relies on the quality of the estimation of the propensity \hat{p} . If \hat{p} is not accurately estimated, the loss function may still exhibit bias. Improving the estimation method could lead to enhanced performance.

Second, our experiments showed that the optimal feature embedding size varies across datasets, with different sizes yielding the best results. Identifying the appropriate embedding size is important for effectively separating irrelevant features from causal features.

Additionally, we did not explore alternative disentanglement methods, which may improve feature representation. Better disentanglement could lead to more effective causal feature selection and ultimately enhance model performance.

Acknowledgements

This work was partially supported by resources provided by the Office of Research Computing at George Mason University (URL: <https://orc.gmu.edu>) and funded in part by grants from the Commonwealth Cyber Initiative.

Ethics Statement

Our research complies with ethical standards, utilizing publicly available datasets. All contents in

the datasets do NOT represent the authors' views.

References

- Yujia Bao, Shiyu Chang, Mo Yu, and Regina Barzilay. 2018. Deriving machine attention from human rationales. *arXiv preprint arXiv:1808.09367*.
- Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion proceedings of the 2019 world wide web conference*, pages 491–500.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR.
- Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, and Dhruv Batra. 2015. Reducing overfitting in deep networks by decorrelating representations. *arXiv preprint arXiv:1511.06068*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Karan Goel, Albert Gu, Yixuan Li, and Christopher Ré. 2020. Model patching: Closing the subgroup performance gap with data augmentation. *arXiv preprint arXiv:2008.06775*.
- Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. 2018. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, pages 1929–1938. PMLR.
- Dirk Hovy and Anders Søgaard. 2015. Tagging performance correlates with author age. In *Proceedings of the 53rd annual meeting of the Association for Computational Linguistics and the 7th international joint conference on natural language processing (volume 2: Short papers)*, pages 483–488.
- Badr Youbi Idrissi, Martin Arjovsky, Mohammad Pezeshki, and David Lopez-Paz. 2022. Simple data balancing achieves competitive worst-group-accuracy. In *Conference on Causal Learning and Reasoning*, pages 336–351. PMLR.
- Pavel Izmailov, Polina Kirichenko, Nate Gruver, and Andrew G Wilson. 2022. On feature learning in the presence of spurious correlations. *Advances in Neural Information Processing Systems*, 35:38516–38532.
- Polina Kirichenko, Pavel Izmailov, and Andrew Gordon Wilson. 2022. Last layer re-training is sufficient for robustness to spurious correlations. *arXiv preprint arXiv:2204.02937*.

- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. 2021. Wilds: A benchmark of in-the-wild distribution shifts. In *International conference on machine learning*, pages 5637–5664. PMLR.
- Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. 2021. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, pages 6781–6792. PMLR.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- J Pearl. 2009. *Causality*. Cambridge university press.
- Shikai Qiu, Andres Potapczynski, Pavel Izmailov, and Andrew Gordon Wilson. 2023. Simple and fast group robustness by automatic feature reweighting. In *International Conference on Machine Learning*, pages 28448–28467. PMLR.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. 2019. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*.
- Rachael Tatman. 2017. Gender and dialect bias in youtube’s automatic captions. In *Proceedings of the first ACL workshop on ethics in natural language processing*, pages 53–59.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Jingzhao Zhang, Aditya Menon, Andreas Veit, Srinadh Bhojanapalli, Sanjiv Kumar, and Suvrit Sra. 2020. Coping with label shift via distributionally robust optimisation. *arXiv preprint arXiv:2010.12230*.
- Michael Zhang, Nimit S Sohoni, Hongyang R Zhang, Chelsea Finn, and Christopher Ré. 2022. Correct-n-contrast: A contrastive approach for improving robustness to spurious correlations. *arXiv preprint arXiv:2203.01517*.
- Wenbo Zhang, Tong Wu, Yunlong Wang, Yong Cai, and Hengrui Cai. 2023. Towards trustworthy explanation: On causal rationalization. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 41715–41736. PMLR.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Yuqing Zhou, Ruixiang Tang, Ziyu Yao, and Ziwei Zhu. 2024. Navigating the shortcut maze: A comprehensive analysis of shortcut learning in text classification by language models. *arXiv preprint arXiv:2409.17455*.

A Appendix

A.1 Experiment Setup

The reported performance was achieved using the settings outlined in Table 5. All experiments were conducted once, with a fixed random seed to ensure reproducibility. The pre-trained model, 'bert-base-uncased,' consists of approximately 110M parameters. The statistics for each dataset are provided in Table 6. Each experiment can be completed within 24 hours on a single A100 GPU.

A.2 Explainability Analysis

Figure 2 presents a review of beer palate labeled as 0 (corresponding to a 0.6 palate rating), along with the behaviors of models trained using standard ERM and CCR in predicting this label. In the ERM model, there is a strong reliance on the words that describe "appearance" (e.g., "appearance" and "golden"). These words show high SHAP values (highlighted in red), indicating that the model attributes significant importance to these terms when predicting a label of 0. In contrast, the CCR model exhibits much lower SHAP values for the same appearance-related terms, indicating that CCR has successfully reduced the model’s reliance on the spurious correlation between appearance-related features and label 0. The SHAP values in the CCR model are closer to neutral (around zero), showing that the model is less biased toward irrelevant features and more robust to spurious correlations compared to the standard ERM model.

Datasets	CivilComments	MultiNLI	Yelp-Author-Style	Beer-Concept-Occurrence
Coefficient for disentanglement	0.5	0.5	0.5	0.5
Coefficient for causal constraints	0.1	2	3	2
Learning rate	2.00E-02	2.00E-01	3.00E-03	3.00E-02
Weight decay	0	0	1.00E-04	1.00E-04
Batch size	32	32	32	32
Size of Feature Embedding	128	128	128	150
Epoch	15	10	30	30
Seed	42	42	42	42

Table 5: The hyperparameter settings for each dataset.

	Group IDs	Class Label	Group (spurious feature, class)	# Training Data
CivilComments	0	0	(no identities, non-toxic)	148486
	1	0	(has identities, non-toxic)	90337
	2	1	(no identities, toxic)	12731
	3	1	(hasidentities, toxic)	17784
MultiNLI	0	0	(no negations, contradiction)	57498
	1	0	(has negations, contradiction)	11158
	2	1	(no negations, entailment)	67376
	3	1	(has negations, entailment)	1521
	4	2	(no negations, neutral)	66630
Yelp-Author-Style	0	0	(Hemingway’s Style, Rating 2)	684
	1	0	(Shakespeare’s Style, Rating 2)	214
	2	1	(Hemingway’s Style, Rating 4)	252
	3	1	(Shakespeare’s Style, Rating 4)	660
Beer-Concept-Style	0	0	(Comments on Appearance , Rating 0.6)	477
	1	0	(Comments on Aroma, Rating 0.6)	261
	2	1	(Comments on Appearance, Rating 1.0)	65
	3	1	(Comments on Aroma, Rating 1.0)	432

Table 6: The statistics of each dataset. The statistics for CivilComments and MultiNLI are sourced from (Qiu et al., 2023).

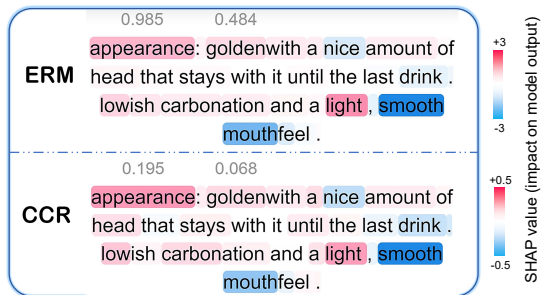


Figure 2: SHAP Analysis for BERT with ERM and CCR