Collaborative Knowledge Infusion for Low-resource Stance Detection

Ming Yan, Joey Tianyi Zhou Member, IEEE, and Ivor W. Tsang Fellow, IEEE.

Abstract—Stance detection is the view towards a specific target by a given context (e.g. tweets, commercial reviews). Targetrelated knowledge is often needed to assist stance detection models in understanding the target well and making detection correctly. However, prevailing works for knowledge-infused stance detection predominantly incorporate target knowledge from a singular source that lacks knowledge verification in limited domain knowledge. The low-resource training data further increases the challenge for the data-driven large models in this task. To address those challenges, we propose a collaborative knowledge infusion approach for low-resource stance detection tasks, employing a combination of aligned knowledge enhancement and efficient parameter learning techniques. Specifically, our stance detection approach leverages target background knowledge collaboratively from different knowledge sources with the help of knowledge alignment. Additionally, we also introduce the parameter-efficient collaborative adaptor with a staged optimization algorithm, which collaboratively addresses the challenges associated with low-resource stance detection tasks from both network structure and learning perspectives. To assess the effectiveness of our method, we conduct extensive experiments on three public stance detection datasets, including low-resource and cross-target settings. The results demonstrate significant performance improvements compared to the existing stance detection approaches.

I. INTRODUCTION

C TANCE detection is the view towards a specific target with a given context, such as tweets or commercial reviews. Typically, those given contexts in stance detection tasks are mostly short-length contexts, which makes it challenging to predict the target's stance for the data-driven detection models with such limited information. Large pretrained language models (PLMs) are becoming the default backbone to enhance the stance detection model with learned commonsense knowledge, leading to great success in this field [1], [2]. To further enrich the knowledge of targets, the straightforward approach is to incorporate the target-related background knowledge as extra supplementary knowledge for the pretrained stance detection model, which has been shown to substantially improve model performance [3], [4]. In detail, those works infuse explicitly knowledge individually through knowledge graph [5], [6], Wikipedia [7], [8], generative knowledge [9], [10], leveraging PLMs' knowledge feature learning and representation capability by fine-tuning entire models' parameters. However, those knowledge-infuse solutions are quite inefficient in finetuning large PLM backbones on the limited training data. For instance, the few/zero-shot stance detection dataset VAST [11] has very limited training data or even no training data for each target. Besides the low-resource challenge, unbalanced dataset distribution is another challenge for the stance detection task, leading the training trajectory to fall into the local minima. Last but not least, we find some background knowledge is not always infused correctly in the knowledge infusion process. This is because a single knowledge source in previous works can not fully cover and support enough knowledge for diverse targets [4]. For example, the target "breaking the law" in Wikipedia is erroneously linked to a heavy metal music song rather than its ground truth definition of engaging in activities contrary to the law.

1

To address the aforementioned challenges, we propose a novel collaborative knowledge-infused stance detection method for training the large detection model in the lowresource setting efficiently. Specifically, we introduce a retrieval-based knowledge verifier that mitigates incorrect knowledge infusion by selecting the high-semantic background knowledge from different knowledge sources, rather than relying on a single knowledge source. Furthermore, we present a trainable collaborative adaptor integrated into PLMs to enable efficient parameter learning in low-resource stance detection tasks. Concretely, the collaborative adaptor freezes the parameter weights of large PLM and fine-tunes the parameter-efficient adaptor only, which alleviates the over-fitting effects on large PLM in low-resource scenarios. However, we empirically find that intuitively adding adaptors into PLM may lead to unstable training in the new stance detection tasks. We think the initialized weights of the collaborative adaptor can not work well with the pretrained PLMs in the early fine-tuning stage of new tasks. Moreover, the unbalanced data distribution further impacts the stable training. So we design a staged optimization algorithm for the adaptive model training in unbalanced distributions. The primary objective of the first optimization stage is to prevent the training trajectory from converging to a local minimum leading to unexpected performance. In the second stage, our model introduces a weighted cross-entropy loss to balance the biased stance categories and further improve the model performance in low-resource stance detection tasks. In other words, we progressively used label smooth (stage-1) and weighted loss (stage-2) separately to reduce the over-fitting effects in our low-resource stance detection tasks, which is different from traditional optimization paradigms using those two in the whole training process.

We conducted extensive experiments on three public stance detection datasets, encompassing the low-resource stance detection, and cross-target stance detection tasks. Experimental

Ming Yan, Joey Tianyi Zhou, and Ivor W. Tsang are with Centre for Frontier AI Research (CFAR), and Institute of High Performance Computing (IHPC), Agency for Science Technology and Research (A*STAR), 1 Fusionopolis Way, Singapore, 138632, Singapore e-mail: Yan_Ming,Joey_Zhou, Ivor_Tsang@cfar.a-star.edu.sg

results demonstrate the superior performance of our method compared to state-of-the-art approaches across all stance detection tasks. The contributions of our work are summarised as follows,

- We introduce collaborative knowledge verification to assist the detection model in selecting more semanticrelated knowledge from different knowledge sources. To our knowledge, this is the first work to infuse verified knowledge into the knowledge enhancement stance detection task.
- We introduce a collaborative adaptor to selective learning features in an efficient way for the low-resource setting. It contains three sub-components which are architecturally located in different positions of the backbone model, learning different features collaboratively.
- To alleviate the unbalanced effects of low-resource stance detection tasks, we also provided a staged optimization algorithm to improve the training efficiency in large PLMs. Experiments showed the superiority of our method in different low-resource settings and outperformed state-of-the-art approaches on three public stance detection datasets.

II. RELATED WORKS

A. Knowledge Enhancement

Knowledge enhancement increases the capabilities in thinking, understanding, and reasoning for the data-driven models beyond the original training data. In recent years, there has been a growing trend in infusing external-specific knowledge as complementary knowledge to the large pretrained models [12]. Depending on the infused knowledge, knowledge infusion methods can be broadly categorized into structuredknowledge infusion (*e.g.*, knowledge graph) and unstructuredknowledge infusion (*e.g.*, Wikipedia, text corpus).

Domain-specific experts typically collect structured knowledge and encompass well-organized and rich knowledge. For instance, CKE-Net [5] utilizes the structured knowledge base (ConceptNet) to enhance its model's common sense knowledge in zero/few-shot stance detection tasks. Similarly, K-BERT [13] incorporates domain knowledge through entity triplets obtained from the knowledge graph. Other methods like JAKET [14], ERNIE [15] and Entity-as Experts [16] also infuse knowledge from knowledge bases through grounding knowledge with entity linking technologies. Structured knowledge provides well-organized and domain-specific knowledge for specific targets in stance detection tasks. However, its utility is limited by the pre-defined scope of available knowledge, which may not cover all targets encountered in practical scenarios.

In contrast, unstructured knowledge offers more flexibility and can be easily collected from a wide range of diverse domains. For instance, the VAST [11] dataset introduces thousands of diverse targets that mostly can not be found in the well-constructed structured knowledge. To incorporate unstructured knowledge into the stance detection models, WS-BERT [4] directly infuses external knowledge from Wikipedia as its inputs to pre-trained models for stance detection in the VAST dataset. Another knowledge infusion paradigm is that finetune PLMs on the specific domain corpus to embed the domain-specific knowledge, as demonstrated by Sci-BERT [17], Bio-BERT [18], BERTweet [19]. In addition to domain-specific finetuning, Self-talk [20] offers another interesting solution by exploring knowledge from its own training corpus with hand-crafted prompts, enhancing language model learning with task-related knowledge. Furthermore, DDP [21] and K-Former [22] present the retrieval-based knowledge infusion methods by retrieving knowledge from feature pools and online websites, respectively. Nevertheless, more efforts are still needed to collaborate structured and unstructured knowledge together in a correct and efficient manner for large PLM-based models, particularly in low-resource tasks.

B. Stance Detection

Stance detection refers to the identification of attitudes toward a specific context or topic, typically framed as a stance classification problem (pros, cons, neutral) for the neural network-based models. Stance detection encompasses various tasks depending on the specific topics involved, including rummer stance detection [23], fake news stance detection [24], disinformation/misinformation stance detection [25], multi/cross-language stance detection [26], [27] and zeroshot stance detection [11], etc. In this study, we focus on in stance detection tasks of background knowledge infusion and low-resource training. To infuse background knowledge, existing approaches try to incorporate knowledge from different sources. For instance, CKE-Net [5] introduces targetrelated knowledge from ConceptNet, which is trained on the common sense knowledge graph. Similarly, BS-GGCN [28] simplifies the whole Concept-Net graph to a compact sentencerelated graph, enabling more efficient knowledge embedding for stance detection. Moreover, WS-BERT [4] leverages the background knowledge from Wikipedia pages as its additional input to improve the model performance in stance detection. Regarding the low-resource challenge, STCC [29] employs contrastive learning to enhance target representation in lowresource stance detection tasks. Different from STCC building contrastive examples in the existence of the target, Joint-CL [30] builds contrastive examples from the prototype graph representation of the target's link, which further improves the model performance on the unseen targets. While most of those works solve the knowledge infusion and low-resource separately, it is essential to consider two challenges together to improve the performance of stance detection models.

III. METHODOLOGY

Before delving into our methodology, we define the notations of the stance detection task as follows: Given a context set \mathbb{C} with elements c_i where $i = 1, 2, ..., n_c$, and a target set \mathbb{T} with element t_j where $j = 1, 2, ..., n_t$. The stance detection task is formulated to predict the stance ythat maximizes $P(y|\mathbb{C},\mathbb{T}\})$, where the stance set is Y ={pros, cons, neutral}. Regarding the knowledge enhancement stance detection task, its objective is defined as maximizing $P(y|\{\mathbb{C},\mathbb{T},\mathbb{K}\})$, where \mathbb{K} denotes the infused knowledge that assists in the stance detection task.



Fig. 1. Overview of our stance detection architecture: knowledge alignment, parameter-efficient learning, and staged optimization. Knowledge alignment collaboratively selects semantic similar knowledge of the target from different knowledge sources. Parameter-efficient learning introduces the collaborative adaptor and knowledge augmentation into the stance detection model to perform low-resource learning. Staged optimization algorithm optimizes the classifier with label smoothing, then pushes the classifier edge aligning to data distribution with the help of weighted loss.

The overview of our proposed methodology, as illustrated in Figure 1, contains three modules: (1) Knowledge alignment, which aims to collaborative select semantic target knowledge from structured and unstructured knowledge sources. (2) Parameter-efficient learning, which involves collaborative adaptor and knowledge augmentation to enhance model performance in low-resource settings. (3) Staged optimization algorithm, which further refines the adaptive model through strategies such as label smoothing and weighted loss.

A. Knowledge Alignment

We introduce a knowledge alignment module to assist the stance detection model infuse target-related background knowledge correctly, particularly for the targets that lack matching items in singular knowledge source (*e.g.*, Wikipedia). To help the unmatched targets infuse knowledge out of Wikipedia, we incorporate retrieval knowledge from the Internet, specifically through Google search, as an additional knowledge source. Thus, our approach adopts a multi-source knowledge infusion paradigm across structured Wikipedia and unstructured Internet, which selects the semantic similar knowledge as the collaborative knowledge \mathbb{K} to align to the target \mathbb{T} from multiple knowledge sources. Figure 2 illustrates our knowledge alignment module, as follows,

In the paradigms of knowledge enhancement stance detection, detection models mostly infuse extra knowledge through the target (\mathbb{T}) rather than the given context (\mathbb{C}). There are two reasons. Firstly, the stance detection task aims to identify the stance (Y) of the target, which may not be explicitly mentioned in the given context. Consequently, the target contains more information compared to the given context. Secondly, the context is typically long and complex, making it challenging to locate target-related information for infusing background knowledge (\mathbb{K}). In our proposed collaborative knowledge infusion approach (Figure 2), we collaboratively incorporate the background knowledge into the detection model by retrieving the target-related knowledge from Wikipedia and the Internet.

For structured Wikipedia knowledge, we utilize the target as the keyword to retrieve background knowledge through Wikipedia's API¹. This API returns a summary of the matched



Fig. 2. Knowledge alignment. The target's collaborative knowledge is the knowledge with a higher semantic similarity score from Wikipedia or the Internet. The target's Wikipedia knowledge is obtained by Wikipedia's API. The target's knowledge from the Internet is obtained by Google retrieval.

Wikipedia page. In cases where is no match for a target, we follow the setting of [4] and consider the target itself as the knowledge without introducing additional information.

For unstructured Internet knowledge, we retrieve the targetrelated web pages by using the searching prompt "What is the meaning/definition of TARGET (\mathbb{T})"? as the search term for the Google search engine. Subsequently, we select the top three pages from the Google search results and employ BeautifulSoup² to parse the *HTML* contents of these pages into candidate passage lists (\mathbb{D}). The next step involves filtering out unrelated contexts from the candidate passage lists, as web pages often contain noise and extraneous information. To accomplish this, we utilize *TF-IDF* ranker to identify and exclude noisy passages from a long list of candidate passages.

$$\operatorname{TFIDF}(\mathbb{T},\mathbb{D}) = \operatorname{TF}(\mathbb{T},\mathbb{D}) \cdot \operatorname{IDF}(\mathbb{T}),$$
 (1)

Once the knowledge related to the target has been collected from Wikipedia and the Internet, we introduce the knowledge verifier to select more accurate knowledge from multiple

²https://git.launchpad.net/beautifulsoup

sources as the infused knowledge. Knowledge verification involves feature encoding and feature similarity comparison, that selects semantic similar knowledge among different knowledge sources. Concretely, we employ Sentence-BERT [31] to encode the target \mathbb{T} and its corresponding knowledge (Wikipedia knowledge \mathbb{K}_w and Internet knowledge \mathbb{K}_g) into embedding features ($\mathbb{T}^{em}, \mathbb{K}_q^{em}, \mathbb{K}_q^{em}$), as follows,

$$\left\{\mathbb{T}^{em}, \mathbb{K}_{g}^{em}, \mathbb{K}_{w}^{em}\right\} = Stance(\{\mathbb{T}, \mathbb{K}_{w}, \mathbb{K}_{g}\}), \qquad (2)$$

Subsequently, we compute the semantic similarity between the stance target and different knowledge using the classical cosine similarity:

$$S(\mathbb{T}^{em}, \mathbb{K}^{em}_w) = \frac{\mathbb{T}^{em}, \mathbb{K}^{em}_w}{||\mathbb{T}^{em}|| \times ||\mathbb{K}^{em}_w||},$$
(3)

$$S(\mathbb{T}^{em}, \mathbb{K}_g^{em}) = \frac{\mathbb{T}^{em}, \mathbb{K}_g^{em}}{||\mathbb{T}^{em}|| \times ||\mathbb{K}_g^{em}||},$$
(4)

Finally, we select the knowledge with the highest semantic similarity as the collaborative knowledge \mathbb{K} to be infused into our model, which is expressed as follows:

$$\mathbb{K} = argmax\{S(\mathbb{T}^{em}, \mathbb{K}^{em}_w), S(\mathbb{T}^{em}, \mathbb{K}^{em}_g)\}, \qquad (5)$$

By collaborative integration of this verified knowledge, our stance detection model allows for the inclusion of more reliable knowledge in stance detection tasks. This knowledge infusion manner expands the scope of target knowledge by incorporating information from both structured and unstructured knowledge sources. As a result, it addresses the limitations associated with relying on a single knowledge source, including the issues of out-of-scope knowledge and false infusions.

B. Efficient Parameter Learning

To enable efficient parameter learning for low-resource stance detection tasks, our approach introduces collaborative adaptor and knowledge augmentation into our PLM-based stance detection model. Collaborative adaptor significantly reduces parameters compared to fine-tuning the entire model for low-resource stance detection tasks. Additionally, our collaborative adaptor learns diverse feature representations by leveraging the collaboration of multiple adaptors. To address the input-length limitation of stance models when incorporating collaborative knowledge contexts, we also introduce knowledge augmentation, which helps overcome the constraints imposed by the PLM's input length.

Suppose our stance detection model is parameterized by the fixed pretrained PLM backbone W and collaborative trainable adaptors ΔW . The backbone model can be a generic language model BERT or RoBERTa with Transformer architecture. In our low-resource stance detection task, the objective of the knowledge-infused stance detection task is formulated as follows:

$$J = \max_{\Delta W} \sum_{X} \sum_{Y} \sum_{V} \log(P_{W+\Delta W}(Y|\{\mathbb{C},\mathbb{T},\mathbb{K}))).$$
(6)

4

During the training process, our approach keeps the parameters of the pretrained model fixed and trains the collaborative adaptor on the low-resource stance detection dataset. The fixed-weight setting prevents the catastrophic forgetting problem and mitigates the challenges associated with training large models on limited training data (\mathbb{C}). Moreover, the collaborative adaptor has significantly fewer parameters than the PLMs $|\Delta W| << |W|$ (e.g., the prefix-tuning adaptor of our collaborative adaptor only has approximately 0.01% of the parameters W in BERT). This reduction in parameters greatly alleviates the data dependency for large model training.



Fig. 3. Overview of collaborative adaptor in efficient-parameter learning. Q, K, V are the query, key, and value of the Transformer module respectively. LoRA denotes the low-rank adaptor, and the gate is a controller for LoRA.

The motivation of our collaborative adaptor is to introduce multiple adaptive modules that collaboratively work together to provide a more powerful feature representation capability than the individual adaptors. Figure 3 presents the overview of our collaborative adaptor, which consists of gated lowrank adaptation, prefix-tuning, and attentive fusion modules. Those adaptive modules are hierarchically incorporated into different levels of the Transformer architecture. In detail, the gated low-rank adaptor (W_l) is inserted after the bottom hidden state layer of transformer architecture. It maps the selective embedding features from previous layers. In the intermediate level of the Transformer architecture, the prefix-tuning (W_p) introduces additional trainable prefix tokens before the key K and value V to incorporate new task-specific information into the PLMs. At the top of Transformer architecture, we introduce the attentive fusion module (W_f) to further select task-related features from value V. Therefore, all adaptive modules collaborate to learn task-specific information and feature representations in a hierarchical manner.

More specifically, the gated low-rank adaptation not only maintains the parameter efficiency of the low-rank adaptor (LoRA) [32], but also introduces a gate function Φ to selectively incorporate the learned features from LoRA. This gate function empowers the vanilla LoRA with a similar attentive capability to the Transformers module, by passing through the learned features selectively. Here, we default set the gated function Φ to *Sigmoid* function. Mathematically, if we denote the output of fully connection layer with normalization as the hidden state *H*, the low-rank downscale metric W_{down} and upscale metric W_{up} map features in an efficient-computing manner. Our gated LoRA is defined as follows:

$$I^{l} = \Phi(W_{uv}^{T}W_{down}H) + W_{I}H, I \in \{K, V, Q\},$$
(7)

Gated LoRA selectively passes the learned features from LoRA with the gated function Φ , which captures task-specific representation. The LoRA features are then collaborative with the Transformer's features to form the key K, value V, and query Q through an additive operation correspondingly. Before feeding the mapped embedding $\langle K, V, Q \rangle$ into multi-head computation, we introduce prefix tokens (T) into query and key ($\langle K, V \rangle$) to further enrich the feature learning capability of our detection model in the efficient-parameter learning.

$$Z_i = \operatorname{Attention}_i(Q, < \mathsf{T}_k; K >, < \mathsf{T}_v; V >), \qquad (8)$$

 Z_i is the i^{th} output of the multi-head attention computation. The prefix-tuning adaptors are introduced at the intermediate level of every Transformer, with significantly fewer trainable parameters compared to full parameter fine-tuning. To leverage the attentive mechanism, we introduce attentive fusion at the top of the Transformer, selectively activating the prefix-tuning features. As an attentive network, the key (K') and value (V') are derived from the outputs (Z_i) of the multi-head attention layer, and the query (Q') is obtained from the previous layer's query (Q) through a residual connection. The attentive fusion is performed using dot production \bigotimes . All computations of attentive fusion are listed as follows:

$$Z_{af} = \text{Softmax}(Q' \bigotimes K'),$$

$$Z = Z_{af} \bigotimes V'.$$
(9)

In this way, our collaborative adaptor performs efficientparameter learning architecturally across bottom embedding layers, middle of Transformer, and top feature fusion. All the efficient modules collaborate with each other to learn a generic representation for the low-resource stance detection tasks.

Besides the collaborative adaptor in efficient-parameter learning, we also introduce knowledge augmentation to facilitate efficient-parameter learning. In the knowledge infusion module, the crawled knowledge obtained from the Internet often exceeds the maximum input length of backbone PLMs. Our knowledge augmentation approach involves slicing the lengthy knowledge content into properly segmented parts to help the detection model capture the complete semantics of the infused knowledge.

Unlike previous approaches in knowledge-infused stance detection that infuses knowledge (\mathbb{K}) following the paradigm

$$[\mathbb{T}, \mathbb{C}, <\texttt{SEP} >, \mathbb{K}, <\texttt{CLS} >]$$
$$[\mathbb{T} < \texttt{SEP} >, \mathbb{C}, <\texttt{SEP} >, \mathbb{K}, <\texttt{CLS} >]$$

where \mathbb{T} , \mathbb{C} , and \mathbb{K} denote the target, given context, and knowledge, respectively. $\langle SEP \rangle$ and $\langle CLS \rangle$ are separate token and ending token for the PLMs. we reformulate our knowledge infusion paradigm into

$$[\mathbb{P}_t, \mathbb{C}, < \mathtt{SEP} >, \mathbb{K}_{sub}, < \mathtt{CLS} >]$$

Our input paradigm employs the prompt \mathbb{P}_t :

"What's the stance of
$$\mathbb{T}$$
 in following context?"

instead of the target \mathbb{T} to fully leverage the capability of PLMs, which matches the pre-training input format of two sentences split by $\langle \text{SEP} \rangle$, as well as keeping the semantic integrity.

In detail, we conduct the knowledge augmentation by slicing the long collaborative knowledge content into sub-knowledge segments, each of which fits the maximum length requirement. This manner helps the stance detection model capture the entire background knowledge instead of the cropped knowledge with missing information, as in the previous knowledge enhancement paradigms. The sub-knowledge segment \mathbb{K}_{sub} is sampled from the collaborative knowledge \mathbb{K} as follows:

$$\mathbb{K}_{sub} = [\mathbb{K}_{i*l/2}, \mathbb{K}_{(i+1)*l/2}],$$

s.t. $i \in (0, 1, ..., \lfloor \operatorname{Len}(\mathbb{K})/l \rfloor).$ (10)

The collaborative adaptor and knowledge augmentation work together to optimize efficient-parameter learning by reducing trainable parameters and addressing data limitations in low-resource stance detection tasks. The collaborative adaptor reduces the data consumption in training large-scale PLMs, while knowledge augmentation expands the training data to further improve training efficiency.

C. Staged Optimization Algorithm

To address the challenges of data discrepancy and domain gap between training data and pretrained models in the lowresource stance detection task, we propose a staged optimization algorithm that combines collaborative knowledge infusion and efficient parameter learning. However, the collaborative adaptor weights are initialized randomly, which may hinder its cooperation with the pretrained backbone PLMs during the initial training phase. Another issue is the unbalanced data distribution that is often overlooked in stance detection tasks. Our algorithm aims to mitigate these challenges in lowresource stance detection.

Algorithm 1 presents the algorithm details for the lowresource stance detection task. Once the collaborative knowledge \mathbb{K} and adaptor initialization are prepared, we set the first stage step for label smooth training, and set the second stage with the weighted loss for unbalanced stance categories. Before training begins, we prepare the prompt \mathbb{P}_t and augmented knowledge \mathbb{K}_{sub} with the given stance target \mathbb{T} as the training input triplet $\langle \mathbb{C}, \mathbb{P}_t, \mathbb{K}_{sub} \rangle$. Our two-stage learning algorithm incorporates weighted loss and label smoothing to improve the adaptive model's performance in unbalanced stance detection. Our algorithm aims to enhance the model's capability to handle ambiguous and diverse inputs by adjusting the loss function weights and introducing label smoothing

Require:
Initialize the target's collaborative knowledge \mathbb{K} .
Initialize the weights of collaborative adaptor $W_{l,p,f}$.
Ensure:
Set the input $\langle \mathbb{C}, \mathbb{T}, \mathbb{K} \rangle$ and stance label Y.
Set the total steps N, stage-1 step N_{s1} , stage-2 steps N_{s2} .
Set the initial learning rate α_0 , Set the loss weights
$[heta_a, heta_b].$
1: repeat
2: for $it = 1$ to N do
3: Fetch training data batch $< \mathbb{C}, \mathbb{T}, Y >$.
4: Set knowledge augmentation $\mathbb{K}_{sub} \leftarrow < \mathbb{T}, \mathbb{K} >$.
5: Set Prompt $\mathbb{P} \leftarrow \mathbb{P}_t$.
6: Stage-1: label smooth
7: if $it \leq N_{s1}$ then
8: Set smooth factor ε
9: else
10: $\varepsilon \leftarrow 0$
11: end if
12: Stage-2: weighted loss
13: if $it \ge N_{s2}$ then
14: $\theta \leftarrow \theta_a$
15: else
16: $\theta \leftarrow \theta_b$
17: end if
18: Feed Forward computing:
19: $\tilde{Y} \leftarrow model(\mathbb{C}, \mathbb{P}_t, \mathbb{K}_{sub} W, W_{l,p,f})$
20: Compute cost function:
21: $CE(Y, Y) \leftarrow -\sum_{c=1}^{c} \theta \cdot Y_c log(Y_c)$
22: $J(Y,\tilde{Y}) \leftarrow -(1 - \varepsilon)log(1 - CE(Y,\tilde{Y}) - \varepsilon)log(1 - CE(Y,\tilde{Y}))$
arepsilon log(CE(Y,Y))
23: Compute gradient:
24: $\Delta W_{l,p,f} \leftarrow \frac{\partial J(Y,Y)}{\partial W_{l,p,f}}$
25: Update gradient:
26: $W_{l,p,f} \leftarrow W_{l,p,f} + \alpha \Delta W_{l,p,f}$
27: end for
28: until convergence

during the training process. In the first stage, label smoothing is applied to soften the training targets, allowing the model to better handle uncertain data instances. Meanwhile, label smoothing helps the collaborative adapter convergence with newly initialed parameters. In the second stage, a weighted loss function is employed to assign different weights to different classes, thereby improving the model's ability to handle unbalanced datasets. Our algorithm provides a promising solution for tackling the challenges of data discrepancy and unbalanced data distribution in low-resource stance tasks.

IV. EXPERIMENTS

To evaluate the effectiveness of our proposed method, we conducted extensive experiments on three publicly available stance detection datasets and different low-resource settings. In this section, we provide a brief description of the datasets and the compared methods used in our stance detection task. Finally, we summarize the results using F1 as the default evaluation metric.

A. Datasets

VAST [11] is a typical zero-shot/few-shot stance detection dataset that covers a wide range of over 6000 targets across various themes, including politics, sports, education, immigration, and public health, etc. The *VAST* dataset consists of 13447, 2062, 3006 examples in its training, validation, and test sets, respectively. Notably, the majority of targets in *VAST* are designed for zero-shot setting. It has an average of approximately 2.4 examples per target. This characteristic makes *VAST* particularly suitable for zero-shot/few-shot stance detection tasks.

P-Stance [33] is stance detection dataset specific to the political domain. It contains in-target and cross-target settings with 21, 574 labeled tweets on three specific targets: "Biden", "Sanders" and "Trump". In the in-target setting, the target and classifier are the same in both the training and evaluation sets. Conversely, in the cross-target setting, the targets are entirely different, allowing for the evaluation of the generalization performance.

COVID-19-Stance [34] is stance detection dataset constructed from COVID-19-related tweets. It contains 6,133 tweets with respect to four specific targets: "Anthony S. Fauci, M.D. (Fauci)", "Keep School Closed (School)", "Stay at Home Order (Home)" and "Wearing a Face Mask (Mask)". *COVID-19-Stance* is also an unbalanced dataset in terms of class distribution.

B. Compared Methods

TAN [35] is a classical attention-based method for the stance detection task. It contains a target-specific attention extractor and a long short-term memory network.

BERT [36] is a well-known Transformer-based pretrained language model widely used for various downstream tasks. We employ BERT as our baseline for reference in the stance detection task.

WS-BERT-Dual [4] infuses target-related knowledge from extra Wikipedia to enhance background knowledge of PLMs in stance detection tasks. It utilizes two pretrained encoders to encode tweets and knowledge separately.

In addition to the shared baselines mentioned above, we also introduce other strong specific baselines for different stance datasets. For VAST task, we introduce graph convolution networks-based methods *BERT-GCN*, *CKE-NET* [9], and *BSRGCN* [28]. Those methods joined large pretrained models with graphic convolution networks to leverage the learning capability of the models on heterogeneous data with structured graphic representation. For the zero-shot setting in *VAST* task, we select BSRGCN and contrastive learning-based *Joint-CL* [30] as the compared methods.

For *PStance* task, we choose the bi-recurrent neural networks (*BiCE*) [37] and gated convolutional neural networks (*GCAE*) [38] and PGCNN [39] as the baselines. Specifically, the GCAE uses a TanH as the gate function to selectively output the sentiment futures according to the given aspect.

Similarly, PGCNN also used the parameterized filters as the gate function to effectively capture the aspect-specific features. Moreover, we also include *BERT-Tweet* [33] as the compared method, which is pretrained on the target Tweet domain data.

For *COVID-19-Stance* task, we choose *CT-BERT* [34] as the baseline, which pretrained on the COVID-19-related tweet corpus to enhance task-specific domain knowledge. We also include *CT-BERT-NS* and *CT-BERT-DAN* [40], which incorporate self-training and domain adaption into CT-BERT to further improve model representation capability and reduce domain gap in the stance detection task.

C. Implementation Details

We utilize RoBERTa as the backbone model for both VASE and PStance tasks. The batch size is set to 16, and the learning rates range from 1e - 5 to 5e - 5 in our experiments. Since COVID-19-Stance is a task related to the COVID-19 pandemic, we employ CT-BERT and BERT as the backbone models to leverage COVID-19-related knowledge from pretrained models. Due to our GPU memory limitations, the batch size is set to 8. Regarding the hyperparameters of the collaborative adaptor, we set the rank to r = 8 for the low-rank adaptor. The prefix-taken is set to 100 with a dropout rate 0.2. The reducing factor for the feature fusion module is 16, and all gates are the ReLU function. The models are implemented using Pytorch, and the maximum input length is default set to 512 tokens. We trained the models for a maximum of 30 epochs, and applied stopping with a patience of 10 epochs. The optimizer used is AdamW with a weight decay of 1e-5. All the experiments are conducted with the same random seed on four NVIDIA RTX A5000 GPU cards.

D. Results

To verify the effectiveness of our proposed method, we evaluate its performance on three public stance detection tasks: *VAST*, *PStance*, and *COVID-19-Stance*. Firstly, we evaluate the proposed method on the *VAST*, a low-resource dataset with a significantly larger number of targets than the other two datasets. Additionally, we evaluate the method's performance on the *PStance* and *COVID-19-Stance* datasets. Following previous work [4], all the datasets are evaluated using the macro-average F1-score as the standard metric. The overall performance is calculated as the average across all stances.

VAST dataset officially splits into two sub-tasks: zero-shot stance detection (600 targets) and few-shot stance detection (159 targets). Zero-shot setting does not include any targets in its training set, while the few-shot setting has very limited training samples (approximately 14.8 examples per target) in its training set. In contrast, the *PStance* and *COVID-19-Stance* datasets have over hundreds of training examples per target. Table I summarizes the evaluation results on the VAST dataset.

From the numbers presented in Table I, we can observe that the baseline method BERT achieves clear improvements around 2% in the overall performance compared to the none pretrained baseline TAN, which indicates pretrained models have a stronger feature representation capability than none

TABLE I EXPERIMENTAL RESULTS ON VAST

Method	Zero-shot	Few-shot	Overall
TAN	66.6	66.3	66.5
BERT	68.5	68.4	68.4
BERT-GCN	68.6	69.7	69.2
CKE-Net	70.2	70.1	70.1
BSRGCN	72.6	70.2	71.3
Joint-CL	72.3	71.6	72.3
WS-BERT-Dual	75.3	73.6	74.5
Ours	81.9	79.6	80.7

pretrained TAN. Building upon BERT, BERT-GCN incorporates graphic convolution networks (GCN) with BERT further improving the overall F1 score to 69.2. Similarly, the GCNbased methods CKE-Net and BSRGCN demonstrate progressive improvements by leveraging graph convolution networks, achieving an F1 score of 71.3. Specifically, BSRGCN performs better in the zero-shot setting, benefiting from the unsupervised training on the domain-specific corpus. Joint-CL further enhances model performance through contrastive learning, achieving an overall F1-score of 72.3. To enhance background knowledge, WS-BERT-Dual introduces target-related knowledge from Wikipedia, resulting in signification improvements compared to previous methods.

However, all those solutions overlooked the fact that VAST is a low-resource task, particularly for large pretrained models. Our method addresses this issue by incorporating efficientparameter learning, and staged optimization for the lowresource task. Another neglect point in those solutions is that the infused target's knowledge should be the corrected knowledge. Our method addresses this issue by incorporating the collaborative knowledge infusion that introduces knowledge from multiple knowledge sources in a more accurate way. As a result, our method achieves new state-of-the-art (SOTA) performance on VAST, achieving an overall F1-score of 80.7. Interestingly, we find that the zero-shot settings achieve higher scores than the few-shot settings, especially in the pretrainbased methods. This difference can be attributed to the fact that the zero-shot and few-shot sets are two distinct subsets with completely different targets in the test set. Consequently, we can treat these two settings as two separate datasets.

TABLE II EXPERIMENTAL RESULTS ON PSTANCE

Method	Trump	Biden	Sanders	Avg.
TAN	77.1	77.6	71.6	75.1
BiCE	77.2	77.7	71.2	75.4
PGCNN	76.9	76.6	72.1	75.2
GCAE	79.0	78.0	71.8	76.3
BERT	78.3	78.7	72.5	76.5
BERT-Tweet	82.5	81.0	78.1	80.5
WS-BERT-Dual	85.8	83.5	79.0	82.8
Ours	86.2	84.1	80.5	83.6

Different from VAST with a large number of targets in a low-resource setting, *PStance* contains only 3 targets, and *COVID-19-Stance* contains 4 targets. Table II presents the evaluation results of compared methods on *PStance*. The classical recurrent neural network-based TAN and BiCE obtain comparable performance to the GCN-based PGCNN and GCAE, yielding an F1-score around 75 76. Those results approach the performance of pretrained BERT. This similarity in performances suggests that rich training sources can benefit different types of models, and even achieve comparable performance to pretrain-based BERT. Different from BERT pretrained on the generic corpus, BERT-Tweet enhances domain-specific knowledge by being pretrained on the Twitter corpus, resulting in significant improvements of 4% in the overall F1-score. Building upon BERT-Tweet, WS-BERT-Dual further infused the target background knowledge from Wikipedia, attaining an overall F1-score of 82.8. In light of WS-BERT-Dual, our method further optimizes knowledge augmentation through parameter-efficient learning, achieving the best performance across all the targets than the compared methods in *PStance* task.

 TABLE III

 EXPERIMENTAL RESULTS ON COVID-19-STANCE

Method	Fauci	Home	Mask	School	Avg.
TAN	54.7	53.6	54.6	53.4	54.1
ATRGU	61.2	52.1	59.9	52.7	56.5
GCAE	64.0	64.5	63.3	49.0	60.2
CT-BERT	81.8	80.0	80.3	75.3	79.8
CT-BERT-NS	82.1	78.4	83.3	75.3	79.8
CT-BERT-DAN	83.2	78.7	82.5	71.7	79.0
WS-BERT-Dual	83.6	85.0	86.6	82.2	84.4
Ours	86.05	86.76	86.91	83.33	85.76

We also conducted evaluations of the proposed method on the domain-specific COVID stance detection and present the results in Table III. From the comparison of the results, we can clearly observe the performance gap between traditional gated-based methods (TAN, ATGRU, GCAE) and pretrainbased models (CT-BERT and its variants). The gated-based methods, which only conduct finetuning on its rich training set, obtain low average F1-scores below 60.2, lacking any background knowledge specific to the target domain.

In contrast, pretrained models trained on the COVID-related Twitters data exhibit good background knowledge and feature representation for COVID-19-Stance, resulting in a high average F1-score above 79.0. Meanwhile, self-training and domain adaptation techniques applied to CT-BERT led to performance improvements in stance detection for "Fauci" and "Mask", but no substantial improvements in the overall F1score. With the help of dual pretrained model encoders, WS-BERT-Dual further elevates the overall performance to 84.4. Similarly, our method achieves the best performance among the compared methods in COVID-19-Stance by leveraging collaborative adaptor and staged optimization. Based on the extensive experimental comparisons, we can conclude that our proposed method performs well not only in low-resource VAST stance detection task but also in rich-resource PStance and COVID-19-Stance tasks.

V. DISCUSSION

A. Ablation Study

We performed an ablation study on the main modules of our method, namely collaborative knowledge infusion, efficient parameter learning, and staged optimization, using the VAST and *PStance* datasets. In addition to reporting the overall F1-score for zero-shot and few-shot settings in *VAST*, we also provide the detailed results for three specific stances: pros, cons, and neutral. For *PStance*, we report the average performance across the different targets.

Table IV presents the results of the ablation study conducted on VAST, where RoBERTa serves as the backbone model. We study the impact of each module on the backbone performance. All individual modules that work with the backbone outperformed the finetuning of the vanilla backbone. The collaborative knowledge infusion (KI) module, which includes knowledge verification and augmentation, facilitated the learning of target-specific background knowledge and achieved remarkable improvements of 75.1%. Likewise, the efficient parameter learning module (EP) proved beneficial for the stance detection model on the low-resource VAST data with the help of collaborative adaptors. When comparing the performance across different stances, we observed that the neutral stance exhibited significantly higher scores compared to the pros and cons stances. Our staged optimization (SO) module tries to address this bias by incorporating label smooth and weighted loss, resulting in overall performance improvements. Furthermore, extensive ablation studies were conducted to assess different combinations of the modules. We can observe that two module combinations further improve model performance by 1% - 2%. Similarly, our method incorporating all three modules achieved the best overall F1-score of 80.7%on VAST, highlighting the effectiveness of each module in addressing the challenges of the low-resource VAST task.

We also conducted an ablation study on PStance, which benefits from a relatively rich training source than the lowresource VAST dataset. Table V presents a summary of the ablation study on PStance, focusing on the efficiency of three modules with the same backbone RoBERTa. Notably, the ablation study results differ from the results obtained for VAST. Interestingly, we observed a slight decrease in model performance with the collaborative knowledge infusion (KI) module than the vanilla backbone. This performance drop may be attributed to the specific dataset, as PStance only contains four targets compared to diverse targets in VAST. In other words, only four background long-sequential knowledge content \mathbb{K} is infused into the model training pipeline, which may hinder feature learning on short-sequential raw tweets content \mathbb{C} . In the single-modular settings, we found backbone incorporating knowledge infusion (KI) or staged optimization (SO) exhibited superior performance compared to the efficient parameter learning (EP) module in the rich-source PStance task. This suggests that full parameter finetuning is more effective than the adaptor-based solution in data-rich tasks. Similarly, in the two-modular settings, we observed that the setting with EP module (setting 4&6) performed worse than the setting without EP modular (setting 5). We also observe that the backbone with two-modular settings yielded more improvements in the overall F1-score than the single-modular settings. In the three-modular setting, we find the performance of the 'Trump' and 'Sanders' targets could be further improved compared to the two-modular settings. However, the 'Biden' category experienced a significant drop compared to its results

Backbones	KI	EP	SO	Cons(ZS)	Pros(ZS)	Neu(ZS)	Zero-Shot	Cons(FS)	Pros(FS)	Neu(FS)	Few-Shot	Avg.
RoBERTa-L				65.7	59.0	95.0	73.3	65.6	60.5	97.4	74.5	73.9
RoBERTa-L	\checkmark			66.4	64.9	93.5	74.9	66.4	63.2	96.3	75.3	75.1
RoBERTa-L		\checkmark		71.4	72.3	90.7	78.1	67.7	68.6	88.5	74.9	76.5
RoBERTa-L			\checkmark	74.5	72.0	92.0	79.5	69.2	67.8	87.0	74.7	77.0
RoBERTa-L	\checkmark	\checkmark		70.7	70.4	95.5	78.9	68.9	70.1	95.9	78.3	78.6
RoBERTa-L		\checkmark	\checkmark	69.1	73.5	94.9	79.2	68.5	72.8	96.7	79.3	79.3
Ours	\checkmark	\checkmark	\checkmark	75.2	75.3	95.1	81.9	71.5	72.6	94.7	79.6	80.7

TABLE IV Ablation Study on VAST

'Neu' is short for the neutral stance.

'ZS' and 'FS' are shorts of Zero-shot and Few-shot. 'L' denotes the large size backbone.

'KI', 'EP', and 'SO' denote knowledge infusion, collaborative adaptor, and staged optimization.

TABLE V Ablation Study on PStance

Setting	KI	EP	SO	Trump	Biden	Sanders	Avg.
0				85.5	82.7	76.9	81.7
1	\checkmark			84.6	82.7	76.9	81.4
2		\checkmark		76.3	80.5	78.5	78.4
3			\checkmark	84.4	83.6	78.4	82.1
4	\checkmark	\checkmark		84.8	82.7	73.7	80.4
5		\checkmark	\checkmark	86.1	83.3	80.5	83.3
6	\checkmark		\checkmark	85.0	85.5	80.5	83.7
7	\checkmark	\checkmark	\checkmark	86.2	84.1	80.5	83.6

'KI', 'EP', and 'SO' denote collaborative knowledge infusion, efficient parameter learning, and staged optimization.

in the setting-6. We attribute this to the negative impact of the 'EP' module in the three-modular setting, which slightly decreased the overall performance compared to the twomodular setting 6. Thus, we can conclude that the adaptorbased solution does not always perform well in rich-resource tasks.

B. Cross-target Stance Detection on PStance

We evaluate the model's generalization performance on the rich-source PStance, we conducted cross-target stance detection, training model on one target, and evaluated on another target (e.g., training on Trump and testing on Biden). We employed BERT-Tweet [33] and WS-BERT-Dual [4] as the strong baselines. BERT-Tweet is pretrained on the Twitter corpus, benefiting from domain-specific knowledge. WS-BERT-Dual is a dual-path architecture using BERT and BERT-Tweet as feature encoders to incorporate target-specific Wikipedia knowledge. We follow the experimental settings of WS-BERT-Dual, testing on three targets: Trump, Biden, and Sanders. We employ the knowledge infusion and the staged optimization modules for cross-target stance detection tasks, as our ablation study demonstrated the two-modular setting performs best in rich-resource PStance. All the experimental results are summarized in the following table:

From the results of Table VI, we observed that the WS-BERT-D achieved significant improvements (average 5% F1-score) in all six cross-target pairs compared to the baseline BERT-Tweet, benefiting from the additional BERT branch for encoding extra Wikipedia knowledge. In contrast, our method only uses the RoBERTa as the backbone for cross-target

TABLE VI CROSS-TARGET STANCE DETECTION ON PSTANCE

Cross-Targets	BERT-Tweet	WS-BERT-Dual	Ours
$\texttt{Trump} \rightarrow \texttt{Biden}$	58.9	68.3	68.2
$\mathtt{Trump} o \mathtt{Sanders}$	56.5	64.4	67.8
$\mathtt{Biden} o \mathtt{Trump}$	63.6	67.7	72.1
$\texttt{Biden} \to \texttt{Sanders}$	67.0	69.0	74.8
$\mathtt{Sanders} o \mathtt{Trump}$	58.7	63.6	63.4
$\mathtt{Sanders} o \mathtt{Biden}$	73.0	76.8	78.9
Avg.	63.0	68.3	70.9

* All the reported results are F1 score.

* In cross-target setting, the stance detection model is trained on lefttarget data and tested on the right-target data.

stance detection. We further improved four of six cross-target pairs by introducing the staged optimization, resulting in a slight performance drop in "Trump \rightarrow Biden" and "Sanders \rightarrow Trump" than the WB-BERT-Dual. Additionally, we noticed that the two pairs' target results are not symmetric to each other. Overall, our proposed solution achieved a new stateof-the-art performance in cross-target stance detection on the *PStance* dataset.

C. Low-resource Stance Detection

In this section, we further evaluated our method's performance in the low-resource settings of stance detection on the subsets of PStance and COVID-19-Stance datasets. Specifically, we randomly sampled 5%, 10%, 15%, 20% data from their training sets as our low-resource data settings, and kept its test sets for evaluation.

TABLE VII Low-resource Stance Detection on PStance

Settings		Baseline			Ours		Avg.(%)
~ 8-	Trump	Biden	Sanders	Trump	Biden	Sanders	
5%	69.3	70.5	63.5	71.2	73.9	68.2	3.4 ↑
10%	70.5	72.4	71.0	73.8	75.0	72.5	2.5 ↑
15%	72.9	76.1	72.3	74.2	78.1	73.9	1.7 ↑
20%	74.0	78.4	74.5	76.9	79.4	75.1	1.7 ↑

'Settings' means different low-resource settings with different sample percentages from PStance.

'Avg.' denotes the average performance improved between our method and the baseline.

Following the setting of the ablation study, we set the BERT as the baseline and backbone of our method in all the low-

 TABLE VIII

 LOW-RESOURCE STANCE DETECTION ON COVID-19-STANCE

Settings	ngs Baseline					Ours			
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	Fauci	Home	Mask	School	Fauci	Home	Mask	School	8-(/-0)
5%	34.6	44.0	37.0	22.7	40.8	44.0	37.0	22.7	6.6 ↑
10%	57.8	59.9	42.3	26.3	72.9	59.9	42.3	26.3	12.7 ↑
15%	69.6	62.1	49.0	40.3	77.2	62.1	49.0	40.3	6.4 ↑
20%	70.8	74.2	54.0	45.4	80.6	74.2	54.0	45.4	12.8 ↑

resource stance detection evaluations. Table VII and Table VIII summarized the comparison results on different low-resource settings of PStance and COVID-19-Stance.

From the reported results of Table II on different lowresource settings, we can observe that our method surpasses the baseline compared to the baseline in all low-resource data settings, which shows our method presents the effectiveness of our proposed method. Overview the whole performance across different settings, we can obviously find the progressive increasing trends with the training data adding in both the baseline and our method. Last but not least, our method benefits more with less training data, and the average improvements are reduced with more data introduced from 5% to 20%.

Table VIII summarized the low-resource setting evaluation results on the COVID-19-Stance dataset. Similar to the lowresource setting results of PStance, our method surpassed the baseline with large margins in all low-resource settings. The average performance can achieve more than 6%. Different from the performance trends in PStance, the performance improvement trends do not keep consistently changing with increasing of training data. We think the main reason causing this is the domain gap and diverse data distributions in different stance topics of COVID-19-Stance.

## D. Efficient Parameter Learning

In this section, we compare the performance of our efficient parameter learning paradigm with the entire model finetuning paradigm in the low-resource stance detection task. We selected the classic BERT and RoBERTa for entire model finetuning, using both basic (B) and large (L) model sizes. Our efficient parameter learning approach utilizes the largesize RoBERTa as the backbone. All the models are evaluated in the zero-shot and few-shot settings as defined by *VAST* dataset. Note, the few-shot and zero-shot in traditionally computer vision tasks are evaluated with same test data. However, in our VAST NLP dataset, the few-shot and zero-shot settings are evaluated with different test data.

Table IX presents the performance of different models on the zero-shot stance detection task with respect to three stances: pros, cons, and neutral. In the entire model finetuning setting, we observe that the basic-sized models outperformed the large-sized models in terms of average F1 score. Specifically, the BERT models experienced a 1% drop in average F1-score from basic-size to large-size, while the RoBERTa models even encountered a decrease of more than 4%. The large model's zero-shot performance decay may attribute to combined training few-shot on the low-resource VAST dataset.

TABLE IX Zero-shot Performance on VAST

Methods	Param*	Pros	Cons	Neutral	Avg.	
BERT-B	110M	64.0	63.2	94.2	73.8	
BERT-L	340M	60.0	66.8	94.1	72.4	
RoBERTa-B	110M	67.4	72.3	93.7	77.8	
RoBERTa-L	340M	65.7	59.0	95.0	73.3	
Ours	3K	75.2	75.3	95.1	81.9	

Param* denotes the total trainable parameters.

'B' and 'L' denote the basic and large model sizes, respectively.

The large-size models exhibit reduced generalization capability with limited training data samples, resulting in performance decay in zero-shot stance detection tasks. Consistent with the findings of the ablation study, the neutral stance achieved significantly higher scores (above 90%) than the pros and cons stances in the zero-shot setting. In contrast, our efficient parameter learning method, which maintains the pretrained model's generalization capability by freezing its parameter, achieved the best performance in the zero-shot stance detection task.

TABLE X Few-shot Performance on VAST

Methods	Param*	Pros	Cons	Neutral	Avg.
BERT-B	110M	64.2	65.1	95.1	74.8
BERT-L	340M	64.2	61.3	91.8	73.8
RoBERTa-B	110M	64.6	70.8	95.1	76.8
RoBERTa-L	340M	65.6	60.5	97.4	74.5
Ours	3K	71.5	72.6	94.7	79.6

Param* denotes the total trainable parameters.

'B' and 'L' denote the basic and large model sizes, respectively.

Table X presents the results of different models in the fewshot VAST stance detection task. Similar to the zero-shot setting, we observe a decline in overall performance as the model size increases. Additionally, the neutral stance detection performance exhibited significant superiority over the pros and cons stances. For the entire model finetuning paradigm, there have been no notable variations between basic and large-size models in the pros and neutral stances. However, in the cons stance, the large-size models experienced a sharp drop. This can be attributed to the unbalanced data distribution, which leads to model performance decay when trained on limited samples. Therefore, the entire model finetuning using a large pretrained language model is not an optional solution for zero/few-shot stance detection tasks. The reason for zero-shot setting (Table IX) achieved better performance than the fewshot setting (Table X) is mainly depends on the special test data setting in the VAST dataset that the few-shot and zeroshot settings have total different test data samples, which is different from traditional few-shot and zero-shot setting with the same test data. In other words, the settings in Table IX and X can be seen as two datasets. This divergence from traditional evaluation methods is a key factor in understanding the experimental results. Moreover, this pattern of zero-shot settings outperforming few-shot settings is not unique to our study but is also observed in comparable works, such as TAN [35], BERT [36], CKE-NET [9], BSRGCN [8], and WS-BERT-Dual [4], which report similar anomalies in Table 1. Our proposed method offers a solution for the low-resource stance detection task, which has only 3K trainable, significantly fewer than the millions of trainable parameters required for full model finetuning. Despite the parameter reduction, our method still can achieve superior performance with a considerable margin.

## E. Collaborative Adaptor Analysis

Collaborative adaptor is an essential part of efficient parameter learning in the low-resource detection task, which consists of three modules: gated LoRA, prefix-tuning and attentive fusion. To assess the importance of each module, we conducted evaluations by removing individual modules from our collaborative adaptor.

TABLE XI Collaborative Adaptor Analysis on VAST

Settings	Pros	Cons	Neutral	F1
w/o Gated LoRA	71.9	70.3	95.8	79.3
w/o Prefix-tuning	70.6	73.3	94.3	79.4
w/o Attentive Fusion	66.1	62.1	92.5	73.6
All	73.4	73.9	94.9	80.7
<i>w/o</i> denotes without	ut.			

Table XI presents the performance of different settings of the collaborative adaptor. For instance, the setting 'w/o gated LoRA' indicates the removal of the 'gated LoRA' module. All settings were evaluated on the low-resource VAST using the same hyperparameters. We observed the 'gated LoRA' and 'prefix-tuning' modules exhibited similar drops of approximately 1% with slight variations across different stances. Surprisingly, in the 'w/o attentive fusion' setting, all stance scores experience a sharp decline of approximately 7% on average. Through the performance comparisons, we discovered that the attentive fusion module had a more significant impact on down-steam stance detection tasks than the gated LoRA and Prefix-tuning modules in the efficient-parameter learning paradigm. One possible explanation for this observation is that attentive fusion is more closely connected to the stance prediction classifier than the other two models, which serve as feature extractors with less impact on the final stance prediction.

# F. Comparison with ChatGPT

ChatGPT ³ attracts lots of attention in the natural language processing community due to its impressive performance

on conversational tasks, leading to its utilization in various downstream NLP tasks. In this section, we aim to evaluate the performance of ChatGPT on *VAST*, which is a varied stance topics dataset with over a thousand targets. To adapt ChatGPT for the stance detection task, we constructed the prompt as follows:

"Please choose one stance from cons, pros, neutral for < TARGET  $> (\mathbb{T})$  on following content: < TWEET  $> (\mathbb{C})$ ?"

We sequentially selected and evaluated 100 samples, comprising 33 cons, 33 pros, and 34 neutral instances. Regarding ChatGPT is an evolving system⁴, all the evaluation results are reported as follows,

TABLE XII VAST SAMPLES EVALUATION ON CHATGPT

GT/Pred	Cons	Pros	Neutral	Recall
Cons (33)	18	4	9	58.1
<b>Pros</b> (33)	3	18	14	51.4
Neutral (34)	12	11	11	33.3
Precision	54.5	54.5	32.4	/

The row is prediction (Pred) and the column is ground truth (GT).

The evaluation results presented in Table XII indicate that ChatGPT's performance on *VAST* stance detection dataset is not as impressive as anticipated, which is similar to the findings of the work directly using the chain-of-thought [41] in ChatGPT for stance detection on *VAST* dataset with only 62.3 F1 performance. From the result analysis, we observe that ChatGPT often predicts the pros or cons stances to the neutral stance, resulting in the neutral stance in a low recall score. This tendency might stem from ChatGPT's inclination to produce mild and friendly responses [42], leading to a bias toward predicting neutral stances. Furthermore, we observed that ChatGPT trends to output a neutral stance for sensitive topics, such as voting, humanity, and elections.

#### VI. CONCLUSION

In this paper, we propose a method for low-resource stance detection that collaborative infuses verified target knowledge with efficient parameter learning. Firstly, we enhanced the infusion of target-related knowledge by extending it beyond structured Wikipedia to encompass a broader range of unstructured information from the entire Internet. To ensure the selection of relevant semantic background knowledge, a knowledge verifier is employed. Secondly, we introduce efficient-parameter learning through collaborative adaptors, which involve a minimal number of trainable parameters by freezing the weights of large PLM-based models. This manner not only facilitates efficient model training in low-resource stance detection tasks but also retains the rich prior knowledge encoded in pretrained models. Thirdly, a staged optimization algorithm is proposed to mitigate the impact of unbalanced data. Additionally, knowledge augmentation and prompting techniques are integrated into our efficient parameter learning framework for low-resource stance detection. Experimental results demonstrate the effectiveness of our method on three

⁴The ChatGPT results were evaluated on 04 June 2023.

³https://openai.com/blog/chatgpt

public datasets with state-of-the-art performance. In future work, we plan to further explore efficient-parameter learning in the context of multi-modal stance detection tasks.

## REFERENCES

- [1] S. Ghosh, P. Singhania, S. Singh, K. Rudra, and S. Ghosh, "Stance detection in web and social media: A comparative study," in *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, F. Crestani, M. Braschler, J. Savoy, A. Rauber, H. Müller, D. E. Losada, G. Heinatz Bürki, L. Cappellato, and N. Ferro, Eds. Cham: Springer International Publishing, 2019, pp. 75–87.
- [2] A. Sen, M. Sinha, S. Mannarswamy, and S. Roy, "Stance classification of multi-perspective consumer health information," in *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, ser. CODS-COMAD '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 273–281. [Online]. Available: https://doi.org/10.1145/3152494.3152518
- [3] K. Kawintiranon and L. Singh, "Knowledge enhanced masked language model for stance detection," in *Proceedings of the 2021 Conference* of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Online: Association for Computational Linguistics, Jun. 2021, pp. 4725–4735. [Online]. Available: https://aclanthology.org/2021.naacl-main.376
- [4] Z. He, N. Mokhberian, and K. Lerman, "Infusing knowledge from Wikipedia to enhance stance detection," in *Proceedings of the 12th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis.* Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 71–77. [Online]. Available: https://aclanthology.org/2022.wassa-1.7
- [5] R. Liu, Z. Lin, Y. Tan, and W. Wang, "Enhancing zero-shot and few-shot stance detection with commonsense knowledge graph," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Online: Association for Computational Linguistics, Aug. 2021, pp. 3152–3157. [Online]. Available: https://aclanthology. org/2021.findings-acl.278
- [6] O. Agarwal, H. Ge, S. Shakeri, and R. Al-Rfou, "Knowledge graph based synthetic corpus generation for knowledge-enhanced language model pre-training," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, Jun. 2021, pp. 3554–3565. [Online]. Available: https://aclanthology.org/2021.naacl-main.278
- [7] Y. Xu, C. Zhu, S. Wang, S. Sun, H. Cheng, X. Liu, J. Gao, P. He, M. Zeng, and X. Huang, "Human parity on commonsenseqa: Augmenting self-attention with external attention," in *IJCAI 2022*, July 2022, human parity result on CommonsenseQA.
- [8] X. Wang, T. Gao, Z. Zhu, Z. Zhang, Z. Liu, J. Li, and J. Tang, "KEPLER: A Unified Model for Knowledge Embedding and Pretrained Language Representation," *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 176–194, 03 2021. [Online]. Available: https://doi.org/10.1162/tacl_a_00360
- [9] Y. Lin, Y. Meng, X. Sun, Q. Han, K. Kuang, J. Li, and F. Wu, "BertGCN: Transductive text classification by combining GNN and BERT," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Online: Association for Computational Linguistics, Aug. 2021, pp. 1456–1462. [Online]. Available: https: //aclanthology.org/2021.findings-acl.126
- [10] J. Wei, X. Wang, D. Schuurmans, M. Bosma, brian ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou, "Chain of thought prompting elicits reasoning in large language models," in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [11] E. Allaway and K. McKeown, "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).* Online: Association for Computational Linguistics, Nov. 2020, pp. 8913–8931. [Online]. Available: https://aclanthology.org/2020.emnlp-main.717
- [12] C. Zhu, Y. Xu, X. Ren, B. Y. Lin, M. Jiang, and W. Yu, "Knowledgeaugmented methods for natural language processing," ACL 2022, p. 12, 2022.
- [13] W. Liu, P. Zhou, Z. Zhao, Z. Wang, Q. Ju, H. Deng, and P. Wang, "K-bert: Enabling language representation with knowledge graph," *Proceedings of the AAAI Conference on Artificial Intelligence*,

- [14] D. Yu, C. Zhu, Y. Yang, and M. Zeng, "Jaket: Joint pre-training of knowledge graph and language understanding," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, pp. 11 630–11 638, Jun. 2022. [Online]. Available: https://ojs.aaai.org/index. php/AAAI/article/view/21417
- [15] Z. Zhang, X. Han, Z. Liu, X. Jiang, M. Sun, and Q. Liu, "ERNIE: Enhanced language representation with informative entities," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 1441–1451. [Online]. Available: https://aclanthology.org/P19-1139
- [16] T. Févry, L. Baldini Soares, N. FitzGerald, E. Choi, and T. Kwiatkowski, "Entities as experts: Sparse memory access with entity supervision," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 4937–4951. [Online]. Available: https://aclanthology.org/2020.emnlp-main.400
- [17] I. Beltagy, K. Lo, and A. Cohan, "SciBERT: A pretrained language model for scientific text," 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)*. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3615–3620. [Online]. Available: https://aclanthology.org/D19-1371
- [18] K. r. Kanakarajan, S. Ramamoorthy, V. Archana, S. Chatterjee, and M. Sankarasubbu, "Saama research at MEDIQA 2019: Pretrained BioBERT with attention visualisation for medical natural language inference," in *Proceedings of the 18th BioNLP Workshop* and Shared Task. Florence, Italy: Association for Computational Linguistics, Aug. 2019, pp. 510–516. [Online]. Available: https: //aclanthology.org/W19-5055
- [19] D. Q. Nguyen, T. Vu, and A. Tuan Nguyen, "BERTweet: A pre-trained language model for English tweets," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 9–14. [Online]. Available: https:// aclanthology.org/2020.emnlp-demos.2
- [20] V. Shwartz, P. West, R. Le Bras, C. Bhagavatula, and Y. Choi, "Unsupervised commonsense question answering with self-talk," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 4615–4629. [Online]. Available: https://aclanthology.org/2020.emnlp-main.373
- [21] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, and W.-t. Yih, "Dense passage retrieval for opendomain question answering," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 6769–6781. [Online]. Available: https://aclanthology.org/2020.emnlp-main.550
- [22] Y. Yao, S. Huang, L. Dong, F. Wei, H. Chen, and N. Zhang, "Kformer: Knowledge injection in transformer feed-forward layers," in *Natural Language Processing and Chinese Computing*, W. Lu, S. Huang, Y. Hong, and X. Zhou, Eds. Cham: Springer International Publishing, 2022, pp. 131–143.
- [23] D. Küçük and F. Can, "Stance detection: A survey," ACM Comput. Surv., vol. 53, no. 1, feb 2020. [Online]. Available: https://doi.org/10.1145/3369026
- [24] A. ALDayel and W. Magdy, "Stance detection on social media: State of the art and trends," *Information Processing and Management*, vol. 58, no. 4, p. 102597, 2021. [Online]. Available: https: //www.sciencedirect.com/science/article/pii/S0306457321000960
- [25] M. Hardalov, A. Arora, P. Nakov, and I. Augenstein, "A survey on stance detection for mis- and disinformation identification," in *Findings of the Association for Computational Linguistics:* NAACL 2022. Seattle, United States: Association for Computational Linguistics, Jul. 2022, pp. 1259–1277. [Online]. Available: https: //aclanthology.org/2022.findings-naacl.94
- [26] M. Mohtarami, J. Glass, and P. Nakov, "Contrastive language adaptation for cross-lingual stance detection," 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)*. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 4442–4452. [Online]. Available: https://aclanthology.org/D19-1452

- [27] E. Zotova, R. Agerri, M. Nuñez, and G. Rigau, "Multilingual stance detection in tweets: The Catalonia independence corpus," in *Proceedings* of the Twelfth Language Resources and Evaluation Conference. Marseille, France: European Language Resources Association, May 2020, pp. 1368–1375. [Online]. Available: https://aclanthology.org/ 2020.lrec-1.171
- [28] Y. Luo, Z. Liu, Y. Shi, S. Z. Li, and Y. Zhang, "Exploiting sentiment and common sense for zero-shot stance detection," in *Proceedings of the 29th International Conference on Computational Linguistics*. Gyeongju, Republic of Korea: International Committee on Computational Linguistics, Oct. 2022, pp. 7112–7123. [Online]. Available: https://aclanthology.org/2022.coling-1.621
- [29] R. Liu, Z. Lin, H. Ji, J. Li, P. Fu, and W. Wang, "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*. Gyeongju, Republic of Korea: International Committee on Computational Linguistics, Oct. 2022, pp. 6944–6954. [Online]. Available: https: //aclanthology.org/2022.coling-1.605
- [30] B. Liang, Q. Zhu, X. Li, M. Yang, L. Gui, Y. He, and R. Xu, "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)*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 81–91. [Online]. Available: https://aclanthology.org/2022.acl-long.7
- [31] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," 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). Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3982–3992. [Online]. Available: https: //aclanthology.org/D19-1410
- [32] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *International Conference on Learning Representations*, 2022. [Online]. Available: https://openreview.net/forum?id=nZeVKeeFYf9
- [33] Y. Li, T. Sosea, A. Sawant, A. J. Nair, D. Inkpen, and C. Caragea, "P-stance: A large dataset for stance detection in political domain," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021.* Online: Association for Computational Linguistics, Aug. 2021, pp. 2355–2365. [Online]. Available: https://aclanthology. org/2021.findings-acl.208
- [34] K. Glandt, S. Khanal, Y. Li, D. Caragea, and C. Caragea, "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).* Online: Association for Computational Linguistics, Aug. 2021, pp. 1596–1611. [Online]. Available: https://aclanthology.org/2021.acl-long.127
- [35] J. Du, R. Xu, Y. He, and L. Gui, "Stance classification with target-specific neural attention," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2017, pp. 3988–3994. [Online]. Available: https://doi.org/10.24963/ ijcai.2017/557
- [36] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "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). Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186. [Online]. Available: https://aclanthology.org/N19-1423
- [37] I. Augenstein, T. Rocktäschel, A. Vlachos, and K. Bontcheva, "Stance detection with bidirectional conditional encoding," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, Nov. 2016, pp. 876–885. [Online]. Available: https://aclanthology.org/ D16-1084
- [38] W. Xue and T. Li, "Aspect based sentiment analysis with gated convolutional networks," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 2514–2523. [Online]. Available: https://aclanthology.org/ P18-1234
- [39] B. Huang and K. Carley, "Parameterized convolutional neural networks for aspect level sentiment classification," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing.* Brussels, Belgium: Association for Computational Linguistics, Oct.-

Nov. 2018, pp. 1091–1096. [Online]. Available: https://aclanthology. org/D18-1136

- [40] B. Zhang, M. Yang, X. Li, Y. Ye, X. Xu, and K. Dai, "Enhancing cross-target stance detection with transferable semanticemotion knowledge," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, Jul. 2020, pp. 3188–3197. [Online]. Available: https://aclanthology.org/2020.acl-main.291
- [41] B. Zhang, X. Fu, D. Ding, H. Huang, Y. Li, and L. Jing, "Investigating chain-of-thought with chatgpt for stance detection on social media," *arXiv preprint arXiv:2304.03087*, 2023.
- [42] J. Wei, D. Huang, Y. Lu, D. Zhou, and Q. V. Le, "Simple synthetic data reduces sycophancy in large language models," *arXiv preprint* arXiv:2308.03958, 2023.