

Topic-Aware Causal Intervention for Counterfactual Detection

Thong Nguyen

National University of Singapore
e0998147@u.nus.edu

Truc-My Nguyen

Ho Chi Minh city University of Technology
nguyenmy3399@gmail.com

Abstract

Counterfactual statements, which describe events that did not or cannot take place, are beneficial to numerous NLP applications. Hence, we consider the problem of counterfactual detection (CFD) and seek to enhance the CFD models. Previous models are reliant on clue phrases to predict counterfactuality, so they suffer from significant performance drop when clue phrase hints do not exist during testing. Moreover, these models tend to predict non-counterfactuals over counterfactuals. To address these issues, we propose to integrate neural topic model into the CFD model to capture the global semantics of the input statement. We continue to causally intervene the hidden representations of the CFD model to balance the effect of the class labels. Extensive experiments show that our approach outperforms previous state-of-the-art CFD and bias-resolving methods in both the CFD and other bias-sensitive tasks.

1 Introduction

Counterfactual statements describe an event that may not, did not, or cannot occur, and the consequence(s) that did not occur as well (O’Neill et al., 2021). For example, consider the statement — *I would purchase this physics book, but I really want that my brain has a tiny amount of interest in science!*. We can partition the statement into two components: a component about the event (*my brain has a tiny amount of interest in science*) as the antecedent, and the consequence of the event (*I would purchase this physics book*) as the consequent. Both the antecedent and the consequent did not take place (*neither the speaker has purchased the book nor he is interested in science*). Accurate detection of such counterfactual statements is beneficial to various NLP applications, such as in social media or psychology. In social media, counterfactual detection (CFD) can be helpful by eliminating irrelevant content (O’Neill et al., 2021).

Scenarios	Examples	mBERT Predictions	Labels
Clue phrase Anomaly	It doesn’t work as well as I was hoping it would, it is a waste of money.	Negative	Positive
	I don’t like to go into the plot a lot. The blurb represents the book fairly.	Negative	
	Who would have thought a pillow could make such a difference.	Positive	Negative
	The girlfriend was annoying, and it made me wonder if any man in his right mind would have put up with her behavior as long as he did.	Positive	
Cross-lingual input	It would have been, people would say, worse than Watergate.	Positive	Positive
	ウォーターゲート事件よりもひどかったかもしれない、と人々は言うだろう。	Negative	

Table 1: Examples of counterfactual detection from the Amazon-2021 dataset. We denote mBERT predictions of *positive* (counterfactual) and *negative* (non-counterfactual) classes.

For example, in the previous statement, we should not return science or physics content because the user is not interested. Detecting counterfactuality can also give useful features to perform psychology assessment of huge populations (Son et al., 2017).

Previous development of monolingual and multilingual CFD methods depend on extensive labelled datasets (O’Neill et al., 2021). However, in CFD datasets, the percentage of counterfactual examples is heavily low, even approaching 1 – 2% (Son et al., 2017). This class imbalance has two weaknesses. First, because counterfactual hints are so limited for the CFD model to learn, it tends to rely on clue phrases, e.g. *if*, *I wish*, etc., to detect counterfactuality. When the existence of such clue phrases does not correlate with the counterfactuality, the model might be led to false predictions. As illustrated in Table 1, the mBERT baseline predicts all incorrect classes for both counterfactual examples, which do not include clue phrases, and non-counterfactual ones, which include clue phrases. Moreover, the performance might substantially drop if the model is tested upon a language different from the training language. As shown in Table 1, the multilin-




Figure 1: For each topic, we count the percentage of inputs in which the topic has the largest probability in the topic representation. Topic 1, 2, and 9 refer to three top topics of the input document, in descending order of probability.

gual mBERT predicts the correct class for the English statement, but misclassifies the Japanese one of similar meaning. Second, the class imbalance causes the CFD model to bias towards the non-counterfactual class over the counterfactual one, thus resulting in sub-optimal performance.

To address the first issue, we propose to incorporate neural topic model (NTM) into the CFD module. Particularly, we aim to approximate the global semantics of the input statement learned from the posterior distribution of the NTM. The posterior distribution generates the global semantics in terms of the topic representation to guide the CFD model towards semantics of the input instead of the clue phrases. However, a challenge exists that the NTM tends to repetitively assign large weights to a certain small group of topics. In Figure 1, even though the input statement is about *ear buds* and *iPhone*, the NTM still infers it to be highly related to *stories*, *book*, and *reviews*. To cope with this challenge, we propose to adapt backdoor adjustment that adjusts the behavior of neural topic model to make it consider all topics fairly. To the best of our knowledge, no study has explored the benefit of backdoor adjusted NTM for counterfactual detection.

To address the second issue, we view the CFD problem from a causal perspective. Our perspective gives rise to a causal graph where the class imbalance plays a confounder role in influencing hidden representations of the input statement. Based on the graph, we propose to perform causal intervention on these representations to remove the confounding effect of the imbalance phenomenon and enhance the model prediction.

To sum up, our contributions are as follows:

- We propose a novel neural topic model equipped with the backdoor adjustment to produce effective topic representations for benefiting counterfactual detection.
- We propose causal intervention upon hidden representations to ameliorate the confounding effect of the class imbalance in counterfactual detection datasets.
- Extensive experiments demonstrate that we significantly outperform state-of-the-art CFD and bias-resolving approaches. Our method is also applicable to other bias-sensitive natural language understanding tasks.

2 Related Work

2.1 Representational Intervention for Deep Learning

Representational intervention has been popularly adopted in deep learning applications. Some include document summarization (Nguyen et al., 2021; Nguyen and Luu, 2022), topic modeling (Wu et al., 2024c, 2023a, 2024d, 2023b, 2024a), document ranking (Nguyen et al., 2023c, 2022), sentiment analysis (Nguyen et al., 2023b, 2024a, 2023a), video moment retrieval (Nguyen et al., 2023d, 2024d), and video question answering (Nguyen et al., 2024c,e). As one notable approach for representational intervention, causal inference has attracted myriad attention as a method to interpret adversarial attacks (Zhao et al., 2022) and eradicate spurious confounding factors in SGD optimizer (Tang et al., 2020).

2.2 Predictive Biases in Deep Learning

Research community has long searched for objective-based and augmentation-based countermeasures against biases that drive deep learning models to ignore the input content when making predictions (Wu et al., 2024b; Nguyen et al., 2024b; Nguyen and Luu, 2021; Nguyen et al., 2024f). For the objective-based direction, Karimi Mahabadi et al. (2020) propose to increase the loss weight of rare examples and subtract the gradients of the biased model from the main one to mitigate their spurious influences. In the second direction, Wang et al. (2022) perturb words to prevent the confounding effect of language bias. Wang and Culotta (2021) suggest augmenting the original training

set with samples containing antonyms of high coefficient terms and reverse label. However, their method demands human supervision and solely involves sentiment classification. Focusing on Counterfactual Detection, O’Neill et al. (2021) decide to mask clue phrases and populate counterfactual examples through backtranslation. Nevertheless, they find that these methods suffer from deficiency since counterfactuality also depends on the context. Contrast to them, we decide to causally intervene the contextualized representations to reduce the confounding effect of the biases.

3 Methodology

In this section, we sequentially formulate the preliminaries of counterfactual detection and neural topic model, introduce our proposed causal perspective for the task, and then articulate the implementation details of our framework.

3.1 Problem Formulation

Given an input sentence $S = \{w_1, w_2, \dots, w_N\}$ and its bag-of-words (BOW) representation \mathbf{x}_{bow} , we aim to train a model function f that maps S and \mathbf{x}_{bow} to a probability scalar $y \in [0, 1]$. The probability magnitude will denote whether the input sentence is counterfactual or not.

3.2 Neural Topic Model (NTM)

Our neural topic model possesses the Variational AutoEncoder architecture (Miao et al., 2017; Kingma and Welling, 2013). It consists of an encoder to produce topic representation and a decoder to reconstruct the original input based upon the representation.

Topic Encoder. Its function is to encode the input \mathbf{x}_{bow} into the topic representation θ . In the beginning, \mathbf{x}_{bow} is forwarded to both non-linear and linear layers to estimate the mean μ and standard deviation σ of the variational distribution $q(\mathbf{z}|\mathbf{x})$:

$$\pi = f_0(\mathbf{x}_{\text{bow}}), \quad \mu = f_\mu(\pi), \quad \log \sigma = f_\sigma(\pi), \quad (1)$$

where we implement f_0 as a non-linear layer with the softplus activation function; f_μ and f_σ are two linear layers. Subsequently, to lessen the gradient variance, we adapt the reparameterization trick (Kingma and Welling, 2013) to draw the latent vector \mathbf{z} :

$$\mathbf{z} = \mu + \sigma \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2)$$

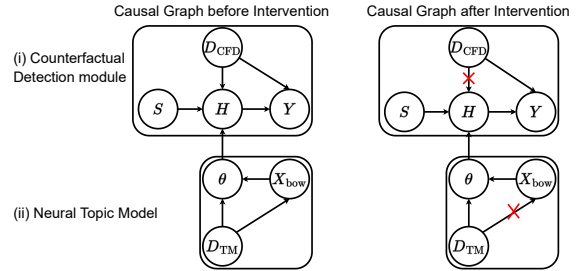


Figure 2: (left) Our proposed causal model for counterfactual detection. (right) The causal graph after removing arrows from D_{CFD} to H and D_{TM} to X_{bow} , eliminating spurious effects of the label and topic biases.

Then, we normalize \mathbf{z} with the softmax function to attain the topic representation θ as:

$$\theta = \text{softmax}(\mathbf{z}). \quad (3)$$

Topic Decoder. Given the topic representation θ , the decoder’s task is to reconstruct the original input \mathbf{x}_{bow} as \mathbf{x}'_{bow} . It performs the sampling process to extract the word distribution:

- For each word $w \in \mathbf{x}_{\text{bow}}$, draw $w \sim \text{softmax}(f_\phi(\theta))$,

where f_ϕ denotes a ReLU-activated non-linear transformation. In the ensuing sections, we designate the weight matrix $\phi = (\phi_1, \phi_2, \dots, \phi_K) \in \mathbb{R}^{V \times K}$ of f_ϕ as the topic-word distribution, in which V and K denote the vocabulary size and the number of topics, respectively. We also leverage the topic representation θ as global semantics to enhance the counterfactual detection module.

3.3 Causal Perspective into Counterfactual Detection

To investigate the relation among factors in our counterfactual detection system, we propose a structural causal graph (SCG) in Figure 2. Our graph includes vertices, direct edges, and two sub-graphs for denoting random variables, causal effects, with the pre-intervened and post-intervened states, respectively.

SCG for Topic Modeling. In this component, the topic bias D_{TM} is the confounder that influences variables θ and X_{bow} in the neural topic model.

- $X_{\text{bow}} \leftarrow D_{\text{TM}} \rightarrow \theta$: This backdoor path elicits the spurious correlation between \mathbf{x}_{bow} and θ instances. In topic modeling, neural topic models have a tendency to align documents with a repetitive set of topics.

- $\theta \rightarrow H$: Because of the confounder D_{TM} , the inferred global semantics might comprise irrelevant entries that do not represent the true semantics of the document. Therefore, the fallacious semantics could become detrimental noise to adulterate the hidden representations \mathbf{h} , which is the direct input to the counterfactual classifier.

SCG for Counterfactual Detection. This component delineates causalities among four variables in counterfactual detection: input sequence S , encoded content H , output prediction Y , and the imbalanced label bias D_{CFD} . In detail, the imbalanced label distribution confounds both the predicted output Y and variable H , leading to erroneous correlation between H and Y .

- $H \leftarrow D_{\text{CFD}} \rightarrow Y$ specifies the effect of D_{CFD} on hidden representations. In practice, the overwhelming population of the negative label in counterfactual datasets might result in learned representations that mostly express non-counterfactual features, thus driving the detection model towards the non-counterfactual response during prediction.

Causal Intervention on Textual Representations. We now present the method to remove the confounding effects. To obtain the deconfounded representations, we capture the causal effect from X_{bow} to θ and from H to Y via the Causal Intervention technique, i.e. Backdoor Adjustment (Pearl, 2009), with the following theorem to remove the arrow from D_{TM} to X_{bow} and D_{CFD} to H .

Theorem 1. (Backdoor Adjustment (Pearl, 2009)) Let $o \in \{y, \theta\}$, $i \in \{\mathbf{x}_{\text{bow}}, \mathbf{h}\}$, and $n \in \{d_{\text{TM}}, d_{\text{CFD}}\}$. Then,

$$p(o|\text{do}(i)) = p^{N \rightarrow I}(o|\text{do}(i)) = \sum_n p(o|i, n) \cdot p(n). \quad (4)$$

This theorem shows that we can model the deconfounded likelihood $p(o|\text{do}(i))$ through estimating $p^{N \rightarrow I}(o|i, n)$ and $p(n)$. We will expound the implementation of $p(o|\text{do}(i))$ to assist the model in predicting counterfactuality in Section 3.4 and deconfound neural topic model in Section 3.5.

3.4 Model Implementation

Our overall framework is illustrated in Figure 3. **NTM for Text Encoder.** To address the issue of model reliance on clue phrases in counterfactual

detection, we propose to condition contextualized representations on global semantics yielded by the neural topic model.

Initially, we append the special token [CLS] to the beginning of the input sequence. Then, the text encoder converts each discrete token w_i into the hidden vector \mathbf{h}_i as:

$$\mathbf{h}_{[\text{CLS}]}, \mathbf{h}_1, \dots, \mathbf{h}_{|S|} = \text{TextEncoder}([\text{CLS}], w_1, \dots, w_{|S|}). \quad (5)$$

We insert global semantics θ of the input S into the encoded information:

$$\mathbf{h}_i^{\text{topic}} = \tanh(\text{Linear}([\mathbf{h}_i, \theta])), \quad (6)$$

where $[\cdot]$ denotes the concatenation operator.

Causal Intervention for Predicting Counterfactuality. As mentioned in Section 3.3, we propose to debias hidden vectors from the imbalanced label bias. To this end, we set $p(d_{\text{CFD}}) = \frac{1}{|\mathcal{Y}|}$, where \mathcal{Y} is the set of groundtruth labels. Formally, the Eq. (4) becomes:

$$p(y|\text{do}(c)) = \frac{1}{|\mathcal{Y}|} \sum_{d_{\text{CFD}}} p(y|c, d_{\text{CFD}}). \quad (7)$$

Because of d_{CFD} , we need to incorporate the label information into $p(y|c, d_{\text{CFD}})$. We propose that information of each label exists in the hidden vectors of the inputs belonging to that label and denote such set of inputs for each label l as \mathcal{D}_l . Inspired by the prototypical network (Snell et al., 2017), we extract the information as follows:

$$\mathbf{h}_{[\text{CLS}]}^l = \frac{1}{|\mathcal{D}_l|} \sum_{S_j \in \mathcal{D}_l} \mathbf{h}_{j, [\text{CLS}]}. \quad (8)$$

Hereafter, we forward both the topic-oriented representation $\mathbf{h}_i^{\text{topic}}$ and the label information in $\mathbf{h}_{[\text{CLS}]}^l$ to the non-linear layer to classify the counterfactuality as:

$$\mathbf{h}'_{[\text{CLS}]} = \text{Linear} \left(\left[\mathbf{h}_{[\text{CLS}]}^{\text{topic}}, \text{Linear} \left[\{\mathbf{h}_{[\text{CLS}]}^l\}_{l \in \mathcal{Y}} \right] \right] \right), \quad (9)$$

$$p_{\text{CFD}} = p(y|\text{do}(c)) = \frac{1}{|\mathcal{Y}|} \sum_l \varphi(\mathbf{h}'_{[\text{CLS}]}) \quad (10)$$

where φ denotes the sigmoid function.

3.5 Training Strategy

Deconfounding NTM. To deconfound the NTM, we derive the Eq. (4) as:

$$p(\theta|\text{do}(\mathbf{x}_{\text{bow}})) = \sum_{d_{\text{TM}}} \frac{p(\theta, \mathbf{x}_{\text{bow}}|d_{\text{TM}}) \cdot p(d_{\text{TM}})}{p(\mathbf{x}_{\text{bow}}|d_{\text{TM}})}. \quad (11)$$




Figure 3: Illustration of the Topic-aware Causal Intervention Framework for Counterfactual Detection. Here the green component denotes the neural topic model, the purple component the text encoder, and the orange component our causal intervention operation for counterfactuality prediction.

In NTMs, topics are parameterized as word distributions (Blum and Haghtalab, 2016; Austin, 2011), similar to \mathbf{x}_{bow} . Hence, we conjecture that topic representation is a decomposed variant of each \mathbf{x}_{bow} , and we can only fully observe the distribution of the decompositions as in Figure 1 with the same number of times we retrieve \mathbf{x}_{bow} . Furthermore, as the training progresses, the output \mathbf{x}'_{bow} will converge to \mathbf{x}_{bow} . As such, we propose to approximate Eq. (11) following the propensity score modeling approach (Rosenbaum and Rubin, 1983):

$$\begin{aligned}
 p(\theta | \text{do}(\mathbf{x}_{\text{bow}})) &\approx \sum_{d_{\text{TM}}} p(\theta, \mathbf{x}'_{\text{bow}}, d_{\text{TM}}) \\
 &= \prod_i \frac{\phi_i \cdot \theta}{\|\phi_i\| \cdot \|\theta\|},
 \end{aligned} \quad (12)$$

where i refers to a word in \mathbf{x}_{bow} , and we empirically add the magnitude of θ . The denominator works as a normalizer to balance the magnitude of the variables.

Training Objective. Our framework jointly optimizes the Neural Topic Model and Counterfactual Detection (CFD) module. To train the CFD module, we employ the binary cross-entropy loss as:

$$\begin{aligned}
 \mathcal{L}_{\text{CFD}}(S, \mathbf{x}_{\text{bow}}, y) &= \\
 &- y \log p_{\text{CFD}} - (1 - y) \log(1 - p_{\text{CFD}}).
 \end{aligned} \quad (13)$$

For the NTM, with the Eq. (12), we obtain the deconfounded evidence lower bound as:

$$\begin{aligned}
 \mathcal{L}_{\text{NTM}}(\mathbf{x}_{\text{bow}}) &= \text{KL}(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \\
 &- \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\log p_{\phi}(\mathbf{x}_{\text{bow}}|\theta)] \\
 &- \gamma \cdot \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\sum_{i=1}^V \log \frac{\phi_i \cdot \theta}{\|\phi_i\| \cdot \|\theta\|} \right],
 \end{aligned} \quad (14)$$

where the first term denotes the Kullback-Leibler divergence between the prior and posterior distribution, the second term the reconstruction error

of the output compared with the input, the third term the deconfounded objective in Eq. (12), V the vocabulary size, and γ the hyperparameter to control the deconfounding effect upon the training, respectively.

To conclude, our entire architecture is optimized with the linear combination of the loss functions \mathcal{L}_{NTM} and \mathcal{L}_{CFD} as:

$$\mathcal{L} = \mathcal{L}_{\text{CFD}} + \lambda_{\text{NTM}} \mathcal{L}_{\text{NTM}}, \quad (15)$$

where λ denotes the hyperparameter weight to scale the topic modeling component.

4 Experiments

4.1 Datasets

We evaluate on two prevalent datasets for the counterfactual detection task, SemEval-2020 (Yang et al., 2020) and Amazon-2021 (O’Neill et al., 2021). While SemEval-2020 comprises English documents, Amazon-2021 covers statements in three languages, English, Japanese, and German. For our experiments, we inherit the original train/val/test splits. To verify the generalizability of our methods, we measure our performance on two other bias-sensitive document analysis tasks, Paraphrase Identification with the MRPC dataset (Dolan and Brockett, 2005), and Implicit Sentiment Analysis (ISA) with CLIPeval from SemEval 2015 task 9 (Russo et al., 2015). These two tasks have been shown to sustain syntactic phrase and label biases (Li et al., 2020; Wang et al., 2022). The statistics of the datasets are provided in the Appendix. For evaluation metrics, we report Matthew’s Correlation Coefficient (MCC) (Boughorbel et al., 2017), the Accuracy (Acc), and F1 score.

4.2 Implementation Details

For the topic model, we select the topic number $T = 15$ based on the validation performance. Because at the beginning of the training process, the reconstructed output \mathbf{x}'_{bow} does not resemble the input \mathbf{x}_{bow} , we decide to adapt the linear warm-up strategy (Gilmer et al., 2021) with the number of warm-up steps $N_{\text{wp}} = 1000$ for the value of γ before fixing it to 0.25. We finetune two pretrained multilingual language models, mBERT (Devlin et al., 2018) and XLM-R (Conneau et al., 2020) for the CFD task, and the monolingual BERT (Devlin et al., 2018) and RoBERTa (Liu et al., 2019) for the PI and ISA tasks. All variants are equipped with a linear layer on top of the pretrained language model. Our entire architecture is trained end-to-end on the A100 GPU, with the batch size of 16 and λ_{NTM} of 0.5 for 50 epochs, adopting Adam optimizer for the learning rate of 10^{-5} and L2 regularization equal 10^{-6} . For the counterfactual detection and paraphrase identification tasks, $\mathcal{Y} = \{0, 1\}$, meanwhile for the implicit sentiment analysis task, $\mathcal{Y} = \{-1, 0, 1\}$.

4.3 Baselines

As baselines, we compare our work against a wide variety of recent state-of-the-art bias-resolving causal intervention and data augmentation approaches for Counterfactual Detection: (i) **Stochastic Perturbation (SP)** (Wang et al., 2022), leveraging word perturbation to causally intervene the spurious effect of the language bias confounder; (ii) **Masking** (O’Neill et al., 2021), masking clue phrases in counterfactual detection to eliminate their effect upon the training; (iii) **Debiased Focal Loss (DFL)** (Karimi Mahabadi et al., 2020), de-emphasizing the loss contribution of easy biased examples and direct the model towards hard but less biased ones; (iv) **Product of Experts** (Karimi Mahabadi et al., 2020), aggregating the predictions of two models, one trained with the biased and the other with both biased and unbiased examples; (v) **Backtranslation** (O’Neill et al., 2021), a data augmentation method on the input level to increase the number of rare-label samples.

4.4 Comparison with State-of-the-arts

Results on Original Test Sets. We train and test the baselines and our model on the original test sets in Table 2. In the English variant of Amazon-2021 dataset, with mBERT we achieve an improved

accuracy of 1.5 points, and MCC score of 2.4 points with XLM-R. For German documents, our XLM-R outperforms the baseline using Backtranslation with 1.5 points, while our method adopted on mBERT enhances the SP approach with 1.0 point in MCC. On the Japanese version, where the language upholds syntactic and morphological features separate from English and German, our mBERT-based and XLM-R-based models accomplish absolute enhancements of 5.1 and 1.2 points in F1 metric, respectively, compared with DFL and SP, which are the best previous approaches.

On the SemEval-2020 dataset, which is at a larger scale and concerns diverse domains (O’Neill et al., 2021), our general performance is also auspicious. In particular, our mBERT system surpasses the DFL model by a mean MCC of 2.1 points. In addition, our XLM-R polishes the SP approach by 1.1 points of the F1 score. These results corroborate that our counterfactual detection model is able to cope with harmful confounding impacts of different biases, thus producing more generalizable representations to attain better performance.

Results on Balanced Test Sets. We randomly sample 500 samples from each class in the test set of SemEval-2020, then evaluate our method in Table 3. As can be seen, our method surpasses the best previous baseline, i.e. Backtranslation, with a significant margin of 1.5 points of accuracy for the mBERT variant, and surpasses DFL with 2.8 points of MCC for the XLM-R variant. These results verify that our causal intervention technique is able to mitigate the confounding effect of the class imbalance and makes the CFD model impartially consider the counterfactual and non-counterfactual choices when making prediction.

4.5 Zero-Shot Cross-lingual Evaluation

To clearly confirm whether our methods have the ability to deal with the bias of clue phrases, we conduct the zero-shot cross-lingual evaluation. In particular, we proceed to finetune the standard mBERT, XLM-R, and our counterfactual detection architectures on the Japanese portion of Amazon-2021 dataset, then directly validate the performance on the English portion, and similarly we finetune the models on the German training set and test them on the Japanese test set. We indicate the results in Table 4 and 5.

As can be observed, our model is capable of enhancing zero-shot cross-lingual counterfactual

Methods	Amazon-2021 (CD)									SemEval-2020 (CD)			MRPC (PI)			CLIPeVal (ISA)		
	En			De			Jp			En			En					
	Acc	MCC	F1	Acc	MCC	F1	Acc	MCC	F1	Acc	MCC	F1	Acc	MCC	F1	Acc	MCC	F1
mBERT/BERT	91.79	72.29	79.19	90.79	77.02	93.00	92.93	60.87	58.93	94.39	68.87	71.83	91.75	83.86	91.35	83.10	73.79	80.67
w/ DFL	93.88	81.70	81.30	91.11	79.58	93.47	<u>94.00</u>	<u>66.61</u>	<u>69.89</u>	<u>96.63</u>	<u>81.23</u>	<u>82.80</u>	92.15	84.62	91.69	<u>85.25</u>	<u>77.22</u>	<u>83.76</u>
w/ PoE	<u>94.03</u>	82.72	81.54	90.90	78.53	93.32	93.79	66.23	69.74	95.33	77.66	80.00	92.23	85.61	91.73	84.18	76.16	82.37
w/ Backtranslation	<u>94.03</u>	<u>83.07</u>	<u>81.89</u>	90.26	73.51	92.47	93.25	62.50	60.27	95.94	79.18	80.99	92.30	85.95	91.75	83.10	74.22	82.33
w/ Masking	93.43	78.60	81.01	91.43	79.79	93.89	93.68	64.28	68.02	95.81	78.73	80.84	—	—	—	—	—	—
w/ SP	93.63	81.19	81.21	<u>91.65</u>	<u>81.06</u>	<u>93.86</u>	93.61	62.76	66.83	95.08	77.35	79.02	<u>92.38</u>	<u>86.53</u>	<u>91.76</u>	84.44	76.97	83.13
Our Model	95.52	86.37	83.05	92.29	82.08	94.40	95.29	73.79	75.00	96.97	83.31	84.81	93.65	87.38	93.20	86.79	79.28	85.00
XLM-R/roBERTa	92.63	73.16	82.89	90.55	80.18	93.37	92.96	64.70	67.25	94.43	83.68	85.05	91.25	88.46	92.75	88.16	81.31	85.78
w/ DFL	94.66	85.22	83.46	90.84	80.32	93.53	94.40	74.21	75.68	96.41	84.43	85.56	93.92	88.99	94.05	<u>88.41</u>	81.64	85.89
w/ PoE	94.52	84.90	83.09	90.86	80.58	93.74	94.12	73.23	75.59	96.43	84.55	85.87	93.50	88.91	93.75	87.87	80.80	84.99
w/ Backtranslation	94.95	85.85	83.81	<u>91.79</u>	<u>81.02</u>	<u>93.99</u>	94.33	73.75	75.67	95.86	84.36	85.43	93.34	88.78	93.62	87.60	80.48	84.86
w/ Masking	94.20	83.75	82.96	90.21	79.75	93.25	94.57	74.44	75.73	95.79	84.22	85.11	—	—	—	—	—	—
w/ SP	<u>95.40</u>	<u>86.35</u>	<u>83.87</u>	91.18	80.76	93.87	<u>94.87</u>	<u>75.03</u>	<u>76.85</u>	<u>96.77</u>	<u>85.18</u>	<u>86.34</u>	<u>94.60</u>	<u>89.22</u>	<u>94.17</u>	<u>88.41</u>	<u>81.69</u>	<u>87.03</u>
Our Model	96.85	88.74	84.44	92.51	82.49	94.57	95.82	76.01	77.97	97.44	86.09	87.44	95.55	91.05	95.14	89.49	83.55	88.07

Table 2: Numerical results on original test sets of the Counterfactual Detection (CD), Paraphrase Identification (PI), and Implicit Sentiment Analysis (ISA) tasks. We respectively bold and underline the best and second-to-best results.

Methods	SemEval-2020		
	Acc	MCC	F1
mBERT	85.75	74.21	83.55
w/ DFL	91.48	83.53	90.94
w/ PoE	90.45	82.01	89.60
w/ Backtranslation	91.68	83.21	91.54
w/ Masking	89.53	80.28	88.52
w/ SP	90.60	81.92	89.94
Our model	93.13	86.52	92.84
XLM-R	88.95	79.69	88.04
w/ DFL	92.78	86.01	92.38
w/ PoE	91.33	83.50	90.66
w/ Backtranslation	92.48	85.94	92.49
w/ Masking	90.60	82.15	89.83
w/ SP	89.40	80.13	88.25
Our model	94.33	88.83	94.14

Table 3: Numerical results on balanced test sets of the CFD task on the SemEval-2020 dataset.

Models	Jp → En			De → Jp		
	Acc	MCC	F1	Acc	MCC	F1
mBERT	91.83	49.78	52.81	80.84	40.54	44.92
Our Model	93.40	59.41	60.71	91.54	50.10	54.34

Table 4: Cross-lingual Zero-shot mBERT results on the Amazon-2021 dataset.

detection capacity of both mBERT and XLM-R, surpassing mBERT with a large margin of 1.6 and XLM-R with 1.1 points in accuracy for the English test set. For the Japanese test set, we outperform mBERT and XLM-R with 9.6 points of F1, and 9.7 points of MCC, respectively. These results substantiate that our methods can mitigate the clue phrase bias in the language models.

4.6 Adaptability to Other Bias-sensitive Tasks

Experimental results in PI and ISA tasks are given in Table 2. For the MRPC dataset, our BERT model performance exceeds the one of the SP method by 1.3 points in accuracy, and 1.4 points in F1. With respect to the RoBERTa backbone, we also surpass

Models	Jp → En			De → Jp		
	Acc	MCC	F1	Acc	MCC	F1
XLM-R	92.70	61.55	62.83	87.87	45.82	50.49
Our Model	93.85	62.79	64.35	89.19	55.56	51.67

Table 5: Cross-lingual Zero-shot XLM-R results on the Amazon-2021 dataset.

Methods	Acc	MCC	F1
Our Model	95.52	86.37	83.05
w/o Debaised CFD Objective	94.63	83.75	82.96
w/o Deconfounded Topic Model Objective	94.33	83.29	82.52
w/o Neural Topic Model	93.43	80.59	82.40

Table 6: Results from ablating different deconfounding components on the English Amazon-2021 dataset.

the SP method by 1.8 points of MCC, and 1.0 point of F1. Regarding the CLIPeVal dataset, integrating our approaches into BERT and RoBERTa extends the performance with 2.1 points in MCC, and 1.1 points in accuracy, respectively.

Those aforementioned results have shown that our methods have the capability of tackling biases in not only counterfactual detection but also other natural language understanding tasks.

4.7 Ablation Study

Effect of Deconfounding Components. In this ablation, we experiment with removing each component that helps the model deconfound. Particularly, we train and test the ablated mBERT on the English portion of the Amazon-2021 dataset. As shown in Table 6, solely employing one of the elements does enhance the overall counterfactual recognition, but being less effective than the joint approach. Without combining the deconfounding mechanisms, the model might not be able to cope with multiple biases.

Effect of Global Semantics. Here, we investigate the performance of our method when utiliz-

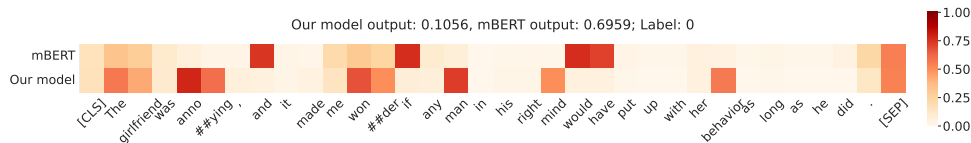


Figure 4: Attention weights of the [CLS] token to all other words, and output scores of mBERT and Our Model. The score is in range [0, 1]. The input: “The girlfriend was annoying, and it made me wonder if any man in his right mind would have put up with her behavior as long as he did.”

Methods	Acc	MCC	F1
XLM-R + NTM	92.51	82.49	94.57
XLM-R + PFA	91.86	81.09	94.09
XLM-R + LDA	91.97	81.83	94.25

Table 7: Ablation results with various types of global semantics on the German Amazon-2021 dataset.




Figure 5: Topic Percentages and inferred Top Topics from Figure 1 after Causal Intervention.

ing conventional topic models. We consider two choices, i.e. Poisson Factor Analysis (PFA) and Latent Dirichlet Allocation (LDA), and finetune the XML-R model on the German subset in Amazon-2021. As can be seen in Table 7, NTM burnishes the counterfactual detector more effectively than traditional topic models. More results on the two ablation experiments can be found in the Appendix.

4.8 Case Study

Impact of Causal Intervention on Attention

Logits. Here we randomly select one example from Table 1 and visualize the average attention scores in the heads of all layers of the [CLS] token to the remaining words. As shown in Figure 4, whereas mBERT’s [CLS] strongly pays attention to clue phrases “if” and “would have”, our model distributes the attention impartially and emphasizes content words, such as “annoying”, “man”, and “behavior”. This could help to explain our more reasonable prediction than mBERT. We provide attention visualizations of other examples in the Appendix. These visualizations demonstrate that our approach can resolve the confounding influence of clue phrases and improve model prediction.

Impact of Causal Intervention on Topic Distri-

bution. In Figure 5, we obtain topic representation from our neural topic model for the document of the Amazon-2021 dataset in Figure 1, and then count the percentage of documents sharing the top topic, i.e. possessing the largest likelihood. Different from Figure 1, our deconfounded topic model does not lean towards a subset of topics to assign top probabilities. Moreover, all three leading topics reveal the semantics of the document, which concerns *headset*, *shopping*, and *phone*. These results demonstrate that our approach is capable of resolving the topic bias phenomenon to produce faithful global semantics for counterfactual detection.

5 Conclusion

In this paper, we propose a causal intervention framework that discovers biases in the counterfactual detection problem. In order to cope with clue phrase, topic, and label biases, we propose to utilize global semantics and extend the training strategy with deconfounding training objectives. Comprehensive experiments demonstrate that our model can ameliorate detrimental influences of biases to polish previous state-of-the-art baselines for not only the counterfactual detection but also other bias-sensitive NLU tasks.

6 Limitations

We consider the following two limitations as our future work: (1) Extend the problem to circumstances with multiple observable confounding variables. The problem will become more complex if additional confounding factors are explicitly taken into account. Studying such complex scenario is potential to enhance the applicability and our understanding towards the proposed debiasing technique; (2) Explore the impact of causal intervention on generative tasks. We have only verified the effectiveness of causal intervention in discriminative language models. Whether the effectiveness applies for generative tasks such as machine translation, document summarization, etc., remains an open problem and interesting research direction.

References

- Peter C Austin. 2011. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*, 46(3):399–424.
- Avrim Blum and Nika Haghtalab. 2016. Generalized topic modeling. *arXiv preprint arXiv:1611.01259*.
- Sabri Boughorbel, Fethi Jarray, and Mohammed El-Anbari. 2017. Optimal classifier for imbalanced data using matthews correlation coefficient metric. *PLoS one*, 12(6):e0177678.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *ACL*.
- 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*.
- Bill Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Third International Workshop on Paraphrasing (IWP2005)*.
- Justin Gilmer, Behrooz Ghorbani, Ankush Garg, Sneha Kudugunta, Behnam Neyshabur, David Cardoze, George Dahl, Zachary Nado, and Orhan Firat. 2021. A loss curvature perspective on training instability in deep learning. *arXiv preprint arXiv:2110.04369*.
- Rabeeh Karimi Mahabadi, James Henderson, et al. 2020. End-to-end bias mitigation by modelling biases in corpora. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, CONF*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Xiaoya Li, Xiaofei Sun, Yuxian Meng, Junjun Liang, Fei Wu, and Jiwei Li. 2020. Dice loss for data-imbalanced nlp tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 465–476.
- 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*.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *International Conference on Machine Learning*, pages 2410–2419. PMLR.
- Cong-Duy Nguyen, Thong Nguyen, Duc Vu, and Anh Luu. 2023a. Improving multimodal sentiment analysis: Supervised angular margin-based contrastive learning for enhanced fusion representation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14714–14724.
- Cong-Duy Nguyen, Thong Nguyen, Xiaobao Wu, and Anh Tuan Luu. 2024a. Kdmcse: Knowledge distillation multimodal sentence embeddings with adaptive angular margin contrastive learning. *arXiv preprint arXiv:2403.17486*.
- Cong-Duy Nguyen, The-Anh Vu-Le, Thong Nguyen, Tho Quan, and Anh-Tuan Luu. 2023b. Expand bert representation with visual information via grounded language learning with multimodal partial alignment. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5665–5673.
- Thong Nguyen, Yi Bin, Xiaobao Wu, Xinshuai Dong, Zhiyuan Hu, Khoi Le, Cong-Duy Nguyen, See-Kiong Ng, and Luu Anh Tuan. 2024b. Meta-optimized angular margin contrastive framework for video-language representation learning. *arXiv preprint arXiv:2407.03788*.
- Thong Nguyen, Yi Bin, Junbin Xiao, Leigang Qu, Yicong Li, Jay Zhangjie Wu, Cong-Duy Nguyen, See-Kiong Ng, and Luu Anh Tuan. 2024c. Video-language understanding: A survey from model architecture, model training, and data perspectives. *arXiv preprint arXiv:2406.05615*.
- Thong Nguyen and Anh Tuan Luu. 2021. Contrastive learning for neural topic model. *Advances in neural information processing systems*, 34:11974–11986.
- Thong Nguyen, Anh Tuan Luu, Truc Lu, and Tho Quan. 2021. Enriching and controlling global semantics for text summarization. *arXiv preprint arXiv:2109.10616*.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Khoi M Le, Zhiyuan Hu, Cong-Duy Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024d. Read-pvla: Recurrent adapter with partial video-language alignment for parameter-efficient transfer learning in low-resource video-language modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18824–18832.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Anh Tuan Luu, Cong-Duy Nguyen, Zhen Hai, and Lidong Bing. 2023c. Gradient-boosted decision tree for listwise context model in multimodal review helpfulness prediction. *arXiv preprint arXiv:2305.12678*.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy Nguyen, See Kiong Ng, and Anh Luu. 2023d. Demaformer: Damped exponential moving average transformer with energy-based modeling for temporal language grounding. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3635–3649.

- Thong Nguyen, Xiaobao Wu, Anh-Tuan Luu, Cong-Duy Nguyen, Zhen Hai, and Lidong Bing. 2022. Adaptive contrastive learning on multimodal transformer for review helpfulness predictions. *arXiv preprint arXiv:2211.03524*.
- Thong Thanh Nguyen, Zhiyuan Hu, Xiaobao Wu, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024e. Encoding and controlling global semantics for long-form video question answering. *arXiv preprint arXiv:2405.19723*.
- Thong Thanh Nguyen and Anh Tuan Luu. 2022. Improving neural cross-lingual abstractive summarization via employing optimal transport distance for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11103–11111.
- Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024f. Topic modeling as multi-objective optimization with setwise contrastive learning. In *The Twelfth International Conference on Learning Representations*.
- James O’Neill, Polina Rozenshtein, Ryuichi Kiryo, Motoko Kubota, and Danushka Bollegala. 2021. I wish i would have loved this one, but i didn’t—a multilingual dataset for counterfactual detection in product review. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7092–7108.
- Judea Pearl. 2009. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146.
- Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.
- Irene Russo, Tommaso Caselli, and Carlo Strapparava. 2015. Semeval-2015 task 9: Clipeval implicit polarity of events. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 443–450.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Youngseo Son, Anneke Buffone, Joe Raso, Allegra Larche, Anthony Janocko, Kevin Zembroski, H Andrew Schwartz, and Lyle Ungar. 2017. Recognizing counterfactual thinking in social media texts. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 654–658.
- Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. 2020. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *Advances in Neural Information Processing Systems*, 33:1513–1524.
- Siyin Wang, Jie Zhou, Changzhi Sun, Junjie Ye, Tao Gui, Qi Zhang, and Xuanjing Huang. 2022. Causal intervention improves implicit sentiment analysis. *arXiv preprint arXiv:2208.09329*.
- Zhao Wang and Aron Culotta. 2021. Robustness to spurious correlations in text classification via automatically generated counterfactuals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14024–14031.
- Xiaobao Wu, Xinshuai Dong, Thong Nguyen, Chaoqun Liu, Liang-Ming Pan, and Anh Tuan Luu. 2023a. Infotm: A mutual information maximization perspective of cross-lingual topic modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13763–13771.
- Xiaobao Wu, Xinshuai Dong, Thong Thanh Nguyen, and Anh Tuan Luu. 2023b. Effective neural topic modeling with embedding clustering regularization. In *International Conference on Machine Learning*, pages 37335–37357. PMLR.
- Xiaobao Wu, Xinshuai Dong, Liangming Pan, Thong Nguyen, and Anh Tuan Luu. 2024a. Modeling dynamic topics in chain-free fashion by evolution-tracking contrastive learning and unassociated word exclusion. *arXiv preprint arXiv:2405.17957*.
- Xiaobao Wu, Thong Nguyen, and Anh Tuan Luu. 2024b. A survey on neural topic models: methods, applications, and challenges. *Artificial Intelligence Review*, 57(2):18.
- Xiaobao Wu, Thong Nguyen, Delvin Ce Zhang, William Yang Wang, and Anh Tuan Luu. 2024c. Fastopic: A fast, adaptive, stable, and transferable topic modeling paradigm. *arXiv preprint arXiv:2405.17978*.
- Xiaobao Wu, Fengjun Pan, Thong Nguyen, Yichao Feng, Chaoqun Liu, Cong-Duy Nguyen, and Anh Tuan Luu. 2024d. On the affinity, rationality, and diversity of hierarchical topic modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19261–19269.
- Xiaoyu Yang, Stephen Obadinma, Huasha Zhao, Qiong Zhang, Stan Matwin, and Xiaodan Zhu. 2020. Semeval-2020 task 5: Counterfactual recognition. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 322–335.
- Haiteng Zhao, Chang Ma, Xinshuai Dong, Anh Tuan Luu, Zhi-Hong Deng, and Hanwang Zhang. 2022. Certified robustness against natural language attacks by causal intervention. *arXiv preprint arXiv:2205.12331*.

A Attention Visualization

In this section, we visualize the attention weights of the [CLS] token to the words of the examples in Table 1.




Figure 6: Attention weights of the [CLS] token to all other words, and output scores of mBERT and our model. The score is in range [0, 1]. The input: “It doesn’t work as well as I was hoping it would, it is a waste of money.”




Figure 7: Attention weights of the [CLS] token to all other words, and output scores of mBERT and our model. The score is in range [0, 1]. The input: “I don’t like to go into the plot a lot. The blurb represents the book fairly.”




Figure 8: Attention weights of the [CLS] token to all other words, and output scores of mBERT and our model. The score is in range [0, 1]. The input: “Who would have thought a pillow could make such a difference.”




Figure 9: Attention weights of the [CLS] token to all other words, and output scores of mBERT and our model. The score is in range [0, 1]. The input: “It would have been, people would say, worse than Watergate.”

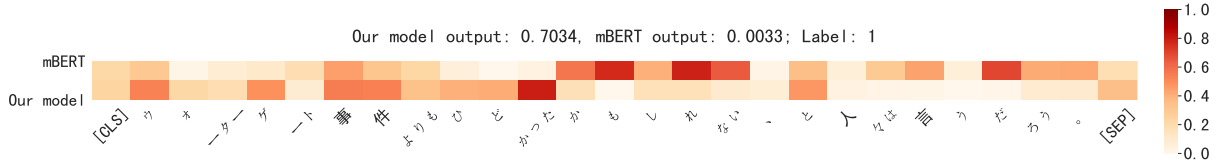


Figure 10: Attention weights of the [CLS] token to all other words, and output scores of mBERT and our model. The score is in range [0, 1]. The input: “ウォーターゲート事件よりもひどかったかもしれない、と人々は言うだろう。”

B Dataset Statistics

In this section, we present the statistics of all the datasets pertaining to Counterfactual Detection, Paraphrase Identification, and Implicit Sentiment Analysis tasks.

Dataset	Variant	Train	Val	Test	#Pos	#Neg	#Neutral	#Pos in Test	#Neg in Test	Total
Amazon-2021	En	4018	335	670	954	4069	-	131	539	5023
	De	5600	466	934	4840	2160	-	650	284	7000
	Jp	5600	466	934	667	6333	-	96	838	7000
SemEval-2020	En	13000	-	7000	2192	17808	-	738	6262	20000
MRPC	En	49184	2000	2000	23493	29691	-	907	1093	53184
CLIPeval	En	1347	-	371	580	796	342	216	155	1718

Table 8: Statistics of the Datasets.

C Additional Ablation Studies

Impact of Deconfounding Components. We compare our model with its ablated variants in all subsets of the Amazon-2021 dataset. As can be observed in Table 9, jointly utilizing deconfounded neural topic model and debiased objective can tackle the clue phrase, label, and topic biases, leading to the largest overall improvement.

Methods	En			De			Jp		
	Acc	MCC	F1	Acc	MCC	F1	Acc	MCC	F1
Our Model	95.52	86.37	83.05	92.29	82.08	94.40	95.29	73.79	75.00
w/o Debiased CFD objective	94.63	83.75	82.96	92.15	81.49	94.36	94.87	73.13	73.61
w/o Deconfounded Topic Model objective	94.33	83.29	82.52	91.94	81.31	94.19	94.80	72.78	71.66
w/o Neural Topic Model	93.43	80.59	82.40	91.76	81.20	94.05	94.72	72.60	71.09

Table 9: Results of subsequently pruning deconfounding components on the Amazon-2021 dataset.

Impact of Global Semantics. In addition to the results in Table 7, we execute our model with different topic models on other languages of the Amazon-2021 dataset. The results are shown in Table 10.

Methods	En			De			Jp		
	Acc	MCC	F1	Acc	MCC	F1	Acc	MCC	F1
XLM-R + NTM	96.85	88.74	84.44	92.51	82.49	94.57	95.82	76.01	77.97
XLM-R + PFA	96.63	87.19	83.97	91.86	81.09	94.09	95.72	75.46	77.68
XLM-R + LDA	96.40	87.11	83.10	91.97	81.83	94.09	95.18	75.07	77.53

Table 10: Results with heterogeneous topic models on the Amazon-2021 dataset.