

# EDDA: An Encoder-Decoder Data Augmentation Framework for Zero-Shot Stance Detection

Daijun Ding<sup>1,2</sup>, Li Dong<sup>2</sup>, Zhichao Huang<sup>3</sup>, Guangning Xu<sup>3</sup>,  
Xu Huang<sup>3</sup>, Bo Liu<sup>4</sup>, Liwen Jing<sup>5</sup>, Bowen Zhang<sup>2\*</sup>

<sup>1</sup>Shenzhen University    <sup>2</sup>Shenzhen Technology University

<sup>3</sup>Harbin Institute of Technology, Shenzhen

<sup>4</sup>Xi'an Jiaotong University,    <sup>5</sup>Shenzhen X-Institute

ding\_dai\_jun@outlook.com, dongli@sztu.edu.cn, iceshzc@gmail.com,

guangningxu35@gmail.com, huangxuu@outlook.com,

raymanliu@126.com, ljing@connect.ust.hk, zhang\_bo\_wen@foxmail.com

## Abstract

Stance detection aims to determine the attitude expressed in text towards a given target. Zero-shot stance detection (ZSSD) has emerged to classify stances towards unseen targets during inference. Recent data augmentation techniques for ZSSD increase transferable knowledge between targets through text or target augmentation. However, these methods exhibit limitations. Target augmentation lacks logical connections between generated targets and source text, while text augmentation relies solely on training data, resulting in insufficient generalization. To address these issues, we propose an encoder-decoder data augmentation (EDDA) framework. The encoder leverages large language models and chain-of-thought prompting to summarize texts into target-specific if-then rationales, establishing logical relationships. The decoder generates new samples based on these expressions using a semantic correlation word replacement strategy to increase syntactic diversity. We also analyze the generated expressions to develop a rationale-enhanced network that fully utilizes the augmented data. Experiments on benchmark datasets demonstrate our approach substantially improves over state-of-the-art ZSSD techniques. The proposed EDDA framework increases semantic relevance and syntactic variety in augmented texts while enabling interpretable rationale-based learning.

**Keywords:** zero-shot stance detection, data augmentation, chain-of-thought

## 1. Introduction

Stance detection aims to automatically determine the attitude (i.e., *favor*, *against*, or *neutral*) expressed in opinionated text towards a given target (Du et al., 2017). Conventionally, stance detection has relied on designing target-specific classifiers to enable predictions on a single topic. Subsequently, cross-target stance detection has emerged as a subclass of generic stance detection, where the classifier is adapted from different but topically related domains (Zhang et al., 2020). In practice, exhaustively enumerating all potential in-target or associated cross-target entities for training is infeasible. Consequently, Zero-shot stance detection (ZSSD) has gained traction as a promising approach focused on accurately classifying stance towards unseen targets during inference (Allaway and Mckeown, 2020).

Previous works have utilized attention-based methods (Allaway and Mckeown, 2020), graph network approaches (Liu et al., 2021), and external background knowledge (Zhu et al., 2022) for ZSSD. However, most of these methods can only extract information from seen targets present in the training data. Recently, data augmentation techniques

have gained research attention by increasing transferable knowledge between targets through augmentation of the textual content or targets. These data augmentation approaches can be categorized into target augmentation and text augmentation methods. Target augmentation focuses on generating or retrieving new, unseen targets from the training corpus (Li et al., 2023). Text augmentation leverages large pre-trained models to produce new text-label pairs via prompt-based learning (Xu et al., 2022; Zhang et al., 2023c).

Despite promising results, prior data augmentation approaches exhibit two key limitations when applied to ZSSD. First, current target augmentation methods lack inherent logical and semantic connections between the generated targets and source text, which can yield uninterpretable and counterintuitive predictions. For example, in sentence “We believe that Hillary’s *emails* have significant issues”, the original target is “Hillary”, while the target augmentation method may introduce “email” as a new target. However, the original text does not semantically relate to this new target, which can lead to confusion. Second, existing text augmentation techniques rely exclusively on the training data. As more samples are synthesized, the distribution of augmented text converges towards the original training distribution, resulting in insufficient

---

\* Corresponding author.

generalization. Empirically, we observe that current text augmentation methods only provide marginal accuracy improvements compared to conventional techniques on this task (see TDDA and TarBK in Table 1).

To address the aforementioned issues, we propose a simple yet effective data augmentation method for ZSSD called the encoder-decoder data augmentation (EDDA) framework, as shown in Figure 1. Unlike prior techniques, EDDA focuses on increasing syntactic diversity in the augmented text while maintaining logical and semantic relevance between the text and the target. Specifically, the encoder leverages a one-shot chain of thought prompting method to let large language models (LLMs) understand and summarize the training texts into target-specific prediction rationales expressed in an *if-then* expression. *If-then* expression can clearly describe the logical connections between text and target. Moreover, representing predictions as rationales has been shown to improve model comprehension and performance for ZSSD (Jayaram and Allaway, 2021). Next, the decoder aims to generate augmented samples based on the *if-then* expression. To increase textual syntactic variety, inspired by (Huang et al., 2023), we propose a straightforward semantic correlation word replacement strategy for the decoder input. This increases the diversity of generated syntactic structures while retaining semantic logic.

Additionally, we analyze the *if-then* expression generated by EDDA and propose a rationale-enhanced network (REN) to fully exploit the augmented data. Experimental results on multiple widely adopted benchmark datasets demonstrate that our proposed approach substantially outperforms current state-of-the-art methods.

The main contributions of our work can be summarized as follows:

- We propose a novel EDDA framework<sup>1</sup> that maintains semantic relevance between the augmented text and target while increasing syntactic diversity. This substantially improves ZSSD performance.
- We introduce a novel *if-then* expression representation derived from LLMs via chain-of-thought prompting. This effectively encodes the stance prediction process in an interpretable manner. Moreover, this representation is model-agnostic. Through simple integration, it can augment existing stance detection models by incorporating prior knowledge.
- We conduct extensive experiments on several widely used benchmarks to verify the effectiveness of our model for ZSSD. The experimental

results show that our model consistently outperforms the compared methods.

## 2. Related Work

### 2.1. Zero-shot Stance Detection

ZSSD is to recognize the attitude towards unseen targets by leveraging stance features learned from known targets. Earlier works have primarily focused on model architecture, including attention mechanisms, graph convolutional networks, and contrastive learning frameworks. For instance, Allaway and Mckeown (2020) constructed a large-scale dataset called VAST for ZSSD encompassing diverse topics in domains such as politics and education. Liu et al. (2021) proposed an attention-based model to better capture relational patterns between texts and topics. Liang et al. (2022a,b) introduced a contrastive learning approach that relates target-invariant and target-specific representations. Luo et al. (2022); Zhu et al. (2022); He et al. (2022); Zhang et al. (2023a) incorporated background knowledge to better generalizing what they learned about known topics to new unseen topics.

Since these prior works have focused solely on seen targets, data augmentation methods for ZSSD have recently gained significant research attention. Earlier data augmentation studies primarily followed traditional text augmentation techniques to increase transferable knowledge between topics (Liang et al., 2022a). With the advent of LLMs such as the GPT series, recent works have proposed prompt-based or instruction-based data augmentation approaches for ZSSD. These methods effectively leverage the data generation and weak supervision capabilities of large models, yielding substantial performance gains on ZSSD. These augmentation techniques can be categorized into target augmentation and text augmentation. Li et al. (2023) focused on enhancing target diversity during training, allowing models trained on more informative targets to improve generalization to unseen targets. For text augmentation, Zhang et al. (2023c) proposed augmenting target-relevant text segments, while Xu et al. (2022) further augmented text samples based on prompt learning. Additionally, following (Dai et al., 2023), we replace GPT-3 in Xu et al. (2022) with GPT-3.5 and construct text-driven data augmentation (TDDA) method.

### 2.2. Text Data Augmentation

**Traditional Text Data Augmentation** Zhang et al. (2015) replaced local words with their synonyms in external semantic lexicons. Wei and Zou (2019) presented EDA that consists of four operations

<sup>1</sup><https://github.com/Szu-Ddj/EDDA>

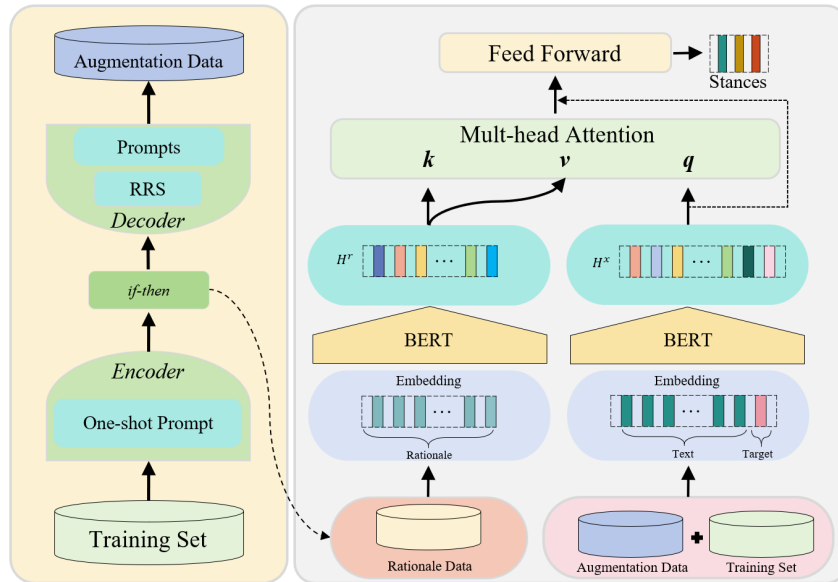


Figure 1: The architecture of our proposed encoder-decoder data augmentation (EDDA) framework.

(e.g., random deletion or insertion), which boost the performance on text classification tasks. For stance detection, Li and Caragea (2021) constructed auxiliary sentences that contain target and label information to generate target-aware sentences.

**Data Augmentation with LLM** With the development of LLM, parameter-free prompt-based data augmentation methods that have rapidly advanced across natural language processing. For example, Edwards et al. (2021) proposed a prompt-based approach that generates the training texts based on human-selected seed samples. Wang et al. (2022) further extended this by automating seed selection, achieving improved performance. Yoo et al. (2021); Dai et al. (2023) proposed an iterative data generation method, mixing training samples with newly generated samples at each step.

### 2.3. Chain-of-Thought model (CoT)

Recent works have explored enhancing chain-of-thought (CoT) prompting to elicit impressive multi-hop reasoning from LLMs (Wei et al.; Zhou et al., 2022; Zhang et al., 2022). For example, Cai et al. (2023) proposed a human-in-the-loop system augmented with CoT prompting, investigating how manual correction of sub-logic in rationales can refine LLMs reasoning. Fei et al. (2023) introduced a multi-step CoT approach that decomposes downstream tasks into multiple stages to improve prediction effectiveness. Ling et al. (2023) presented a new CoT technique, which iteratively infers tasks via deductive reasoning and verification. Ding et al. (2024) utilized multi-step CoT to acquire the positional perspective of targets and applied it to stance detection.

Inspired by multi-step CoT techniques, we pro-

pose a novel encoder-decoder CoT framework that effectively generates prediction rationales in *if-then* format.

## 3. Methodology

### 3.1. Task Description

We use  $D^{train} = \{x_i^{train}, p_i^{train}\}_{i=1}^{N^{train}}$  to denote the collection of labeled data, where  $x$  and  $p$  denote the input text and the corresponding target, respectively. Each  $(x, p)$  pair in  $D^{train}$  has a stance label  $y$ . Given an input sentence  $x^u$  and a corresponding target  $p^u$  as a test set (unseen target), this study aims to predict a stance label  $\hat{y}$  for the input sentence  $x^u$  towards the given target  $p^u$ .

### 3.2. EDDA

**Encoder.** The encoder aims to leverage the reasoning capabilities of LLMs to transform input text data into *if-then* representations encapsulating predictive logic. Inspired by prompt-based instruction learning methods, we conduct a one-shot prompt for LLMs to enable comprehension of input text and generation of *if-then* expression. Specifically, we define the one-shot prompt  $P'1$  to the LLMs to acquire the *if-then* expression. Based on the  $P'1$  template, the LLMs can effectively transform the given input into *if-then* expressions, as exemplified by the generation results presented in Figure 2.

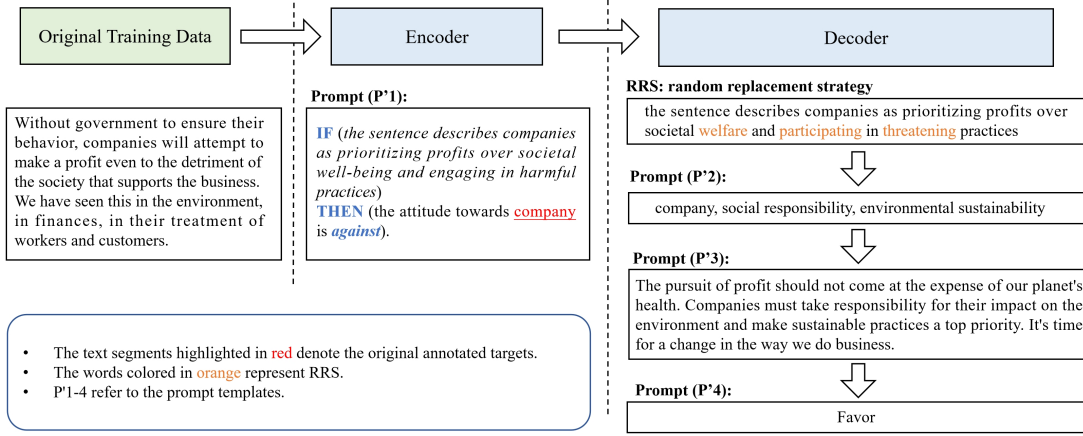


Figure 2: An example of the entire EDDA process.

P1: Your task is to add calls to a question-answering API to a piece of text. The questions should help you get information required to complete the text. You can call the API by writing “[RULE: If (A) then (B)]” where “A” is the reason why “B”. Here are some examples of API calls:  
**Input:** What’s the attitude of the sentence “[given text]” to the target “[given target]”? Select an answer from (favor, against, neutral).  
**Output:** If (reason) then (attitude is [stance label]).

P2: What topics do you think the following inferences [given if-then] are most likely to involve? Please answer with one to three topics.

P3: Please generate two tweets expressing attitude with the following requirements: Tweet expresses attitude to the [given target] with the following reason: [given if-then]

P4: What is the attitude of the sentence: [given text] to the [given target] select from “favor, against or neutral”.

**Decoder.** The aim of the decoder is to enable the LLMs to automatically generate augmented data based on the *if-then* expressions. Notably, inspired by Huang et al. (2023), we propose an additional step of randomly replacing words, termed as random replacement strategy (RRS), in the *if-then* expressions before feeding them into the model. This word replacement prevents the model from generating samples with similar syntax from identical *if-then* expressions, thereby mitigating the problem of producing overly homogeneous data. Specifically, we utilize the SenticNet (Cambria et al., 2022) semantic dictionary to randomly substitute emotion words in the “reason” component of the *if-then* expressions with a fixed probability. Since there is a logical relationship between the reason and decision components in the *if-then* expressions, randomly replacing semantically associated emotion words preserves this relationship while avoiding overly consistent input samples.

Next, we design a three-step prompting method to generate augmented data. First, potential targets are generated based on *if-then* expressions following prompt template P2. Second, text data satisfying the *if-then* logic and conditioned on the given targets is generated adhering to prompt template P3. Third, pseudo-labels are produced based on the generated text-target pairs with prompt P4.

### 3.3. Rationale-Enhanced Network

In the data augmentation process for stance detection, we introduce additional *if-then* expressions. Inspired by Jayaram and Allaway (2021), which shows that prediction rationales can effectively improve stance detection performance, we construct an effective transformer-based stance detection network for ZSSD, termed the REN.

REN comprises a text representation layer and a rationale-guided attention layer. The text representation layer utilizes BERT as the backbone to generate a contextualized representation for each token in the input sample. The rationale-guided attention layer learns sentence representations via attention mechanisms.

**Text representation layer.** First, the text and target are treated as a sentence pair, following the format: [CLS] text [SEP] target [SEP]. This input is then fed into BERT to obtain the hidden state representations  $h_x = \text{BERT}([\text{CLS}] \text{ text } [\text{SEP}] \text{ target } [\text{SEP}])$ . Additionally, the *if-then* expressions are encoded into hidden states  $h_r$  in a similar manner.

**Rationale-guided attention layer.** After acquiring  $h_x$  and  $h_r$ , we feed them into the attention mechanism to obtain rationale-text relation representation:

$$\text{Att} = \text{softmax}\left(\frac{[W_q h_x][W_k h_r]^T}{\sqrt{d_k}}\right)[W_v h_r], \quad (1)$$



where  $W_q$ ,  $W_k$  and  $W_v$  are trainable parameters. Subsequently, we integrate these rationale-text relation representations with the text representation vectors, enabling the rationale knowledge to effectively augment the final text representations, as follows:

$$\hat{y} = \text{softmax}(W_o(\lambda \text{Att} + h_x)). \quad (2)$$

where  $W_o$  denotes trainable parameters.

Finally, we utilize the cross-entropy between the predicted stance  $\hat{y}$  and the ground-truth stance  $y$  as our loss function for stance detection:

$$\mathcal{L} = - \sum_{i=1}^N \sum_{j=1}^C y_{ij} \log \hat{y}_{ij}, \quad (3)$$

where  $N$  represents the number of instances in the training set.  $C$  denotes the number of possible stance categories.

## 4. Experimentation Setup

### 4.1. Experimental Data

To evaluate the effectiveness of our method, we conduct extensive experiments on SemEval-2016 Task 6 (SEM16) (Mohammad et al., 2016) and VAST (Allaway and Mckeown, 2020). SEM16 contains 6 pre-defined targets across multiple domains, including Donald Trump (DT), Hillary Clinton (HC), Feminist Movement (FM), Legalization of Abortion (LA), Atheism (A), and Climate Change (CC). Each instance could be classified as ‘‘Favor’’, ‘‘Against’’ or ‘‘Neutral’’. Following Wei and Mao (2019), we remove targets A and CC due to data quality issues, and regard one target as the zero-shot testing target while training on the other five. randomly select 15% of the training set as the development data to tune hyperparameters. VAST contains a large number of diverse targets. Each data instance is comprised of a sentence, a target, and a stance polarity towards the target, which can be ‘‘Pro’’, ‘‘Con’’, or ‘‘Neutral’’. Following Li et al. (2023), we evaluate our model’s performance on zero-shot topics with 100% and 10% training sizes, respectively.

### 4.2. Experimental Implementation

**Training Settings.** In this paper, we conduct experiments with gpt-3.5-turbo (GPT) and LLaMA2-70B (LLaMA) as LLMs. We utilize Bert-base (Devlin et al., 2019) as the backbone of the rationale-enhanced model. The AdamW optimizer is applied to train the model, with a mini-batch size of 16, dropout of 0.3 and a learning rate set to  $5e^{-6}$ . Following Li et al. (2023), we run the method four times and report the average score for our method.

**Evaluation Metric.** For the SEM16 dataset, following Liang et al. (2022a), we utilize Macro-avg,

the average F1 score on the *Favor* and *Against* classes, as the evaluation metric. For the VAST dataset, following Li et al. (2023), we employ the Macro F1 score, defined as the average F1 score across all labels, to assess model stability. To detect text similarity in the generated content, we utilize SimCSE (Gao et al., 2021) to evaluate semantic similarity, ROUGE (Lin, 2004) to measure content similarity, and the Levenshtein distance (Miller et al., 2009) to quantify structural similarity.

### 4.3. Compared Baseline Methods

To evaluate the effectiveness of our proposed EDDA framework, we compare EDDA with a series of strong baselines, including statistics-based methods: BiLSTM & Bicond (Augenstein et al., 2016), CrossNet (Xu et al., 2018), SEKT (Zhang et al., 2020), TPDG (Liang et al., 2021) and TOAD (Allaway et al., 2021). Fine-tuning based methods: Bert-Joint (Devlin et al., 2019), JointCL (Liang et al., 2022b) and Bert-GCN (Liu et al., 2021). Prompt-tuning based methods: MPT (Hu et al., 2022) and KEPrompt (Huang et al., 2023). Knowledge-based method: TarBK (Zhu et al., 2022). Data augmentation methods: PT-HCL (Liang et al., 2022a), SDAgu (Zhang et al., 2023c), TTS (Li et al., 2023) and OpenStance (Xu et al., 2022). For LLMs, we compare LLMs-EDDA with LLMs, the latter employing a zero-shot prompt for ZSSD (Zhang et al., 2023b).

**Statistics-based methods.** BiLSTM utilized a bidirectional-LSTM to encode the underlying sentence and the corresponding target. Bicond (Augenstein et al., 2016) utilized two BiLSTM models to separately encode the underlying sentences and their corresponding targets. CrossNet (Xu et al., 2018) is a variant of BiCond, which leverages a self-attention layer to capture informative words. TPDG (Liang et al., 2021) proposed a target-adaptive graph convolutional network. TOAD (Allaway et al., 2021) conducted adversarial learning to generalize across topics.

**Fine-tuning based methods.** Bert-Joint (Devlin et al., 2019) employed Bert for encoding both text and topics, followed by two successive fully connected layers. JointCL (Liang et al., 2022b) proposed a contrastive learning method to leverage the stance features of known targets. Bert-GCN (Liu et al., 2021) utilized node information aggregation to create Graph Convolutional Networks (GCN) and incorporated it into the model. TarBK (Zhu et al., 2022) incorporated the targeted background knowledge from Wikipedia into the model for stance detection. SEKT (Zhang et al., 2020) introduced semantic knowledge as the transferable knowledge between domains.

**Prompt-tuning based methods.** MPT (Hu et al., 2022) utilized a prompt-tuning method for stance detection, which employs a verbalizer defined by

	Model	SEM16				VAST (100%)			VAST (10%)		
		HC	FM	LA	DT	Pro	Con	All	Pro	Con	All
Sta.	BiLSTM	31.6	40.3	33.6	30.8	43.7	43.6	37.5	39.2	24.6	32.9
	Bicond	32.7 <sup>‡</sup>	40.6 <sup>‡</sup>	34.4 <sup>‡</sup>	30.5 <sup>‡</sup>	44.6 <sup>‡</sup>	47.4 <sup>‡</sup>	42.8 <sup>‡</sup>	29.8 <sup>◇</sup>	40.1 <sup>◇</sup>	34.8 <sup>◇</sup>
	CrossNet	38.3 <sup>‡</sup>	41.7 <sup>‡</sup>	38.5 <sup>‡</sup>	35.6 <sup>‡</sup>	46.2 <sup>‡</sup>	43.4 <sup>‡</sup>	43.4 <sup>‡</sup>	37.3 <sup>◇</sup>	32.9 <sup>◇</sup>	36.2 <sup>◇</sup>
	SEKT	50.1	44.2	44.6	46.8	50.4 <sup>‡</sup>	44.2 <sup>‡</sup>	41.8 <sup>‡</sup>	-	-	-
	TPDG	50.9 <sup>‡</sup>	53.6 <sup>‡</sup>	46.5 <sup>‡</sup>	47.3 <sup>‡</sup>	53.7 <sup>‡</sup>	49.6 <sup>‡</sup>	51.9 <sup>‡</sup>	-	-	-
	TOAD	51.2 <sup>‡</sup>	54.1 <sup>‡</sup>	46.2 <sup>‡</sup>	49.5 <sup>‡</sup>	42.6 <sup>‡</sup>	36.7 <sup>‡</sup>	41.0 <sup>‡</sup>	-	-	-
Bert	TGA Net	49.3 <sup>‡</sup>	46.6 <sup>‡</sup>	45.2 <sup>‡</sup>	40.7 <sup>‡</sup>	53.7	62.0	67.4	47.6 <sup>◇</sup>	58.2 <sup>◇</sup>	64.1 <sup>◇</sup>
	Bert-Joint	50.1	42.1	44.8	41.0	56.8	67.6	71.0	50.7	56.3	64.9
	Bert-GCN	50.0 <sup>‡</sup>	44.3 <sup>‡</sup>	44.2 <sup>‡</sup>	42.3 <sup>‡</sup>	58.3 <sup>‡</sup>	60.6 <sup>‡</sup>	68.6 <sup>‡</sup>	-	-	-
	JointCL	54.4	54.0	50.0	50.5	64.9	63.2	71.2	53.8 <sup>◇</sup>	57.1 <sup>◇</sup>	65.5 <sup>◇</sup>
	TarBK	55.1 <sup>ℵ</sup>	53.8 <sup>ℵ</sup>	48.7 <sup>ℵ</sup>	50.8 <sup>ℵ</sup>	65.7 <sup>ℵ</sup>	63.9 <sup>ℵ</sup>	73.6 <sup>ℵ</sup>	-	-	-
	PT-HCL	54.5 <sup>‡</sup>	54.6 <sup>‡</sup>	50.9 <sup>‡</sup>	50.1 <sup>‡</sup>	61.7 <sup>‡</sup>	63.5 <sup>‡</sup>	71.6 <sup>‡</sup>	-	-	-
	SDAgu	-	-	-	-	60.1	65.2	71.3	-	-	-
	Openstance*	-	-	-	-	63.1	66.4	73.1	61.4	60.1	70.2
	TTS*	-	-	-	-	59.5	65.2	71.4	55.8	65.5	70.3
	TDDA-GPT	53.3	54.5	54.2	51.5	61.8	68.7	73.4	56.5	66.0	70.3
	TDDA-LLaMA	-	-	-	-	63.4	67.5	73.4	61.0	63.9	70.4
	EDDA-GPT	<b>77.4<sup>†</sup></b>	<b>69.7<sup>†</sup></b>	<b>62.7<sup>†</sup></b>	<b>69.8<sup>†</sup></b>	66.9 <sup>†</sup>	68.2 <sup>†</sup>	75.1 <sup>†</sup>	62.6 <sup>†</sup>	66.9 <sup>†</sup>	72.7 <sup>†</sup>
EDDA-LLaMA	-	-	-	-	<b>68.3<sup>†</sup></b>	<b>70.7<sup>†</sup></b>	<b>76.3<sup>†</sup></b>	<b>61.7<sup>†</sup></b>	<b>67.8<sup>†</sup></b>	<b>73.2<sup>†</sup></b>	
LLMs	GPT	78.9	68.3	62.3	68.6	63.8	56.8	65.1	63.8	56.8	65.1
	GPT-EDDA	<b>80.1</b>	<b>69.2</b>	<b>62.4</b>	<b>69.5</b>	65.3	64.2	68.5	65.3	64.2	68.5
	LLaMA	-	-	-	-	66.8	63.4	69.0	66.8	63.4	69.0
	LLaMA-EDDA	-	-	-	-	<b>67.7</b>	<b>63.1</b>	<b>70.3</b>	<b>67.7</b>	<b>63.1</b>	<b>70.3</b>

Table 1: Experimental results on two zero-shot stance detection (ZSSD) datasets. The results with <sup>‡</sup>, <sup>◇</sup> and <sup>ℵ</sup> are retrieved from Liang et al. (2022a), Li et al. (2023) and Zhang et al. (2023c), respectively. The \* mark indicates that we utilize Bert as the backbone classifier for a fair comparison. The <sup>†</sup> mark refers to a  $p$ -value  $< 0.05$ . The best scores are in bold. Due to policy restrictions, LLaMA cannot be applied to controversial targets, so experiments with LLaMA are limited to the non-controversial VAST dataset.

human experts. KEPrompt (Huang et al., 2023) introduced external lexicons to define the verbalizer for the prompt framework.

**Data Augmentation methods.** PT-HCL (Liang et al., 2022a) developed a novel approach to cross-target and zero-shot stance detection using contrastive learning. SDAgu (Zhang et al., 2023c) proposed a self-supervised data augment approach based on coreference resolution. TTS (Li et al., 2023) proposed to augment the training set with different diverse targets. TDDA is our proposed text-based data augmentation technique. It entails the random extraction of seven instances from the original dataset, with the substitution of three instances at every third iteration.

	Models	F→L	L→F	H→D	D→H
Sta.	BiCond	45.0	41.6	29.7	35.8
	CrossNet	45.4	43.3	43.1	36.2
	SEKT	53.6	51.3	47.7	42.0
	TPDG	58.3	54.1	50.4	52.9
Bert	Bert-Joint	48.2	34.4	45.9	42.7
	MPT	42.1	47.6	47.1	58.7
	KEPROMPT	49.1	54.2	54.6	60.9
	JointCL	58.8	54.5	52.8	54.3
	PT-HCL	59.3	54.6	53.7	55.3
	TarBK	59.1	54.6	53.1	54.2
	EDDA-GPT	<b>61.7<sup>†</sup></b>	<b>65.4<sup>†</sup></b>	<b>64.5<sup>†</sup></b>	<b>76.2<sup>†</sup></b>

Table 2: Performance comparison of cross-target stance detection.

## 5. Experimental Results

### 5.1. Result

The main experimental results of ZSSD on two benchmark datasets are reported in Table 1. We observe that our method consistently outperforms all baselines on all datasets, which verifies the effectiveness of our proposed approach in ZSSD. Furthermore, compared to previous models, our method shows statistically significant improvements ( $p$ -value  $< 0.05$ ) in most evaluation metrics, validating the effectiveness of our proposed approach in ZSSD.

Specifically, we first observe that most neural-network based methods (statistics-based and fine-tuning based methods) perform poorly under zero-shot settings, owing to their reliance on seen targets and samples. Second, contrastive learning methods achieve moderate performance gains compared to traditional neural-network based methods. This is attributed to their ability to learn target-related and target-agnostic features separately, thereby improving model performance. Furthermore, data augmentation methods significantly outperform other contrastive approaches, especially with limited training data (VAST 10%). This demonstrates that data augmentation enhances the model’s generalization to unseen targets by

TDDA	EDDA
I think it's important to focus on investing in renewable energy sources, but we also can't abandon gas and oil completely right away. We need to transition thoughtfully.	The alarming consequences of relying on fossil fuels are staring us in the face. From air pollution to climate change, it's clear that we need to shift gears towards sustainable energy options. Let's leave the fossil fuel era behind and embrace a greener tomorrow.
I strongly believe in the power of art and culture to unite communities and promote understanding.	Absolutely love how studying language helps unravel the intricate layers of culture! It's like a key that unlocks a whole new world of understanding.
I believe that the minimum wage should be raised to provide workers with fair compensation for their labor and to decrease income inequality.	The federal minimum wage is a joke. How can anyone survive on such a meager income? And let's not forget the tipped employees who are being exploited with a subminimum wage. This needs to change NOW!

Table 3: Examples of text generated by TDDA and EDDA.

expanding the data.

The proposed EDDA method yields superior performance over all baselines in most tasks. For example, our method improves 2.7% on VAST and 2.8% on VAST (10%) over the best competitors. Notably, our model shows substantial improvements compared to target augmentation methods like TTS and SDAug. This indicates that solely applying data augmentation techniques tailored for seen targets does not generalize well to unseen targets. Furthermore, our generative strategy considerably outperforms TDDA, owing to TDDA's reliance on training text. This over-dependence causes TDDA's generated samples to exhibit increasingly consistent syntactic patterns as the sample size grows, resulting in limited diversity. By contrast, our proposed strategy encourages greater syntactic variety by incorporating *if-then* augmentation. We posit this promotes the model's ability to generalize across diverse unseen targets.

The results indicate that the EDDA chain of thought enhances the performance of LLMs. This enhancement can be attributed to the ability of *if-then* prompts to better align stance explanation and stance classification, thereby strengthening the reasoning capabilities of LLMs. Additionally, models fine-tuned using EDDA demonstrate an enhanced comprehension of stance knowledge through task-specific supervision, thereby bolstering stance detection performance built upon the foundation of LLMs. To evaluate the generalizability of our EDDA framework, we also evaluate EDDA in the cross-target condition on SEM16. From the experimental results shown in Table 2, we can see that EDDA performs overall better than all the comparison methods under the cross-target condition. This verifies the effectiveness and generalizability of EDDA in dealing with stance detection.

## 5.2. Analysis

**Analysis of Data Distributions** To qualitatively demonstrate how the proposed EDDA generative framework improves predictive performance, we present t-SNE (Van der Maaten and Hinton, 2008) visualizations of the intermediate vectors learned by TTS, OpenStance, TDDA and EDDA on the Vast

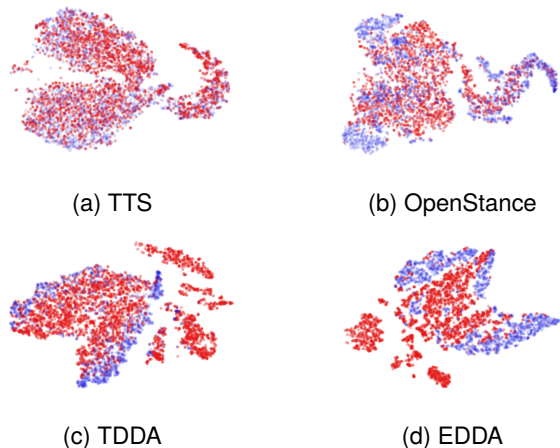


Figure 3: Visualization of intermediate vectors in Bert, where red represents the original training set and blue represents the augmented dataset.

dataset. The results are shown in Figure 3. Specifically, comparing EDDA (Figure 3.d) with TTS (Figure 3.a) reveals that the data augmented by the target augmentation method highly overlaps with the training samples, making it difficult to effectively generalize to an unseen domain. Second, compared to the text augmentation methods of OpenStance (Figure 3.b) and TDDA (Figure 3.c), the data distribution generated by the EDDA method differs substantially more from the original training samples. This implies that by summarizing from *if-then* expressions, EDDA can produce more generalized data, thereby better extending to unseen targets and enhancing ZSSD performance.

	SimCSE <sup>aug</sup> ↓	SimCSE <sup>test</sup> ↑	ROUGE ↓	Levenshtein ↑
TDDA	0.466	0.403	0.195	190.6
EDDA	0.443	0.432	0.152	245.1

Table 4: Text similarity comparison. <sup>aug</sup> denotes the similarity between the augmented texts, <sup>test</sup> indicates the similarity between the augmented texts and the actual test samples.

**Analysis of Textual Features** To further understand the specific differences between TDDA and

EDDA, we provide example texts generated by both methods in Table 3. Despite prompts to increase diversity, TDDA produced simple texts with limited linguistic variance, making them less realistic. In contrast, the texts from EDDA contain rich sentiments, discuss relevant issues, and use more sentence structure. Consequently, EDDA leads to more significant performance gains.

Quantitative analysis further revealed the advantages of EDDA over TDDA. The results are presented in Table 4. Here, the smaller the SimCSE<sup>aug</sup> and ROUGE, the better performance a model delivers. We add the downward arrow symbol “↓” to indicate this. EDDA texts exhibited greater diversity, as evidenced by lower similarity scores. Specifically, we utilized SIMCSE, ROUGE and Levenshtein metrics to compute similarity scores. From the dataset, 300 instances were randomly sampled, and each metric was calculated over 10 iterations to obtain average values. From the results, EDDA also demonstrated better generalization and covered wider content with richer syntactic variations. In summary, both qualitative and quantitative analyses indicate EDDA’s superiority in producing varied, naturalistic texts compared to TDDA.

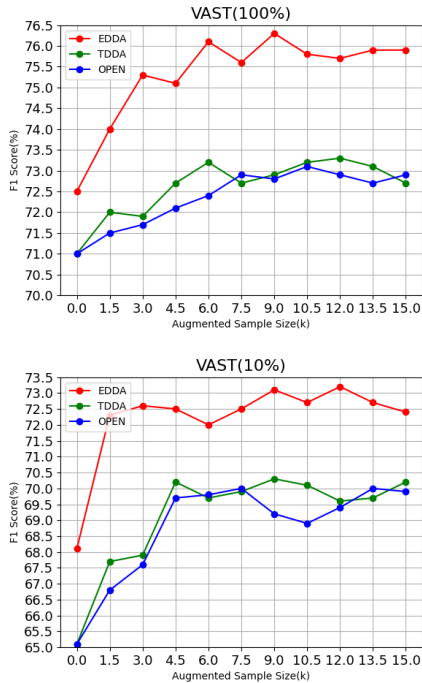


Figure 4: The experimental results with respect to varying augmentation data size.

**Analysis of Sample Size** To analyze how increasing the number of text samples affects model performance, we conduct experiments using the VAST dataset with 100% and 10% training data. Figure 4 shows the performance changes after augmenting the samples. With the full training set,

accuracy improved steadily as more augmented data was added, peaking at around 6000 samples. It then declined and stabilized. Compared to the baseline data generation methods, EDDA consistently improves accuracy. Moreover, under the 10% sample setting, EDDA leads to even larger gains. This demonstrates the efficacy of our proposed method, especially when training data is scarce.

### 5.3. ZSSD Baselines with EDDA

We evaluated whether our proposed EDDA method could improve other existing ZSSD models, including BiLSTM, CrossNet, TGA-Net, Bert-Joint and JointCL. Table 5 shows the results, where “+EDDA” signifies the integration of our enhanced data, and “+ L” denotes the incorporation of our generative *if-then* expressions. The results demonstrate that our approach can effectively improve the performance of other baseline models for ZSSD. This indicates that our approach is not limited to a specific model but presents a generalizable solution for this task.




Figure 5: Comparison of EDDA with data augmentation baselines for previous ZSSD task.

### 5.4. Comparison with Data Augmentation Baselines

To further demonstrate the effectiveness of EDDA, following Li et al. (2023), we compare EDDA with other data augmentation baselines. We use the same base model(Bart) for all comparisons and conduct experiments on the VAST dataset. The results are shown in Figure 5, where TDDA exhibited only a slight improvement, while EDDA significantly outperformed other data augmentation methods. The red dashed line represents the base model.

First, our experiments demonstrate that traditional data augmentation techniques (e.g. +BT, EDA, +CBert and +ASDA) are ineffective for the ZSSD task and can even degrade performance. A potential explanation is that word-level perturbations may disrupt the semantic structure of the text. Second, we find that target-based augmentation methods relying on the LLMs are underper-



formed by text-based techniques. The inferior results could arise because extracting target words directly from text can cause logical and semantic inconsistencies between texts and targets, limiting improvements. Third, EDDA and TDDA show significant improvements over the baseline, indicating the effectiveness of our proposed encoder-decoder augmentation as a novel augmentation for ZSSD.

Model	Pro	Con	All
BiLSTM	43.7	43.6	37.5
BiLSTM + L	60.0	64.4	65.3
BiLSTM + EDDA	47.2	45.8	41.7
BiLSTM + EDDA + L	61.0	62.6	68.1
CrossNet	40.5	44.4	42.9
CrossNet + L	53.3	61.8	66.3
CrossNet + EDDA	46.0	39.0	43.5
CrossNet + EDDA + L	58.7	62.6	68.2
TGA-Net	53.7	62.0	67.4
TGA-Net + L	55.6	62.5	68.9
TGA-Net + EDDA	58.7	62.9	69.6
TGA-Net + EDDA + L	60.9	62.5	70.2
Bert-Joint	56.8	67.6	71.0
Bert-Joint + L	58.3	68.9	72.1
Bert-Joint + EDDA	62.8	67.8	73.8
Bert-Joint + EDDA + L	64.9	66.8	74.3
JointCL	64.9	63.2	71.2
JointCL + L	62.8	66.2	72.5
JointCL + EDDA	62.3	70.1	73.9
JointCL + EDDA + L	65.2	69.5	74.4
Ours	61.5	64.2	70.8
Ours + L	62.5	67.4	73.1
Ours + EDDA	62.9	69.8	73.9
Ours + EDDA + L	<b>68.3</b>	<b>70.7</b>	<b>76.3</b>

Table 5: Performance of our EDDA applied to other zero-shot baselines.

## 6. Conclusion

ZSSD determines the text’s attitude towards an unseen target. Recent data augmentation techniques exhibit limitations due to insufficient generalization. To address this, we propose an encoder-decoder data augmentation framework (EDDA). The Encoder generates target-specific if-then rationales using LLMs, establishing logical relationships. The Decoder increases syntactic diversity of new samples based on these expressions using word replacement. We also develop a rationale-enhanced network to fully utilize augmented data. Experiments show EDDA substantially outperforms state-of-the-art by increasing semantic relevance and syntactic variety while enabling interpretable rationale-based learning. In future work, we intend to establish a filtering framework to eliminate low-quality or erroneous generated instances. Additionally, we aim to integrate the results from multiple large language models to construct a high-quality augmented dataset.

## 7. Acknowledgements

Bowen Zhang was supported by National Nature Science Foundation of China (No.62306184), Natural Science Foundation of Top Talent of SZTU (grant no. GDRC202320). Xiaojiang Peng was supported by the National Natural Science Foundation of China (62176165), the Stable Support Projects for Shenzhen Higher Education Institutions (20220718110918001), the Natural Science Foundation of Top Talent of SZTU(GDRC202131).

## 8. Bibliographical References

- Emily Allaway and Kathleen Mckeown. 2020. Zero-shot stance detection: A dataset and model using generalized topic representations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8913–8931.
- Emily Allaway, Malavika Srikanth, and Kathleen Mckeown. 2021. Adversarial learning for zero-shot stance detection on social media. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4756–4767.
- I Augenstein, T Rocktaeschel, A Vlachos, and K Bontcheva. 2016. Stance detection with bidirectional conditional encoding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Sheffield.
- Zefan Cai, Baobao Chang, and Wenjuan Han. 2023. Human-in-the-loop through chain-of-thought. *arXiv preprint arXiv:2306.07932*.
- Erik Cambria, Qian Liu, Sergio Decherchi, Frank Xing, and Kenneth Kwok. 2022. Senticnet 7: a commonsense-based neurosymbolic ai framework for explainable sentiment analysis. *Proceedings of LREC 2022*.
- Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Yihan Cao, Zihao Wu, Lin Zhao, Shaochen Xu, Wei Liu, Ninghao Liu, et al. 2023. Auggpt: Leveraging chatgpt for text data augmentation. *arXiv preprint arXiv:2302.13007*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics.
- Daijun Ding, Rong Chen, Liwen Jing, Bowen Zhang, Xu Huang, Li Dong, Xiaowen Zhao, and Ge Song.

2024. Cross-target stance detection by exploiting target analytical perspectives. *arXiv preprint arXiv:2401.01761*.
- Jiachen Du, Ruifeng Xu, Yulan He, and Lin Gui. 2017. Stance classification with target-specific neural attention networks. *International Joint Conferences on Artificial Intelligence*.
- Aleksandra Edwards, Asahi Ushio, Jose Camacho-Collados, H el ene de Ribaupierre, and Alun Preece. 2021. Guiding generative language models for data augmentation in few-shot text classification. *arXiv preprint arXiv:2111.09064*.
- Hao Fei, Bobo Li, Qian Liu, Lidong Bing, Fei Li, and Tat-Seng Chua. 2023. Reasoning implicit sentiment with chain-of-thought prompting. *arXiv preprint arXiv:2305.11255*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Zihao He, Negar Mokherian, and Kristina Lerman. 2022. Infusing wikipedia knowledge to enhance stance detection. *arXiv preprint arXiv:2204.03839*.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. 2022. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2225–2240.
- Hu Huang, Bowen Zhang, Yangyang Li, Baoquan Zhang, Yuxi Sun, Chuyao Luo, and Cheng Peng. 2023. Knowledge-enhanced prompt-tuning for stance detection. *ACM Transactions on Asian and Low-Resource Language Information Processing*.
- Sahil Jayaram and Emily Allaway. 2021. Human rationales as attribution priors for explainable stance detection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5540–5554.
- Yingjie Li and Cornelia Caragea. 2021. Target-aware data augmentation for stance detection. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1850–1860.
- Yingjie Li, Chenye Zhao, and Cornelia Caragea. 2023. Tts: A target-based teacher-student framework for zero-shot stance detection. In *Proceedings of the ACM Web Conference 2023*, pages 1500–1509.
- Bin Liang, Zixiao Chen, Lin Gui, Yulan He, Min Yang, and Ruifeng Xu. 2022a. Zero-shot stance detection via contrastive learning. In *Proceedings of the ACM Web Conference 2022*, pages 2738–2747.
- Bin Liang, Yonghao Fu, Lin Gui, Min Yang, Jiachen Du, Yulan He, and Ruifeng Xu. 2021. Target-adaptive graph for cross-target stance detection. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23*, pages 3453–3464.
- Bin Liang, Qinlin Zhu, Xiang Li, Min Yang, Lin Gui, Yulan He, and Ruifeng Xu. 2022b. Jointcl: A joint contrastive learning framework for zero-shot stance detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 81–91. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Zhan Ling, Yunhao Fang, Xuanlin Li, Zhiao Huang, Mingu Lee, Roland Memisevic, and Hao Su. 2023. Deductive verification of chain-of-thought reasoning. *arXiv preprint arXiv:2306.03872*.
- Rui Liu, Zheng Lin, Yutong Tan, and Weiping Wang. 2021. Enhancing zero-shot and few-shot stance detection with commonsense knowledge graph. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3152–3157.
- Yun Luo, Zihan Liu, Yuefeng Shi, Stan Z Li, and Yue Zhang. 2022. Exploiting sentiment and common sense for zero-shot stance detection. *arXiv preprint arXiv:2208.08797*.
- Frederic P Miller, Agnes F Vandome, and John McBrewhster. 2009. Levenshtein distance: Information theory, computer science, string (computer science), string metric, damerau? levenshtein distance, spell checker, hamming distance.
- Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. Semeval-2016 task 6: Detecting stance in tweets. In *Proceedings of the International Workshop on Semantic Evaluation, SemEval '16, San Diego, California*.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).

- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed H Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.
- Jason Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6382–6388.
- Penghui Wei and Wenji Mao. 2019. Modeling transferable topics for cross-target stance detection. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1173–1176. ACM.
- Chang Xu, Cecile Paris, Surya Nepal, and Ross Sparks. 2018. Cross-target stance classification with self-attention networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 778–783.
- Hanzi Xu, Slobodan Vucetic, and Wenpeng Yin. 2022. Openstance: Real-world zero-shot stance detection. *arXiv preprint arXiv:2210.14299*.
- Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyeong Park. 2021. Gpt3mix: Leveraging large-scale language models for text augmentation. *arXiv preprint arXiv:2104.08826*.
- Bowen Zhang, Daijun Ding, Guangning Xu, Jinjin Guo, Zhichao Huang, and Xu Huang. 2023a. Twitter stance detection via neural production systems. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Bowen Zhang, Xianghua Fu, Daijun Ding, Hu Huang, Yangyang Li, and Liwen Jing. 2023b. Investigating chain-of-thought with chatgpt for stance detection on social media. *arXiv preprint arXiv:2304.03087*.
- Bowen Zhang, Min Yang, Xutao Li, Yunming Ye, Xiaofei Xu, and Kuai Dai. 2020. Enhancing cross-target stance detection with transferable semantic-emotion knowledge. Association for Computational Linguistics.
- Jiarui Zhang, Shaojuan Wu, Xiaowang Zhang, and Zhiyong Feng. 2023c. Task-specific data augmentation for zero-shot and few-shot stance detection. In *Companion Proceedings of the ACM Web Conference 2023*, pages 160–163.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.
- Qinglin Zhu, Bin Liang, Jingyi Sun, Jiachen Du, Lanjun Zhou, and Ruifeng Xu. 2022. Enhancing zero-shot stance detection via targeted background knowledge. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2070–2075.