PairCFR: Enhancing Model Training on Paired Counterfactually Augmented Data through Contrastive Learning

Anonymous ACL submission

Abstract

Counterfactually Augmented Data (CAD) involves creating new data samples by applying minimal yet sufficient modifications to flip the label of existing data samples to the other classes. Training with CAD enhances model robustness against spurious features that happen to correlate with labels by spreading the casual relationships across different classes. Yet, recent research reveals that CAD may lead models to overly focus on modified features while ignoring other important contextual information, inadvertently introducing biases that may impair performance on out-of-distribution (OOD) datasets. To mitigate this issue, we employ contrastive learning to promote global feature alignment in addition to counterfactual clues. We theoretically prove that contrastive loss can encourage models to leverage a broader range of features beyond those modified ones. Comprehensive experiments on two human-edited CAD datasets demonstrate that our propose method outperformed the state-of-the-art on OOD datasets.

1 Introduction

In the field of Natural Language Processing (NLP), a significant body of research (McCoy et al., 2019; Wang and Culotta, 2020; Poliak et al., 2018; Gururangan et al., 2018) has raised the concern that deep learning models can overfit spurious correlations, such as dataset-specific artifacts and biases, rather than focusing on the more complex, generalizable task-related features. For example, studies in Gururangan et al. (2018) and Poliak et al. (2018) demonstrate that classifiers trained exclusively on hypotheses can still achieve decent results on some Natural Language Inference (NLI) datasets, which ideally requires comparing hypotheses with premises to determine the labels. The existence of biases in training datasets can severely degrade the performance of deep learning models on out-of-distribution (OOD) datasets.

Counterfactually Augmented Data (CAD) has emerged as a promising approach to mitigate this issue by making minimal modifications to existing data samples such that the corresponding labels are switched to the other classes (Kaushik et al., 2020; Wen et al., 2022; Pryzant et al., 2023). This technique aims to establish direct causal relationships for models to learn more effectively and enhance generalization across different datasets (Teney et al., 2020; Kaushik et al., 2021).

However, the efficacy of CAD is not always guaranteed, particularly when both contexts and the modified information should be considered together to make predictions (Joshi and He, 2022; Huang et al., 2020). For instance, in sentiment analysis, simply replacing positive adjectives such as "good" or "excellent" with negative counterparts like "terrible" or "bad" will potentially risk models to overemphasize these changes and even assign zero weights to the broader unmodified context (Joshi and He, 2022). Consequently, the trained models may fail to understand more nuanced expressions like irony or negation, exemplified by sentences such as "Is it a good movie ????" or "This movie is not that good."

To solve the above risks of CAD training, an intuitive solution is to increase the diversity of counterfactual samples (Joshi and He, 2022; Sen et al., 2023), thereby disentangling the suspicious correlations between edited features and labels. Nonetheless, this kind of method often relies on human knowledge to steer the diversification, bearing high expenditure and time consumption (Huang et al., 2020). Others attempt to design additional constraints to align the model gradient with the straight line from CFE and the original input or minimize invariant risk, failing to exploit the complex effects of augmented feature components (Teney et al., 2020; Fan et al., 2024).

In this paper, we introduce a simple yet effective learning strategy to mitigate the overfitting problem

associated with CAD. Inspired by the recent success of contrastive learning (CL) in feature alignment (Gao et al., 2021; Wang et al., 2022b; Liu et al., 2023a,b) and its strengths in capturing global relationships (Park et al., 2023), we propose to employ a contrastive learning objective to complement the standard cross-entropy (CE) loss. While CL compels the model to extract complementary effects among counterfactually augmented data to alleviate the feature degeneration, CE ensures the induced feature representations are effectively used for classification. Our mathematical proof further corroborates the advantage of combining the two losses in training models on CAD, resulting in enhanced generation capability.

In summary, our contributions are as follows:

- We introduce a contrastive learning-based framework, namely Pairwisely Counterfactual Learning with Contrastive Loss Regularization (PairCFR), for training models on CAD, which prevents overfitting to non-robust minor edits, thereby ensuring a better generalization performance.
- We offer theoretical proof for understanding the synergistic benefits of combining CE and CL loss, revealing their complementary effects in preventing models from relying solely on counterfactual edits for classification.
- We conducted comprehensive experiments to demonstrate that the models trained under our learning framework achieve superior OOD generalization performance on two human-edited CAD datasets.

2 Related work

Counterfactually Augmented Data. Counterfactual examples (CFEs) suggest the minimal modifications required in an input instance to elicit a different outcome (Wachter et al., 2017; Barocas et al., 2020). This property has inspired researchers (Kaushik et al., 2020; Wu et al., 2021) to adopt CFEs as a meaningful data augmentation in NLP, aiming to mitigate spurious correlations and improve causal learning. Early efforts (Kaushik et al., 2020) involved creating CAD datasets with manual sentence edits for label reversal. To ease the high cost of manual annotation, subsequent works adopt large language models (LLMs) for cost-effective generation of CAD (Wu et al., 2021; Madaan et al., 2021; Wen et al., 2022; Pryzant et al., 2023; Chen et al., 2023). However, findings from various investigations have indicated that training on CAD does not always ensure improved generalization on OOD tasks (Huang et al., 2020; Joshi and He, 2022; Fan et al., 2024). Consequently, our emphasis in this work is not on generating CAD, but rather on the exploration of methodologies to effectively utilize the inherent prior knowledge within CAD.

Contrastive Learning. Contrastive learning is initially proposed to learn a better embedding space by clustering similar samples closely while pushing dissimilar ones far apart (Schroff et al., 2015; Sohn, 2016; Oord et al., 2018; Wang and Isola, 2020a). For example, the triplet loss (Schroff et al., 2015) minimizes the distance between an anchor point and its positive sample while maximizing the distance from a negative sample. The N-pair loss (Sohn, 2016) maximizes the distance between an anchor point with multiple negative points. Meanwhile, InfoNCE (Oord et al., 2018) separates positive samples from multiple noise samples with cross-entropy loss. Enhanced by other efficient techniques, e.g., data augmentation (Chen et al., 2020), hard negative sampling (Schroff et al., 2015), and memory bank (Wu et al., 2018), CL has propelled significant advancements in various domains, under both supervised and unsupervised settings. In this section, we explore the untapped potential of CL to enhance the OOD generalization of models trained on CAD.

Training with CAD. The task of effectively training a robust model with CAD has received limited attention. One simple approach is to directly use cross-entropy loss (Kaushik et al., 2020). Teney et al. (2020) introduces gradient supervision over pairs of original data and their counterfactual examples, ensuring the model gradient aligns with the straight line between the original and counterfactual points. Fan et al. (2024) considers original and counterfactual distribution as two different environments and proposes a dataset-level constraint using invariant risk minimization. Following these works, we introduce a learning framework employing contrastive loss as a regularizer to notably enhance the generalization of fine-tuned models.

3 Methodology

3.1 Motivation

Recent studies have empirically shown that while perturbed features in CAD are robust, they may inhibit the model's ability to learn other robust fea-

tures that remain unperturbed (Kaushik et al., 2020; Joshi and He, 2022). In this section, we mathematically show that the standard cross-entropy loss, which is commonly used for training models on CAD, can exacerbate this tendency.

Given an instance $\mathbf{x} \in \mathbb{R}^{m\times 1}$, we train a single-layer non-linear function $f_W(x) = \sigma(W^T\mathbf{x})$, where $W \in \mathbb{R}^{m\times 1}$ and σ is the sigmoid function, to predict the label $y \in \{0, 1\}$. We expand $\mathbf{x}$, whose label $y = 1$, as $\mathbf{x} = [x_r, x_c]^T$, where x_r denotes the features to be revised (perturbed) and x_c denotes the constant (unperturbed) features. The counterfactual example of $\mathbf{x}$ can be written as $\mathbf{c} = [c_r, x_c]^T$, with label $y = 0$. We expand the weights $W = [w_r, w_c]^T$ and take it into the function f_W to obtain $f_W(x) = \sigma(w_r x_r + w_c x_c)$ and $f_W(c) = \sigma(w_r c_r + w_c x_c)$. The CE loss on $\mathbf{x}$ and its counterfactual $\mathbf{c}$ is calculated as

$$\begin{aligned}\mathcal{L}_{CE}(\mathbf{x}, \mathbf{c}) = &-\log(f_W(\mathbf{x})) \\ &- \log(1 - f_W(\mathbf{c})).\end{aligned} \tag{1}$$

By minimizing the CE loss, we enforce $f_W(\mathbf{x})$ to be as large as 1 and $f_W(\mathbf{c})$ to be as small as 0. Hence, we have

$$W = \underset{W}{\operatorname{argmax}} \frac{f_W(\mathbf{x})}{f_W(\mathbf{c})} = \underset{W}{\operatorname{argmax}} \frac{w_r}{w_c}, \tag{2}$$

which results in a higher weight w_r for the revised features and a lower weight w_c for the constant or unperturbed features. The derivation process can be found in the Appendix A. Similar phenomenon is observed in both least square loss (Joshi and He, 2022) and Fisher's Linear Discriminant on CAD (Fan et al., 2024).

The above observations indicate that the CE loss alone can lead the model to overfit the revised features in CAD, which necessitates incorporating a regularization that compels the model to consider a broader range of features.

3.2 The Role of Contrastive Loss

Recent research findings have empirically shown that models trained under contrastive loss mainly focus on capturing global relationships (Park et al., 2023) compared with negative log-likelihood losses such as masked language modeling. Inspired by this, we propose to employ CL to complement standard CE loss for training models on CAD. In the following, we start from the introduction of CL loss and then mathematically show how CL encourages the model to select a broader range of features beyond the edited ones in the counterfactual data.

Given an anchor sample $\mathbf{x}_i$ from a data batch $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$, $\forall \mathbf{x}_i \in \mathcal{D}$, we have its positive samples in $\mathcal{P}_i \equiv \{\mathbf{x}_p | y_p = y_i, p \neq i\}$ and negative samples in $\mathcal{N}_i \equiv \{\mathbf{x}_n | y_n \neq y_i, n \neq i\}$, where $\mathcal{N}_i$ contains the counterfactual samples $\mathbf{c}$ for every $\mathbf{x}_i$. The contrastive loss of an anchor $\mathbf{x}_i$ is

$$\mathcal{L}_{CL} = -\underset{\mathbf{x}_p \in \mathcal{P}_i}{\mathbb{E}} \left[\log \frac{e^{s_{ip}/\tau}}{e^{s_{ip}/\tau} + \sum_{\mathbf{x}_n \in \mathcal{N}_i} e^{s_{in}/\tau}} \right], \tag{3}$$

where $s_{xy} = \frac{\mathbf{z}_x \cdot \mathbf{z}_y}{\|\mathbf{z}_x\| \|\mathbf{z}_y\|}$ measures the cosine similarity between the hidden representations of a pair of samples, and τ is a temperature scaling factor for controlling the extent to which we separate positive and negative pairs (Wang and Isola, 2020b).

Without loss of generality, we assume $\mathbf{W} \in \mathbb{R}^{m\times d}$ that directly maps the input instance into a d-dimensional embedding space, $\mathbf{z}_i = \mathbf{W}^T\mathbf{x}_i$. To obtain the gradient of the CL loss coming from negative samples, we have

$$\begin{aligned}\left.\frac{\partial \mathcal{L}_{CL}}{\partial \mathbf{W}}\right|_{s_{in}} &= \frac{\partial \mathcal{L}_{CL}}{\partial s_{in}} \times \frac{\partial s_{in}}{\partial \mathbf{W}} \\ &= \frac{1}{\tau} P_{in} \times \mathbf{A}_{in}\mathbf{W}.\end{aligned} \tag{4}$$

The full derivation process can be found in the appendix B. Here, we have

$$P_{in} = \underset{\mathbf{x}_p \in \mathcal{P}_i}{\mathbb{E}} \left[\frac{e^{s_{in}/\tau}}{e^{s_{ip}/\tau} + \sum\limits_{\mathbf{x}_n \in \mathcal{N}_i} e^{s_{in}/\tau}} \right], \tag{5}$$

which indicates the probability of $\mathbf{x}_i$ being recognized as $\mathbf{x}_n$. $\mathbf{A}_{in} = \mathbf{x}_i\mathbf{x}_n^T + \mathbf{x}_n\mathbf{x}_i^T \in \mathbb{R}^{m\times m}$ is a symmetric matrix derived from the outer product of $\mathbf{x}_i$ and $\mathbf{x}_n$. Each element of $\mathbf{A}_{i,n}$ indicates the digit-level dot product between the features of $\mathbf{x}_i$ and $\mathbf{x}_n$, which provides a full view of the entire feature space when comparing a pair of samples. A higher value leads to a larger gradient update and the weights $\mathbf{W}$ are optimized by considering the whole feature sets.

The above analysis implies that the CL loss has the capability of capturing global features beyond those being edited. When learning on CAD under CL, we pair each instance $\mathbf{x}$ with its CFE, $\mathbf{c}$ to compel the model to disparate $\mathbf{x}$ from all negative samples, including its counterfactual example $\mathbf{c}$:

$$\sum_{\mathbf{x}_n \in \mathcal{N}_i} e^{s_{in}/\tau} = e^{s_{ic}/\tau} + \sum_{\mathbf{x}_n \in \mathcal{N}_i \setminus \mathbf{c}} e^{s_{in}/\tau}, \tag{6}$$

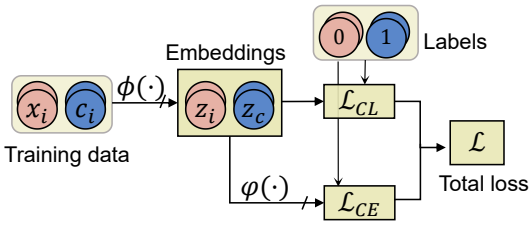

Figure 1: The overall learning framework.

where the non-bold c is the index of CFE. Let us revisit the toy example with $\mathbf{x} = [x_r, x_c]^T$ and $\mathbf{c} = [c_r, x_c]^T$. Though minimizing the similarity between $\mathbf{x}$ and $\mathbf{c}$ encourages the model to focus on features x_r, the other negative samples in the batch, e.g., $\mathbf{x}' = [x'_r, x'_c]^T$, will enforce the model to use both w_r and w_c to compare the difference. Hence, the existence of real negative samples could help the model capture the relationships between x_r and its context x_c.

As all s_{in} equally contribute to updating the model weights, the number of non-CFE negatives moderates the learning from local CAD and global patterns. A smaller batch size will manifest the influence of edited features, whereas a larger batch size may dilute the local differences in CAD, as discussed in the experiments 5.4.

3.3 Overall Learning Framework

Next, we introduce our proposed learning framework, Pairwisely Counterfactual Learning with Contrastive Loss Regularization, named PairCFR for short. As shown in Figure 1, a model f can be decomposed into two modules, $\phi(\cdot)$ and $\varphi(\cdot)$, i.e., $f = \varphi \circ \phi$, where $\phi(\cdot)$ encodes the input sentence into a hidden embedding, and $\varphi(\cdot)$ maps $\phi(\mathbf{x})$ for classification. For transformer-based models, we instantiate $\phi(\mathbf{x})$ using the [CLS] representation, denoted as $\mathbf{z}$. We explicitly pair the original sentences $\mathbf{x}$ and their CFEs, $\mathbf{c}$, in the same batch to provide additional training signals indicative of the underlying causal relationships.

The standard cross-entropy loss is computed on the logits vector projected from $\varphi(z)$. Optimizing CE loss enforces φ to identify a small set of features from $\mathbf{z}$ and assign them higher weights to quickly reach a local minimum while optimizing CL loss compels ϕ to consider the entire feature space of $\mathbf{z}$ to meet the distance constraints. Overall, we combine the two losses as follows.

$$\mathcal{L} = \lambda \mathcal{L}_{CL} + (1 - \lambda)\mathcal{L}_{CE}, \tag{7}$$

where λ is the trade-off factor to balance the two losses. To compute CL on a batch, we sample positive pairs that have the same label while all the negative samples including the CFE of the anchor sample are considered.

4 Experimental Setup

In the following, we introduce experimental settings, which include benchmark tasks, evaluation measurements, competitive baselines and implementation details.

4.1 Benchmark Tasks & Evaluations

We evaluate our learning framework on two NLP tasks, sentiment analysis (SA) and natural language inference (NLI). We use two human-edited CAD datasets (Kaushik et al., 2020), which ensures good-quality counterfactual data (Sen et al., 2023), to train all the models. The IMDb augmented dataset contains 4880 data samples with an original to CFE ratio of 1:1. The SNLI dataset contains 11330 data samples with an original to CFE ratio of 1:4. More details are reported in Appendix C.1.

To eliminate the random effect, we train each model for multiple runs (10 runs for SA and 7 runs for NLI) using different random seeds. We report the average test **accuracy** and **standard deviation** for both in-domain (ID) datasets and several out-of-domain (OOD) datasets. More details about the datasets and evaluation are summarized in the Appendix C.2.

4.2 Implementation Details

We finetune the BERT base (Devlin et al., 2019), RoBERTa base (Liu et al., 2019), Sentences-BERT (SBERT, multi-qa-distilbert-cos) (Reimers and Gurevych, 2019) and T5 base (Raffel et al., 2020) models with the original or CAD datasets on HuggingFace platform (Wolf et al., 2020). Following the common practices of transformers (Devlin et al., 2019), we take the embedding of the "[CLS]" token as sentence representation and finetune the whole model. We set the maximum token length to 350 for SA and 64 for NLI.

We employ the original dataset splits (Kaushik et al., 2020), where the train, validation, and test sets are balanced and follow a ratio of 7:1:2. We finetune all models up to 20 epochs using the AdamW optimizer and a linear learning rate scheduler with the warmup ratio as 0.05. The best learning rate is manually selected from $\{1e^{-4}, 1e^{-5}, 3e^{-5}, 5e^{-5}, 5e^{-6}, 1e^{-6}\}$. We apply

the early stopping strategy with a patience of 5 and the best model is selected according to the lowest validation loss. To determine the trade-off factor λ and temperature τ, we conducted a grid search in the range $[0, 1]$ with a step size of 0.1. Our code is released on GitHub repository [1].

4.3 Baselines

We compare our method PairCFR with the following baselines. For a fair comparison, we employ other forms of augmentation or increase the sampling number for the first three baselines without counterfactual augmentation, to ensure all approaches have the same number of training data.

Vanilla (Devlin et al., 2019). This method refers to a general model fine-tuning with original sentences. We include this baseline to verify the improvement of our method result from both the introduction of CAD and the novel learning framework.

BTSCL (Gunel et al., 2020). This approach employs the supervised contrastive loss (Khosla et al., 2020) into the model training where augmented positive samples are obtained through back-translating a given sentence (Ng et al., 2020).

CouCL (Wang et al., 2022a). As counterfactual examples (CFEs) are rare in a mini-batch, CouCL samples CFEs from the entire training dataset, where an example with lower confidence corresponds to a higher likelihood of being selected. Subsequently, it adopts the self-supervised contrastive loss to push representations of positive CFEs and negative CFEs far apart.

The following approaches study how to train a robust model with annotated CAD:

HCAD (Kaushik et al., 2020). It collects two human-edited CAD datasets and fine-tunes a pre-trained model on CAD with the cross-entropy loss.

CFGSL (Teney et al., 2020). As domain priors in CAD may be lost due to random shuffling in preprocessing (Kaushik et al., 2020), CFGSL pairs original data and its counterfactual example in the same batch and introduces a gradient supervision loss (GSL) alongside the cross-entropy loss. The GSL enforces the model gradient to align with the straight line from the original point to CFE.

ECF (Fan et al., 2024). It introduces two additional losses to mine the causal structures of CAD. The first loss extracts the dataset-level invariance through Invariant Risk Minimization (IRM) while the second loss is applied to pairs of original sentences and CFEs, preventing the model from relying on correlated features.

[1] https://anonymous.4open.science/r/PairCAD-BDA3/

5 Results and Analysis

5.1 Overall Performance Comparison

Table 1 reports the overall performance comparisons. Our proposed PairCFR method outperforms all the baseline models on three out of four OOD datasets for both SA and NLI tasks across four different backbone models. We make the following observations.

Firstly, we observe similar findings as reported in (Joshi and He, 2022): CAD-based methods may perform worse than non-CAD methods on OOD tasks, e.g., HCAD, CFGSL, and ECF results in lower accuracy on some OOD datasets when fine-tuning the RoBERTa, and HCAD always lags behind CouCL on SA task using fine-tuned T5 model. These could be due to the failure in extracting complementary features between CFEs and the original data; Secondly, the introduction of CAD alters the original data distribution, hence, as anticipated, the CAD-based methods underperform non-CAD methods on ID datasets. Thirdly, our proposed PairCFR exhibits superior OOD performance compared to the baselines, achieving the highest accuracy on all OOD datasets, with the sole exceptions being the Yelp and Negation datasets. We postulate that the noted exceptions may be attributed to Yelp and Negative datasets having distributions similar to the ID datasets. The above results validate that PairCFR possesses a heightened capability to learn prior knowledge in CAD.

5.2 Few-shot Learning Performance

Data augmentation, including counterfactual augmentation, is frequently utilized to enhance few-shot learning performance. In this section, we investigate the effectiveness of our proposed PairCFR in few-shot learning scenarios.

We experimented with the finetuned BERT model on SNLI dataset, gradually increasing the number of training samples from 50 to 4000 (with the proportion of original to counterfactuals, 1:4, remains unchanged). Few-shot experiment results on SNLI are reported in Figure 2 and more experiments on other datasets are reported in the Appendix C.5. Our PairCFR model exhibits enhanced accuracy and decreased standard deviation across OOD datasets, particularly evident in scenarios with small training samples. For instance,

Table 1: Average performance of various fine-tuned models on ID and OOD test sets. $\overline{Acc}$ denotes the average of all the OOD performance. Highest values are highlighted in **bold**.

Methods	Sentiment Analysis						Natural Language Inference						
	In-Domain	Out-of-Dimain					In-Domain	Out-of-Dimain					
	IMDb	Amazon	Yelp	Twitter	SST-2	$\overline{Acc}$	SNLI	MNLI-m	MNLI-mm	Negation	Spelling-e	Word-o	$\overline{Acc}$
BERT-base-uncased													
Vanilla	90.15±1.66	86.38±0.39	91.03±0.83	81.66±0.27	82.59±1.00	85.42	78.85±0.44	57.43±0.92	59.36±0.80	40.96±4.32	53.56±1.54	50.75±6.65	52.41
BTSCL	**90.43±1.47**	85.45±0.71	**91.97±0.31**	81.79±1.28	83.80±1.17	85.75	**79.02±0.49**	57.28±1.30	59.10±1.42	**43.10±3.65**	53.51±1.74	49.20±4.51	52.44
CouCL	85.67±1.13	86.75±0.22	89.53±0.55	84.41±0.23	85.01±0.43	86.43	71.90±0.95	51.99±1.75	52.20±1.86	38.70±4.69	49.82±2.01	44.03±4.02	47.35
HCAD	88.16±2.70	86.40±0.77	89.94±0.99	83.29±2.71	85.74±1.04	86.34	73.49±1.37	58.53±1.59	60.77±1.46	35.43±3.06	54.01±2.70	54.72±3.29	52.69
CFGSL	88.51±3.29	85.52±1.05	89.58±1.83	84.56±1.53	86.77±0.79	86.61	77.16±0.41	60.11±1.07	62.25±0.66	33.81±1.89	56.37±0.74	58.45±0.97	54.20
ECF	87.71±0.29	86.43±0.10	89.30±0.16	83.05±0.69	86.23±0.18	86.25	73.23±1.52	58.95±0.15	61.19±1.34	42.40±1.07	54.15±0.53	57.10±0.92	54.76
Ours	89.63±1.36	**86.79±0.72**	91.78±0.44	**85.27±0.39**	**86.81±0.97**	**87.66**	75.38±0.21	**60.46±0.38**	**62.27±0.39**	39.21±3.61	**56.84±0.54**	**59.16±0.88**	**55.59**
RoBERTa-base													
Vanilla	92.68±1.15	87.08±1.39	94.00±0.77	81.43±2.82	86.04±2.76	87.14	85.16±0.39	70.35±1.29	71.25±1.59	52.47±5.55	67.36±1.36	61.82±4.54	64.65
BTSCL	**93.09±0.61**	89.46±0.21	**94.74±0.36**	85.72±1.22	87.16±0.87	89.27	**85.72±0.44**	70.83±1.38	72.10±1.32	**56.89±3.78**	67.61±1.32	62.22±3.55	65.93
CouCL	91.22±0.83	89.48±0.19	93.04±0.58	87.40±0.77	88.07±0.66	89.50	82.37±0.52	70.86±1.32	71.38±1.23	51.83±2.71	68.08±1.23	64.68±1.82	65.37
HCAD	90.12±1.74	88.50±0.57	92.18±0.94	83.43±1.75	86.48±0.98	87.65	80.91±0.69	70.35±1.08	70.77±0.76	45.79±4.16	67.37±1.28	64.83±1.47	63.82
CFGSL	90.69±0.92	88.32±0.41	93.48±0.48	83.90±1.78	86.89±0.80	88.15	82.45±0.35	71.59±0.90	71.25±1.06	51.40±1.47	68.86±1.07	62.22±1.99	65.06
ECF	91.05±0.44	88.56±0.32	93.79±0.19	85.82±0.43	87.84±0.59	89.00	81.88±0.17	70.45±1.03	71.18±0.93	51.70±2.38	66.60±0.94	63.76±1.98	64.74
Ours	91.74±0.88	**89.60±0.26**	93.35±0.34	**87.90±0.45**	**88.61±0.41**	**89.87**	82.13±0.51	**71.80±0.53**	**72.12±0.79**	55.19±1.97	**68.88±0.36**	**65.91±1.35**	**66.78**
SBERT-multi-qa-distilbert-cos													
Vanilla	87.61±1.86	80.65±0.67	89.74±0.77	83.95±1.12	82.01±1.59	84.09	76.96±0.53	53.90±2.03	55.90±2.22	45.20±4.18	51.23±2.72	48.27±5.00	50.90
BTSCL	**88.84±2.41**	81.21±0.76	**90.49±0.37**	84.20±0.61	83.62±0.64	84.88	**77.16±0.38**	54.42±1.31	56.14±1.36	**45.40±2.78**	52.44±1.83	49.80±2.63	51.64
CouCL	87.96±0.67	83.92±0.13	89.15±0.18	85.40±0.31	83.48±0.37	85.49	70.61±1.54	55.29±1.45	57.90±1.81	35.86±1.87	52.01±2.26	54.89±1.91	51.19
HCAD	86.09±1.74	83.94±0.39	87.87±0.66	85.91±0.66	82.83±0.90	85.14	71.64±1.04	55.93±1.61	58.70±1.96	35.05±1.22	53.33±1.06	54.86±2.08	51.57
CFGSL	86.05±1.07	82.71±0.73	87.59±0.75	83.36±0.55	83.70±0.49	84.34	70.72±1.06	55.84±0.88	58.52±1.15	36.07±3.38	52.60±1.27	55.57±1.68	51.72
ECF	87.83±0.46	84.51±0.34	88.44±0.20	84.60±0.70	84.27±0.56	85.46	64.55±1.23	49.95±1.84	51.49±1.82	38.59±2.32	48.31±1.67	49.55±2.27	47.58
Ours	87.28±0.75	**84.58±0.22**	88.52±0.30	**86.32±0.35**	**84.31±0.78**	**85.93**	71.48±0.40	**57.19±0.84**	**60.76±0.46**	37.27±2.35	**54.36±0.67**	**56.78±1.24**	**53.27**
T5-base													
Vanilla	92.15±1.49	88.24±0.85	94.44±0.67	83.40±1.38	86.17±2.60	88.06	83.28±0.57	62.62±2.59	65.18±2.10	41.00±2.46	58.76±2.61	48.30±3.27	55.17
BTSCL	**92.78±1.08**	88.50±0.81	**94.89±0.42**	83.37±1.09	87.17±1.07	88.48	**83.66±0.46**	64.01±2.57	66.47±2.24	42.16±2.90	60.01±3.43	50.16±5.69	56.56
CouCL	91.74±0.88	88.91±0.47	93.35±0.34	87.03±0.70	88.61±0.41	89.48	79.81±0.54	70.19±0.58	71.84±0.76	39.82±3.23	66.35±0.68	64.29±1.58	62.50
HCAD	90.09±1.95	88.72±0.85	92.60±0.87	85.63±1.15	85.54±1.28	88.12	80.09±0.73	70.19±0.72	71.60±0.83	45.05±3.94	66.57±0.73	65.30±1.51	63.74
CFGSL	89.48±5.17	88.27±1.05	92.77±1.45	81.56±2.49	82.11±2.50	86.18	80.71±0.64	69.08±0.97	69.85±1.12	45.59±3.74	65.58±1.18	65.80±1.55	63.18
ECF	90.85±0.37	89.27±0.25	92.65±0.44	87.66±0.26	88.57±0.54	89.54	78.93±0.51	69.57±1.14	70.30±1.45	46.14±3.12	64.19±1.08	65.79±1.71	63.20
Ours	91.47±0.89	**89.18±0.21**	93.45±0.63	**87.90±0.45**	**88.64±1.04**	**89.79**	80.87±0.77	**71.38±0.13**	**72.46±0.57**	**46.31±0.50**	**67.37±0.12**	**67.39±0.33**	**64.98**

Table 2: Ablation study for the pairing strategy and the CL loss on various transformer-based models.

Variants		Sentiment Analysis					Natural Language Inference					
		In-Domain	Out-of-Dimain				In-Domain	Out-of-Dimain				
Train Data	Loss	IMDb	Amazon	Yelp	Twitter	SST-2	SNLI	MNLI-m	MNLI-mm	Negation	Spelling-e	Word-o
Bert-base-uncased												
ShuffCAD	CE	88.16±2.70	86.40±0.77	89.94±0.99	83.29±2.71	85.74±1.04	73.49±1.37	58.53±1.59	60.77±1.46	35.43±3.06	54.01±2.70	54.72±3.29
PairCAD	CE	88.23±3.11	86.56±0.34	89.97±1.85	84.15±1.20	85.84±0.85	74.27±0.72	59.13±0.65	60.85±0.88	36.10±1.92	56.14±1.34	55.40±2.83
ShuffCAD	CE+CL	89.18±1.33	86.77±0.65	91.45±0.53	84.14±1.82	86.26±0.99	73.77±1.11	59.39±0.64	61.85±0.86	36.80±4.04	55.62±0.87	57.09±2.45
PairCAD	CE+CL	**89.63±1.36**	**86.79±0.72**	**91.78±0.44**	**85.27±0.39**	**86.81±0.97**	**75.38±0.21**	**60.46±0.38**	**62.27±0.39**	**39.21±3.61**	**56.84±0.54**	**59.16±0.88**
Roberta-base												
ShuffCAD	CE	90.12±1.74	88.50±0.57	92.18±0.94	83.43±1.75	86.48±0.98	80.91±0.69	70.35±1.08	70.77±0.76	45.79±4.16	67.37±1.28	64.83±1.47
PairCAD	CE	90.95±0.84	88.77±0.74	92.77±0.95	83.45±2.53	86.37±1.06	81.69±0.90	70.77±0.49	71.33±0.45	54.38±1.67	67.90±0.63	65.43±0.99
ShuffCAD	CE+CL	91.42±1.01	89.44±0.27	92.91±0.64	86.67±1.05	87.25±0.68	81.95±0.39	71.16±0.60	71.79±0.79	51.43±2.91	68.20±0.57	64.12±1.03
PairCAD	CE+CL	**91.74±0.88**	**89.60±0.26**	**93.35±0.34**	**87.90±0.45**	**88.61±0.41**	**82.13±0.51**	**71.80±0.53**	**72.12±0.79**	**55.19±1.97**	**68.88±0.36**	**65.91±1.35**
Sbert-multi-qa-distilbert-cos												
ShuffCAD	CE	86.09±1.74	83.94±0.39	87.87±0.66	85.91±0.66	82.83±0.90	71.64±1.04	55.93±1.61	58.70±1.96	35.05±1.22	53.33±1.06	54.86±2.08
PairCAD	CE	86.78±1.41	83.55±0.39	88.51±0.77	85.95±0.40	83.20±0.63	70.90±1.02	56.50±0.58	59.03±0.57	35.89±1.98	53.03±1.17	55.04±1.03
ShuffCAD	CE+CL	87.68±1.05	84.23±0.37	88.66±0.77	85.45±0.28	83.60±0.38	71.38±0.62	57.08±0.53	60.01±0.35	35.11±1.64	54.15±0.53	55.59±1.89
PairCAD	CE+CL	**87.05±0.97**	**84.36±0.49**	**88.30±0.38**	**86.10±0.45**	**83.76±0.46**	**71.48±0.40**	**57.19±0.84**	**60.76±0.46**	**37.27±2.35**	**54.36±0.67**	**56.78±1.24**
T5-base												
ShuffCAD	CE	90.09±1.95	88.72±0.85	92.60±0.87	85.63±1.15	85.54±1.28	80.09±0.73	70.19±0.72	71.60±0.83	45.05±3.94	66.57±0.73	65.30±1.51
PairCAD	CE	90.03±1.35	89.02±0.41	92.76±0.99	86.46±1.00	86.59±1.37	79.55±0.66	68.86±0.52	70.75±0.77	45.18±3.49	65.56±0.67	65.64±1.50
ShuffCAD	CE+CL	90.38±1.80	89.03±0.46	93.06±1.29	85.75±0.96	87.24±2.12	80.21±0.10	70.43±0.11	71.78±0.37	45.41±2.08	66.59±0.56	66.28±0.93
PairCAD	CE+CL	**91.47±0.89**	**89.18±0.21**	**93.45±0.63**	**87.03±0.70**	**88.64±1.04**	**80.14-0.96**	**71.02±0.39**	**71.84±0.76**	**45.73±0.71**	**66.87±0.51**	**67.11±0.39**

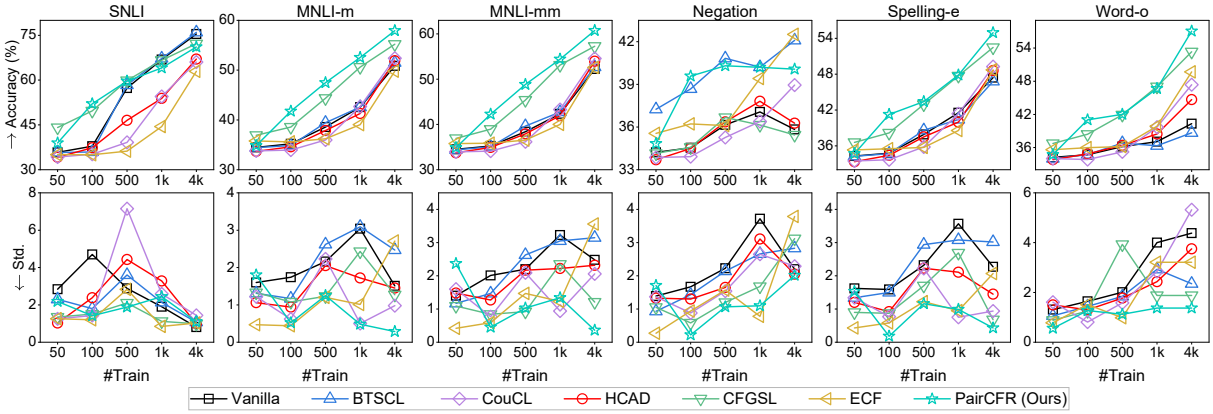

Figure 2: Few-shot learning results of $\text{BERT}_{\text{base}}$ on NLI. x-axis represents the number of training samples (original: augmented data = 1:4) and y-axis represents the averaged accuracy and standard deviation on ID and OODs.

Table 3: The influence of the neutral samples during fine-tuning $\text{BERT}_{\text{base}}$ on SNLI. *w* and *w/o* is the abbreviation of *with* and *without*.

Train Data	netural samples	In-Domain	Out-of-Dimain					
		SNLI	MNLI-m	MNLI-mm	Negation	Spelling-e	Word-o	$\overline{Acc}$
PairCAD	w	73.29±1.09	59.41±0.91	61.66±0.85	35.96±2.81	56.42±1.10	56.14±2.60	53.92
PairCAD	w/o	**75.38±0.21**	**60.46±0.38**	**62.27±0.39**	**39.21±3.61**	**56.84±0.54**	**59.16±0.88**	**55.59**

Table 4: The influence of the counterfactual diversity during fine-tuning T5_{base} on SNLI.

#Train CE+SCL	R:O	In-Domain	Out-of-Domain					
		SNLI	MNLI-m	MNLI-mm	Negation	Spelling-e	Word-o	$\overline{Acc}$
Original (20k)	-	**85.09±0.27**	69.53±1.38	71.62±1.04	45.65±3.53	66.43±1.49	52.89±5.22	61.22
PairCAD (3.3k)	1	74.50±2.51	65.24±1.63	67.61±1.36	38.38±3.42	61.24±1.86	60.61±2.33	58.62
PairCAD (4.9k)	2	76.12±1.58	66.62±1.05	69.31±0.87	42.33±7.31	62.91±1.60	62.61±1.58	60.76
PairCAD (6.4k)	3	77.98±0.82	68.36±1.48	70.00±1.44	43.13±1.17	64.60±1.98	64.45±2.15	62.11
PairCAD (8.3k)	4	80.14±0.96	**71.02±0.39**	**71.84±0.76**	**45.73±0.70**	**66.87±0.51**	**67.11±0.39**	**64.51**

PairCFR significantly outperforms other methods by around 6% on Spelling-e when trained with 100 counterfactually augmented samples.

5.3 Ablation Study

We conduct ablation experiments to verify the efficacy of two important components of our proposed learning framework: (1) the pairing strategy: the integration of original data with their CFEs within the same batch, denoted *PairCAD*, versus *ShuffCAD* where randomly shuffle CFEs and origianls. (2) the CL loss: the incorporation of CL and CE loss versus CE loss alone.

Results in Table 2 offer several insights: 1) The strategy of pairing original data with their CFEs in the same batch improves OOD performance for both SA and NLI tasks. This can be attributed to the preservation of prior knowledge regarding causal relations that might be lost during random shuffling; 2) The efficacy of PairCAD with a CE-only learning framework is not guaranteed. For example, within the T5 model framework, PairCAD underperforms ShuffCAD on the SNLI, MNLI, and Spelling-e datasets when only CE loss is adopted. This underscores the critical role of the CL component in augmenting features within batch CFEs and original scenarios; 3) Integrating the CL term consistently improves model performance in both ID and OODs. Particularly, combining CL with PairCAD yields the best performance across various model assessments, highlighting the effectiveness of contrastive learning and the pairing strategy in leveraging causal relations of CFEs.

5.4 Impact of Batch Size

In this study, we investigate the effect of batch size on learning performance. We conduct experiments on the fine-tuned BERT model for SA and the fine-

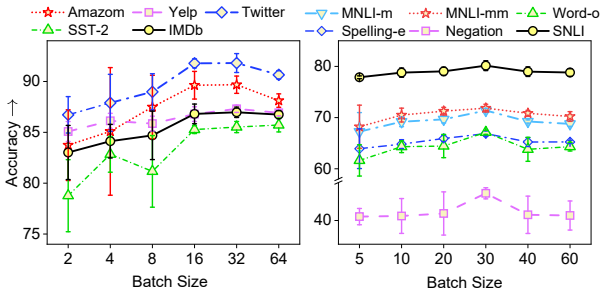

Figure 3: Test results for fine-tuning $\text{BERT}_{\text{base}}$ on IMDb augmented data (left) and T5_{base} on SNLI augmented data (right) with respect to the batch size.

tuned T5 model for NLI, incrementally increasing the batch size while maintaining the original augmentation ratio for each respective task. From Figure 3, we observe that the model performance on most evaluation metrics initially improves with increasing batch size, but eventually reaches a plateau or experiences a slight decline. We contend that the inclusion of negative samples in the CL function provides additional regularization, forcing the model to rely on a broader array of features beyond those edited. However, an excessively large batch size introduces an overwhelming number of negative samples, which may dilute the human priors in CAD, leading to diminished performance. This trend is consistent across both SA and NLI tasks, highlighting the balance required in batch size selection.

5.5 Contribution of Neutral Class in NLI

Do all counterfactual examples unequivocally contribute to enhancing model generality? To address this, we specifically experimented with the fine-tuned BERT model on the NLI task, comparing performance with and without the inclusion of CFEs generated from neutral class samples.

Results in Table 3 reveal that removing CFEs generated from neutral samples significantly enhances the OOD generalization by approximately 2% when training the model on CAD with our proposed learning framework. We attribute this performance difference to the distinct nature of neutral samples. In NLI tasks, entailment and contradiction judgments are often easily made based on semantic alignment or its absence. Conversely, neutral samples represent scenarios where the hypothesis and premise lack any clear relationship, encompassing a vast array of potential expressions. This diversity poses a challenge for models to dentify universal patterns within the neutral class through human annotation, explaining why adding CFEs generated from neutral samples detriment the performance of the model.

This investigation highlights the necessity of contemplating the practical value of adding additional counterfactual examples for specific classes.

5.6 Effect of Counterfactual Diversity

We also investigated the role of CFE diversity in improving model performance. In SNLI, each sentence is associated with 4 CFEs, due to the existence of two opposite targets and modifications made to both the hypothesis and premise. As each edited sentence is obtained by specific perturbation type (e.g.quantifier, lexical, delete), the more CFEs generated from the same example will potentially introduce more diverse counterfactuals (Joshi and He, 2022). We fine-tuned the T5_{base} model by incrementally introducing more CFEs in a batch, ranging from 1 to 4. The results in Table 4, reveal a direct correlation between the number of CFEs and the model's generalization capabilities. Notably, the OOD performance of the model trained on CAD is even better than that trained on a 3 times larger dataset with only original data. We conclude that enhancing counterfactual diversity proves to be an efficient strategy, which is similar to the findings reported in (Joshi and He, 2022).

6 Conclusion

Counterfactually Augmented Data (CAD) can enhance model robustness by explicitly identifying causal features. However, recent research found that CAD may fall behind non-CAD methods on generality. In this work, we introduce PairCFR to overcome this challenge. PairCFR pairs original and counterfactual data during training and includes both contrastive and cross-entropy losses for learning discriminative representations. We prove that contrastive loss aids models in capturing sufficient relationships not represented in CAD, thus improving generality. Extensive experiments demonstrate that our PairCFR achieves superior accuracy and robustness in various scenarios. Our findings emphasize the potential of carefully designed training paradigms in utilization of CAD.

7 Limitations

While PairCFR has been shown to effectively improve models' OOD generalization with human-edited CAD datasets, its effectiveness on LLM-

generated CAD, which may potentially suffer from lower quality due to noisy and insufficient perturbations, has yet to be explored. Our future work will seek to broaden the scope of PairCFR by also applying it to LLM-generated CAD datasets. Furthermore, PairCFR currently utilizes only one form of contrastive loss, i.e. InfoNCE. In the future, we aim to investigate alternative contrastive loss variants and assess their potential to further enhance OOD generalization capabilities.

8 Ethics Statement

This work focuses on reducing shortcut learning in models trained on CAD, thereby improving their robustness and generalization. Similar to other methods designed to mitigate learning from spurious correlations, our proposed PairCFR could help elicit trust in NLP models. It assists models in better-considering context (see Section 3 for details), preventing decision-making based on incomplete or biased information, such as solely on the edited words in CAD. Nonetheless, ensuring absolute fairness in model decisions in complex real-world contexts remains a formidable challenge solely from a model design standpoint. For instance, models could be compromised by low-quality or erroneous counterfactual data, leading to the learning of false relationships and resulting in erroneous or biased real-world decisions. Consequently, it is crucial for practitioners to consider the quality of counterfactual data alongside model design.

References

Solon Barocas, Andrew D Selbst, and Manish Raghavan. 2020. The hidden assumptions behind counterfactual explanations and principal reasons. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 80–89.

Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.

Zeming Chen, Qiyue Gao, Antoine Bosselut, Ashish Sabharwal, and Kyle Richardson. 2023. DISCO: Distilling counterfactuals with large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5514–5528, Toronto, Canada. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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. Association for Computational Linguistics.

Caoyun Fan, Wenqing Chen, Jidong Tian, Yitian Li, Hao He, and Yaohui Jin. 2024. Unlock the potential of counterfactually-augmented data in out-of-distribution generalization. *Expert Systems with Applications*, 238:122066.

Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. In *International Conference on Learning Representations*.

Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.

William Huang, Haokun Liu, and Samuel R. Bowman. 2020. Counterfactually-augmented SNLI training data does not yield better generalization than unaugmented data. In *Proceedings of the First Workshop on Insights from Negative Results in NLP*, pages 82–87, Online. Association for Computational Linguistics.

Nitish Joshi and He He. 2022. An investigation of the (in)effectiveness of counterfactually augmented data. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3668–3681, Dublin, Ireland. Association for Computational Linguistics.

Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2020. Learning the difference that makes a difference with counterfactually augmented data. *International Conference on Learning Representations (ICLR)*.

Divyansh Kaushik, Amrith Setlur, Eduard Hovy, and Zachary C Lipton. 2021. Explaining the efficacy of counterfactually augmented data. *International Conference on Learning Representations (ICLR)*.

Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673.

Meizhen Liu, Xu Guo, He Jiakai, Jianye Chen, Fengyu Zhou, and Siu Hui. 2023a. InteMATs: Integrating granularity-specific multilingual adapters for cross-lingual transfer. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5035–5049, Singapore. Association for Computational Linguistics.

Meizhen Liu, Jiakai He, Xu Guo, Jianye Chen, Siu Cheung Hui, and Fengyu Zhou. 2023b. Grancats: Cross-lingual enhancement through granularity-specific contrastive adapters. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, page 1461–1471, New York, NY, USA. Association for Computing Machinery.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.

Nishtha Madaan, Inkit Padhi, Naveen Panwar, and Diptikalyan Saha. 2021. Generate your counterfactuals: Towards controlled counterfactual generation for text. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13516–13524.

Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.

Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. *arXiv preprint arXiv:1806.00692*.

Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2020. Facebook fair's wmt19 news translation task submission. In *Proc. of WMT*.

Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 188–197.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.

Namuk Park, Wonjae Kim, Byeongho Heo, Taekyung Kim, and Sangdoo Yun. 2023. What do self-supervised vision transformers learn? In *The Eleventh International Conference on Learning Representations*.

Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis only baselines in natural language inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.

Reid Pryzant, Dan Iter, Jerry Li, Yin Lee, Chenguang Zhu, and Michael Zeng. 2023. Automatic prompt optimization with "gradient descent" and beam search. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968, Singapore. Association for Computational Linguistics.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

Sara Rosenthal, Noura Farra, and Preslav Nakov. 2019. Semeval-2017 task 4: Sentiment analysis in twitter. *arXiv preprint arXiv:1912.00741*.

Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823.

Indira Sen, Dennis Assenmacher, Mattia Samory, Isabelle Augenstein, Wil Aalst, and Claudia Wagner. 2023. People make better edits: Measuring the efficacy of llm-generated counterfactually augmented data for harmful language detection. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10480–10504.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.

Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29.

Damien Teney, Ehsan Abbasnedjad, and Anton van den Hengel. 2020. Learning what makes a difference from counterfactual examples and gradient supervision. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 580–599. Springer.

Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841.

Jinqiang Wang, Rui Hu, Chaoquan Jiang, Rui Hu, and Jitao Sang. 2022a. Counterexample contrastive learning for spurious correlation elimination. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4930–4938.

Tongzhou Wang and Phillip Isola. 2020a. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR.

Tongzhou Wang and Phillip Isola. 2020b. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR.

Yaushian Wang, Ashley Wu, and Graham Neubig. 2022b. English contrastive learning can learn universal cross-lingual sentence embeddings. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9122–9133, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Zhao Wang and Aron Culotta. 2020. Identifying spurious correlations for robust text classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3431–3440, Online. Association for Computational Linguistics.

Jiaxin Wen, Yeshuang Zhu, Jinchao Zhang, Jie Zhou, and Minlie Huang. 2022. AutoCAD: Automatically generate counterfactuals for mitigating shortcut learning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2302–2317, Abu Dhabi, United Arab Emirates. 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*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel S Weld. 2021. Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. *arXiv preprint arXiv:2101.00288*.

Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level Convolutional Networks for Text Classification. *arXiv:1509.01626 [cs]*.

A The trap in the CE loss

Given the sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$, we have $f_W(\mathbf{x}) = \frac{1}{1+e^{-(w_r x_r + w_c x_c)}}$ and $f_W(\mathbf{c}) = \frac{1}{1+e^{-(w_r c_r + w_c x_c)}}$. The goals of $f_W(\mathbf{x})$ and $f_W(\mathbf{c})$ can be combined as $\max \frac{f_W(\mathbf{x})}{f_W(\mathbf{c})}$. Hence, we have

$$\begin{aligned}\frac{f_W(\mathbf{x})}{f_W(\mathbf{c})} &= \frac{1+e^{-(w_r c_r + w_c x_c)}}{1+e^{-(w_r x_r + w_c x_c)}} \\ &= \frac{e^{(w_r c_r + w_c x_c)}}{e^{(w_r c_r + w_c x_c)} + e^{(w_r c_r - w_r x_r)}} \\ &= \frac{e^{w_c x_c}}{e^{w_c x_c} + e^{-w_r x_r}} \\ &= \frac{1}{1+e^{-\frac{w_r}{w_c}\cdot\frac{x_r}{x_c}}}. \end{aligned} \tag{8}$$

Here, $\frac{x_r}{x_c}$ is a constant and $\frac{w_r}{w_c}$ is to be optimized. Maximizing $\frac{f_W(\mathbf{x})}{f_W(\mathbf{c})}$ is equivalent to maximizing $\frac{w_r}{w_c}$, which means the model tends to assign higher weights w_r to revised features and lower weights w_c to constant (unperturbed) features.

B Gradient analysis of CL

In this section, we introduce the details of gradient of CL with respect to the weight $\mathbf{W}$ through the negative branches $s_{i,n}$. Before talking details, we rewrite the CL term for convenience,

$$\mathcal{L}_{CL} = -\mathop{\mathbb{E}}_{\mathbf{x}_i \in \mathcal{P}_i}\left[\log \frac{e^{s_{ip}/\tau}}{e^{s_{ip}/\tau} + \sum_{\mathbf{x}_n \in \mathcal{N}_i} e^{s_{in}/\tau}}\right]. \tag{9}$$

The total derivative of $\mathcal{L}_{CL}$ over the weights be calculated through the chain rule as

$$\frac{\partial \mathcal{L}_{CL}}{\partial \mathbf{W}} = \frac{\partial \mathcal{L}_{CL}}{\partial s_{in}} \times \frac{\partial s_{in}}{\partial \mathbf{W}} + \frac{\partial \mathcal{L}_{CL}}{\partial s_{ip}} \times \frac{\partial s_{ip}}{\partial \mathbf{W}}, \tag{10}$$

where the gradient coming from the branch s_{in} is

$$\left.\frac{\partial \mathcal{L}_{CL}}{\partial \mathbf{W}}\right|_{s_{in}} = \frac{\partial \mathcal{L}_{CL}}{\partial s_{in}} \frac{\partial s_{in}}{\partial \mathbf{W}}. \tag{11}$$

For simplicity, we let $s_{in} = \mathbf{z}_i^T \mathbf{z}_n$ and drop the the denominator, $\| \mathbf{z}_i \| \| \mathbf{z}_n \|$, which is eliminated in the product of partial derivatives. $\mathbf{z}_i = \mathbf{W}^T \mathbf{x}_i$ and $\mathbf{z}_j = \mathbf{W}^T \mathbf{x}_n$, and then we have

$$\begin{aligned}\frac{\partial s_{in}}{\partial \mathbf{W}} &= \frac{\partial (\mathbf{W}^T \mathbf{x}_i)^T (\mathbf{W}^T \mathbf{x}_n)}{\partial \mathbf{W}} \\ &= \frac{\partial (\mathbf{x}_i^T \mathbf{W})(\mathbf{W}^T \mathbf{x}_n)}{\partial \mathbf{W}} \\ &= \mathbf{x}_i \mathbf{x}_n^T \mathbf{W} + \mathbf{x}_n \mathbf{x}_i^T \mathbf{W} \\ &= \mathbf{A}_{in} \mathbf{W}. \end{aligned} \tag{12}$$

Here, $\mathbf{A}_{in} = \mathbf{x}_i \mathbf{x}_n^T + \mathbf{x}_n \mathbf{x}_i^T$. The CL can be further written by,

$$\begin{aligned}\mathcal{L}(\mathbf{x}_p) &= \mathop{\mathbb{E}}_{\mathbf{x}_p \in \mathcal{P}_i}\left[-log \frac{exp(s_{ip}/\tau)}{exp(s_{ip}/\tau) + \sum_{\mathbf{x}_n \in \mathcal{N}_i} exp(s_{in}/\tau)}\right] \\ &= \mathop{\mathbb{E}}_{\mathbf{x}_p \in \mathcal{P}_i}\left[log(exp(s_{ip}/\tau) + \sum_{\mathbf{x}_n \in \mathcal{N}_i} exp(s_{in}/\tau))\right] \\ &\quad - \mathop{\mathbb{E}}_{\mathbf{x}_p \in \mathcal{P}_i}(s_{ip}/\tau). \end{aligned} \tag{13}$$

Then, we can easily obtain the gradient of CL from a negative sample $s_{i,n}$ and a specific positive sample $s_{i,p}$ by,

$$\begin{aligned}\frac{\partial \mathcal{L}_i}{\partial s_{in}} &= \frac{1}{\tau} \mathop{\mathbb{E}}_{\mathbf{x}_p \in \mathcal{P}_i}\left[\frac{exp(s_{in}/\tau)}{exp(s_{ip}/\tau + \sum_{\mathbf{x}_n \in \mathcal{N}_i} exp(s_{in}/\tau)}\right] \\ &= \frac{1}{\tau} \mathop{\mathbb{E}}_{\mathbf{x}_p \in \mathcal{P}_i}\left[\frac{\exp(s_{in}/\tau)}{\exp(s_{ip}/\tau) + \sum_{\mathbf{x}_n \in \mathcal{N}_i} \exp(s_{in}/\tau)}\right] \\ &= \frac{1}{\tau} P_{in} \qquad \text{(written as } P_{in}\text{)}. \end{aligned} \tag{14}$$

Combining Eq (12) and Eq (14) gives the final gradient from a negatives sample,

$$\begin{aligned}\left.\frac{\partial \mathcal{L}(\mathbf{x}_i)}{\partial \mathbf{W}}\right|_{s_{in}} &= \frac{\partial \mathcal{L}(\mathbf{x}_i)}{\partial s_{in}} \frac{\partial s_{in}}{\partial \mathbf{W}} \\ &= \frac{1}{\tau} P_{in} \mathbf{A}_{in} \mathbf{W} \end{aligned} \tag{15}$$

Summing up gradients in Eqn 15 from all negative samples, we can derive,

$$\begin{aligned}\left.\frac{\partial \mathcal{L}(\mathbf{x}_i)}{\partial \mathbf{W}}\right|_{\mathcal{N}_i} &= \frac{\partial \mathcal{L}(\mathbf{x}_i)}{\partial s_{in}} \frac{\partial s_{in}}{\partial \mathbf{W}} \\ &= \frac{1}{\tau} \sum_{\mathbf{x}_n \in \mathcal{N}_i} P_{in} \mathbf{A}_{in} \mathbf{W} \end{aligned} \tag{16}$$

As the gradient contains pair-wise outer products between the anchor point and all its negative samples, it fully captures the overview of the feature space rather than focusing on a local perspective on edited words.

C Experimental Details

C.1 Training Data

We introduce more details of the CAD data used in model training in our experiments. We adopt two counterfactually augmented datasets from IMDb (Maas et al., 2011) and SNLI (Bowman et al., 2015) in (Kaushik et al., 2021). The counterfactually augmented IMDb dataset contains 2440 original sentences, with each sentence having a corresponding

revised counterfactual example. In SNLI, annotators can revise both the hypothesis and the premise for each of two opposite classes, and each sentence has 4 counterfactual examples. After another round of human filtering, the counterfactual augmented SNLI dataset consists of 9064 counterfactuals and 2266 original examples. During training, we split two CAD datasets into train, validation, test sets as shown in Table 5.

Table 5: Statistic of human-edited CAD datasets.

Dataset	#Train	#Val	#Test	Total No.
Sentiment Analysis: IMDb				
Original	1707	245	488	2440
Revised	1707	245	488	2440
CAD	3414	490	976	4880
Natural Language Inference: SNLI				
Original	1666	200	400	2266
Revised	6664	800	1600	9064
CAD	8330	1000	2000	11330

C.2 ID and OOD datasets

Here, we provide statistics of in-domain (ID) and out-of-domain (OOD) datasets used to evaluate the generalization of models in Table 6.

Since CADs in our experiments are manually revised on samples from IMDb (Maas et al., 2011) and SNLI (Bowman et al., 2015), we include their test datasets for ID evaluation. As for OOD evaluation, we evaluate our sentiment models on Amazon reviews (Ni et al., 2019), Topic-based Tweets sentiment data (Rosenthal et al., 2019), Yelp reviews (Zhang et al., 2015) and SST-2 movie reviews (Socher et al., 2013). On NLI task, we report on the genre-matched (MNLI-m) and genre-mismatched (MNLI-mm) test set of MNLI (Williams et al., 2017), which are more challenging than SNLI due to multiple genres. In addition, We additionally employ the diagnostic datasets Negation, Spelling error, and Word-Overlap provided by Naik et al. (2018) to evaluate models' reasoning abilities across lexical semantics and grammaticality.

C.3 Implementation setting

Next, we list the learning rate lr, batch size bts, trade-off factor λ, and temperature τ of our PairCFR on full CAD datasets in Table 7. In the few-shot setting, these hyperparameters are adjusted accordingly, shown in Table 8. All our experiments run on an A100 GPU (40G). Each implementation with a single seed can be completed in less than 1 hours.

Table 6: Datasets description. ♯ refers to ID datasets.

Dataset	Domain	#Test
Sentiment Analysis #class=2		
IMDb (Maas et al., 2011)♯	movie reviews	67k
Amazon (Ni et al., 2019)	service feedback	207k
Yelp (Zhang et al., 2015)	purchase reviews	38k
Twitter (Rosenthal et al., 2019)	social microblogs	10.3k
SST-2 (Socher et al., 2013)	movie reviews	1.82k
Natural Language Inference #class=3		
SNLI (Bowman et al., 2015)♯	written text	9.82k
MNLI-m (Williams et al., 2017)	mismatched genres	9.83k
MNLI-mm (Williams et al., 2017)	matched genres	9.82k
Negation (Naik et al., 2018)	strong negation	9.83k
Spelling-e (Naik et al., 2018)	spelling errors	9.14k
Word-o (Naik et al., 2018)	large word-overlap	9.83k

Table 7: Hyperparameters of our PairCFR on full CAD training data.

Model	lr	bts	λ	τ
Sentiment Analysis				
$\text{BERT}_{\text{base}}$	3e-5	16	0.7	0.3
$\text{RoBERTa}_{\text{base}}$	3e-6	16	0.9	0.07
SBERT	5e-6	16	0.7	0.7
T5_{base}	1e-4	16	0.8	0.07
Natural Language Inference				
$\text{BERT}_{\text{base}}$	3e-5	30	0.4	0.7
$\text{RoBERTa}_{\text{base}}$	1e-5	30	0.3	0.8
SBERT	5e-5	30	0.2	0.9
T5_{base}	1e-4	30	0.4	0.7

Table 8: Hyperparameters of our PairCFR in few shot setting.

Model	#Train	lr	bts
Sentiment Analysis			
$\text{BERT}_{\text{base}}$	32	1e-4	4
	64	1e-5	8
	128	1e-5	8
	512	1e-5	16
	1024	1e-5	16
Natural Language Inference			
$\text{BERT}_{\text{base}}$	50	1e-5	5
	100	1e-5	5
	500	1e-5	10
	1k	1e-5	10
	4k	1e-5	20

C.4 Hyperparemeter analysis: λ and τ

In this study, we investigate the influence of trade-off factor λ and temperature τ on model generalization. Specifically, we incrementally increase λ or τ from 0.1 to 0.9 by 0.1 and fix other hyperparameters determined by grid search. The exper-

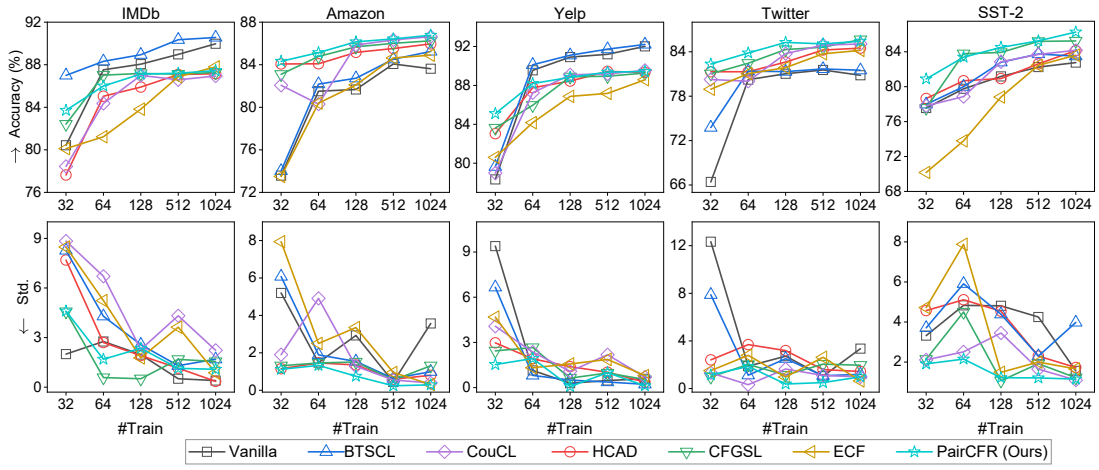

Figure 4: Few-shot learning results of $\text{BERT}_{\text{base}}$ on SA. x-axis represents the number of training samples and y-axis represents the averaged accuracy and standard deviation on ID and OODs.

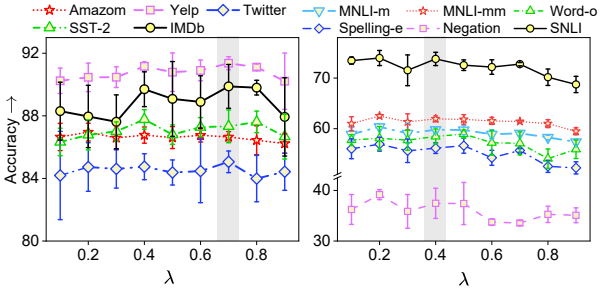

(a) The impact of trade-off term λ in CL. We keep $\tau = 0.3$ for SA (left) and $\tau = 0.7$ for NLI (right), respect to increase the λ.

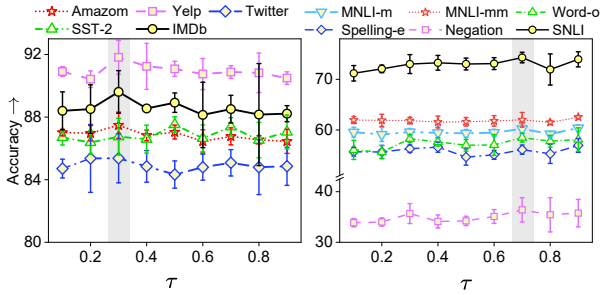

(b) The impact of temperature τ in CL. We keep $\lambda = 0.7$ for SA (left) and $\lambda = 0.4$ for NLI (right), respect to increase the τ.

Figure 5: The performance the BERT_{base} trained on IMDb and SNLI augmented data over ID and OODs. Grey is the overall optimal point.

imental results on ID and OODs are reported in Figure 5. We observe that with λ or τ increasing from 0.1, the model performance initially increases and then declines. In SA, the model perform better for a larger λ and a lower temperature 0.3 (i.e., $\lambda = 0.7, \tau = 0.3$), while in NLI, a larger temperature and smaller λ is favored (i.e., $\lambda = 0.4, \tau = 0.7$). We hypothesize that in SA, the model may overly depend on perturbed words for predictions, as revision patterns are relatively smaller than in NLI. Therefore, we incorporate a smaller temperature λ to enforce uniform distribution in representations and introduce a higher regularization from contrastive learning. More insights will be explored in future work.

C.5 Few-shot learning on NLI

Here we show the few-shot learning result of BERT model on SA task with progressively increasing IMDb augmented data from 32 to 1024. Similar to the trend of few-shot learning in the NLI task, discussed in Section 5.2, our approach exhibits notable significant performance even with limited data in the SA task.