

Uncovering Main Causalities for Long-tailed Information Extraction

Guoshun Nan^{1*}, Jiaqi Zeng^{2*}, Rui Qiao^{3*}, Zhijiang Guo⁴ and Wei Lu¹

¹StatNLP Research Group, Singapore University of Technology and Design

²Carnegie Mellon University ³National University of Singapore ⁴University of Cambridge

nanguoshun@gmail.com, jiaqizen@andrew.cmu.edu

qiaorui@comp.nus.edu.sg, zg283@cam.ac.uk

luwei@sutd.edu.sg

Abstract

Information Extraction (IE) aims to extract structural information from unstructured texts. In practice, long-tailed distributions caused by the selection bias of a dataset, may lead to incorrect correlations, also known as spurious correlations, between entities and labels in the conventional likelihood models. This motivates us to propose counterfactual IE (CFIE), a novel framework that aims to uncover the main causalities behind data in the view of causal inference. Specifically, 1) we first introduce a unified structural causal model (SCM) for various IE tasks, describing the relationships among variables; 2) with our SCM, we then generate counterfactuals based on an explicit language structure to better calculate the direct causal effect during the inference stage; 3) we further propose a novel debiasing approach to yield more robust predictions. Experiments on three IE tasks across five public datasets show the effectiveness of our CFIE model in mitigating the spurious correlation issues.

1 Introduction

The goal of Information Extraction (IE) is to detect the structured information from unstructured texts. Previous deep learning models for IE tasks, such as named entity recognition (NER; Lample et al. 2016), relation extraction (RE; Peng et al. 2017) and event detection (ED; Nguyen and Grishman 2015), are largely proposed for learning under some reasonably balanced label distributions. However, in practice, these labels usually follow a long-tailed distribution (Doddington et al., 2004). Figure 1 shows such an unbalanced distribution on the ACE2005 (Doddington et al., 2004) dataset. As a result, performance on the instance-scarce (tail) classes may drop significantly. For example, on an existing model for NER (Jie and Lu, 2019), the macro F1 score of instance-rich (head) classes




Figure 1: Class distribution of the ACE2005 dataset.

can be 71.6, while the score of tail classes sharply decreases to 41.7.

The underlying causes for the above issue are the biased statistical dependencies between entities¹ and classes, known as spurious correlations (Srivastava et al., 2020). For example, an entity *Gardens* appears 13 times in the training set of OntoNotes5.0, with the NER tag location *LOC*, and only 2 times as organization *ORG*. A classifier trained on this dataset tends to build spurious correlations between *Gardens* and *LOC*, although *Gardens* itself does not indicate a location. Most existing works on addressing spurious correlations focus on images, such as re-balanced training (Lin et al., 2017), transfer learning (Liu et al., 2019) and decoupling (Kang et al., 2019). However, these approaches may not be suitable for natural language inputs. Recent efforts on information extraction (Han et al., 2018; Zhang et al., 2019) incorporate prior knowledge, which requires data-specific designs.

Causal inference (Pearl et al., 2016) is promising in tackling the above spurious correlation issues caused by unbalanced data distribution. Along this line, various causal models have been proposed for

¹For the NER task, a model may also build the spurious correlations between the part-of-speech (POS) tags of entities and class labels. For RE and ED tasks, a model may also learn incorrect correlations between features like NER tags and labels. We consider all above issues in our proposed causal diagram.

*Contributed equally.

visual tasks (Abbasnejad et al., 2020; Tang et al., 2020b). Despite their success, these methods may be unsatisfactory on textual inputs. Unlike images, which can be easily disentangled with detection or segmentation methods for causal manipulation, texts rely more on the context involving complex syntactic and semantic structures. Hence it is impractical to apply the methods used in images to disentangle tokens’ representations. Recent causal models (Zeng et al., 2020; Wang and Culotta, 2020, 2021) on text classification eliminate biases by replacing target entities or antonyms. These methods do not consider structural information, which has proven effective for various IE tasks as they are able to capture non-local interactions (Zhang et al., 2018; Jie and Lu, 2019). This motivates us to propose a novel framework termed as counterfactual information extraction (CFIE). Different from previous efforts, our CFIE model alleviates the spurious correlations by generating counterfactuals (Bottou et al., 2013; Abbasnejad et al., 2020) based on the syntactic structure (Zhang et al., 2018).

From a causal perspective, counterfactuals state the results of the outcome if certain factors had been different. This concept entails a hypothetical scenario where the values in the causal graph can be altered to study the effect of the factor. Intuitively, the factor that yields the most significant changes in model predictions has the greatest impact and is therefore considered as the main effect. Other factors with minor changes are categorized as side effects. In the context of IE with language structures, counterfactual analysis answers the question on “which tokens in the text would be the key clues for RE, NER or ED that could change the prediction result”. With that in mind, our CFIE model is proposed to explore the language structure to eliminate the bias caused by the side effects and maintain the main effect for prediction. We show the effectiveness of CFIE on three representative IE tasks including NER, RE and ED. Our code and the supplementary materials are available at <https://github.com/HeyyyyyyG/CFIE>.

Specifically, our major contributions are:

- To the best of our knowledge, CFIE is the first study that marries the counterfactual analysis and syntactic structure to address the spurious correlation issue for long-tailed IE. We build different structural causal models (SCM; Pearl et al. 2016) for various IE tasks to better capture the underlying main causalities.

- To alleviate spurious corrections, we generate counterfactuals based on syntactic structures. To achieve more robust predictions, we further propose a novel debiasing approach, which maintains a better balance between the direct effect and counterfactual representations.
- Extensive quantitative and qualitative experiments on various IE tasks across five datasets show the effectiveness of our approach.

2 Model

Figure 2 demonstrates the proposed CFIE method using an example from the ACE2005 dataset (Doddington et al., 2004) on the ED task. As shown in Figure 2 (a), two event types “*Life:Die*” and “*SW:Quit*” of the trigger *killed* have 511 and 19 training instances, respectively. Such an unbalanced distribution may mislead a model to build spurious correlations between the trigger word *killed* and the type “*Life:Die*”. The goal of CFIE is to alleviate such incorrect correlations. CFIE employs SCM (Pearl et al., 2016) as causal diagram as it clearly describes relationships among variables. We give the formulation of SCM as follows.

SCM: Without loss of generality, we express SCM as a directed acyclic graph (DAG) $\mathcal{G} = \{\mathbb{V}, \mathbb{F}, \mathbf{U}\}$, where the set of observables (vertices) are denoted as $\mathbb{V} = \{V_1, \dots, V_n\}$, the set of functions (directed edges) as $\mathbb{F} = \{f_1, \dots, f_n\}$, and the set of exogenous variables (e.g., noise) (Pearl et al., 2016) as $\mathbf{U} = \{U_1, \dots, U_n\}$ for each vertice. Here n is the number of nodes in \mathcal{G} . In the deterministic case where \mathbf{U} is given, the values of all variables in SCM are uniquely determined. Each observable V_i can be derived from:

$$V_i := f_i(\mathbf{PA}_i, U_i), (i = 1, \dots, n), \quad (1)$$

where $\mathbf{PA}_i \subseteq \mathbf{V} \setminus V_i$ is the set of parents of V_i and “ \setminus ” is an operator that excludes V_i from \mathbf{V} , and f_i refers to the direct causation from \mathbf{PA}_i to its child variable V_i . Next we show how our SCM-based CFIE works.

2.1 Causal Representation Learning

Figure 3 presents our unified SCM \mathcal{G}_{ie} for IE tasks based on our prior knowledge. The variable S indicates the contextualized representations of an input sentence, where the representations are the output from a BiLSTM (Schuster and Paliwal, 1997) or a pre-trained BERT encoder (Devlin et al., 2019). Z_j

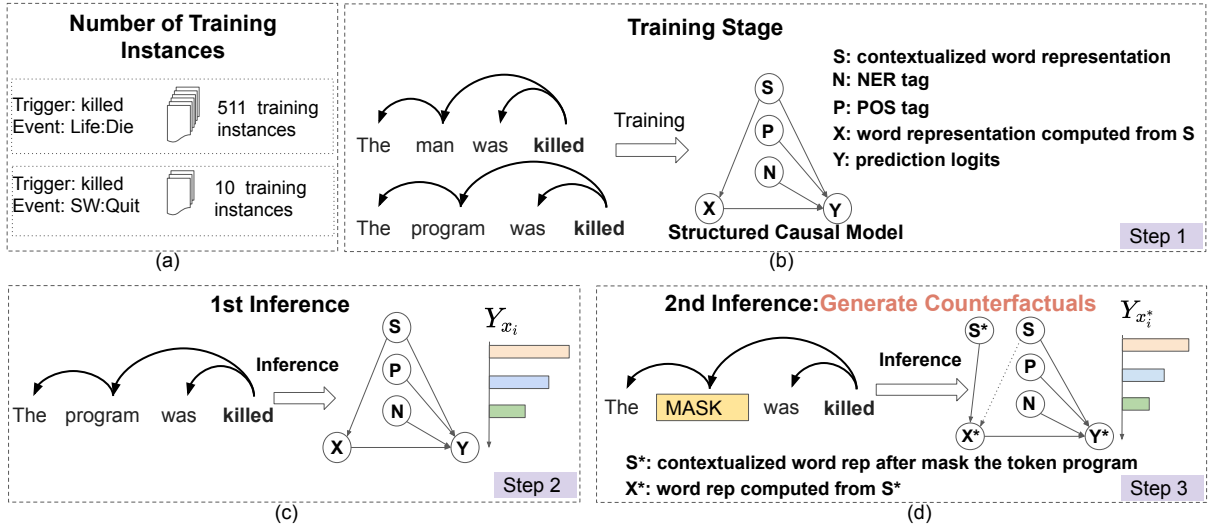


Figure 2: Training and inference stages of CFIE for ED, which detects the triggers and then predict the corresponding event types for a given sentence. (a) The number of training instances for the trigger *killed* labeled with two different event types. (b) Step 1 builds the SCM and trains the model. (c) Step 2 obtains prediction results for each token, e.g., Y_{x_i} for the token *killed*. (d) Step 3 generates the counterfactuals of each token by masking the tokens along the 1st hop of a syntactic tree, and then achieves counterfactual prediction, e.g., $Y_{x_i^*}$ for *killed*.




Figure 3: (a) a unified structured causal model (SCM) for IE tasks. (b) causal interventions on X .

($j \in \{1, 2, \dots, m\}$) represents features such as the NER tags and part-of-speech (POS) tags, where m is the number of features. The variable X is the representation of a relation, an entity and a trigger for RE, NER and ED, respectively, and Y indicates the output logits for classification.

For \mathcal{G}_{ie} , we denote the parents of Y as $\mathcal{E} = \{S, X, Z_1, \dots, Z_m\}$. The direct causal effects towards Y are linear transformations. Transformation for each edge $i \rightarrow Y$ is denoted as $\mathbf{W}_{iY} \in \mathbb{R}^{c \times d}$, where $i \in \mathcal{E}$, c is the number of classes, and d is the dimensional size. We let $\mathbf{H}_i \in \mathbb{R}^{d \times k}$ denote k representations for the node i . Then, the prediction can be obtained by summation $Y_x = \sum_{i \in \mathcal{E}} \mathbf{W}_{iY} \mathbf{H}_i$ or gated mechanism $Y_x = \mathbf{W}_g \mathbf{H}_X \odot \sigma(\sum_{i \in \mathcal{E}} \mathbf{W}_{iY} \mathbf{H}_i)$, where \odot refers to element-wise product, \mathbf{H}_X is the representation of the node X , and $\mathbf{W}_g \in \mathbb{R}^{c \times d}$ and $\sigma(\cdot)$ indicate a linear transformation and the sigmoid function, respectively.

To avoid any single edge dominating the generation of the logits Y_x , we introduce a cross-entropy loss \mathcal{L}_{iY} , $i \in \mathcal{E}$ for each edge. Let \mathcal{L}_Y denote the loss for Y_x , the overall loss \mathcal{L} can be:

$$\mathcal{L} = \mathcal{L}_Y + \sum_{i \in \mathcal{E}} \mathcal{L}_{iY} \quad (2)$$

Step 1 in Figure 2 (b) trains the above causal model, aiming to teach the model to identify the main cause (main effect) and the spurious correlations (side effect) for classification. Our proposed SCM is encoder neutral and it can be equipped with various encoders, such as BiLSTM and BERT.

Fusing Syntactic Structures Into SCM: So far we have built our unified SCM for IE tasks. On the edge $S \rightarrow X$, we adopt different neural network architectures for RE, NER and ED. For RE, we use dependency trees to aggregate long-range relations with graph convolution networks (GCN; Kipf and Welling 2017) to obtain \mathbf{H}_X . For NER and ED, we adopt the dependency-guided concatenation approach (Jie and Lu, 2019) to obtain \mathbf{H}_X .

2.2 Inference and Counterfactual Generation

Given the above SCM, we train our neural model designed for a specific task such as ED. Step 2 in Figure 2 (c) performs inference with our proposed SCM, and Step 3 in Figure 2(d) generates syntax-based counterfactuals to obtain better main effect.

Interventions: For \mathcal{G}_{ie} , an intervention indicates an operation that modifies a subset of variables $\mathbf{V} \subseteq \mathbb{V}$ to new values where each variable $V_i \in \mathbf{V}$ is generated by manual manipulations. Thus, the causal dependency between V_i and its parents $\{\mathbf{PA}_i, U_i\}$ will be cut off, as shown in Figure 3(b). Such an intervention for one variable $X \in \mathbb{V}$ can be expressed by the *do*-notation $do(X = x^*)$ where x^* is the given value (Pearl, 2009).

Counterfactuals: Unlike interventions, the concept of counterfactual reflects an imaginary scenario for “what would the outcome be had the variable(s) been different”. Let $Y \in \mathbb{V}$ denote the outcome variable, and let $X \in \mathbb{V} \setminus \{Y\}$ denote the variable of study. The counterfactual is obtained by setting $X = x^*$ and formally estimated as:

$$Y_{x^*}(u) = Y_{\mathcal{G}_{x^*}}(u) \quad (3)$$

where \mathcal{G}_{x^*} means assigning $X = x^*$ for all equations in the SCM \mathcal{G} . We slightly abuse the notation and use Y_{x^*} as a short form for $Y_{x^*}(u)$, since the exogenous variable u is not explicitly required here². For SCM \mathcal{G} , the counterfactual Y_{x^*} of the original instance-level prediction Y_x is computed as:

$$\begin{aligned} Y_{x^*} &= f_Y(do(X = x^*), S = s, Z = z) \\ &= \sum_{i \in \mathcal{E} \setminus \{X\}} \mathbf{W}_{iY} \mathbf{H}_i + \mathbf{W}_{XY} \mathbf{H}_{x^*} \end{aligned} \quad (4)$$

where f_Y is the function that computes Y . Compared to the vanilla formula for Y_x , we only replace its feature \mathbf{H}_X with \mathbf{H}_{x^*} .

Counterfactual Generation: There are many language structures such as dependency and constituency trees (Marcus et al., 1993), semantic role labels (Palmer et al., 2005), and abstract meaning representations (Banarescu et al., 2013). We choose the dependency tree in our case as it can capture rich relational information and complex long-distance interactions that have proven effective on IE tasks. Counterfactuals lead us to think about: what are the key clues that determine the relations of two entities for RE, and a certain span of a sentence to be an entity or an event trigger for NER and ED respectively? As demonstrated in Figure 2 (d), we mask entities, or the tokens in the scope of 1 hop on the dependency tree. Then this masked sequence is fed to a BiLSTM or BERT encoder to output new contextualized representations

²Derivations are given in the supplementary materials.

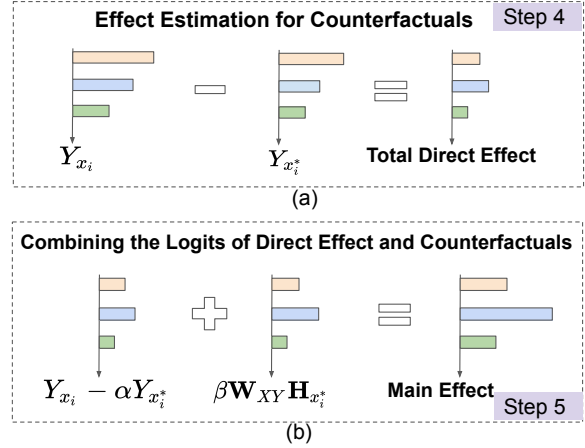


Figure 4: Causal effect estimation: (a) Step 4 computes the TDE by subtraction of outputs for each token in Step 2 and 3, e.g., $Y_{x_i} - \alpha Y_{x_i^*}$. (b) Step 5 obtains more robust predictions by highlighting each token’s counterfactual representations, e.g., $\mathbf{W}_{XY} \mathbf{H}_{x_i^*}$ for *killed*.

S^* , as shown in Figure 2 (d). Then we feed S^* to the function of the edge $S \rightarrow X$ to get X^* . This operation also aligns with a recent finding (Zeng et al., 2020) that the entity itself may be more important than context in NER. By doing so, the key clues are expected to be wiped off in the representations X^* of counterfactuals, strengthening the main effect while reducing the spurious correlations.

2.3 Causal Effect Estimation

As shown in Figure 4, we estimate the causal effect in the Step 4 and use the representation of counterfactuals for a more robust prediction in Step 5. Inspired by SGG-TDE (Tang et al., 2020b), we compare the original outcome Y_x and its counterfactual Y_{x^*} to estimate the main effect so that the side effect can be alleviated with Total Direct Effect (TDE) (Pearl, 2009)³:

$$\text{TDE} = Y_x - Y_{x^*} \quad (5)$$

As both of the context and entity (or trigger) play important roles for the classification in NER, ED, and RE, we propose a novel approach to further alleviate the spurious correlations caused by side effects, while strengthening the main effect at the same time. The interventional causal effect of the i -th entity in a sequence can be described as:

$$\text{Main Effect} = Y_{x_i} - \alpha Y_{x_i^*} + \beta \mathbf{W}_{XY} \mathbf{H}_{x_i^*} \quad (6)$$

where α, β are the hyperparameters that balance the importance of context and entity (or trigger) for

³Derivations are given in the supplementary materials.

NER, ED, and RE. The first part $Y_{x_i} - \alpha Y_{x_i^*}$ indicates the main effect, which reflects more about the debiased context, while the second part $\mathbf{W}_{XY}\mathbf{H}_{x_i^*}$ reflects more about the entity (or trigger) itself. Combining them yields more robust prediction by better distinguishing the main and side effect.

As shown in Step 4 of Figure 4(a), the sentence “The program was killed” produces a biased high score for the event “*Life:Die*” in Y_x and results in wrong prediction due to the word “*killed*”. By computing the counterfactual Y_{x^*} with “program” masked, the score for “*Life:Die*” remains high but the score for “*SW:Quit*” drops. The difference computed by $Y_x - \alpha Y_{x^*}$ may help us to correct the prediction while understanding the important role of the word “program”. However, we may not only rely on the context since the entity (trigger) itself is also an important clue. To magnify the difference and obtain more robust predictions, we strengthen the impact of entity (trigger) on the final results by $\mathbf{W}_{XY}\mathbf{H}_{x_i^*}$ as shown in Step 5 of Figure 4(b). Such a design differs from SGG-TDE (Tang et al., 2020a) by providing more flexible adjustment and effect estimation with hyperparameters α and β . We will show that our approach is more suitable for long-tailed IE tasks in experiments.

3 Experiments

3.1 Datasets and Settings

We use five datasets in our experiments including OntoNotes5.0 (Pradhan et al., 2013) and ATIS (Tur et al., 2010) for the NER task, ACE2005 (Dodington et al., 2004) and MAVEN (Wang et al., 2020b) for ED, and NYT24 (Gardent et al., 2017) for the RE task. The labels in the above datasets follow long-tailed distributions. We categorize the classes into three splits based on the number of training instances per class, including Few, Medium, and Many, and also report the results on the whole dataset with the Overall setting. We focus more on Mean Recall (MR; Tang et al. 2020b) and Macro F1 (MF1), two more balanced metrics to measure the performance of long-tailed IE tasks. MR can better reflect the capability in identifying the tail classes, and MF1 can better represent the model’s ability for each class, whereas the conventional Micro F1 score highly depends on the head classes and pays less attention to the tail classes. The hyperparameter α in Equation (6) is set as 1 for NER and ED tasks, and 0 for the RE task. We tune the optimal

α on the development sets⁴.

3.2 Baselines

We categorize the baselines used in our experiments into three groups and outline them as follows. **Conventional models** include BiLSTM (Chiu and Nichols, 2016), BiLSTM+CRF (Ma and Hovy, 2016), C-GCN (Zhang et al., 2018), Dep-Guided LSTM (Jie and Lu, 2019), and BERT (Devlin et al., 2019). These neural models do not explicitly take the long-tailed issues into consideration.

Re-weighting/Decoupling models refer to loss re-weighting approaches including Focal Loss (Lin et al., 2017), and two-stage decoupled learning methods (Kang et al., 2019) that include τ -normalization, classifier retraining (cRT) and learnable weight scaling (LWS).

Causal model includes SGG-TDE (Tang et al., 2020b). There are also recent studies based on the deconfounded methodology (Tang et al., 2020a; Yang et al., 2020) for images, which however seem not applicable to be selected as a causal baseline in our case for text. We ran some of the baseline methods by ourselves since they may have not been reported on NLP datasets.

3.3 Task Definitions

We show the definition of the IE sub-tasks used in our experiments as follows, including named entity recognition (NER), event detection (ED) and relation extraction (RE).

Named Entity Recognition: NER is a sequence labeling task that seeks to locate and classify named entities in unstructured text into pre-defined categories such as person, location, etc.

Event Detection: ED aims to detect the occurrences of predefined events and categorize them as triggers from unstructured text. An event trigger is defined as the words or phase that most clearly expresses an event occurrence. Taking the sentence “a cameraman died in the Palestine Hotel” as an example, the word “died” is considered as the trigger with a “*Life:Die*” event.

Relation Extraction: The goal of RE is to identify semantic relationships from text, given two or more entities. For example, “Paris is in France” states a “is in” relationship between two entities “Paris” to “France”. Their relation can be denoted by the triple (Paris, is in, France).

⁴More details are attached in the supplementary materials.

Model	OntoNotes5.0								ATIS							
	Few		Medium		Many		Overall		Few		Medium		Many		Overall	
	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1
BiLSTM (Chiu and Nichols, 2016)	67.5	69.9	72.6	75.3	88.1	85.4	76.4	76.8	66.2	69.0	89.8	85.9	93.0	92.2	84.2	83.1
BiLSTM+CRF (Ma and Hovy, 2016)	60.7	63.6	65.3	69.1	86.9	86.9	71.6	73.5	58.1	60.4	87.4	83.5	93.0	93.5	81.0	80.2
C-GCN (Zhang et al., 2018)	68.3	69.8	69.1	72.9	90.9	86.6	77.3	76.8	63.2	65.3	87.3	83.2	91.8	89.3	82.0	80.1
Dep-Guided LSTM (Jie and Lu, 2019)	61.8	69.3	70.2	73.7	89.8	84.3	74.1	75.8	60.6	65.4	94.1	90.4	93.2	92.9	84.6	84.3
Focal Loss (Lin et al., 2017)	64.1	65.5	69.9	71.2	87.7	84.7	74.2	73.9	48.9	49.8	89.3	84.6	91.1	89.8	78.7	76.6
cRT (Kang et al., 2019)	64.1	68.5	73.9	75.3	88.0	85.2	75.0	76.1	68.1	71.7	92.3	88.0	92.8	92.2	85.7	84.8
τ - Normalization (Kang et al., 2019)	61.1	66.7	72.8	76.4	88.0	85.7	73.5	75.7	64.8	68.0	89.9	86.2	93.0	92.5	83.9	83.1
LWS (Kang et al., 2019)	58.7	64.9	71.6	76.1	87.6	85.2	72.1	74.7	66.2	69.1	89.9	85.9	93.0	92.2	84.3	83.2
SGG-TDE (Tang et al., 2020b)	71.9	68.8	77.9	74.8	91.2	86.7	80.4	76.7	67.5	67.1	95.3	89.4	93.7	93.5	87.1	84.5
Ours (Glove)	76.7	68.9	83.6	76.2	92.0	87.6	83.8	77.3	71.8	73.1	95.6	91.4	94.3	93.5	88.6	87.0
BERT (Devlin et al., 2019)	77.7	76.5	81.4	78.6	94.0	90.7	84.6	82.4	60.7	65.9	97.2	89.1	93.9	93.5	86.1	84.0
Ours (BERT)	80.6	79.1	85.1	80.4	94.5	91.4	86.7	84.1	69.9	71.2	97.2	91.5	93.5	93.1	88.5	86.4

Table 1: Evaluation results on the OntoNotes5.0 dataset and ATIS datasets for the NER task.

Model	ACE2005								MAVEN							
	Few		Medium		Many		Overall		Few		Medium		Many		Overall	
	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1	MR	MF1
BiLSTM (Chiu and Nichols, 2016)	34.2	35.6	55.1	58.2	64.9	67.0	52.3	54.8	36.5	40.7	78.3	79.9	80.4	82.3	67.1	69.5
BiLSTM+CRF (Ma and Hovy, 2016)	41.4	45.1	49.8	52.2	70.1	70.5	51.8	54.1	43.4	46.8	79.0	79.8	82.3	83.0	69.6	71.1
C-GCN (Zhang et al., 2018)	41.4	44.1	51.2	55.8	66.4	71.2	52.0	56.1	49.7	51.7	81.8	80.8	82.6	82.1	73.1	73.0
Dep-Guided LSTM (Jie and Lu, 2019)	42.8	41.7	49.8	56.0	71.1	71.6	52.4	55.8	44.7	45.4	76.5	78.2	75.9	78.9	67.8	69.3
Focal Loss (Lin et al., 2017)	38.6	42.9	50.7	58.8	74.6	76.0	52.6	58.5	45.4	51.5	78.6	81.3	85.4	87.2	70.3	73.8
cRT (Kang et al., 2019)	44.8	47.4	58.8	60.1	68.8	68.5	57.6	58.9	49.7	55.4	78.4	81.3	82.1	85.0	71.0	74.6
τ - Normalization (Kang et al., 2019)	34.3	35.6	50.9	53.8	82.7	68.3	53.3	52.5	21.1	26.7	60.0	68.5	74.4	80.0	51.0	58.4
LWS (Kang et al., 2019)	34.3	35.6	61.2	60.2	76.8	71.7	58.2	56.9	33.3	38.7	77.6	79.7	81.6	81.7	65.9	68.7
SGG-TDE (Tang et al., 2020b)	34.3	33.9	61.5	59.7	77.4	73.3	58.5	56.5	39.8	36.2	83.3	78.0	87.8	85.2	71.9	67.4
Ours (Glove)	47.1	49.7	64.3	59.9	80.5	73.3	63.5	60.2	60.4	57.4	86.8	82.2	89.1	86.6	79.8	76.0
BERT(Devlin et al., 2019)	47.6	48.9	67.8	67.5	84.5	76.8	66.5	65.1	61.1	61.8	86.1	84.6	90.3	89.4	79.8	78.9
Ours (BERT)	61.9	63.2	76.5	76.6	85.3	80.7	74.9	74.4	61.8	62.7	86.6	84.5	90.0	89.4	80.2	79.1

Table 2: Evaluation results on the ACE2005 and MAVEN datasets for event detection.

3.4 Main Results

NER: Table 1 shows the comparison results on both OntoNotes5.0 and ATIS datasets. Our models perform best or achieve comparable results under most settings, including Few, Medium, Many and Overall. For example, our model achieves more than 8 points higher MR comparing with the C-GCN model under the Few setting with Glove embeddings on the both of the two benchmarks. The results show the superiority of CFIE in handling the instance-scarce classes for the long-tailed NER. Comparing with a causal baseline SGG-TDE, our model consistently performs better in terms of the two metrics. The results confirm our hypothesis that language structure can help a causal model to better distinguish main effect from the side effect. CFIE also obtains large performance gains with the BERT-based encoder under most of the settings, showing the effectiveness of our approach in mitigating the bias issue with a pre-trained model. It is interesting that BERT-based models perform worse than Glove-based ones on ATIS. The reason is probably that BERT, which is trained on Wikipedia, may not perform well on a small dataset collected from a very different domain.

ED: Table 2 shows comparison results on both of the ACE2005 and MAVEN datasets. Overall, our model significantly outperforms the previous causal baseline SGG-TDE under the Few setting by a large margin, specifically, 12.8 and 15.8 points higher in terms of MR and MF1 respectively on ACE2005 dataset, 20.6 and 21.2 points higher in terms of the two metrics on MAVEN dataset with Glove embeddings. Meanwhile, our model is able to achieve better or comparable results under other settings, such as Medium and Many. The results further confirm the robustness of our model for tail classes with few training instances available. Our model also performs better or comparable than BERT baselines under Few, Medium, Many and Overall settings, indicating that BERT models still suffer from bias issues on the long-tailed IE tasks.

RE: As shown in Table 3, we further evaluate CFIE on the NYT24 dataset. Our method significantly outperforms all other methods in MR and MF1 for tail classes. The overall performance is also competitive. Although Focal Loss achieves the best overall scores, its ability to handle the classes with very few data points drops significantly, which is the main focus of our work. The results for rela-

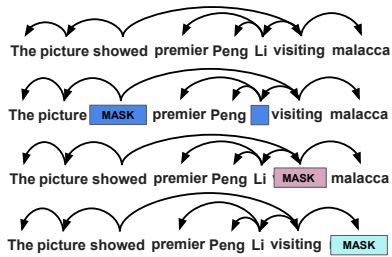


Figure 5: Masking operations.

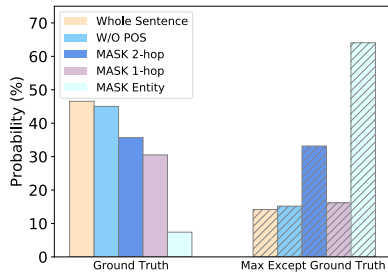


Figure 6: Prediction distributions.

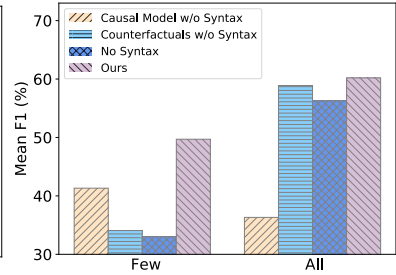


Figure 7: Syntax contribution.

Model	NYT24			
	Few		Overall	
	MR	MF1	MR	MF1
C-GCN (Zhang et al., 2018)	24.0	26.7	51.2	52.6
Focal Loss (Lin et al., 2017)	56.0	54.6	65.7	65.5
cRT (Kang et al., 2019)	66.0	24.2	65.6	50.5
τ - Normalization (Kang et al., 2019)	40.0	40.0	53.5	54.6
LWS (Kang et al., 2019)	40.0	40.0	53.5	54.6
SGG-TDE (Tang et al., 2020b)	60.0	57.1	61.0	60.2
Ours (Glove)	68.0	68.6	65.3	63.6

Table 3: Results on the NYT24 dataset for RE.

tion extraction further confirm our hypothesis that the proposed CFIE is able to alleviate spurious correlations caused by unbalanced dataset by learning to distinguish the main effect from the side effect. We also observe that CFIE outperforms the previously proposed SGG-TDE by a large margin for both Few and Overall settings, i.e., 11.5 points and 3.4 points improvement in terms of MF1. This further proves our claim that properly exploring language structure on causal models will boost the performance of IE tasks.

3.5 Discussion

What are the key factors for NER? We have hypothesised that the factors, such as 2-hop and 1-hop context on the dependency tree, the entity itself, and POS feature, may hold the potential to be the key clues for NER predictions. To evaluate the impact of these factors, we first generate new sequences by masking these factors. Then we feed the generated sequences to the proposed SCM to obtain the predictions. Figure 5 illustrates how we mask the context based on a dependency tree. Figure 6 shows a qualitative example for predicting the NER tag for the entity “malacca”. It visualizes the variances of the predictions, where the histograms in the left refer to prediction probabilities for the ground truth class, while the histograms in the right are the max predictions except the results of ground truth class. For example, the “Mask 2-hop” operation with a blue rectangle in Figure 5 masks tokens

“showed” and “Li” on the dependency tree, and the corresponding prediction probability distribution is given in Figure 6, which is expressed as the blue bar. We observe that masking the entity, i.e., “malacca”, will lead to the most significant performance drop, indicating that entity itself plays a key role for the NER task. This also inspires us to design a more robust debiasing method as shown in Step 5 in our framework.

Does the syntax structure matter? To answer this question, we design three baselines including: 1) Causal Models w/o Syntax that doesn’t employ dependency trees during the training stage, and only uses it for generating counterfactuals, 2) Counterfactuals w/o Syntax that employs dependency structures for training but uses a null input as the intervention during the inference state. We refer such a setting from the previous study (Tang et al., 2020a), and 3) No Syntax that is the same as the previous work SGG-TDE (Tang et al., 2020b) which doesn’t involve dependency structures in both training and inference stages. As shown in Figure 7, our model outperforms all three baselines on the ACE2005 dataset under both Few and All settings, demonstrating the effectiveness of dependency structure in improving the causal models for the long-tailed IE tasks both in the training and inference stages.

How do various interventions and SCMs affect performance? We study this question on ACE2005 dataset for ED task. We design three interventional methods including 1) Intervene X & NER, 2) Intervene X & POS, 3) Intervene X & NER & POS. Figure 8 shows that introducing interventions solely on X is able to achieve the best performance under both Few and All settings. We also introduce three variants of our proposed SCM : 1) SCM w/o NER, 2) SCM w/o POS, 3) SCM w/o NER and POS. Figure 9 shows that removing the NER node will significantly decrease the




Figure 8: Various interventions.




Figure 9: Various SCMs.

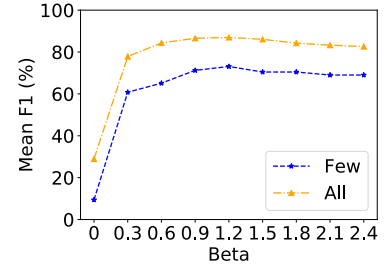


Figure 10: Various β .

ED performance, especially over the Few setting. The results prove the superiority of our proposed SCM that explicitly involves linguistic features to calculate main effect.

How does the hyper-parameter β impact the performance? To evaluate the impact of β on the performance, we tuned the parameter on four datasets including OntoNotes5.0, ATIS, ACE2005, and MAVEN. As shown in Figure 10, when increasing β from 0 to 2.4 on ATIS dataset, the F1 scores increase dramatically then decrease slowly. The F1 scores reach the peak when β is 1.2. As the value of β represents the importance of entity for classifications, we therefore draw a conclusion that, for NER task, an entity plays a relatively more important role than the context (Zeng et al., 2020). We observe that the performance significantly drops when β is 0. This suggests that directly applying previous causal approach (Tang et al., 2020b) may not yield good performance. The result further confirms the effectiveness of Step 5 in CFIE.

3.6 Case Study

Figure 11 shows two cases to visualize the predictions of baseline models and our CFIE model for long-tailed NER and ED, respectively. We use the “BIOES” tagging scheme for both cases and choose Dep-Guided LSTM (Jie and Lu, 2019) and SGG-TDE (Tang et al., 2020b) as baselines. In the first case for NER, the baseline assigns “chinese” with the label S-NORP, which indicates “nationalities or religious or political groups”, while the corrected annotation is “S-LANGUAGE”. This is caused by the spurious correlations between “chinese” and S-NORP learned from unbalanced data. For this case, there are 568 and 20 training instances for S-NORP and S-LANGUAGE, respectively. The numbers of training instances for each type are indicated in the third column of Figure 11. The numbers in the 4-th to 6-th columns indicate the probability of the token “chinese” predicted as a certain label.

For example, in the 6-th column “Predictions of CFIE”, the prediction probability is 35.17% for the label S-LANGUAGE. In the second case for ED, we demonstrate a similar issue for the trigger word “attack”, and compare it with the two baselines. For both cases, previous SGG-TDE outputs relatively unbiased predictions compared with Dep-Guided LSTM, although the predictions are also incorrect. Our CFIE model can obtain correct results for both instances, showing the effectiveness of our novel debiasing approach. Compared to CFIE, the inferior performance of SGG-TDE is due to ignoring the importance of entity (trigger) for the NER and ED tasks.

4 Related Work

Long-tailed IE: RE (Zeng et al., 2014; Peng et al., 2017; Quirk and Poon, 2017; Song et al., 2018; Lin et al., 2019; Peng et al., 2020; Nan et al., 2020; Qu et al., 2020; Guo et al., 2020; Zhang et al., 2021c; Zheng et al., 2021; Zhang et al., 2020b; Ye et al., 2021; Bai et al., 2021; Nan et al., 2021a), NER (Lample et al., 2016; Chiu and Nichols, 2016; Xu et al., 2021b), and ED (Nguyen and Grishman, 2015; Huang et al., 2018) are mainstream IE tasks in NLP. For the long-tailed IE, recent models (Lei et al., 2018; Zhang et al., 2019) leverage external rules or transfer knowledge from data-rich classes to the tail classes. Plenty of re-balancing models are also proposed, including re-sampling strategies (Mahajan et al., 2018; Wang et al., 2020a) that aim to alleviate statistical bias from head classes, and re-weighting approaches (Milletari et al., 2016; Lin et al., 2017) which assign balanced weights to the losses of training samples from each class to boost the discriminability via robust classifier decision boundaries. Another line is decoupling approaches (Kang et al., 2019) that decouple the representation learning and the classifier by direct re-sampling. Different from the above works that are based on conventional approaches, we tackle

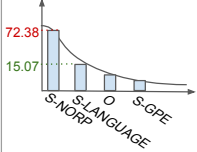
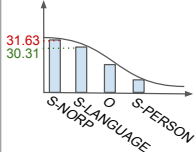
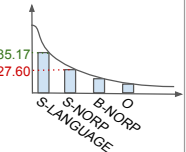
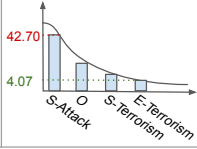
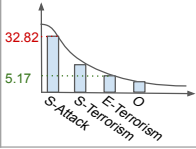
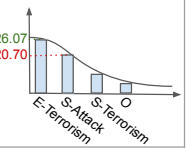
Task	Instances	#Instances/Class in Training set	Predictions of Dep-Guided LSTM	Predictions of SGG-TDE	Predictions of CFIE
NER	The wangma computer company was named for the wubi font, the coding method of chinese characters invented by the famous computer expert - yongmin wang.	S-NORP : 568 O : 243 I-ORG : 42 S-LANGUAGE : 20			
ED	The Camp Chapman attack was a suicide attack by Human Khalil Abu-Mulal al-Balawi against the Central Intelligence Agency facility inside Forward Operating Base Chapman on December 30, 2009.	S-Attack : 545 O : 230 S-Terrorism : 32 E-Terrorism : 12			

Figure 11: Two cases selected from OntoNote5.0 and MAVEN for NER and ED tasks respectively, with unbalanced distributions for the targeting entity and event trigger. The two baseline models Dep-Guided LSTM and SGG-TDE tend to predict incorrect results caused by the spurious correlations, while our proposed CFIE model is able to yield better predictions.

the long-tailed IE problem from the perspective of causal inference.

Causal Inference: Causal inference (Pearl et al., 2016; Rubin, 2019) has been applied in many areas, including visual tasks (Tang et al., 2020b; Abbasnejad et al., 2020; Niu et al., 2021; Yang et al., 2020; Zhang et al., 2020a; Yue et al., 2020; Yang et al., 2021; Nan et al., 2021b), model robustness and stable learning (Srivastava et al., 2020; Zhang et al., 2020a; Shen et al., 2020; Yu et al., 2020; Dong et al., 2020), generation (Wu et al., 2020), language understanding (Feng et al., 2021b), and recommendation systems (Jesson et al., 2020; Feng et al., 2021a; Zhang et al., 2021d; Wei et al., 2021; Wang et al., 2021a; Tan et al., 2021; Wang et al., 2021b; Ding et al., 2021). Works most related to ours are (Zeng et al., 2020; Wang and Culotta, 2021) that generates counterfactuals for weakly-supervised NER and text classifications, respectively. Our method is remotely related to (Tang et al., 2020b) proposed for image classifications. The key differences between our methods and previous ones: 1) counterfactuals in our method are generated by a task-specific pruned dependency structure on various IE tasks. While in previous works, counterfactuals are generated by replacing the target entity with another entity or their antonyms (Zeng et al., 2020; Wang and Culotta, 2021), or simply masking the targeting objects in an image (Tang et al., 2020b). These method do not consider the complex language structure that has been proven useful for IE tasks. 2) compared with previous method SGG-TDE (Tang et al., 2020b), our inference mechanism is more robust for various IE tasks, simultaneously mitigating the spurious correlations and strength-

ening salient context.

5 Concluding Remarks

This paper presents CFIE, a novel framework for tackling the long-tailed IE issues in the view of causal inference. Extensive experiments on three popular IE tasks, named entity recognition, event detection, and relation extraction, show the effectiveness of our method. Our CFIE model provides a new perspective on tackling spurious correlations by exploring language structures based on structured causal models. We believe that our models may also find applications in other NLP tasks that suffer from spurious correlation issues caused by unbalanced data distributions. Our future work includes developing more powerful causal models for the long-tailed distribution problems using the task-specific language structures learned from the data. We are also interested in addressing the spurious correlations in various vision and language tasks (Nan et al., 2021b; Li et al., 2021; Xu et al., 2021a; Fan et al., 2020; Liu et al., 2021; Zhang et al., 2021a,b; Chen et al., 2020, 2021).

Acknowledgments

We would like to thank the anonymous reviewers for their thoughtful and constructive comments. This work is done when Jiaqi Zeng was working as a research intern at SUTD and when Rui Qiao was working as a research assistant at SUTD. This project/research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-RP-2019-012).

References

- Ehsan Abbasnejad, Damien Teney, Amin Parvaneh, Javen Shi, and Anton van den Hengel. 2020. Counterfactual vision and language learning. In *Proc. of CVPR*.
- Xuefeng Bai, Yulong Chen, Linfeng Song, and Yue Zhang. 2021. Semantic representation for dialogue modeling. In *Proc. of ACL*.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*.
- Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual reasoning and learning systems: The example of computational advertising. *The Journal of Machine Learning Research*, 14(1):3207–3260.
- Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. 2020. Scanrefer: 3d object localization in rgb-d scans using natural language. In *Proc. of ECCV*.
- Zhenyu Chen, Ali Gholami, Matthias Nießner, and Angel X Chang. 2021. Scan2cap: Context-aware dense captioning in rgb-d scans. In *Proc. of CVPR*.
- Jason P.C. Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional LSTM-CNNs. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL*.
- Sihao Ding, Fuli Feng, Xiangnan He, Yong Liao, Jun Shi, and Yongdong Zhang. 2021. Causal incremental graph convolution for recommender system re-training. *arXiv preprint arXiv:2108.06889*.
- George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. The automatic content extraction (ACE) program – tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Zhenhua Dong, Hong Zhu, Pengxiang Cheng, Xinhua Feng, Guohao Cai, Xiuqiang He, Jun Xu, and Jirong Wen. 2020. Counterfactual learning for recommender system. In *ACM RecSys*.
- Hehe Fan, Linchao Zhu, Yi Yang, and Fei Wu. 2020. Recurrent attention network with reinforced generator for visual dialog. *ACM Transactions on Multimedia Computing, Communications, and Applications TOMM*, 16(3):78:1–78:16.
- Fuli Feng, Weiran Huang, Xiangnan He, Xin Xin, Qifan Wang, and Tat-Seng Chua. 2021a. Should graph convolution trust neighbors? a simple causal inference method. In *Proc. of SigIR*.
- Fuli Feng, Jizhi Zhang, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. 2021b. Empowering language understanding with counterfactual reasoning. In *Proc. of ACL*.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. The WebNLG challenge: Generating text from RDF data. In *Proceedings of the 10th International Conference on Natural Language Generation*. Association for Computational Linguistics.
- Zhijiang Guo, Guoshun Nan, Wei Lu, and Shay B Cohen. 2020. Learning latent forests for medical relation extraction. In *Proc. of IJCAI*.
- Xu Han, Pengfei Yu, Zhiyuan Liu, Maosong Sun, and Peng Li. 2018. Hierarchical relation extraction with coarse-to-fine grained attention. In *Proc. of EMNLP*.
- Lifu Huang, Heng Ji, Kyunghyun Cho, Ido Dagan, Sebastian Riedel, and Clare Voss. 2018. Zero-shot transfer learning for event extraction. In *Proc. of ACL*.
- Andrew Jesson, Sören Mindermann, Uri Shalit, and Yarin Gal. 2020. Identifying causal-effect inference failure with uncertainty-aware models. In *Proc. of NeurIPS*.
- Zhanming Jie and Wei Lu. 2019. Dependency-guided lstm-crf for named entity recognition. In *Proc. of EMNLP*.
- Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. 2019. Decoupling representation and classifier for long-tailed recognition. In *Proc. of ICLR*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proc. of ICLR*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proc. of NAACL*.
- Kai Lei, Daoyuan Chen, Yaliang Li, Nan Du, Min Yang, Wei Fan, and Ying Shen. 2018. Cooperative denoising for distantly supervised relation extraction. In *Proc. of COLING*.

- Yicong Li, Xun Yang, Xindi Shang, and Tat-Seng Chua. 2021. Interventional video relation detection. In *ACM International Conference on Multimedia*.
- Hongtao Lin, Jun Yan, Meng Qu, and Xiang Ren. 2019. Learning dual retrieval module for semi-supervised relation extraction. In *Proc. of WWW*.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proc. of ICCV*.
- Daizong Liu, Xiaoye Qu, Jianfeng Dong, Pan Zhou, Yu Cheng, Wei Wei, Zichuan Xu, and Yulai Xie. 2021. Context-aware biaffine localizing network for temporal sentence grounding. In *Proc. of CVPR*.
- Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. 2019. Large-scale long-tailed recognition in an open world. In *Proc. of CVPR*.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proc. of ACL*.
- Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. 2018. Exploring the limits of weakly supervised pretraining. In *Proc. of ECCV*.
- M. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank. *Comput. Linguistics*, 19:313–330.
- Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE.
- Guoshun Nan, Zhijiang Guo, Ivan Sekulić, and Wei Lu. 2020. Reasoning with latent structure refinement for document-level relation extraction. In *Proc. of ACL*.
- Guoshun Nan, Guoqing Luo, Sicong Leng, Yao Xiao, and Wei Lu. 2021a. Speaker-oriented latent structures for dialogue-based relation extraction. In *Proc. of EMNLP*.
- Guoshun Nan, Rui Qiao, Yao Xiao, Jun Liu, Sicong Leng, Hao Zhang, and Wei Lu. 2021b. Interventional video grounding with dual contrastive learning. In *Proc. of CVPR*.
- Thien Huu Nguyen and Ralph Grishman. 2015. Event detection and domain adaptation with convolutional neural networks. In *Proc. of ACL*.
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. Counterfactual vqa: A cause-effect look at language bias. In *Proc. of CVPR*.
- Martha Palmer, Paul R. Kingsbury, and D. Gildea. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31:71–106.
- Judea Pearl. 2009. *Causality*. Cambridge university press.
- Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. 2016. *Causal inference in statistics: A primer*. John Wiley & Sons.
- Hao Peng, Tianyu Gao, Xu Han, Yankai Lin, Peng Li, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2020. Learning from context or names? an empirical study on neural relation extraction. In *Proc. of EMNLP*.
- Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen tau Yih. 2017. Cross-sentence n-ary relation extraction with graph lstms. *Transactions of the Association for Computational Linguistics*, 5:101–115.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. Towards robust linguistic analysis using OntoNotes. In *Proc. of CoNLL*.
- Meng Qu, Tianyu Gao, Louis-Pascal Xhonneux, and Jian Tang. 2020. Few-shot relation extraction via bayesian meta-learning on relation graphs. In *Proc. of ICML*.
- Chris Quirk and Hoifung Poon. 2017. Distant supervision for relation extraction beyond the sentence boundary. In *Proc. of EACL*.
- Donald B Rubin. 2019. Essential concepts of causal inference: a remarkable history and an intriguing future. *Biostatistics & Epidemiology*, 3(1):140–155.
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Zheyang Shen, Peng Cui, Jiashuo Liu, Tong Zhang, Bo Li, and Zhitang Chen. 2020. Stable learning via differentiated variable decorrelation. In *ACM SIGKDD*.
- Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. N-ary relation extraction using graph-state lstm. In *Proc. of EMNLP*.
- Megha Srivastava, Tatsunori Hashimoto, and Percy Liang. 2020. Robustness to spurious correlations via human annotations. In *Proc. of ICML*.
- Juntao Tan, Shuyuan Xu, Yingqiang Ge, Yunqi Li, Xu Chen, and Yongfeng Zhang. 2021. Counterfactual explainable recommendation. In *Proc. of CIKM*.
- Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. 2020a. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *Proc. of NeurIPS*.

- Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. 2020b. Unbiased scene graph generation from biased training. In *Proc. of CVPR*.
- Gokhan Tur, Dilek Hakkani-Tür, and Larry Heck. 2010. What is left to be understood in atis? In *Proc. of 2010 IEEE Spoken Language Technology Workshop*.
- Tao Wang, Yu Li, Bingyi Kang, Junnan Li, Junhao Liew, Sheng Tang, Steven Hoi, and Jiashi Feng. 2020a. The devil is in classification: A simple framework for long-tail instance segmentation. In *Proc. of ECCV*.
- Wenjie Wang, Fuli Feng, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. 2021a. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proc. of SigIR*.
- Xiaozhi Wang, Ziqi Wang, Xu Han, Wangyi Jiang, Rong Han, Zhiyuan Liu, Juanzi Li, Peng Li, Yankai Lin, and Jie Zhou. 2020b. Maven: A massive general domain event detection dataset. In *Proc. of EMNLP*.
- Zhao Wang and Aron Culotta. 2020. Identifying spurious correlations for robust text classification. In *Proc. of EMNLP*.
- Zhao Wang and Aron Culotta. 2021. Robustness to spurious correlations in text classification via automatically generated counterfactuals. In *In Proc. of AAAI*.
- Zhenlei Wang, Jingsen Zhang, Hongteng Xu, Xu Chen, Yongfeng Zhang, Wayne Xin Zhao, and Ji-Rong Wen. 2021b. Counterfactual data-augmented sequential recommendation. In *Proc. of SigIR*.
- Tianxin Wei, Fuli Feng, Jiawei Chen, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2021. Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system. In *Proc. of KDD*.
- Yiquan Wu, Kun Kuang, Yating Zhang, Xiaozhong Liu, Changlong Sun, Jun Xiao, Yueting Zhuang, Luo Si, and Fei Wu. 2020. De-biased court’s view generation with causality. In *Proc. of EMNLP*.
- Li Xu, He Huang, and Jun Liu. 2021a. Sutd-trafficqa: A question answering benchmark and an efficient network for video reasoning over traffic events. In *Proc. of CVPR*.
- Lu Xu, Zhanming Jie, Wei Lu, and Lidong Bing. 2021b. Better feature integration for named entity recognition. In *Proc. of NAACL*.
- Xu Yang, Hanwang Zhang, and Jianfei Cai. 2020. Deconfounded image captioning: A causal retrospect. *arXiv preprint arXiv:2003.03923*.
- Xun Yang, Fuli Feng, Wei Ji, Meng Wang, and Tat-Seng Chua. 2021. Deconfounded video moment retrieval with causal intervention. In *Proc. of SigIR*.
- Hongbin Ye, Ningyu Zhang, Shumin Deng, Mosha Chen, Chuanqi Tan, Fei Huang, and Huajun Chen. 2021. Contrastive triple extraction with generative transformer. In *Proc. of AAAI*.
- Xiyu Yu, Tongliang Liu, Mingming Gong, Kun Zhang, Kayhan Batmanghelich, and Dacheng Tao. 2020. Label-noise robust domain adaptation. In *Proc. of ICML*.
- Zhongqi Yue, Hanwang Zhang, Qianru Sun, and Xian-Sheng Hua. 2020. Interventional few-shot learning. In *Proc. of NeurIPS*.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jian Zhao. 2014. Relation classification via convolutional deep neural network. In *Proc. of COLING*.
- Xiangji Zeng, Yunliang Li, Yuchen Zhai, and Yin Zhang. 2020. Counterfactual generator: A weakly-supervised method for named entity recognition. In *Proc. of EMNLP*.
- Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. 2020a. Causal intervention for weakly-supervised semantic segmentation. In *Proc. of NeurIPS*.
- Hao Zhang, Aixin Sun, Wei Jing, Guoshun Nan, Liangli Zhen, Joey Tianyi Zhou, and Rick Siow Mong Goh. 2021a. Video corpus moment retrieval with contrastive learning. In *Proc. of SigIR*.
- Hao Zhang, Aixin Sun, Wei Jing, Liangli Zhen, Joey Tianyi Zhou, and Rick Siow Mong Goh. 2021b. Natural language video localization: A revisit in span-based question answering framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ningyu Zhang, Xiang Chen, Xin Xie, Shumin Deng, Chuanqi Tan, Mosha Chen, Fei Huang, Luo Si, and Huajun Chen. 2021c. Document-level relation extraction as semantic segmentation. In *Proc. of IJCAI*.
- Ningyu Zhang, Shumin Deng, Zhen Bi, Haiyang Yu, Jiacheng Yang, Mosha Chen, Fei Huang, Wei Zhang, and Huajun Chen. 2020b. Openue: An open toolkit of universal extraction from text. In *Proc. of EMNLP*.
- Ningyu Zhang, Shumin Deng, Zhanlin Sun, Guanying Wang, Xi Chen, Wei Zhang, and Huajun Chen. 2019. Long-tail relation extraction via knowledge graph embeddings and graph convolution networks. In *Proc. of NAACL*.
- Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021d. Causal intervention for leveraging popularity bias in recommendation. In *Proc. of SigIR*.

Yuhao Zhang, Peng Qi, and Christopher D Manning. 2018. Graph convolution over pruned dependency trees improves relation extraction. In *Proc. of EMNLP*.

Hengyi Zheng, Rui Wen, Xi Chen, Yifan Yang, Yunyan Zhang, Ziheng Zhang, Ningyu Zhang, Bin Qin, Ming Xu, and Yefeng Zheng. 2021. Prgc: Potential relation and global correspondence based joint relational triple extraction. In *Proc. of ACL*.