

Structural Bias for Aspect Sentiment Triplet Extraction

Chen Zhang¹, Lei Ren², Fang Ma¹, Jingang Wang^{1,2}, Wei Wu², Dawei Song¹

¹Beijing Institute of Technology
{czhang,mfang,dwsong}@bit.edu.cn

²Meituan NLP
{wangjingang02,wuwei30}@meituan.com
renlei_work@163.com

Abstract

Structural bias has recently been exploited for aspect sentiment triplet extraction (ASTE) and led to improved performance. On the other hand, it is recognized that explicitly incorporating structural bias would have a negative impact on efficiency, whereas pretrained language models (PLMs) can already capture implicit structures. Thus, a natural question arises: Is structural bias still a necessity in the context of PLMs? To answer the question, we propose to address the efficiency issues by using an adapter to integrate structural bias in the PLM and using a cheap-to-compute relative position structure in place of the syntactic dependency structure. Benchmarking evaluation is conducted on the SemEval datasets. The results show that our proposed structural adapter is beneficial to PLMs and achieves state-of-the-art performance over a range of strong baselines, yet with a light parameter demand and low latency. Meanwhile, we give rise to the concern that the current evaluation default with data of small scale is under-confident. Consequently, we release a large-scale dataset for ASTE. The results on the new dataset hint that the structural adapter is confidently effective and efficient to a large scale. Overall, we draw the conclusion that structural bias shall still be a necessity even with PLMs.¹

1 Introduction

Aspect sentiment triplet extraction (ASTE) is a task central to fine-grained opinion mining. Compared to aspect sentiment classification that only aims to predict sentiment polarities for various aspects, ASTE instead extracts descriptive opinion units in the form of triplets (i.e., aspect-opinion-sentiment tuples). For example, (food, great, POS) and (service, dreadful, NEG) are aspect sentiment

¹Jingang Wang and Dawei Song are the corresponding authors.

²Code and data are available at <https://github.com/GenZC/StructBias>.

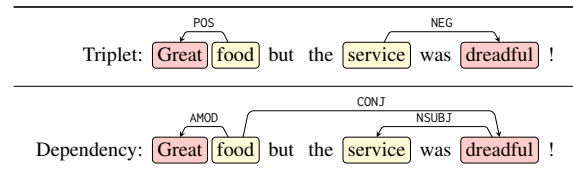


Figure 1: Triplet (top) and dependency structures (bottom) of an illustrative sentence. Spans shaded in yellow are aspects, spans shaded in red are opinions, and arcs indicate either sentiment or structural relations. Irrelevant structural relations are neglected for brevity.

triplets for the sentence in Figure 1 (top), where {POS, NEG, NEU} respectively represent {positive, negative, neutral}.

While ASTE can be generally tackled with neural models in either a pipeline manner (Peng et al., 2020) or a multi-task manner (Xu et al., 2020; Zhang et al., 2020; Wu et al., 2020a; Chen et al., 2021a; Xu et al., 2021), the aspect sentiment triplets can be rather derivable from dependency structures (e.g., syntactic dependency trees) with hand-crafted rules (Wu et al., 2009; Sun et al., 2017). For the example in Figure 1 (bottom), the triplets can be recognized via certain structural dependency relations.² Various studies are motivated by this intuition and exploit dependency bias to enhance neural ASTE models (Chen et al., 2021b), yet without a necessary comparison with pretrained language models (PLMs). On the other hand, recent advances find that using PLMs can already achieve compelling performance (Yan et al., 2021; Zhang et al., 2021; Huang et al., 2021) owing to implicit structures captured by PLMs (Wu et al., 2020b). It signals that, compared with PLMs, explicit structural biases such as dependency bias, may become cumbersome (Dai et al., 2021) due to *parameter inefficiency* and *latency inefficiency*. That is, combining dependencies into models can

²Please see https://downloads.cs.stanford.edu/nlp/software/dependencies_manual.pdf for what these structural relations exactly stand for.

require redundant parameters to achieve structure encoding, while producing structures can also require increased latency to achieve external parsing. Therefore, a critical question naturally arises: Is structural bias still a *necessity* for ASTE in the context of PLMs?

In this paper, we aspire to answer the question from two perspectives: 1) whether structural bias can be incorporated into PLMs in a flexible way in terms of both parameter and latency efficiency; and 2) whether structural bias can enhance PLMs for ASTE.

To boost the parameter efficiency, we develop the idea of *adapter* and put forward a parameter-efficient adapter that can incorporate structural bias. The adapter (Houlsby et al., 2019) was proposed initially to integrate additional modules into PLMs and enable PLMs to leverage inductive bias efficiently. Although feasible, such adapters can be far from lightweight. For example, Liu et al. (2021) introduces a series of linear transformations in their proposed adapter which involve numerous parameters. In contrast, instead of introducing carefully-designed plugins, we propose to use structured attention maps induced with structures, to additively impact the raw attention maps in self-attention, thus requiring only a tiny amount of incremental parameters.

To improve the latency efficiency, we argue that dependency distance is a sufficient simplification of the dependency graph since the simplification has been proven equally powerful in downstream tasks like aspect sentiment classification (Zhang et al., 2019b). On this basis, we further propose to use *relative distance* as an alternative to the dependency distance. The intuition lies in the observation that opinions predominantly locate closely to their corresponding aspects (Xu et al., 2020; Ma et al., 2021), and thus we posit that using relative distance bias would suffice for the purpose of ASTE. In fact, the relative distance is also exhibited to bring merits to the transformer architectures in previous work (Shaw et al., 2018; Raffel et al., 2020). As the relative distance can be obtained with cheap operations in lower latency, the latency efficiency issue is thereby resolved.

We conduct a benchmarking comparative study on the SemEval datasets (Pontiki et al., 2014). The results show that models with the proposed structural adapter achieve the state-of-the-art (SOTA) performance compared with an array of strong

baselines, indicating that incorporating structural bias is beneficial to PLMs. We also conduct a further study on how the relative distance-derived structural adapter overwhelms its alternatives. The results demonstrate that the structural adapter is an appealing choice. Specifically, our structural adapter realizes a $1,000\times$ scale-down in terms of incremental parameters and a $1,000\times$ speed-up of distance derivation.

In summary, structural bias can be flexibly incorporated into PLMs and improve both parameter and latency efficiency. The structural adapter is imposed with only a light parameter demand. The relative distance can be implemented with low latency. Moreover, the structural adapter vastly improves the SOTA performance. Therefore, structural bias is still be a necessity even in the context of PLMs to achieve a better ASTE performance.

Last but not least, in the view that current benchmarks are of small scales, we create a large-scale ASTE dataset termed `LoSTed`. `LoSTed` is collected from one of the largest review platform in China, namely DianPing.³ The dataset will be released to facilitate a more confident evaluation for ASTE and other possible research directions. The results on `LoSTed` hint that structural adapter confidently improves the performance. Furthermore, compared with the results on the SemEval datasets, the model performance generally tends to be lower on this dataset, suggesting that the large-scale deployment of ASTE systems is still challenging.

2 Methodology

2.1 Task Formulation

Given a sequence of tokens $\{t_i\}_{i=1}^n$ as input, ASTE requires a model to output a set of triplets $\{(a, o, s)_i\}_{i=1}^m$, where a , o , s are the aspect, opinion, and sentiment, respectively. Concretely, an aspect a can be decomposed to two elements, i.e. (a_0, a_1) , that separately denote the start and end positions. Likewise, an opinion o can be decomposed similarly.

2.2 PLM with Structural Adapter

When a PLM is employed as the backbone, the tokens are first transformed to embeddings, and then manipulated by subsequent transformer blocks. While the PLM can capture semantic interactions, we additionally present how to include structural interactions with a structural adapter.

³Please see <https://www.dianping.com/>.

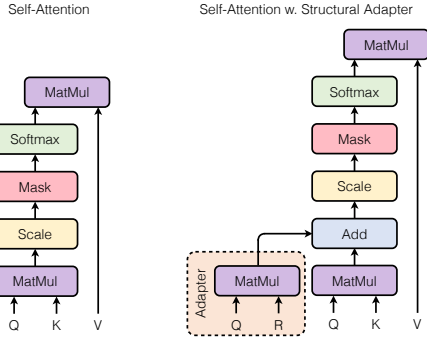


Figure 2: Differences between self-attention (left) and self-attention with the structural adapter applied (right). Q, K, and V separately stand for linear-transformed queries, keys, and values. R stands for distances.

2.2.1 Embedding

The tokens are generally augmented and encoded with the PLM. For example, if the PLM being used is a BERT (Devlin et al., 2019), the tokens should be augmented as:

$$[\text{CLS}] \quad t_1 \quad \dots \quad t_i \quad \dots \quad t_n \quad [\text{SEP}]$$

After that, the augmented tokens are converted to embeddings $\{\mathbf{t}_i\}_{i=0}^{n+1}$.

2.2.2 Transformer Block

The input embeddings are operated by the succeeding transformer blocks (Vaswani et al., 2017), each of which consists of a self-attention module and a feed-forward network module. The self-attention module is typically organized in a query-key-value formulation. Specifically, for any input $\{\mathbf{x}_i\}_{i=1}^n$, the output can be roughly written as:

$$\mathbf{z}_i = \sum_{j=1}^n \alpha_{ij} (\mathbf{x}_j \mathbf{W}_V) \quad \alpha_{ij} = \text{softmax}_j(e_{ij}) \quad (1)$$

$$e_{ij} = \frac{\mathbf{x}_i \mathbf{W}_Q (\mathbf{x}_j \mathbf{W}_K)^\top}{\sqrt{d}}$$

Here, we omit special tokens and multiple heads for simplicity. The parameters \mathbf{W}_Q , \mathbf{W}_K and \mathbf{W}_V are learnable linear transformations for the query, key, and value. d is the head dimensionality.

2.2.3 Structural Adapter

In order to integrate the dependency or relative distance into the self-attention, the structural adapter is imposed to derive structured attention maps to bias the raw attention maps induced with the self-attention additively. The procedure is depicted as

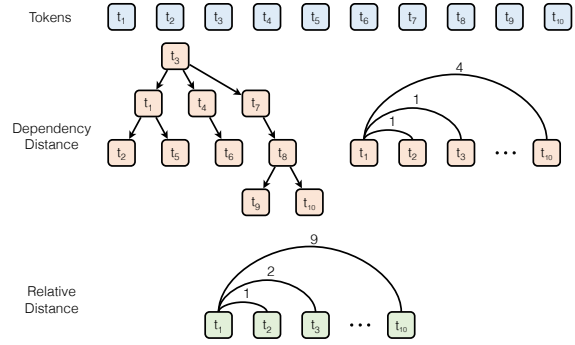


Figure 3: Derivation procedures of dependency distance (middle) and relative distance (bottom) given a sequence of tokens (top). Large distance values like 9 in this example may be clipped by the threshold τ . And distances will be made directional when employed, e.g., distance from t_1 to t_2 is set to 1 while that from t_2 to t_1 is actually set to -1.

below:

$$e_{ij} = \frac{\mathbf{x}_i \mathbf{W}_Q (\mathbf{x}_j \mathbf{W}_K + \mathbb{r}_{ij})^\top}{\sqrt{d}}$$

$$= \underbrace{\frac{\mathbf{x}_i \mathbf{W}_Q (\mathbf{x}_j \mathbf{W}_K)^\top}{\sqrt{d}}}_{\text{raw attention map}} + \underbrace{\frac{\mathbf{x}_i \mathbf{W}_Q \mathbb{r}_{ij}^\top}{\sqrt{d}}}_{\text{structured attention map}} \quad (2)$$

where \mathbb{r}_{ij} indicates the distance embedding between two tokens t_i and t_j . It is also noteworthy that each relation embedding is shared across different heads, but kept independent from one layer to another layer. This behavior is inspired by Shaw et al. (2018), which is originally proposed to encode the relative positions but found to be applicable to encode arbitrary relations (Wang et al., 2020). The differences between self-attention and self-attention with the structural adapter applied are shown in Figure 2.

2.2.4 Distance Derivation

Specifically, the dependency distance between two tokens is obtained by computing the shortest distance on the dependency graph with the networkx toolkit,⁴ and the dependency graph is produced with an off-the-shelf dependency parser stanza (Qi et al., 2020).⁵ The relative distance between two tokens can be yielded by enumerating the number of tokens lying in-between.

⁴Please see <https://networkx.org/> for more information.

⁵Please see <https://stanfordnlp.github.io/stanza/> for more information.

We follow the de facto implementation that treats the distance from t_i to t_j different from that from t_j to t_i . We assign one to positive and the other to negative (Raffel et al., 2020). We also manually set a distance threshold that denotes the maximum distance τ . In doing so, we intend to avoid introducing too many parameters while maintaining as much information as possible. Henceforth, we will refer to the structural adapter with the dependency distance and relative distance as `STRUCTAPT-DEP` and `STRUCTAPT-REL`. The derivation of both dependency distance and relative distance is illustrated in Figure 3.

2.3 Triplet Parser

Learning from two multi-task triplet parsing architectures MTL (Zhang et al., 2020) and GTS (Wu et al., 2020a), we establish a triplet parser that comprises two independent taggers (i.e., one for aspect and the other for opinion tagging), a sentiment scorer, and a triplet decoder. Conceptually, the two taggers are used to uncover continuous tokens that form an aspect or opinion span. The sentiment scorer is used to determine the token-level sentiment relation (if there is one) between two candidate tokens. Moreover, the triplet decoder produces triplets by gathering the information from the taggers and the sentiment scorer.

2.3.1 Aspect and Opinion Taggers

Following MTL, the taggers generate aspect and opinion tags in $\{B, I, O\}$ format, after which the aspect and opinion spans are inferred with $\{B, I, O\}$ tags. Presuming the hidden states of the PLM are $\{\mathbf{h}_i\}_{i=1}^n$ in spite of augmented tokens, the taggers are depicted as:

$$\begin{aligned} \mathbf{r}_i^{(a)} &= \text{ReLU}(\mathbf{W}_1^{(a)}\mathbf{h}_i + \mathbf{b}_1^{(a)}) \\ \mathbf{y}_i^{(a)} &= \text{softmax}(\mathbf{W}_2^{(a)}\mathbf{r}_i^{(a)} + \mathbf{b}_2^{(a)}) \\ \mathbf{r}_i^{(o)} &= \text{ReLU}(\mathbf{W}_1^{(o)}\mathbf{h}_i + \mathbf{b}_1^{(o)}) \\ \mathbf{y}_i^{(o)} &= \text{softmax}(\mathbf{W}_2^{(o)}\mathbf{r}_i^{(o)} + \mathbf{b}_2^{(o)}) \end{aligned} \quad (3)$$

where $\{\mathbf{W}_1^{(a)}, \mathbf{b}_1^{(a)}, \mathbf{W}_2^{(a)}, \mathbf{b}_2^{(a)}\}, \{\mathbf{W}_1^{(o)}, \mathbf{b}_1^{(o)}, \mathbf{W}_2^{(o)}, \mathbf{b}_2^{(o)}\}$ are two sets of weights and biases for two feed-forward networks customized to aspect and opinion tagging.

2.3.2 Sentiment Scorer

The sentiment scorer produces token-level sentiment relations among all tokens. In addition to $\{\text{POS}, \text{NEG}, \text{NEU}\}$, there is also a `NONE` relation to

account for the case of no relation. Unlike the sentiment scorer in MTL and GTS that only predicts uni-directional sentiment relations, we present a sentiment scorer that predicts bi-directional sentiment relations. The uni-directional relation means: a sentiment relation between an aspect token and an opinion token is always directed from the aspect token to the opinion token. In contrast, the bi-directional means: a sentiment relation is both directed from the aspect token to the opinion token and directed from the opinion token to the aspect token. This behavior allows more information to be transduced to the subsequent triplet decoding process to alleviate potential errors. Similarly, the sentiment scorer can be described as:

$$\begin{aligned} \mathbf{r}_i^{(h)} &= \text{ReLU}(\mathbf{W}_1^{(h)}\mathbf{h}_i + \mathbf{b}_1^{(h)}) \\ \mathbf{r}_i^{(d)} &= \text{ReLU}(\mathbf{W}_1^{(d)}\mathbf{h}_i + \mathbf{b}_1^{(d)}) \\ \mathbf{y}_{i,j}^{(s)} &= \text{softmax}(\mathbf{r}_i^{(h)\top}\mathbf{W}_2^{(s)}\mathbf{r}_j^{(d)} + \\ &\quad \mathbf{W}_2^{(h)}\mathbf{r}_i^{(h)} + \mathbf{W}_2^{(d)}\mathbf{r}_j^{(d)} + \mathbf{b}_2^{(s)}) \end{aligned} \quad (4)$$

Here, $\{\mathbf{W}_1^{(h)}, \mathbf{b}_1^{(h)}\}, \{\mathbf{W}_1^{(d)}, \mathbf{b}_1^{(d)}\}$ are weights and biases separately for two feed-forward networks yielding head and dependent representations. These head and dependent representations are then organized in a biaffine manner (Dozat and Manning, 2017), where $\{\mathbf{W}_2^{(h)}, \mathbf{b}_2^{(h)}, \mathbf{W}_2^{(d)}, \mathbf{b}_2^{(d)}\}$ are weights and biases. The biaffine module predicts both aspect-to-opinion and opinion-to-aspect relations at the same time, since either aspects or opinions can be heads or dependents interchangeably.

Additionally, $\mathbf{y}^{(s)}$ refers to a sentiment probability map where $\mathbf{y}_{i,j}^{(s)}$ indicates a probability over 4 sentiment relations from the i -th token to the j -th one. If we apply the `argmax` operation on the probability map, we get a sentiment relation map. An example of the sentiment relation map is given in Figure 4 (left).

2.3.3 Triplet Decoder

Span-level sentiment relations are viable by searching for the most frequent sentiment relation in the set of indexed sentiment relations. Assume there are 2 tokens in a predicted target span and 1 token in a predicted opinion span, then we say there are $2*(2*1)=4$ indexed sentiment relations between the two spans with the bi-directional interplay. Finding the most frequent sentiment relation in inclusive sentiment relations produced with the sentiment relation map gives the sentiment relation between

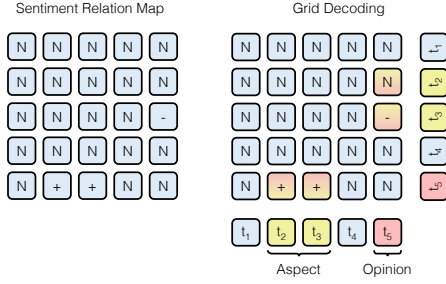


Figure 4: Sentiment relation map (left) and grid decoding algorithm (right). N, +, - respectively are short for NONE, POS, NEG relations. Gradient colors indicate either aspect-to-opinion relations or opinion-to-aspect relations within the bi-directional interplay. The grid decoding algorithm refers to tagged aspects and opinions to index the sentiment relation map.

the two spans. The algorithm is detailed in Figure 4 (right). As the bi-directional interplay is considered, the potential error (i.e., producing a NONE or NEG relation instead of a POS one.) in the example is alleviated.

Since the whole multi-task learning framework is generally borrowed from MTL while the triplet decoding strategy is adapted from the grid decoding in GTS, we name the proposed model as Multi-task learning with Grid decoding (MuG).

2.4 Finetuning

The PLM, the structural adapter, and the triplet parser can be jointly optimized by minimizing an overall objective that contains two sources of losses, i.e., tagging loss \mathcal{L}_t and parsing loss \mathcal{L}_p . Both losses can be measured by the cross-entropy functions. The joint objective is formulated as follows:

$$\min_{\theta} \mathcal{L} = \min_{\theta} \mathcal{L}_t + \mathcal{L}_p \quad (5)$$

where θ stands for all parameters, which might be disassembled to $\{\theta_{\text{PLM}}, \theta_{\text{Adapter}}, \theta_{\text{Parser}}\}$.

3 Benchmarking Evaluation

3.1 Data

We conduct a comparative study on 4 benchmarking datasets from **SemEval** 2014, 2015, and 2016 (Pontiki et al., 2014), in which one contains data from laptop domain (**L**) and the other three contain data from restaurant domain (**R**). The triplet annotations are obtained from Xu et al. (2020). The statistics of these datasets are displayed in Appendix A.

3.2 Models

We compare a wide range of baseline models with varying backbones (e.g., BiLSTM, BERT) and different paradigms (i.e., extractive vs. generative). We list these baselines according to their paradigms as below:

Extractive Paradigm

- KWHW (Peng et al., 2020) is a pipeline system that first extracts aspect-sentiment pairs and opinions, and then pairs them in a binary manner.
- JET^o (Xu et al., 2020) is a position-aware sequence tagging system that jointly extracts triplets.
- MTL (Zhang et al., 2020) is a multi-task learning system which realizes aspect and opinion extraction with tagging while sentiment relation extraction with parsing.
- GTS (Wu et al., 2020a) transforms the triplet extraction problem as a grid tagging problem and achieves the extraction via a grid decoding algorithm.
- Span (Xu et al., 2021) is a span-level triplet extraction system that learns span-level interactions for a more accurate triplet prediction.

Generative Paradigm

- UGF (Yan et al., 2021) is a unified generative system based on BART (Lewis et al., 2020) for all sub-tasks in aspect sentiment analysis.
- GAS (Zhang et al., 2021) likewise is built upon T5 (Raffel et al., 2020) where extractive constraints are applied to the decoding space.

On another note, we have some variants in the comparative study to facilitate the understanding of our adapter. To examine the broad applicability of our structural adapter, we additionally test the structural adapter through the lens of the SOTA extractive systems, namely GTS and Span. To conduct a fair comparison, we initiate the PLM in our model not only with BERT (Devlin et al., 2019) but also with RoBERTa (Liu et al., 2019) to see whether our model with the adapter is competitive with those enhanced by advanced generative PLMs.

3.3 Implementation and Metrics

Typically, the adapter-based finetuning only tunes θ_{Adapter} and θ_{Parser} , and freezes θ_{PLM} . As a randomly initialized adapter can be an unsteady factor

to the PLM, standard finetuning (i.e., tuning all parameters) can result in performance with high variance (Houlsby et al., 2019). However, sub-optimal phenomenon has been observed in the literature (Liu et al., 2021) that such adapter-based finetuning is less promising than standard fine-tuning if the adapter is intended to integrate discrete information (e.g., structural information in our case). Thus, we adapt θ_{Adapter} and θ_{PLM} to the concerned task via finetuning all parameters (i.e., θ_{PLM} , θ_{Adapter} , and θ_{Parser}). Other implementation details are listed in Appendix B.

Following the common practice in the area, we adopt the exact match precision, recall, and F1 scores as the evaluation metrics. Namely, only when the corresponding elements from two triplets exactly match each other, will it be counted as one match. Further, to gain a robust evaluation, we average values over 10 runs and employ the mean value as the final number.

3.4 Performance Analysis

From the results presented in Table 1, we discover two key findings. The first is that the structural adapter incorporated with the relative distance can primarily improve performance across different models and different PLMs, though the improvements over RoBERTa are not as consistent as those over BERT on different datasets. The second is that the previous SOTA models are further boosted by the structural adapter and yield new SOTA results. These findings generally indicate the *effectiveness*, thus *necessity*, of the structural adapter. Conversely, the dependency distance is prone to parsing errors and sometimes underperforms the relative distance.

Moreover, we surprisingly observe that MuG with RoBERTa is a relatively strong baseline even compared with those remarkable generative ASTE models. Concretely, MuG with RoBERTa approximates or outperforms GAS with T5 in terms of F1 scores. This phenomenon encourages some retrospectives on whether generative ASTE models are superior to extractive ones, or the superiority is resulted by the generative PLM.

It can be arguable that the improvements of the structural adapter are marginal; however, we conjecture the inherent reason is that the data for evaluation is of small scale. According to the aforementioned unstable behavior of the adapter when encountering small-scale data in Section 3.3, we think the evaluation is under-confident and there-

fore conduct a large-scale evaluation in Section 4 to verify the guess and to get more confident results.

3.5 Parameter Analysis

We examine the gap between the structural adapter and structural layer (with dependency distance or relative distance, referred to as STRUCTLYR-DEP and STRUCTLYR-REL respectively) from the perspective of incremental parameter scale. The structural layer is exactly a stack of additional transformer layers built upon the PLM, each of which is applied with the structural adapter. The best number of stacked layers is 2 in our pilot study.

The incremental parameters in Table 2 mean that additional parameters are brought to MuG. The structural adapter achieves 1,000 \times scale-down without performance loss compared with the structural layer. Contrarily, it seems that structural layers risk the model on the under-fitting issue due to over-parameterization and get degraded performance compared with MuG. We hereby argue that *parameter efficiency* of the structural adapter is permissible.

3.6 Latency Analysis

To better understand the difference between latency consumed by dependency distance derivation and relative distance derivation. We test the latency caused by the above two derivation procedures.

While dependency distance derivation costs around 4 micro-seconds per token (250 tokens/ms in other words), relative distance derivation only spends 3e-3 micro-seconds per token (333,000 tokens/ms in other words). That is, the relative distance derivation enjoys a 1,000 \times speed-up compared with the dependency distance derivation. Hence, the *latency efficiency* of relative distance derivation is numerically verified.

4 Large-scale Evaluation

4.1 Data

Being aware that the above benchmarking evaluation may be under-confident considering that the data is of small scale, we release a large-scale ASTE dataset, short-named `Losted`. The data is collected from one of the largest review platform in China, namely DianPing. After necessary pre-processing steps, these reviews are manually annotated by 10 proficient assessors. For sanity, double-check on these annotations is carried out by a researcher who has devoted herself to the area for

Model	L14			R14			R15			R16		
	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁
KWHW BiLSTM [*]	37.38	50.38	42.87	43.24	63.66	51.46	48.07	57.51	52.32	46.96	64.24	54.21
JET ^o BiLSTM [*]	53.03	33.89	41.35	61.50	55.13	58.14	64.37	44.33	52.50	70.94	57.00	63.21
MTL BiLSTM [‡]	51.00	40.07	44.81	63.87	54.76	58.90	57.50	42.56	48.73	59.03	54.84	56.73
GTS BiLSTM [‡]	60.32	38.98	47.25	71.08	56.38	62.85	66.60	46.91	55.02	68.75	56.02	61.71
JET ^o BERT [*]	55.39	47.33	51.04	70.56	55.94	62.40	64.45	51.96	57.53	70.42	58.37	63.83
GTS BERT [‡]	57.09	50.33	53.48	69.49	67.75	68.59	61.59	58.21	59.81	65.75	68.32	66.99
w/ STRUCTAPT-REL	57.89	51.57	<u>54.47</u>	68.94	68.26	<u>68.60</u>	62.17	58.63	<u>60.28</u>	66.17	69.79	<u>67.91</u>
Span BERT [‡]	62.57	56.02	59.08	71.77	70.42	71.06	62.06	63.26	<u>62.63</u>	68.57	71.12	69.79
w/ STRUCTAPT-REL	64.72	56.80	60.47	72.53	71.75	72.13	62.80	63.79	63.17	68.94	70.74	69.80
MuG BERT	58.30	52.21	55.06	68.40	67.64	68.00	60.65	54.12	57.10	66.26	67.39	66.74
w/ STRUCTAPT-DEP	59.39	52.95	<u>55.95</u>	67.69	68.90	<u>68.27</u>	60.74	55.77	<u>58.11</u>	64.73	68.33	66.45
w/ STRUCTAPT-REL	59.54	52.56	<u>55.75</u>	68.92	68.12	<u>68.50</u>	59.83	56.78	<u>58.17</u>	65.31	68.83	<u>67.01</u>
UGF BART [†]	61.41	56.19	58.69	65.52	64.99	65.25	59.14	59.38	59.26	66.60	68.68	67.62
GAS T5 [†]	–	–	60.78	–	–	72.16	–	–	62.10	–	–	70.10
MuG RoBERTa	64.18	57.03	60.33	70.47	71.88	71.16	63.78	61.88	62.79	68.61	72.20	70.34
w/ STRUCTAPT-DEP	64.18	56.41	60.03	71.62	71.92	<u>71.72</u>	63.96	61.67	62.70	68.85	71.81	70.28
w/ STRUCTAPT-REL	64.12	57.16	<u>60.53</u>	73.26	71.93	<u>71.17</u>	62.86	63.82	63.12	69.15	74.12	70.44

Table 1: Benchmarking evaluation results. The marker ^{*} indicates results of the model are cited from Xu et al. (2020). The marker [†] indicates results of the model are cited from its original paper. The marker [‡] indicates results of the model are reproduced from its released code. Results are replaced by – to indicate they are not available. F1 scores are underlined to indicate they outperform their adapter-ablated counterparts. F1 scores are **boldfaced** to indicate they are the best-performing ones in their areas.

Model	#Params+	L14	R14
MuG BERT	0.00 M	55.06	68.00
w/ STRUCTLYR-DEP	14.17 M	52.52	67.03
w/ STRUCTLYR-REL	14.17 M	51.57	67.04
w/ STRUCTAPT-DEP	0.01 M	55.95	68.27
w/ STRUCTAPT-REL	0.01 M	55.75	68.50

Table 2: Parameter comparison. F1 scores are reported. #Params+ is short for number of incremental parameters.

years.

For clarity, the pre-processing steps include: 1) removing user identities for privacy consideration; 2) chunking the reviews to shorter examples as they are generally too long (e.g., longer than 512); 3) tokenizing these examples; 4) removing examples without annotations, with less than 4 tokens, or with more than 128 tokens; 5) removing triplets in an example if the triplet has more than 8 tokens in the aspect or has more than 16 tokens in the target, for a too long aspect or opinion indicates the triplet may be not well annotated. Ultimately, these examples are formatted in the format we mentioned in Section 2.1.

We attain the dataset with a total of 27,835 examples. We uniformly split it into train, development,

Dataset	#S	#T	#T/S	#Tk/S	
SemEval R14	train	1266	2336	1.85	17.31
	dev	310	577	1.86	15.81
	test	492	994	2.02	16.34
L _{osted}	train	19485	38050	1.95	34.94
	dev	2783	5334	1.92	34.88
	test	5567	10820	1.94	35.04

Table 3: Statistics of L_{osted}, with a comparison with SemEval R14. #S denotes number of sentences, #T denotes number of triplets, #T/S denotes average number of triplets per sentence, and #Tk/S denotes average number of tokens per sentence.

and test sets with a ratio of 7: 1: 2. The statistics are shown in Table 3, where we also include SemEval R14 for comparison purpose. From the statistics, we can summarize that L_{osted} is a much larger dataset with longer sentences, which sets a more challenging benchmark for models to achieve a high performance.

4.2 Models

We conduct experiments based on GTS and MuG. While we only test GTS with BERT-base, we further test MuG with BERT-base, RoBERTa-base, and tentatively with RoBERTa-large. As we

Model	L _{osted}		
	P	R	F ₁
GTS BERT-base	43.81	46.11	44.92
w/ STRUCTAPT-REL	45.38	46.22	<u>45.79</u>
MuG BERT-base	47.20	45.28	46.22
w/ STRUCTAPT-REL	49.64	45.02	<u>47.22</u>
MuG RoBERTa-base	48.10	44.98	46.49
w/ STRUCTAPT-REL	50.40	44.77	<u>47.42</u>
MuG RoBERTa-large	49.49	46.85	48.13
w/ STRUCTAPT-REL	48.33	47.91	48.13

Table 4: Large-scale evaluation results on L_{osted}. F1 scores are underlined to indicate they outperform their adapter-ablated counterparts.

know that only BERT-base is officially released by Devlin et al. (2019) for Chinese, we retrieve RoBERTa-base and RoBERTa-large released by Cui et al. (2019) on Hugging Face.⁶

4.3 Implementation and Metrics

The implementation and metrics strictly follow those used in the benchmarking evaluation, with exceptions listed in Appendix B.

4.4 Analysis

We can see from Table 4 that the adapter is still promising under large-scale evaluation. With the notice that the evaluation results should be more confident, we hence can safely conclude that the structural adapter is effective and structural bias is *a necessity* for ASTE even in the context of PLMs. However, the metrics on L_{osted} are consistently lower than expected, implying the deployment of ASTE systems is still challenging. The structural adapter does not improve RoBERTa-large, we leave the question of how to combine it with large PLMs for future work.

5 Related Work

5.1 Aspect Sentiment Triplet Extraction

Aspect sentiment triplet extraction is a recently proposed task to extract aspects, opinions, and sentiment relations (Peng et al., 2020), serving as a complete solution to aspect sentiment analysis (Zhang et al., 2019a; Ma et al., 2022). While the first-ever work delving into the task takes a pipeline system, succeeding work shifts their attention from pipeline

⁶Please see <https://huggingface.co/hfl/chinese-roberta-wwm-ext> for more information.

models to joint models. Zhang et al. (2020) and Wu et al. (2020a) share similar spirits to treat three sub-tasks in a multi-task manner. Specifically, Wu et al. (2020a) proposes to consider the extraction of three elements in a unified grid tagging scheme. Later studies exploit inductive biases such as span-level interactions (Xu et al., 2021) and structural bias (Chen et al., 2021b). To our surprise, none of them inspects whether inductive biases, particularly structural bias that is shown beneficial in other PLM-enhanced tasks (Wang et al., 2022), are significant for PLM-enhanced ASTE models. Our work seeks to answer this question by putting forward a flexible adapter and checking whether the adapter is a necessity.

5.2 Adapter for PLM

An adapter is an emergent concept which means an efficient module injected into the PLM so that the PLM can better adapt to downstream tasks (Houlsby et al., 2019). Applications including speed translation (Le et al., 2021), language transfer (He et al., 2021), etc. have been witnessed. Traditionally, parameters of the PLM should not be tuned during fine-tuning once the adapter is armed. Nevertheless, recent work (Liu et al., 2021) finds that when injecting discrete information, unfreezing the parameters of the PLM will bring further performance gain. While previous adapters are modules and thus far from truly lightweight, we propose to leverage the structured attention as a sort of adapter, which is lightweight.

6 Conclusion

In this paper, we are concerned about the parameter and latency inefficiency issues of incorporating structural bias to PLMs for aspect sentiment triplet extraction, and raise the question on whether structural bias is a necessity. To answer the question, we propose to use an adapter to integrate the relative position structure into PLMs for a light parameter demand compared with incremental layers and low latency compared with the syntactic dependency structure. We carry out benchmarking experiments on SemEval benchmarks and large-scale experiments on our newly released L_{osted} dataset as a supplementary. The results in two rounds of evaluations show that the structural adapter is an appealing choice regarding its effectiveness, parameter efficiency, and latency efficiency, implying the structural bias, in the form of the structural adapter,

is a necessity even with PLMs.

Acknowledgements

This research was supported in part by Natural Science Foundation of Beijing (grant number: 4222036) and Huawei Technologies (grant number: TC20201228005).

References

- Shaowei Chen, Yu Wang, Jie Liu, and Yuelin Wang. 2021a. [Bidirectional machine reading comprehension for aspect sentiment triplet extraction](#). In *AAAI*, pages 12666–12674.
- Zhexue Chen, Hong Huang, Bang Liu, Xuanhua Shi, and Hai Jin. 2021b. [Semantic and syntactic enhanced aspect sentiment triplet extraction](#). In *ACL*, pages 1474–1483.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. [Pre-training with whole word masking for Chinese BERT](#). *arXiv*, abs/1906.08101.
- Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu, and Xipeng Qiu. 2021. [Does syntax matter? A strong baseline for aspect-based sentiment analysis with roberta](#). In *NAACL*, pages 1816–1829.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *NAACL*, pages 4171–4186.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *ICLR*.
- Ruidan He, Linlin Liu, Hai Ye, Qingyu Tan, Bosheng Ding, Liying Cheng, Jia-Wei Low, Lidong Bing, and Luo Si. 2021. [On the effectiveness of adapter-based tuning for pretrained language model adaptation](#). In *ACL*, pages 2208–2222.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *ICML*, pages 2790–2799.
- Lianzhe Huang, Peiyi Wang, Sujian Li, Tianyu Liu, Xiaodong Zhang, Zhicong Cheng, Dawei Yin, and Houfeng Wang. 2021. [First target and opinion then polarity: Enhancing target-opinion correlation for aspect sentiment triplet extraction](#). *arXiv*, abs/2102.08549.
- Hang Le, Juan Miguel Pino, Changhan Wang, Jiatao Gu, Didier Schwab, and Laurent Besacier. 2021. [Lightweight adapter tuning for multilingual speech translation](#). In *ACL*, pages 817–824.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *ACL*, pages 7871–7880.
- Wei Liu, Xiyan Fu, Yue Zhang, and Wenming Xiao. 2021. [Lexicon enhanced Chinese sequence labeling using BERT adapter](#). In *ACL*, pages 5847–5858.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *arXiv*, abs/1907.11692.
- Fang Ma, Chen Zhang, and Dawei Song. 2021. [Exploiting position bias for robust aspect sentiment classification](#). In *ACL*, pages 1352–1358.
- Fang Ma, Chen Zhang, Bo Zhang, and Dawei Song. 2022. [Aspect-specific context modeling for aspect-based sentiment analysis](#). *arXiv*, abs/2207.08099.
- Haiyun Peng, Lu Xu, Lidong Bing, Fei Huang, Wei Lu, and Luo Si. 2020. [Knowing what, how and why: A near complete solution for aspect-based sentiment analysis](#). In *AAAI*, pages 8600–8607.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [Semeval-2014 task 4: Aspect based sentiment analysis](#). In *SemEval*, pages 27–35.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *ACL*, pages 101–108.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *JMLR*, 21:140:1–140:67.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *NAACL*, pages 464–468.
- Changzhi Sun, Yuanbin Wu, Man Lan, Shiliang Sun, and Qi Zhang. 2017. [Large-scale opinion relation extraction with distantly supervised neural network](#). In *EACL*, pages 1033–1043.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NeurIPS*, pages 5998–6008.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020. [RAT-SQL: relation-aware schema encoding and linking for text-to-sql parsers](#). In *ACL*, pages 7567–7578.

- Lihan Wang, Bowen Qin, Binyuan Hui, Bowen Li, Min Yang, Bailin Wang, Binhua Li, Jian Sun, Fei Huang, Luo Si, and Yongbin Li. 2022. [Proton: Probing schema linking information from pre-trained language models for text-to-sql parsing](#). In *KDD*, pages 1889–1898.
- Yuanbin Wu, Qi Zhang, Xuanjing Huang, and Lide Wu. 2009. [Phrase dependency parsing for opinion mining](#). In *EMNLP*, pages 1533–1541.
- Zhen Wu, Chengcan Ying, Fei Zhao, Zhifang Fan, Xinyu Dai, and Rui Xia. 2020a. [Grid tagging scheme for end-to-end fine-grained opinion extraction](#). In *EMNLP*, pages 2576–2585.
- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020b. [Perturbed masking: Parameter-free probing for analyzing and interpreting BERT](#). In *ACL*, pages 4166–4176.
- Lu Xu, Yew Ken Chia, and Lidong Bing. 2021. [Learning span-level interactions for aspect sentiment triplet extraction](#). In *ACL*, pages 4755–4766.
- Lu Xu, Hao Li, Wei Lu, and Lidong Bing. 2020. [Position-aware tagging for aspect sentiment triplet extraction](#). In *EMNLP*, pages 2339–2349.
- Hang Yan, Junqi Dai, Tuo Ji, Xipeng Qiu, and Zheng Zhang. 2021. [A unified generative framework for aspect-based sentiment analysis](#). In *ACL*, pages 2416–2429.
- Chen Zhang, Qiuchi Li, and Dawei Song. 2019a. [Aspect-based sentiment classification with aspect-specific graph convolutional networks](#). In *EMNLP*, pages 4567–4577.
- Chen Zhang, Qiuchi Li, and Dawei Song. 2019b. [Syntax-aware aspect-level sentiment classification with proximity-weighted convolution network](#). In *SIGIR*, pages 1145–1148.
- Chen Zhang, Qiuchi Li, Dawei Song, and Benyou Wang. 2020. [A multi-task learning framework for opinion triplet extraction](#). In *EMNLP*, pages 819–828.
- Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2021. [Towards generative aspect-based sentiment analysis](#). In *ACL*, pages 504–510.

A Data Statistics of SemEval

B Full Implementation Details

Our models are implemented with PyTorch and verified on an Nvidia V100, and they are generally trained with following instructions.

For parameter settings in the benchmarking evaluation, the batch size is 8 for models without the adapter, whereas it is 6 for models with the adapter for stability, and the maximum norm for gradients is 1. The learning rate is set hierarchically,

Dataset	#S	#T	#T/S	#Tk/S	
L14	train	906	1460	1.61	19.15
	dev	219	346	1.58	19.06
	test	328	543	1.66	15.77
R14	train	1266	2336	1.85	17.31
	dev	310	577	1.86	15.81
	test	492	994	2.02	16.34
R15	train	605	1013	1.67	14.80
	dev	148	249	1.68	14.34
	test	322	485	1.51	15.63
R16	train	857	1394	1.63	15.15
	dev	210	339	1.61	14.16
	test	326	514	1.58	14.70

Table 5: Statistics of four datasets from **SemEval**. #S denotes number of sentences, #T denotes number of triplets, #T/S denotes average number of triplets per sentence, and #Tk/S denotes average number of tokens per sentence. L denotes laptop domain while R denotes restaurant domain.

where the learning rate for the PLM and adapter is searched with $\{1e-5, 2e-5, 3e-5, 5e-5\}$ while that for the triplet parser is set 10 times of the former. The training procedure is scheduled as such: the number of maximum training epochs is 20, and the number of patience epochs is 5. Learning rates are warmed up for the first 2 epochs and decayed for the rest epochs. The threshold for the maximum distance τ is 8.

For parameter settings in the large-scale evaluation, the batch size is accordingly doubled, since we have data of a much larger scale.