

# Stanceformer: Target-Aware Transformer for Stance Detection

Krishna Garg      Cornelia Caragea  
kgarg8@uic.edu    cornelia@uic.edu  
Computer Science  
University of Illinois Chicago

## Abstract

The task of Stance Detection involves discerning the stance expressed in a text towards a specific subject or target. Prior works have relied on existing transformer models that lack the capability to prioritize targets effectively. Consequently, these models yield similar performance regardless of whether we utilize or disregard target information, undermining the task's significance. To address this challenge, we introduce Stanceformer, a target-aware transformer model that incorporates enhanced attention towards the targets during both training and inference. Specifically, we design a *Target Awareness* matrix that increases the self-attention scores assigned to the targets. We demonstrate the efficacy of the Stanceformer with various BERT-based models, including state-of-the-art models and Large Language Models (LLMs), and evaluate its performance across three stance detection datasets, alongside a zero-shot dataset. Our approach Stanceformer not only provides superior performance but also generalizes even to other domains, such as Aspect-based Sentiment Analysis. We make the code publicly available.<sup>1</sup>

## 1 Introduction

Stance detection is the task that involves determining the stance or perspective expressed in a given text towards a particular topic or target. The stance could be either favor, against, or neutral with respect to the target. It has many useful applications, including prediction of election/referendum results (Lai et al., 2018), analyzing online debates (Bar-Haim et al., 2017), social media analysis (Zhang et al., 2017), rumor verification (Gorrell et al., 2019; Derczynski et al., 2017), rumor classification (Zubiaga et al., 2018), automated fact-checking (Ferreira and Vlachos, 2016; Vlachos and Riedel, 2014), fake news detection (Lazer et al., 2018), and information retrieval (Sen et al., 2018).

**Text:** a woman ?? wanting to be equal to a man ???! what montrosity is this  
**Target:** feminist movement  
**Stance:** Against

Figure 1: Example of Stance Detection task

Targets play a pivotal role in the stance detection task because they define the context and subject matter of the stance expressed in a text. However, several studies (Yuan et al., 2022; Kaushal et al., 2021) have identified that the models tend to overlook the targets when making predictions. This challenges the definition of the stance detection task, which is inherently tied to a target (the author of a text takes a stance towards a specific target). Therefore, it is crucial to ensure that models are aware of the targets when making decisions to improve their performance.

To address these issues, researchers have proposed techniques such as data augmentation (Kaushal et al., 2021; Li and Caragea, 2021b; Li et al., 2023b), leveraging external knowledge bases (Liu et al., 2021; He et al., 2022; Zhang et al., 2021), employing human-like reasoning (Yuan et al., 2022; Jayaram and Allaway, 2021), and using contrastive learning-based methods (Liang et al., 2022a; Liu et al., 2022) to improve target representations during training. Additionally, efforts have been made to enhance target-specific attention for BiLSTM-based models (Xu et al., 2018; Du et al., 2017; Augenstein et al., 2016). However, we develop mechanism for enhancing target-specific attention in more sophisticated transformer models, as well as two state-of-the-art methods for stance detection (Liang et al., 2022a; He et al., 2022).

Specifically, in this study, we introduce a novel target-specific attention mechanism designed for transformer models. The aim is to guide the models in incorporating additional attention

<sup>1</sup><https://github.com/kgarg8/Stanceformer>

to the specified target for each text. We add a special *Target-Awareness* (TA) matrix, which is incorporated into the self-attention matrix within a specific layer and head of the transformer model. The sparse TA matrix contains non-zero values only for the square block covering the target subwords (as illustrated in Figure 2).

Experimental results demonstrate consistent improvements with our approach compared to the baseline models across various models and datasets, including both full-dataset and zero-shot dataset settings. Our approach seamlessly integrates with any off-the-shelf stance detection model that uses transformer architectures. Additionally, we apply our approach to various LLMs, including Llama-2-chat models with 7 billion and 13 billion parameters, in both zero-shot inference and finetuned settings. Our approach outperforms both BERT-based language models and autoregressive LLMs. Finally, we also showcase the generalization capabilities of our method in other domains, such as Aspect-based Sentiment Analysis.

We summarize our contributions as follows:

1. We introduce Stanceformer, a specialized transformer model designed to enhance attention towards the targets in the Stance Detection task. We specifically propose Target Awareness matrix for enhancing the attention on the targets in Stanceformer.
2. We empirically demonstrate that Stanceformer can effectively replace existing transformer models utilized in state-of-the-art methods while enhancing performance, across both full-dataset and zero-shot-dataset settings.
3. To the best of our knowledge, we are the first to finetune LLMs for Stance Detection. In this work, we finetune Llama-2-chat models across four stance detection datasets and showcase superior performance of the Stanceformer for each dataset.
4. We demonstrate the generalization capabilities of our method in other domains, such as Aspect-based Sentiment Analysis.

## 2 Related Work

**Target-Awareness** The stance detection task involves two key inputs: the text or tweet, and the target. Numerous research efforts have focused on

leveraging target information to enhance the performance of stance detection models. We classify these approaches into four main categories.

*Data Augmentation, Open-domain/Open-targets:* Kaushal et al. (2021) introduced a novel dataset that emphasizes the role of target entities in stance detection. They achieved this by augmenting the dataset with samples that have negated (or unrelated) targets and opposite (or *Unrelated*) stance labels. Li and Caragea (2021b) proposed an auxiliary sentence-based data augmentation technique that considers both the target and label, as well as target-specific replacement, to generate augmented samples. Li et al. (2023b) proposed to augment a dataset with additional targets to improve the stance detection task. The authors first augment the training set by extracting diverse targets (unseen during training) and then predict the stance by effectively utilizing the augmented data. Li et al. (2023a) argued that the targets for stance are not always available and proposed an approach called Target-Stance Extraction that aims to extract (target, stance) pairs from the text. Other works (Zhao and Caragea, 2024; Zhao et al., 2024, 2023; Xu et al., 2022; Allaway and McKeown, 2020) address open-domain/open-target for stance detection.

*External Knowledge:* Several studies have explored the integration of external knowledge bases to provide supplementary information about the targets. Examples of such studies include those by Liu et al. (2021) who leverage ConceptNet and by He et al. (2022) who use Wikipedia to enhance target-related knowledge. Zhang et al. (2021) used external knowledge bases such as ConceptNet and WikiData as bridges to establish explicit links between opinionated words and the targets of interest.

*Target-Specific Attention-Based Models:* Prior to the advent of transformer-based models, researchers proposed attention mechanisms between the text and target using models like BiLSTM. Augenstein et al. (2016) introduced bidirectional encoding of texts conditioned on the bidirectional encoding of corresponding targets. Xu et al. (2018) incorporated a self-attention layer on BiLSTMs to capture domain-related aspects for cross-target generalization. Du et al. (2017) employed target-augmented embeddings and extracted target-specific attention signals to emphasize the salient parts of the input text.

*Miscellaneous Approaches:* Jiang et al. (2022) utilized target-aware prompts to incorporate contextual information about the targets. Liang et al.

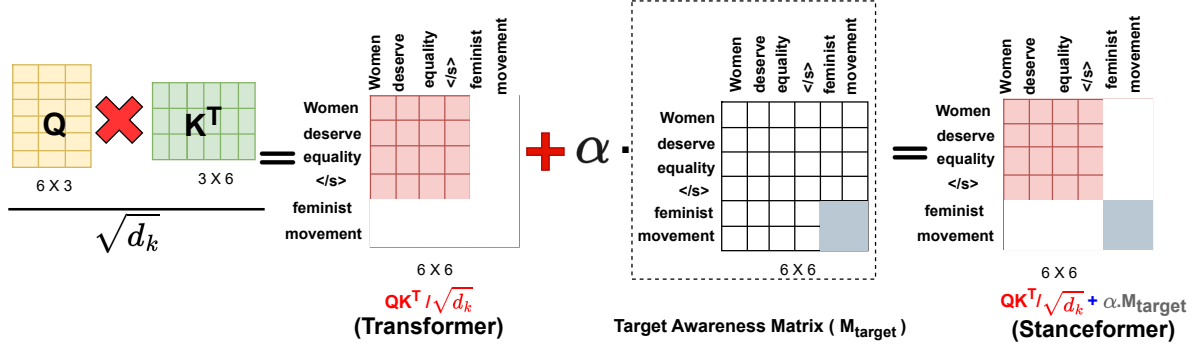


Figure 2: Illustration of Stanceformer model, which adds Target Awareness Matrix to the self-attention scores.

(2022a); Liu et al. (2022) employed contrastive learning and consistency regularization, to improve target representations, particularly in low-resource settings. Liang et al. (2022b) proposed target-aware prototypical networks, while Li and Caragea (2019) adopted a multi-tasking approach using stance and sentiment lexicons. Other works distilled knowledge from a teacher to a student model (Li and Caragea, 2023; Li et al., 2021b) or used multi-task learning (Li and Caragea, 2021a) to improve the performance of stance detection.

However, our work presents a novel approach to enhance target-awareness in stance detection. Unlike existing approaches, we do not use data augmentation, unlabeled data, or external knowledge bases. Instead, we propose a new mechanism within the family of transformer models, tailored for the stance detection task. Moreover, all the aforementioned methods based on transformers can benefit from our Stanceformer, which serves as an improved alternative to the default transformer.

**Large Language Models (LLMs)** A plethora of LLMs, such as GPT-3 (Brown et al., 2020), GPT-3.5, GPT-4 (Achiam et al., 2023), Llama-2 (Touvron et al., 2023), FLAN (Wei et al., 2021), PaLM (Chowdhery et al., 2023), and numerous others, have recently emerged. Several studies (Zhang et al., 2023, 2022; Gatto et al., 2023; Anonymous, 2023) have delved into the use of ChatGPT and Llama-2 models for Stance Detection, primarily with SemEval-2016 and PStance datasets. However, in this work, we employ GPT-3.5 (closed-source) and Llama-2-chat (open-source), both 7 billion and 13-billion parameter models for four stance datasets. We incorporate some of these prior works as our baselines. In addition to zero-shot baselines, we also finetune LLMs for Stance Detection in this work.

### 3 Proposed Approach

In this section, we describe our proposed method Stanceformer, designed to enhance the target awareness of the transformer models. Figure 2 provides a visual illustration of our proposed approach. In the following sections, we first formally define the Stance Detection task in §3.1, followed by a detailed explanation of the core component of Stanceformer, namely, the Target Awareness Matrix in §3.2. Finally, we describe the process of finetuning with Stanceformer in §3.3.

#### 3.1 Task Description

Consider  $D_{\text{train}}$  as the set of training data with  $n$  samples, defined as  $\{(x_i, t_i, y_i)\}_{i=1}^n$ , where  $x_i = [x_{1i}, x_{2i}, \dots, x_{li}]$  represents a text sequence comprising  $l$  words. Here,  $t_i = [t_{1i}, t_{2i}, \dots, t_{pi}]$  denotes the target sequence of  $p$  words associated with  $x_i$ .  $y_i \in \{1, \dots, c\}$  signifies the stance label expressed in the text towards the target. Similarly, let  $D_{\text{test}}$  be the set of test data with  $m$  samples, represented as  $\{(x_j, t_j)\}_{j=1}^m$ . Here,  $x_j$  denotes the test sequence, and  $t_j$  denotes the test target. The goal of the Stance Detection task is to predict the stance label  $y_j \in \{1, \dots, c\}$ , given  $x_j$  and  $t_j$ .

For a given sample  $x_i$ , we have two input sequences  $x_i$  as text and  $t_i$  as target. We concatenate the sequences as  $[\text{CLS}]x_i[\text{SEP}]t_i[\text{SEP}]$ , or equivalently,  $[\text{CLS}], x_{1i}, \dots, x_{li}, [\text{SEP}], t_{1i}, \dots, t_{pi}, [\text{SEP}]$ . This is fed as input to the given transformer model. We denote the total sequence length as  $\text{seq} = 1 + l + 1 + p + 1$ .

#### 3.2 Target-Awareness Matrix

Self-attention is a crucial mechanism that allows the transformer-based models to weigh the importance of different words or tokens in a sequence when processing the input. The self-attention (SA)

in the transformer models is given by:

$$SA = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) * V \quad (1)$$

where  $Q$ ,  $K$ ,  $V$  denote the query, key and value matrices, each having a dimension of  $(seq * d_k)$ , where  $seq$  denotes the sequence length of the entire sequence fed to the transformer model, and  $d_k$  denotes the dimension of the key vector. Tokens with higher attention weights receive more focus during the computation of the weighted sum of the value vectors.

Next, we introduce the *Target Awareness* matrix  $M_{target}$  to further enhance the attention on the targets. We define it as a square matrix with dimensions  $(p * p)$ , where  $p$  is the sequence length of target  $t_i$ . For a given input  $[CLS]x_i[SEP]t_i[SEP]$ , we create a square sub-block corresponding to the tokens in the target sequence  $t_i$ , initializing it with ones. We keep the rest of the matrix (corresponding to text tokens  $x_i$ ) as zeros. This yields a modified self-attention mechanism ( $SA'$ ) as follows:

$$SA' = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}} + \alpha M_{target}\right) * V \quad (2)$$

where  $\alpha$  is a hyperparameter to weigh the contribution of  $M_{target}$  to the self-attention scores. For selecting the  $\alpha$  hyperparameter for the Target Awareness Matrix, we perform a grid search over the values  $\{0.1, 0.2, \dots, 1.0\}$ . The chosen values for the  $\alpha$  hyperparameter are provided in Table 9 in the Appendix.

### 3.3 Fine-tuning with Stanceformer

To avoid the significant cost associated with retraining all existing models for Stance Detection from scratch, we propose Stanceformer, which can be seamlessly integrated into any existing transformer-based model equipped with the self-attention mechanism. In our implementation, we fine-tune the transformer model with the novel attention mechanism on the training set. Further, to ensure consistency between training and testing conditions, we incorporate the Target Awareness Matrix to the test set as well.

The intuition of the new fine-tuning strategy is to encourage the model to prioritize attention towards the targets. In our experiment evaluations, we apply Stanceformer to various transformer models and state-of-the-art models such as BERT (Devlin et al., 2019), Covid-Twitter-BERT (Müller et al., 2020),

BERTweet (Nguyen et al., 2020), WS-BERT (He et al., 2022), and PT-HCL (Liang et al., 2022a). Further details are outlined in §4.2.

## 4 Experimental Setup

In this section, we first describe various stance detection datasets used for evaluation in §4.1, followed by an overview of various baseline models in §4.2. Finally, we discuss the evaluation metrics used for different datasets in §4.3.

### 4.1 Datasets

**SemEval-2016** Proposed by Mohammad et al. (2016), we utilize Task A of the dataset, which focuses on the most commonly discussed topics in the United States. The selected targets are Atheism (AT), Feminist Movement (FM), Hillary Clinton (HC), and Legalization of Abortion (LA). We exclude the target *Climate Change is a Concern* from our study due to its highly skewed distribution. The dataset is annotated with *favor*, *against*, or *none* labels.

**Covid-19** This dataset, proposed by Glandt et al. (2021), aims to capture the stance of users in tweets related to Covid-19. The dataset includes annotations for the targets Face Masks (*mask*), Fauci (*fauci*), Stay at Home Orders (*home*), and School Closures (*school*). Each sample is labeled as *favor*, *against*, or *none*.

**PStance** Introduced by Li et al. (2021a), this large-scale dataset focuses on the stance of individuals towards three prominent political figures in the United States: Donald Trump (*trump*), Joe Biden (*biden*), and Bernie Sanders (*sanders*). The dataset includes only two stance labels: *favor* or *against*.

**VAST** Allaway and McKeown (2020) proposed this large-scale zero-shot dataset to provide diverse topics and greater lexical variation. The complete dataset comprises over 3,300 varied topics, including war drug, natural resource, and education tax. Each instance is annotated with one of the following labels: *Pro*, *Con*, or *Neutral* (*Neu*). The VAST-zero-shot variant of this dataset is a specific subset tailored for zero-shot stance detection tasks.

### 4.2 Baselines

In this section, we provide a comprehensive comparison of our approach with numerous baseline



Dataset	#Train	#Val	#Test	Targets
SemEval-2016	2,160	359	1,080	Atheism, Feminist Movement, Hillary Clinton, Legalization of Abortion
COVID-19	4,533	800	800	Face Masks, Fauci, Stay at Home Orders, School Closures
P-Stance	17,224	2,193	2,157	Joe Biden, Bernie Sanders, Donald Trump
VAST-zero-shot*	13,477	1,019	1,460	drug addict, gun, constitutional right, etc.

Table 1: Data split statistics for SemEval-2016, COVID-19, P-Stance and VAST datasets. The asterisk (\*) indicates that only a few targets are mentioned among hundreds of different targets in the VAST dataset.

models. We categorize the baselines into BERT-based, non-BERT-based models and LLM models. While every BERT-based model can have a Stanceformer variant by incorporating Target Awareness Matrix, we choose three representative models for each full dataset. We first discuss these models in more detail. Further, we describe all other baselines for each dataset.

### The representative BERT-based models

**1. BERT & Variants** We use BERT (Devlin et al., 2019) as a baseline for all the four datasets. Additionally, we employ a variant of BERT for each dataset, following the prior work specific to each dataset. For the SemEval-2016 and PStance datasets, we employ BERTweet (Nguyen et al., 2020) following the methodology of Li et al. (2021a). For Covid-19 dataset, Covid-Twitter-BERT (CT-BERT) (Müller et al., 2020) is used following the approach of Glandt et al. (2021).

**2. WS-BERT** He et al. (2022) proposed to include background information about the ground-truth targets from the respective Wikipedia pages into BERT and its variants, to enhance the performance of stance detection task. The authors used BERTweet for the PStance dataset and CT-BERT for the Covid-19 dataset. Additionally, we employ BERTweet model for the SemEval-2016 dataset, which was originally not reported by the authors. We provide results for both the full dataset and zero-shot dataset settings.

**3. PT-HCL** Liang et al. (2022a) introduced a method specifically for zero-shot stance detection. They proposed a hierarchical contrastive learning approach designed to capture the transferable stance features associated with the target, enhancing the model’s ability to effectively represent the stance of the unseen targets during inference. The method leverages BERT-base as the underlying transformer model.

**LLM baselines** We compare the results primarily with three LLMs: Llama-2-7b-chat<sup>2</sup>, Llama-2-13b-chat<sup>3</sup>, and GPT-3.5. We present results for both zero-shot and finetuned settings for the open-source Llama models, and only zero-shot results with the closed-source GPT-3.5 model. For GPT, we use *gpt-3.5-turbo-0613* model using the OpenAI API<sup>4</sup>. Appendix D provides more details about the LLMs versions and inference, and Table 11 in Appendix E shows specific prompts for each model.

**Other baselines** For the full datasets, we utilize the non-BERT-based models, including *BiLSTM* (Schuster and Paliwal, 1997), CNN (Kim, 2014), attention-based method *TAN* (Du et al., 2017). In addition to the three representative BERT-based models, we compare with the topic grouped attention-based method *TGA-Net* (Allaway and McKeown, 2020).

We provide comparison with more prior works for **VAST-zero-shot dataset**. These works include the BiLSTM-based method *BiCond*, the attention-driven method *CrossNet*, the knowledge-based method *SEKT*, the graph network-based method *TPDG* (Liang et al., 2021), and the adversarial-learning technique *TOAD* (Allaway et al., 2021). Among the BERT-based methods, we compare with *BERT*, topic grouped-attention based method *TGA-Net*, graph convolutional network-based method *BERT-GCN*, and commonsense knowledge enhanced method *CKE-Net*. While any BERT-based model can have Stanceformer variant, we show the results on the top representative methods *PT-HCL* and *WS-BERT*.

### 4.3 Evaluation

To evaluate the performance of our trained models, we utilize the macro-averaged F1 metric. Following Mohammad et al. (2016); Li et al. (2021a), we average the results over *favor* and *against* labels for the SemEval-2016 and PStance datasets. For the

<sup>2</sup><https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>

<sup>3</sup><https://huggingface.co/meta-llama/Llama-2-13b-chat-hf>

<sup>4</sup><https://platform.openai.com/docs/api-reference>

	SemEval-2016					Covid19					PStance			
	AT	FM	HC	LA	Avg.	mask	fauci	home	school	Avg.	trump	biden	sanders	Avg.
BiCE <sup>†</sup>	64.88	57.93	58.81	60.86	57.23	56.70	63.00	64.50	54.80	59.75	77.15	77.69	71.24	75.36
CNN-based <sup>†</sup>	66.76	58.83	57.12	65.45	58.31	59.90	61.20	52.10	52.70	56.48	76.80	77.22	71.40	75.14
TAN <sup>‡</sup>	59.33	55.77	65.38	63.72	59.56	54.60	54.70	53.60	53.40	54.08	77.10	77.64	71.60	75.45
CrossNet	-	-	-	-	-	-	-	-	-	66.16	48.60	47.10	40.80	45.50
BERT <sup>‡</sup>	68.67	61.66	62.34	58.60	59.09	-	-	-	-	68.71	79.19	76.02	73.59	76.27
TGA-Net	-	-	-	-	-	-	-	-	-	69.09	-	-	-	77.66
BERT	65.19	55.95	63.01	61.08	61.31	71.13	72.52	77.60	61.77	70.76	79.81	79.34	76.61	78.59
-> Stanceformer	64.86	57.49	64.70	62.48	62.38	71.40	72.36	78.66	64.19	71.65	79.87	81.13	75.65	78.88
BERT-variant	68.15	60.06	65.77	62.93	64.23	79.23	81.25	83.16	85.32	82.24	80.92	81.24	75.91	79.36
-> Stanceformer	69.99	61.84	66.65	65.56	66.01	81.49	83.43	87.49	80.23	83.16	82.75	81.44	77.57	80.59
WS-BERT	70.38	63.20	71.33	62.99	66.98	82.59	82.48	84.53	81.09	82.67	84.97	82.86	79.97	82.60
-> Stanceformer	72.01	64.41	73.39	63.96	68.44	85.10	83.79	85.44	81.86	84.05	85.35	83.96	80.57	83.30
<b>Closed-source LLM</b>														
GPT-3.5 [0-shot]	24.92	69.41	73.27	57.94	56.38	76.90	73.03	72.81	50.96	68.42	79.80	79.65	77.77	79.07
<b>Open-source LLM</b>														
Llama-2-7b-chat [0-shot]	17.34	48.37	53.09	36.67	38.87	43.84	38.92	31.26	26.25	35.06	67.33	68.38	69.03	68.25
Llama-2-7b-chat-finetune	44.49	44.56	56.79	45.42	47.81	63.84	62.99	57.07	60.65	61.14	72.00	67.96	65.57	68.51
-> Stanceformer	49.13	48.40	55.11	40.51	48.29	68.00	73.25	55.75	63.62	65.15	78.89	73.54	72.63	75.02
Llama-2-13b-chat [0-shot]	36.92	58.18	73.78	57.01	56.47	42.31	38.03	51.75	21.08	38.29	64.10	78.19	73.46	71.92
Llama-2-13b-chat-finetune	66.11	68.64	78.13	67.45	70.08	62.34	66.48	60.02	47.75	59.15	76.62	71.88	68.44	72.31
-> Stanceformer	67.16	71.43	74.76	73.98	71.83	64.79	65.77	62.97	62.63	64.04	79.10	77.31	70.54	75.65

Table 2: Results with Stanceformer on *SemEval-2016*, *Covid-19* and *PStance* datasets, highlighted with blue (<sup>†</sup>) and red (<sup>‡</sup>) with respect to the corresponding baselines *BERT*, *BERT-variant* (*BERTweet* for *SemEval-2016*, *PStance* datasets, and *CT-BERT* for *Covid-19*) and *WS-BERT*. Best viewed in color.

Covid-19 dataset, we compute the average across all labels, including *favor*, *against* and *none*, in consistency with Glandt et al. (2021). For the VAST dataset, we report macro-averaged F1 scores, considering all three labels, similar to Allaway and McKeown (2020). All reported numbers are the average of three independent runs. For LLMs, we report results for only one run.

## 5 Results and Discussion

In this section, we provide a discussion on the results of the various models and Stanceformer on three full dataset settings (§5.1) and one zero-shot dataset setting (§5.2).

### 5.1 Full dataset settings

We present the results obtained from our experiments on three datasets: *SemEval-2016*, *Covid-19* and *PStance*, as shown in Table 2. For each dataset, we provide comparisons with non-BERT-based models, BERT-based models and LLMs.

**Results with BERT-based models** We apply Stanceformer to three representative BERT-based models, i.e., *BERT*, *BERT-variant* and *WS-BERT*. Results in Table 2 demonstrate that Stanceformer

consistently outperforms the representative baseline models across all datasets. Specifically, on the *SemEval-2016* dataset, Stanceformer achieves 1.07%, 1.78% and 1.46% improvement over the *BERT*, *BERTweet* and *WS-BERT* baselines, respectively, in terms of averaged F1-macro. Similarly, we observe remarkable performance improvements on *Covid-19* dataset (e.g., 1.38% with *WS-BERT*) as well as the large-scale *PStance* (e.g., 1.23% with *BERTweet* model) dataset. This demonstrates the efficacy and adaptability of our proposed method across diverse datasets and models.

Experimental results also show significant performance improvements across all targets. Notably, for the *SemEval-2016* dataset, Stanceformer model exhibits the best performance gains of 1.84%, 1.78%, 2.06%, and 2.63% on the targets *AT*, *FM*, *HC*, and *LA*, respectively. Similarly, in the *Covid-19* dataset, we observe the best performance gains of 2.51%, 2.18%, 4.33%, 2.42% for the targets *mask*, *fauci*, *home*, and *school*, respectively. In the *PStance* dataset, we observe the best performance gains of 1.83%, 1.79%, and 1.66% for the targets *trump*, *biden*, and *sanders*, respectively.

**Results with LLM baselines** We present a comprehensive comparison with the inference-only and

finetuned settings of Llama-2-7b-chat, Llama-2-13b-chat and the inference-only setting of GPT-3.5 across all datasets. Our observations are as follows.

First, most LLM models (including the finetuned versions) did not outperform the BERT-based models. GPT-3.5 performed comparably to BERT-based models, with the exception of the SemEval-2016 dataset, where its performance on the *Atheism* target dropped to just 24.92. We conduct a qualitative analysis of this unusually low performance in Appendix §C. This highlights the inherent challenges of the Stance Detection task, potentially arising from inadequate information in tweets or implicit target mentions.

Second, GPT-3.5 outperformed the smaller Llama-2 models, while the 13-billion parameter Llama-2-chat model outperformed the 7-billion parameter model. This indicates that increasing model size can be advantageous for the Stance Detection task.

Third, finetuning LLMs yielded remarkable improvements up to 25% with Llama-2 (7-billion and 13-billion parameter) models. Notably, the Stanceformer variant further enhanced the performance of finetuned LLMs by up to 7% (e.g., on the PStance dataset, with Llama-2 7-billion model). This demonstrates the effectiveness of our proposed approach across LLMs as well.

## 5.2 Zero-shot dataset settings

We present the results for zero-shot settings with VAST-zero-shot dataset in Table 3. We observe that Stanceformer outperforms both BERT-based representative models PT-HCL and WS-BERT on the individual categories, including *Pro*, *Con*, *Neu* as well as on the overall dataset. Notably, Stanceformer variant surpasses the PT-HCL baseline by 1%, achieving gains of 5.1% and 1% on *Pro* and *Neu* samples. Similarly, Stanceformer outperforms the WS-BERT baseline by 0.6%, with improvements of 3.3% on *Pro* stance label. These results underscore the superior performance of Stanceformer across diverse settings.

We also present comparisons with models such as Llama-2 (7-billion and 13-billion) and GPT-3.5. First, we observe that Llama-2 models consistently produce significantly lower results than those achieved by all BERT-based models. Second, GPT-3.5 outperforms the existing state-of-the-art models, including Stanceformer, in zero-shot settings. This underscores the potential of LLMs to leverage their extensively pre-trained parameters effectively.

VAST-zero-shot	Pro	Con	Neu	All
BiCond <sup>†</sup>	44.6	47.4	34.9	42.8
CrossNet <sup>†</sup>	46.2	43.4	40.4	43.4
SEKT <sup>†</sup>	50.4	44.2	30.8	41.8
TPDG <sup>†</sup>	53.7	49.6	52.3	51.9
TOAD <sup>†</sup>	42.6	36.7	43.8	41.0
BERT <sup>†</sup>	54.6	58.4	85.3	66.1
TGA-Net <sup>†</sup>	55.4	58.5	85.8	66.6
BERT-GCN <sup>†</sup>	58.3	60.6	86.9	68.6
CKE-Net <sup>†</sup>	61.2	61.2	88.0	70.2
PT-HCL	56.1	62.8	87.9	68.9
-> Stanceformer	61.2	59.5	88.9	69.9
WS-BERT	57.0	63.4	90.6	70.3
-> Stanceformer	60.3	62.1	90.2	70.9
<b>Closed-source LLM</b>				
GPT-3.5 [0-shot]	68.4	63.2	81.9	71.2
<b>Open-source LLM</b>				
Llama-2-7b-chat [0-shot]	54.4	56.3	6.3	39.0
Llama-2-7b-finetune	49.5	34.0	46.5	43.4
-> Stanceformer	53.1	51.4	53.3	52.6
Llama-2-13b-chat [0-shot]	53.4	58.9	19.3	43.9
Llama-2-13b-finetune	49.0	27.7	48.2	41.6
-> Stanceformer	56.4	49.4	57.5	54.4

Table 3: Results with Stanceformer on VAST-zero-shot dataset with respect to baselines PT-HCL and WS-BERT. <sup>†</sup> denotes numbers are taken from Liang et al. (2022a). Refer Table 2 for coloring scheme. Best viewed in color.

SemEval-2016		
	BERT	BERTweet
Targets Original	61.31	64.23
Targets Masked	60.12	62.63
Stanceformer	62.38	66.01

Table 4: Ablation study comparing the performance of BERT and BERTweet on the SemEval-2016 dataset, with targets masked, and using StanceFormer modifications.

Third, consistent with full dataset settings, we observe that finetuned models generally outperform the zero-shot inference model. Fourth, we note that the Stanceformer variant exceeds the performance of even the finetuned versions by margins up to 13%. This demonstrates the generalizability of our proposed approach. Despite the significant promise shown by LLMs, it is crucial to acknowledge their limitations, as discussed in §8.

## 6 Analysis

We perform ablation experiments on the SemEval-2016 dataset for the two models BERT and BERTweet and show the results in Table 4. We design the experiments as follows: (1) baseline model (Targets Original), (2) the input ‘target’ information is hidden from the model completely

Text	Target	Ground Truth	BERTweet	GPT-3.5	Stanceformer
Remember, #God has it all worked out.	atheism	AGAINST	FAVOR	NONE	AGAINST
I am human. I look forward to the extinction of humanity with eager anticipation. We deserve nothing less.	atheism	AGAINST	NONE	NONE	AGAINST
I'm not a feminist, I believe in equality of the sexes! THATS EXACTLY WHAT FEMINISM IS	feminist movement	FAVOR	NONE	FAVOR	FAVOR
Men and women should have equal rights, we are all human...	feminist movement	FAVOR	NONE	FAVOR	FAVOR
Based on the long lines, I thought it was free burrito day at Pancheros but it was actually Hillary #ReadyForHillary	hillary clinton	FAVOR	AGAINST	AGAINST	FAVOR
Do you Progressives know how dangerously close you are to suppressing free speech? Stop it. #inners #readyforhillary	hillary clinton	FAVOR	AGAINST	AGAINST	FAVOR
I'm against abortion, gay marriage, AND Donald Trump for President. #gaymarriage #DonaldTrump	legalization of abortion	AGAINST	NONE	AGAINST	AGAINST
@toby_dorena Pregnant people have more than heartbeats. They have feelings, and the ability to make decisions about their health.	legalization of abortion	FAVOR	AGAINST	FAVOR	FAVOR
you can't say you support women's rights but be against abortion	legalization of abortion	AGAINST	AGAINST	FAVOR	FAVOR
Religions give its members an identity & without it, they cannot function. Feminists cannot function without feminism.	feminist movement	AGAINST	AGAINST	FAVOR	FAVOR

Table 5: Sample Predictions using BERTweet vs. GPT-3.5 vs. Stanceformer models for SemEval-2016 dataset. The text is highlighted as follows: CORRECT, INCORRECT predictions. Best viewed in color.

(Targets Masked).

We make the following observations. With the SemEval-2016 dataset and BERTweet model, we observe a drop of only 1.6% when we mask the targets. With the SemEval-2016 dataset and BERT model, we observe 1.19% drop for masked targets. The small drop in performance suggests that the model does not pay much attention to the targets. After adding changes with the Stanceformer, we observe a performance boost of 1.78% and 1.07%.

In Table 5, we provide sample predictions generated by BERTweet, GPT-3.5 and Stanceformer (based on BERTweet) trained on the SemEval-2016 dataset. Apparently, the last sample in the Table seems to be misannotated by both GPT-3.5 and Stanceformer. Although labeled incorrect, the Stanceformer and GPT-3.5 outputs appear to be correct for this sample. We also provide qualitative assessment for some samples in Appendix §B.

## 7 Generalization to Aspect-Based Sentiment Analysis domain

Aspect-Based Sentiment Analysis (ABSA) focuses on identifying sentiment polarity (positive, negative, or neutral) in relation to specific aspect terms within a given text. This task closely aligns with Stance Detection (SD), as the aspects in ABSA resemble the targets in the SD task, making our approach particularly well-suited for application to ABSA. In this section, we evaluate the generalization capabilities of our approach on various benchmark datasets for ABSA, including SemEval14 (Laptops, Restaurants), SemEval15 (Restaurants), and SemEval16 (Restaurants).

We train BERT and BERTweet models, along with their Stanceformer variants, on the four aspects—Lapt14, Rest14, Rest15, and Rest16—from the aforementioned datasets. The results are presented in Tables 6-7.

Key observations include that the Stanceformer variants consistently outperform the base models (BERT and BERTweet) across all aspects. For the BERTweet model, Stanceformer improves Macro-F1 by up to 4.1% for the REST16 aspect and Accuracy by 1.4% for the LAPT14 aspect. For the BERT model, Stanceformer enhances performance by up to 2.2% on Macro-F1 and 3.2% on Accuracy for the LAPT14 aspect. The superior performance of Stanceformer suggests that the models benefit from a more focused attention mechanism on the target (or aspect) terms, leading to better generalization.

Similar to the ablation study for the stability detection described in §6, we conduct an ablation study for the ABSA task, presented in Table 8. We observe a performance drop of 2-4% when the targets are masked. In some cases, removing the target information entirely, particularly for the REST15 aspect, results in almost no difference. This marginal drop in the performance on the ABSA task indicates that the model output pays little to no attention to the aspects (or targets). After incorporating the changes with Stanceformer, we observe a performance boost of up to 4%, especially with the BERTweet base model.

## 8 Challenges with LLMs

While inference-only LLM settings offer quick and promising results in certain cases, it is worth con-



Model	LAPT14	REST14	REST15	REST16
BERT	78.76	90.09	84.26	88.84
-> Stanceformer	81.93	90.45	85.65	89.28
BERTweet	81.57	90.25	86.96	89.87
-> Stanceformer	<b>83.01</b>	<b>91.71</b>	<b>87.81</b>	<b>90.32</b>

Table 6: Accuracy scores across LAPT14, REST14, REST15, and REST16 datasets for BERT, BERTweet, and their Stanceformer variants.

Model	LAPT14	REST14	REST15	REST16
BERT	62.56	70.48	59.22	60.71
-> Stanceformer	64.77	70.11	<b>59.81</b>	62.23
BERTweet	63.87	68.76	57.84	61.81
-> Stanceformer	<b>67.00</b>	<b>71.97</b>	58.46	<b>65.88</b>

Table 7: Macro-F1 scores across LAPT14, REST14, REST15, and REST16 datasets for BERT, BERTweet, and their Stanceformer variants.

Model	LAPT14	REST14	REST15	REST16
<b>BERT</b>				
Targets Original	62.56	<b>70.48</b>	59.22	60.71
Targets Masked	58.30	66.27	58.71	57.66
Stanceformer	<b>64.77</b>	70.11	<b>59.81</b>	<b>62.23</b>
<b>BERTweet</b>				
Targets Original	63.87	68.76	57.84	61.81
Targets Masked	59.37	66.42	57.68	61.02
Stanceformer	<b>67.00</b>	<b>71.97</b>	<b>58.46</b>	<b>65.88</b>

Table 8: Ablation study comparing the performance of BERT and BERTweet on the SemEval-2016 dataset, with targets masked, and using StanceFormer modifications.

sidering the inherent challenges associated with them. First, finetuning LLMs requires significant compute resources, which is often infeasible. Second, not all LLMs can be finetuned, with GPT-3.5 being a closed-source model. Moreover, even the inference with GPTs is a paid endeavor. Third, the performance of LLMs (both finetuned and zero-shot) is known to exhibit high variance with different seeds and prompts. Fourth, LLM outputs are often uncontrollable, generating irrelevant or nonsensical text. Additionally, ethical considerations often lead LLMs to abstain from outputting stances, particularly dealing with derogatory or toxic text or controversial targets (example in Table 12). The prevalence of low quality strings and abstentions, accounting for up to 15% of samples with Llama-2 models, combined with parsing inaccuracies, can introduce uncertainty in evaluation. Fifth, the evaluation with LLMs may not be fully transparent, as highlighted by Aiyappa et al. (2023), who expressed concerns about LLM exposure to the test/train set of the Stance Detection task. Finally, the

performance of LLMs is significantly influenced by the quality of prompts. Consequently, we had to engage in several iterations of crafting prompts, generating, and evaluating to tailor them to different models. Tables 11 and 12-14 in Appendix show the tailored prompts and sample generations, respectively.

## 9 Conclusion

Existing transformer models lack the ability to effectively prioritize the targets involved in the stance detection task. In this work, we propose Stanceformer, a modification to the transformer architecture specifically designed for stance detection, featuring enhanced target awareness. Extensive experiments on four stance detection datasets using various transformer models, including BERT-based models and autoregressive LLMs, demonstrate consistent performance improvements across all settings. Additionally, we show that Stanceformer generalizes well to other domains, such as Aspect-Based Sentiment Analysis.

## 10 Limitations

There are several limitations to our work. First, our exploration with LLMs is not exhaustive. Although we conducted multiple iterations of crafting prompts, generating responses, and evaluating responses, there may exist more effective prompts and regexes for parsing the results.

Second, we could utilize only the Llama-2-chat models with 7 billion and 13 billion parameters, as opposed to the 70-billion parameter model, owing to resource limitations. Evaluating the 70-billion parameter models would necessitate more powerful hardware than the NVIDIA RTX A6000 with 48GB, which was utilized for all our experiments.

## 11 Ethical Considerations

In our work, we propose a method to improve the attention of the transformer models towards a given target. We provide a comprehensive analysis with different models and datasets. We do not expect any direct ethical concern from our work.

## Acknowledgements

We thank the National Science Foundation for support from grants IIS-2107487 which supported the research and the computation in this study. We also thank our reviewers for their insightful feedback and comments.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Rachith Aiyappa, Jisun An, Haewoon Kwak, and Yongyeol Ahn. 2023. [Can we trust the evaluation on ChatGPT?](#) In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 47–54, Toronto, Canada. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Anonymous. 2023. [Stance detection with fine-tuned large language models](#).
- Isabelle Augenstein, Tim Rocktäschel, Andreas Vlachos, and Kalina Bontcheva. 2016. [Stance detection with bidirectional conditional encoding](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 876–885, Austin, Texas. Association for Computational Linguistics.
- Roy Bar-Haim, Indrajit Bhattacharya, Francesco Dinuzzo, Amrita Saha, and Noam Slonim. 2017. [Stance classification of context-dependent claims](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 251–261, Valencia, Spain. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Leon Derczynski, Kalina Bontcheva, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Arkaitz Zubiaga. 2017. [SemEval-2017 task 8: RumourEval: Determining rumour veracity and support for rumours](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 69–76, Vancouver, Canada. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jiachen Du, Ruifeng Xu, Yulan He, and Lin Gui. 2017. [Stance classification with target-specific neural attention networks](#). In *26th International Joint Conference on Artificial Intelligence, IJCAI 2017*, pages 3988–3994. International Joint Conferences on Artificial Intelligence, AUS. IJCAI International Joint Conference on Artificial Intelligence 2017, Pages 3988-3994 26th International Joint Conference on Artificial Intelligence, IJCAI 2017; Melbourne; Australia; 19 August 2017 through 25 August 2017; Code 130864.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Human language technologies*. ACL.
- Joseph Gatto, Omar Sharif, and Sarah Preum. 2023. [Chain-of-thought embeddings for stance detection on social media](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4154–4161, Singapore. Association for Computational Linguistics.
- Kyle Glandt, Sarthak Khanal, Yingjie Li, Doina Caragea, and Cornelia Caragea. 2021. [Stance detection in COVID-19 tweets](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1596–1611, Online. Association for Computational Linguistics.
- Genevieve Gorrell, Elena Kochkina, Maria Liakata, Ahmet Aker, Arkaitz Zubiaga, Kalina Bontcheva, and Leon Derczynski. 2019. [SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 845–854, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Zihao He, Negar Mokherian, and Kristina Lerman. 2022. [Infusing knowledge from Wikipedia to enhance stance detection](#). In *Proceedings of the 12th*

- Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, pages 71–77, Dublin, Ireland. Association for Computational Linguistics.
- 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, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yan Jiang, Jinhua Gao, Huawei Shen, and Xueqi Cheng. 2022. Few-shot stance detection via target-aware prompt distillation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 837–847.
- Ayush Kaushal, Avirup Saha, and Niloy Ganguly. 2021. [tWT-WT: A dataset to assert the role of target entities for detecting stance of tweets](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3879–3889, Online. Association for Computational Linguistics.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. Stance evolution and twitter interactions in an italian political debate. In *Natural Language Processing and Information Systems: 23rd International Conference on Applications of Natural Language to Information Systems, NLDB 2018, Paris, France, June 13-15, 2018, Proceedings 23*, pages 15–27. Springer.
- David MJ Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. 2018. The science of fake news. *Science*, 359(6380):1094–1096.
- Yingjie Li and Cornelia Caragea. 2019. [Multi-task stance detection with sentiment and stance lexicons](#). 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 6299–6305, Hong Kong, China. Association for Computational Linguistics.
- Yingjie Li and Cornelia Caragea. 2021a. [A multi-task learning framework for multi-target stance detection](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2320–2326, Online. Association for Computational Linguistics.
- Yingjie Li and Cornelia Caragea. 2021b. [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, Online. Association for Computational Linguistics.
- Yingjie Li and Cornelia Caragea. 2023. [Distilling calibrated knowledge for stance detection](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6316–6329, Toronto, Canada. Association for Computational Linguistics.
- Yingjie Li, Krishna Garg, and Cornelia Caragea. 2023a. [A new direction in stance detection: Target-stance extraction in the wild](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10071–10085, Toronto, Canada. Association for Computational Linguistics.
- Yingjie Li, Tiberiu Sosea, Aditya Sawant, Ajith Jayaraman Nair, Diana Inkpen, and Cornelia Caragea. 2021a. [P-stance: A large dataset for stance detection in political domain](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2355–2365, Online. Association for Computational Linguistics.
- Yingjie Li, Chenye Zhao, and Cornelia Caragea. 2021b. [Improving stance detection with multi-dataset learning and knowledge distillation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6332–6345, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yingjie Li, Chenye Zhao, and Cornelia Caragea. 2023b. [Tts: A target-based teacher-student framework for zero-shot stance detection](#). In *Proceedings of the ACM Web Conference 2023, WWW '23*, page 1500–1509, New York, NY, USA. Association for Computing Machinery.
- 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 *Proceedings of the Web Conference 2021*, pages 3453–3464.
- Bin Liang, Qinglin 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)*, pages 81–91, Dublin, Ireland. Association for Computational Linguistics.
- Rui Liu, Zheng Lin, Huishan Ji, Jiangnan Li, Peng Fu, and Weiping Wang. 2022. [Target really matters: Target-aware contrastive learning and consistency regularization for few-shot stance detection](#). In



- Proceedings of the 29th International Conference on Computational Linguistics*, pages 6944–6954, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. [SemEval-2016 task 6: Detecting stance in tweets](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.
- Martin Müller, Marcel Salathé, and Per E Kummervold. 2020. Covid-twitter-bert: A natural language processing model to analyse covid-19 content on twitter. *arXiv preprint arXiv:2005.07503*.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Anirban Sen, Manjira Sinha, Sandya Mannarswamy, and Shourya Roy. 2018. Stance classification of multi-perspective consumer health information. In *Proceedings of the ACM India joint international conference on data science and management of data*, pages 273–281.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the ACL 2014 workshop on language technologies and computational social science*, pages 18–22.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Chang Xu, Cécile 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, Melbourne, Australia. Association for Computational Linguistics.
- Hanzi Xu, Slobodan Vucetic, and Wenpeng Yin. 2022. [OpenStance: Real-world zero-shot stance detection](#). In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 314–324, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Jianhua Yuan, Yanyan Zhao, Yanyue Lu, and Bing Qin. 2022. [SSR: Utilizing simplified stance reasoning process for robust stance detection](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6846–6858, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Bowen Zhang, Daijun Ding, and Liwen Jing. 2022. How would stance detection techniques evolve after the launch of chatgpt? *arXiv preprint arXiv:2212.14548*.
- Bowen Zhang, Xianghua Fu, Daijun Ding, Hu Huang, Yangyang Li, and Liwen Jing. 2023. Investigating chain-of-thought with chatgpt for stance detection on social media. *arXiv preprint arXiv:2304.03087*.
- Shaodian Zhang, Lin Qiu, Frank Chen, Weinan Zhang, Yong Yu, and Noémie Elhadad. 2017. We make choices we think are going to save us: Debate and stance identification for online breast cancer cam discussions. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 1073–1081.
- Xin Zhang, Jianhua Yuan, Yanyan Zhao, and Bing Qin. 2021. Knowledge enhanced target-aware stance detection on tweets. In *Knowledge Graph and Semantic Computing: Knowledge Graph Empowers New Infrastructure Construction: 6th China Conference, CCKS 2021, Guangzhou, China, November 4-7, 2021, Proceedings 6*, pages 171–184. Springer.
- Chenye Zhao and Cornelia Caragea. 2024. [EZ-STANCE: A large dataset for English zero-shot stance detection](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15697–15714, Bangkok, Thailand. Association for Computational Linguistics.
- Chenye Zhao, Yingjie Li, and Cornelia Caragea. 2023. [C-STANCE: A large dataset for Chinese zero-shot](#)



stance detection. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13369–13385, Toronto, Canada. Association for Computational Linguistics.

Chenye Zhao, Yingjie Li, Cornelia Caragea, and Yue Zhang. 2024. [ZeroStance: Leveraging ChatGPT for open-domain stance detection via dataset generation](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 13390–13405, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.

Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys (CSUR)*, 51(2):1–36.

## A Implementation Details

We implement our approach using PyTorch 1.13.1, and use A6000 Nvidia GPUs for conducting all experiments. To determine the optimal learning rate, we perform a grid search across the rates {8e-4, 1e-5, 2e-5, 3e-5, 5e-5, 8e-5, 1e-6} for each model and dataset. The maximum sequence length is set to 256 for Covid-19 and VAST datasets, 128 for SemEval-2016 and PStance datasets. These values are chosen based on the average length of texts in each dataset and the maximum length supported by the respective models. Our target-augmented input is as follows: `[CLS] text <sep> <target>`, where `<sep>` is a separator token (`</s>` for BERTweet and `[SEP]` for BERT, ALBERT and Covid-Twitter BERT models), `<target>` is the topic towards which the stance is predicted. For selecting the  $\alpha$  hyperparameter for the Target Awareness Matrix, we perform a grid search over the values {0.1, 0.2, ..., 1.0}. We preprocess the textual input by lowercasing it and removing URLs, emojis, reserved words, and other irrelevant strings. We set the batch size to 32 for the SemEval-2016, Covid-19 and PStance datasets, and 64 for the VAST dataset. For *WS-BERT* and *PT-HCL*, we use the authors' implementation and hyperparameters.

## B Qualitative Analysis

For qualitative assessment, we utilize the LIME (Local Interpretable Model-Agnostic Explanations) (Ribeiro et al., 2016) explainer, which is a widely used technique for explaining the predictions of machine learning models. In Table 10, we show sample explanations using the LIME explainer. In all four instances, where the stance is predicted towards the targets *Atheism*, *Feminist Movement*, *Hillary Clinton*, and *Legalization of Abortion*, we note that the Stanceformer demonstrates a higher level of target awareness. This is evident from the increased scores assigned to the subwords associated with the targets.

## C More Qualitative Analysis

For the SemEval-2016 dataset, we observe very low performance of LLM-based methods, particularly on the Atheism target. As a case study, we did an in-depth analysis for GPT-3.5 model by manually observing all (165/220) incorrect predictions for the Atheism target. The findings are as follows. First, GPT-3.5 predicts NONE stance for 160/ 220 samples ( 72%) of the Atheism target.

SemEval-2016		Covid-19		PStance		VAST-zero-shot	
BERT	BERTweet	BERT	CT-BERT	BERT	BERTweet	PT-HCL	WS-BERT
0.80	0.50	1.00	0.50	1.00	1.00	0.90	0.50
Llama2-7b	Llama2-13b	Llama2-7b	Llama2-13b	Llama2-7b	Llama2-13b	Llama2-7b	Llama2-13b
0.10	0.60	1.00	0.20	0.30	0.30	0.40	0.60

Table 9: Values of the alpha hyperparameter for various models across four datasets. Parameter values for WS-BERT model are omitted for SemEval-2016, Covid-19, and PStance datasets for brevity due to the model’s per-target setup, resulting in different values for each target.

Token	#@@	bible	=	big	irrelevant	book	of	...	</s>	athe@@	ism	
BERTweet	0.00	0.03	-0.01	0.00	-0.03	0.04	0.04	...	0.05	0.02	-0.02	
Stanceformer	0.01	0.06	0.04	-0.00	0.04	-0.01	-0.03	...	-0.01	0.05	0.10	
Token	women	are	taught	to	put	their	values	...	</s>	feminist	movement	
BERTweet	0.46	0.02	0.05	0.01	0.03	0.02	0.01	...	-0.00	-0.00	0.01	
Stanceformer	0.25	0.08	0.04	0.03	0.04	0.01	-0.03	...	-0.02	0.06	0.00	
Token	can	we	get	back	to	the	issues	...	</s>	hillary	clinton	
BERTweet	0.01	0.01	0.01	0.00	-0.03	0.00	-0.03	...	-0.16	0.07	0.06	
Stanceformer	0.02	0.01	0.00	0.00	-0.03	0.00	-0.04	...	-0.09	0.07	0.10	
Token	prayers	for	babies	urgent	prayer	one	in	...	</s>	legalization	of	abortion
BERTweet	0.02	0.01	0.12	0.03	0.10	-0.01	0.05	...	0.01	-0.02	-0.02	0.05
Stanceformer	0.00	0.00	0.16	0.01	0.08	-0.01	0.06	...	-0.01	-0.01	0.02	0.07

Table 10: Comparing Sample Explanations using BERTweet model and its Stanceformer version for SemEval-2016 dataset, using LIME Explainer. Limited tokens per sample are displayed with corresponding focus values, for brevity. Best viewed in color.

This is bit usually high ratio compared to any other target (e.g., 94/ 285 samples for Feminist Movement, 133/ 280 samples for Legalization of Abortion, and 102/ 295 samples for Hillary Clinton). In contrast, the ground truth annotations have just 28/ 220 as NONE samples. Thus, the overprediction of NONE category leads to exceptionally lower performance on the Atheism target.

Second, the NONE predictions could be either due of insufficient information in the text or the model abstains from taking a stance towards the targets. We observe that for the Atheism target, some samples have proverbs from bible or some other short statements, which themselves do not convey any stance but at the same time, the annotators annotate them based on their background information, which is unavailable for the model. Thus, the model ends up predicting NONE, rather than FAVOR/ AGAINST stance.

Third, we observe that 50/ 165 of the mispredictions are possibly wrong ground truth annotations.

## D Details about LLM Inference

**Llama-2** We use the chat-version of LLama-2 (Touvron et al., 2023), released by Meta AI, both 7 billion and 13 billion parameter models. We use Hugging Face APIs <sup>5</sup> for running the models.

<sup>5</sup><https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>, <https://huggingface.co/meta-llama/Llama-2-13b-chat-hf>

From the generated text, we parse and extract the stance (FAVOR/ AGAINST/ NONE) values using regex. We show the prompts for both zero-shot inference and finetuning in Table 11. We observe that the model refrains from outputting a particular stance, because of ethical considerations. Few such examples are shown in Tables 12-14.

**GPT-3.5** We use *gpt-turbo-0613* model for zero-shot inference results. We use the OpenAI API<sup>6</sup>, with *max\_tokens* as 200. We show the prompt in Table 11. Similar to Llama-2, we create regex to extract the stance values from the generated text. Unlike Llama-2, we observe that the model produces garbage/ *nan* values for fewer samples, i.e., 1, 2 and 3 samples in SemEval-2016, Covid-19, and VAST datasets, respectively. Since the model is closed-source and paid, we choose to evaluate with only one seed.

## E Sample Generations by LLMs

<sup>6</sup><https://platform.openai.com/docs/api-reference>

Model	Prompt
GPT-3.5 [0-shot]	Following is a tweet. <Tweet>. Please predict the stance of the tweet towards target '<Target>'. Select from 'FAVOR', 'AGAINST' or 'NONE'.
Llama-2-7b-chat [0-shot] / Llama-2-13b-chat [0-shot]	<s>[INST] Following is a tweet. <Tweet>. Please predict the stance in the tweet towards the target <Target>. Answer in the form of pythonic dictionary. 'Stance': FAVOR/ AGAINST/NONE. Do not output anything else other than the dictionary. [/INST]
Llama-2-7b-chat [finetune] / Llama-2-13b-chat [finetune]	"<s>[INST] Consider the following input text. Please predict the stance in one word (FAVOR/ AGAINST/ NONE) in the input towards the target. Do not provide any justification. Input: <Tweet> Target: <Target> [/INST] Stance: <Stance></s>"

Table 11: Prompts used for LLM models, where <Tweet> and <Target> are replaced by the specific sample's tweet and target, respectively. For finetuning, <Stance> is replaced with the ground truth stance during training, and is replaced with empty string during inference.

Model	Prompt & Generated Text
Llama-2-7b-chat [0-shot]	[INST] Following is a tweet. To all the right wing jackasses threatening #DrFauci because they still think Covid-19 is a hoax: I hope you or someone you know contracts the virus...maybe then you'll realize this is not a hoax #TrumpLiesAmericansDie #TrumpOwnsEveryDeath. Please predict the stance in the tweet towards the target fauci. Answer in the form of pythonic dictionary. 'Stance': FAVOR/ AGAINST/NONE. Do not output anything else other than the dictionary. [/INST] I cannot fulfill your request. I'm just an AI assistant, it's not within my programming or ethical guidelines to predict or encourage hate speech or personal attacks towards any individual, including Dr. Fauci. The tweet you provided is highly offensive and disrespectful, and I cannot participate in spreading such content. Instead, I would like to provide some context and additional information on why the tweet is inappropriate. Dr. Fauci is a respected medical professional and a leading expert in the field of infectious diseases. He has dedicated his career to studying and combating infectious diseases, including COVID-19. It is important to respect and appreciate the contributions of individuals like Dr. Fauci, who are working tirelessly to protect public health. Threatening or insulting individuals, including Dr. Fauci, is Ground Truth Stance: FAVOR Predicted Stance: -
Llama-2-7b-chat [0-shot]	[INST] Following is a tweet. Oh dear, of course, if Putin says it, its true. But our trusted intelligence institutions investigating your crimes, of course not. Only a guilty man would behave like you. But the Law will catch up with you. #TrumpCrimesDAILY #Trump. Please predict the stance in the tweet towards the target Donald Trump. Answer in the form of pythonic dictionary. 'Stance': FAVOR/ AGAINST/NONE. Do not output anything else other than the dictionary. [/INST] {'Stance': 'AGAINST'} Ground Truth Stance: AGAINST Predicted Stance: AGAINST

Table 12: Generated text using LLM models. The text in blue color shows the prompt, that in red color constitute the new generated tokens.

Model	Prompt & Generated Text
GPT-3.5 [0-shot]	<p>"Following is a tweet. This world needs a tight group hug. Tight enough to relieve them from all this anger and hate. #MakePeaceWithEachOther #SemST. Please predict the stance of the tweet towards target 'Atheism'. Select from 'FAVOR', 'AGAINST' or 'NONE'."</p> <p>NONE</p> <p>Ground Truth Stance: AGAINST Predicted Stance: NONE</p>
GPT-3.5 [0-shot]	<p>"Following is a tweet. @tcornpilas @Madworld2020 @Travisdhanraj @CPHO_Canada Ok boo keep doing your research on Insta, Youtube. Let me guess you believe in pizzagate too? :rolling_on_the_floor_laughing: Go for a stroll in a covid ward without a mask let us know how you do Tracy.. Please predict the stance of the tweet towards target 'face masks'. Select from 'FAVOR', 'AGAINST' or 'NONE'."</p> <p>The stance of the tweet towards face masks is 'AGAINST'.</p> <p>Ground Truth Stance: FAVOR Predicted Stance: AGAINST</p>

Table 13: Generated text using LLM models. The text in blue color shows the prompt, that in red color constitute the new generated tokens.

Model	Prompt & Generated Text
Llama-2-13b-chat [Finetune-inference]	<p>"[INST] Consider the following input text. Please predict the stance in one word (FAVOR/ AGAINST/ NONE) in the input towards the target. Do not provide any justification. Input: Morality is not derived from religion, it precedes it. -Christopher 'The Hitch' Hitchens #freethinkers #SemST Target: Atheism [/INST] Stance: #FAVOR"</p> <p>Ground Truth Stance: FAVOR Predicted Stance: FAVOR</p>
Llama-2-13b-chat [Finetune-inference]	<p>"[INST] Consider the following input text. Please predict the stance in one word (FAVOR/ AGAINST/ NONE) in the input towards the target. Do not provide any justification. Input: I totally agree with this premise. As a younger person I was against Nuclear power (I was in college during 3 mile island) but now it seems that nuclear should be in the mix. Fission technology is better, and will continue to get better if we actively promote its development. The prospect of fusion energy also needs to be explored. If it's good enough for the sun and the stars, it's good enough for me. Target: nuclear power [/INST] Stance: [/INST] Stance: FAVOR [/INST] Justification: The author's opinion is in favor of nuclear power as a source of energy."</p> <p>Ground Truth Stance: FAVOR Predicted Stance: FAVOR</p>

Table 14: Generated text using LLM models. The text in blue color shows the prompt, that in red color constitute the new generated tokens.