

Adversarial Self-Attention for Language Understanding

Hongqiu Wu^{1,2}, Ruixue Ding⁴, Hai Zhao^{1,2,*}, Pengjun Xie⁴, Fei Huang⁴, Min Zhang³

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University

²Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University

³School of Computer Science and Technology, Soochow University

⁴Damo Academy, Alibaba Group

wuhongqiu@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn,

{ada.drx,chengchen.xpj,f.huang}@alibaba-inc.com, minzhang@suda.edu.cn

Abstract

Deep neural models (e.g. Transformer) naturally learn spurious features, which create a “shortcut” between the labels and inputs, thus impairing the generalization and robustness. This paper advances the self-attention mechanism to its robust variant for Transformer-based pre-trained language models (e.g. BERT). We propose *Adversarial Self-Attention* mechanism (ASA), which adversarially biases the attentions to effectively suppress the model reliance on features (e.g. specific keywords) and encourage its exploration of broader semantics. We conduct a comprehensive evaluation across a wide range of tasks for both pre-training and fine-tuning stages. For pre-training, ASA unfolds remarkable performance gains compared to naive training for longer steps. For fine-tuning, ASA-empowered models outweigh naive models by a large margin considering both generalization and robustness.

Introduction

The emerging pre-trained language models (PrLMs) like BERT (Devlin et al. 2019) have become the backbone of nowadays natural language processing (NLP) systems. It seems to reach the bottleneck in the recent NLP community. This paper rethinks the dilemma from the perspective of self-attention mechanism (SA) (Vaswani et al. 2017), which is broadly chosen as the fundamental architecture of PrLMs and proposes *Adversarial Self-Attention* (ASA)¹.

A large body of empirical evidence (Shi et al. 2021; You, Sun, and Iyyer 2020; Zhang et al. 2020; Wu, Zhao, and Zhang 2021a) indicates that self-attention can benefit from allowing bias, where researchers impose certain priorities (e.g. masking, smoothing) on the original attention structure to compel the model to pay attention to those “proper” tokens. We attribute such phenomena to the nature of deep neural models, which lean to exploit the potential correlations between inputs and labels, even spurious features (Sri-

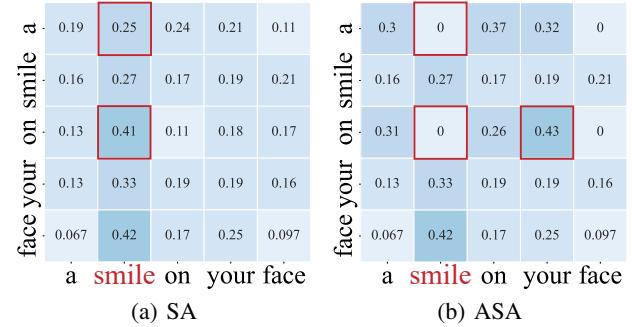


Figure 1: Empirical attention maps from SA and ASA of the first BERT layer. The text is *a smile on your face* from SST-2 dataset. The sum of horizontal scores is equal to 1. We highlight some units receiving strong attention diversions.

vastava, Hashimoto, and Liang 2020). It is harmful for generalizing on test data if the model learns to attend to spurious tokens. Priorly keeping the model from them can hopefully address this. However, crafting priorities is limited to task-specific knowledge and by no means lets the model training be an end-to-end process, whereas generating and storing that knowledge can be even more troublesome.

The idea of ASA is to adversarially bias the self-attention to effectively suppress the model reliance on specific features (e.g. keywords). The biased structures can be learned by maximizing the empirical training risk, which automates the process of crafting specific prior knowledge. Additionally, those biased structures are derived from the input data itself, which sets ASA apart from conventional adversarial training. It is a kind of Meta-Learning (Thrun and Pratt 1998). The learned structures serve as the “meta-knowledge” that facilitates the self-attention process.

We showcase a concrete example in Figure 1. For the word *on*, it never attends to *smile* after being attacked in ASA but strongly attends to *your*. It can be found that *smile* serves as a keyword within the whole sentence, suggesting a positive emotion. The model can predict the right answer based on this single word even without knowing the others.

*Corresponding author; This paper was partially supported by Key Projects of National Natural Science Foundation of China (U1836222 and 61733011).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://github.com/gingasan/adversarialSA>

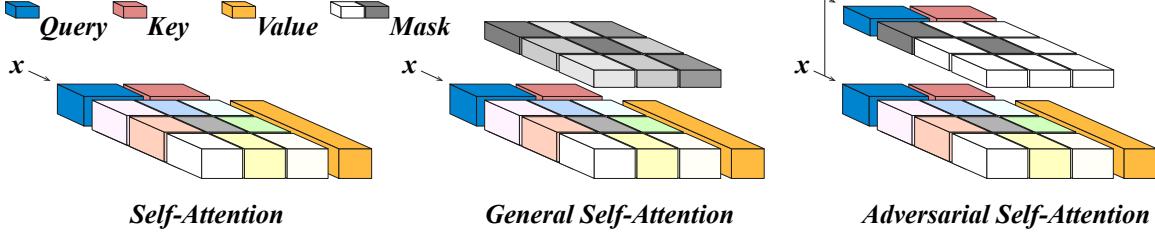


Figure 2: Sketches from self-attention to adversarial self-attention. In self-attention, the attention matrix (the middle colored blocks) is straightforwardly from the query and key components. In general self-attention, the attention matrix is considered to be imposed with another matrix such that its distribution can be biased. In adversarial self-attention, the biasing matrix is learned from the input and is a binary mask (black and white).

Thus, ASA tries to weaken it and let those non-keywords receive more attention. However, in a well-connected structure, the masked linguistic clues can be inferred from their surroundings. ASA prevents the model from shortcut predictions but urges it to learn from contaminated clues, which thus improves the generalization ability.

Another issue of concern is that adversarial training typically results in greater training overhead. In this paper, we design a simple and effective ASA implementation to obtain credible adversarial structures with no training overhead.

This paper is organized as follows. § 2 summarizes the preliminaries of adversarial training and presents a general form of self-attention. § 3 elaborates the methodology of the proposed ASA. § 4 reports the empirical results. § 5 compares ASA with naive masking as well as conventional adversarial training upon the performance and efficiency. § 6 takes a closer look at how ASA works.

Preliminary

In this section, we lay out the background of adversarial training and self-attention mechanism. We begin with a number of notations. Let \mathbf{x} denote the input features, typically a sequence of token indices or embeddings in NLP, and y denote the ground truth. Given a model parameterized by θ , the model prediction can be thus written as $p(y|\mathbf{x}, \theta)$.

Adversarial Training

Adversarial training (AT) (Goodfellow, Shlens, and Szegedy 2015) encourages model robustness by pushing the perturbed model prediction to y :

$$\min_{\theta} \mathcal{D}[y, p(y|\mathbf{x} + \delta^*, \theta)] \quad (1)$$

where $\mathcal{D}[\cdot]$ refers to KL divergence and $p(y|\mathbf{x} + \delta^*, \theta)$ refers to the perturbed model prediction under an adversarial perturbation δ^* . In this paper, we apply the idea of virtual adversarial training (VAT) (Miyato et al. 2019), where y is smoothed by the model prediction $p(y|\mathbf{x}, \theta)$. However, the smoothness can hold only if the amount of training samples is sufficiently large so that $p(y|\mathbf{x}, \theta)$ can be so close to y . This assumption is tenable for PrLMs. Even when fine-tuning on limited samples, the model can not predict so badly in the support of large-scale pre-trained weights.

The adversarial perturbation δ^* is defined to maximize the empirical risk of training:

$$\delta^* = \arg \max_{\delta; \|\delta\| \leq \epsilon} \mathcal{D}[p(y|\mathbf{x}, \theta), p(y|\mathbf{x} + \delta, \theta)] \quad (2)$$

where $\|\delta\| \leq \epsilon$ refers to the decision boundary restricting the adversary δ . We expect that δ^* is so minor but greatly fools the model. Eq. 1 and Eq. 2 make an adversarial “game”, in which the adversary seeks to find out the vulnerability while the model is trained to overcome the malicious attack.

General Self-Attention

Standard self-attention (SA, or vanilla SA) (Vaswani et al. 2017) can be formulated as:

$$\text{SA}(Q, K, V) = \text{Softmax} \left(\frac{Q \cdot K^T}{\sqrt{d}} \right) \cdot V \quad (3)$$

where Q , K , and V refer to the query, key, and value components respectively, and \sqrt{d} is a scaling factor. In this paper, we define the pair-wise matrix $\text{Softmax} \left(\frac{Q \cdot K^T}{\sqrt{d}} \right)$ as the *attention topology* $\mathcal{T}(Q, K)$. In such a topology, each unit refers to the attention score between a token pair, and every single token is allowed to attend to all the others (including itself). The model learns such a topology so that it can focus on particular pieces of text.

However, empirical results show that manually biasing this process (i.e. determining how each token can attend to the others) can lead to better SA convergence and generalization, e.g. enforcing sparsity (Shi et al. 2021), strengthening local correlations (You, Sun, and Iyyer 2020), incorporating structural clues (Wu, Zhao, and Zhang 2021a). Basically, they smooth the output distribution around the attention structure with a certain priority μ . We call μ the *structure bias*. Therefore, it leads us to a more general form of self-attention:

$$\text{SA}(Q, K, V, \mu) = \mathcal{T}(Q, K, \mu) \cdot V \quad (4)$$

where $\mathcal{T}(Q, K, \mu)$ is the biased attention topology $\text{Softmax} \left(\frac{Q \cdot K^T}{\sqrt{d}} + \mu \right)$. In standard SA, μ equals an all-equivalent matrix (all elements are equivalent). That means the attentions on all token pairs are unbiased.

The general form in Eq. 4 indicates that we are able to manipulate the way the attentions unfold between tokens via overlapping a specific structure bias μ . The corresponding sketches are in Figure 2. We focus on the masking case, which is commonly used to mask out the padding positions, where μ refers to a binary matrix with elements in $\{0, -\infty\}$. When an element equals $-\infty$, the attention score of that unit is off ($= 0$). Note that the mask here is different from that in dropout (Srivastava et al. 2014) since the masked units will not be discarded but will be redistributed to other units.

Adversarial Self-Attention Mechanism

This section presents the details of our proposed *Adversarial Self-Attention* mechanism (ASA).

Definition

The idea of ASA is to mask out those attention units to which the model is most vulnerable. Specifically, ASA can be regarded as an instance of general self-attention with an adversarial structure bias μ_x^* . However, those vulnerable units vary from inputs. Thus, we let μ_x^* be a function of \mathbf{x} , denoted as μ_x^* . ASA can be eventually formulated as:

$$\text{ASA}(Q, K, V, \mu_x^*) = \mathcal{T}(Q, K, \mu_x^*) \cdot V \quad (5)$$

where μ_x is parameterized by η namely $\mu_x^* = \mu(\mathbf{x}, \eta^*)$. We also call μ_x the “adversary” in the following. Eq. 5 indicates that μ_x^* acts as “meta-knowledge” learned from the input data itself, which sets ASA apart from other variants where the bias is predefined based on a certain priority.

Optimization

Similar to adversarial training, the model is trained to minimize the following divergence:

$$\min_{\theta} \mathcal{D} [p(y|\mathbf{x}, \theta), p(y|\mathbf{x}, \mu(\mathbf{x}, \eta^*), \theta)] \quad (6)$$

where $p(y|\mathbf{x}, \mu(\mathbf{x}, \eta^*), \theta)$ refers to the model prediction under the adversarial structure bias. We evaluate η^* by maximizing the empirical risk:

$$\eta^* = \arg \max_{\eta: \|\mu(\mathbf{x}, \eta)\| \leq \epsilon} \mathcal{D} [p(y|\mathbf{x}, \theta), p(y|\mathbf{x}, \mu(\mathbf{x}, \eta), \theta)] \quad (7)$$

where $\|\mu(\mathbf{x}, \eta)\| \leq \epsilon$ refers to the new decision boundary for η . The design of this constraint is necessary for keeping ASA from hurting model training.

Generally, researchers use L_2 or L_∞ norm to make the constraint in adversarial training. Considering that $\mu(\mathbf{x}, \eta)$ is in form of a binary mask, it is more reasonable to constrain it by limiting the proportion of the masked units, which comes to L_0 or L_1 norm (since $\mu(\mathbf{x}, \eta)$ is binary, they are the same), namely $\|\mu(\mathbf{x}, \eta)\|_1 \leq \epsilon$. The question is that it is cumbersome to heuristically develop the value of ϵ . As an alternative, we transform the problem with a hard constraint into an unconstrained one with a penalty:

$$\begin{aligned} \eta^* = \arg \max_{\eta} & \mathcal{D} [p(y|\mathbf{x}, \theta), p(y|\mathbf{x}, \mu(\mathbf{x}, \eta), \theta)] \\ & + \tau \|\mu(\mathbf{x}, \eta)\|_1 \end{aligned} \quad (8)$$

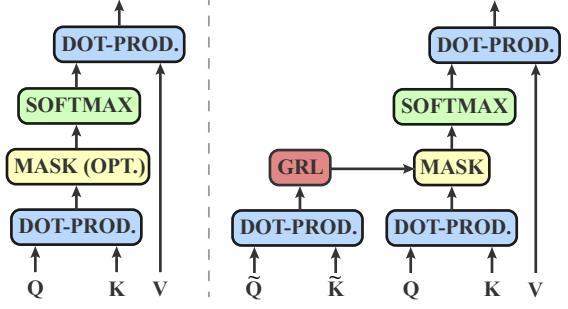


Figure 3: Diagrams of a self-attention (SA) layer (left) and an ASA layer (right). Readers may refer to our supplementary material for the implementation detail.

where we use a temperature coefficient τ to control the intensity of the adversary. Eq. 8 indicates that the adversary needs to maximize the training risk and at the same time mask the least number of units as possible. Our experiments show that it is much easier to adjust τ than to adjust ϵ as in adversarial training. We find good performances when $\tau = 0.1 \sim 0.3$.

Eventually, we generalize Eq. 8 to a model with n self-attention layers (e.g. BERT). Let $\mu(\mathbf{x}, \eta) = \{\mu(\mathbf{x}, \eta)^1, \dots, \mu(\mathbf{x}, \eta)^n\}$, where $\mu(\mathbf{x}, \eta)^i$ refers to the adversary for the i^{th} layer. The penalty term thus becomes the summation $\|\mu(\mathbf{x}, \eta)\|_1 = \sum_{i=1}^n \|\mu(\mathbf{x}, \eta)^i\|_1$.

Fast Implementation

Adversarial training is naturally expensive. To remedy this, we propose a fast and simple implementation of ASA. There are two major points below.

Feature sharing Adversarial training algorithms like K -PGD (Madry et al. 2018) barely avoid multiple inner optimization steps to achieve the high-quality solutions of the adversarial examples (for ASA, they are adversarial structures). Indeed, multiple inner steps could be a disaster for LMs, especially in the pre-training process, which will cost several days or weeks to go through once.

Though there are different ways to implement the ASA adversary $\mu_x = \mu(\mathbf{x}, \eta)$, we are supposed to allow it to obtain a nice solution of μ_x^* in few steps (e.g. only one step). Thus, for the i^{th} self-attention layer, we let the input of $\mu(\mathbf{x}, \eta)^i$ be the input hidden states \mathbf{h}^i of this layer. It does not contradict $\mu(\mathbf{x}, \eta)^i$, since \mathbf{h}^i is encoded from \mathbf{x} , suggesting that the adversary of each layer can access all useful features in the hidden states learned by the model from the lower layers. We apply two linear transformations on \mathbf{h}^i to obtain two components \tilde{Q} and \tilde{K} , and take a dot-product of them to obtain the matrix $\tilde{Q} \cdot \tilde{K}^T / \sqrt{d}$. It is a symmetrical process of computing $Q \cdot K^T / \sqrt{d}$ in vanilla self-attention. The difference is that we will then binarize the matrix using the reparameterization trick (Jang, Gu, and Poole 2017).

Such a design allows us to take only one inner step but obtain a nice μ_x^* . A potential risk is that we can not ensure the performance of ASA in the early training steps since η is randomly initialized. However, it is not naive since we

generally utilize a very small learning rate at the beginning of LM training. The impact of these sacrificed training steps on the model can be negligible.

Gradient reversal Another concern is that adversarial training algorithms always leverage an alternate optimization style, where we temporarily freeze one side of the model and the adversary, and update the other. It requires at least twice backward passes for the inner maximization and outer minimization. To further accelerate training, we adopt Gradient Reversal Layer (GRL) to merge two passes into one. GRL is first introduced in Domain-Adversarial Training (Ganin and Lempitsky 2015), acting as a switch during the backward pass to twist the sign of the gradient from the previous modules. The subsequent modules will be consequently optimized in the opposite direction.

We show a diagram in Figure 3. We see that an ASA layer is composed of five components, where \tilde{Q} and \tilde{K} are competing with Q and K through GRL.

Training

We eventually present the training objective for training an ASA-empowered model. Given a task with labels, we let $\mathcal{L}_e(\theta)$ be the task-specific loss (e.g. classification, regression). The model needs to make the right predictions against the ASA adversary so that we obtain the ASA loss $\mathcal{L}_{asa}(\theta, \eta) = \mathcal{D}[p(y|\mathbf{x}, \theta), p(y|\mathbf{x}, \mu(\mathbf{x}, \eta), \theta)]$. At the same time, the adversary is subject to the penalty term $\mathcal{L}_c(\eta) = \|\mu(\mathbf{x}, \eta)\|_1$. The final training objective thus consists of three components:

$$\mathcal{L}_e(\theta) + \alpha \mathcal{L}_{asa}(\theta, \eta) + \tau \mathcal{L}_c(\eta) \quad (9)$$

where we find ASA performs just well when simply fixing the balancing coefficient α to 1 so that τ is the only hyper-parameter of learning ASA.

Eq. 9 explains how ASA works. One point is that the penalty term is paralleled to the model since it is not associated with θ . We seek to find the optimal model parameters θ to minimize the first two terms. On the other hand, because of GRL, we seek to find the optimal adversary parameters η to maximize the last two terms.

Though Eq. 9 covers all cases when fine-tuning the ASA-empowered model on downstream tasks. We will discuss more on pre-training. ASA is consistent with the current trend of self-supervised language modeling like MLM (Devlin et al. 2019) and RTD (Clark et al. 2020), where we construct the negative samples based on the super large corpus itself without additional labeling.

To be more concrete, we rely on MLM pre-training setting (Devlin et al. 2019) in what follows and the other situations can be easily generalized. MLM intends to recover a number of masked pieces within the sequence and the loss of it is obtained from the cross-entropy on those selected positions, denoted as \mathcal{L}_{mlm} . Thus, we can compute the divergence between the two model predictions before and after biased by ASA on those positions and obtain the token-level ASA loss $\mathcal{L}_{asa}^t(\theta, \eta)$.

Aside from the masked positions, the beginning position is also crucial to PrLMs (e.g. [CLS] in BERT), which is

always used as an indicator for the relationship within the sequence (e.g. sentence order, sentiment). Thus, we pick this position out from the final hidden states and calculate the divergence on it as another part of the ASA loss $\mathcal{L}_{asa}^s(\theta, \eta)$. Finally, pre-training with ASA can be formulated as:

$$\mathcal{L}_{mlm}(\theta) + \mathcal{L}_{asa}^t(\theta, \eta) + \mathcal{L}_{asa}^s(\theta, \eta) + \tau \mathcal{L}_c(\eta) \quad (10)$$

where \mathcal{L}_{asa}^t and \mathcal{L}_{asa}^s refer to the token-level and sentence-level ASA loss respectively.

Based on Eq. 10, we may touch on the idea of ASA pre-training from two perspectives: (a) **Structural loss**: ASA acts as a regularizer on the empirical loss of language modeling (the same in Eq. 9); (b) **Multiple objectives**: ASA can be viewed as two independent self-supervised pre-training objectives in addition to MLM.

Experiments

Our implementations are based on *transformers* (Wolf et al. 2020).

Setup

We experiment on five NLP tasks down to 10 datasets:

- **Sentiment Analysis**: Stanford Sentiment Treebank (SST-2) (Socher et al. 2013), which is a single-sentence binary classification task;
- **Natural Language Inference (NLI)**: Multi-Genre Natural Language Inference (MNLI) (Williams, Nangia, and Bowman 2018) and Question Natural Language Inference (QNLI) (Wang et al. 2019), where we need to predict the relations between two sentences;
- **Semantic Similarity**: Semantic Textual Similarity Benchmark (STS-B) (Cer et al. 2017) and Quora Question Pairs (QQP) (Wang et al. 2019), where we need to predict how similar two sentences are;
- **Named Entity Recognition (NER)**: WNUT-2017 (Aguilar et al. 2017), which contains a large number of rare entities;
- **Machine Reading Comprehension (MRC)**: Dialogue-based Reading Comprehension (DREAM) (Sun et al. 2019), where we need to choose the best answer from the three candidates given a question and a piece of dialogue;
- **Robustness learning**: Adversarial NLI (ANLI) (Nie et al. 2020) for NLI, PAWS-QQP (Zhang, Baldridge, and He 2019) for semantic similarity, and Hell-SWAG (Zellers et al. 2019) for MRC.

We verify the gain of ASA on the top of two different self-attention (SA) designs: vanilla SA in BERT-base (Devlin et al. 2019) and its stronger variant RoBERTa-base (Liu et al. 2019), and disentangled SA in DeBERTa-large (He et al. 2021). In addition, we do the experiments on both pre-training ($\tau = 0.1$, Eq. 10) and fine-tuning ($\tau = 0.3$, Eq. 9) stages (the training details can be found in Appendix). For pre-training, we continue to pre-train BERT based on MLM with ASA on the English Wikipedia corpus. Besides, for fair enough comparison, we train another BERT with vanilla SA (BERT \dagger in Table 1). We set the batch size to 128 and train both models for 20K steps with FP16. Note that we directly fine-tune them without ASA.

Results

Results on generalization Table 1 summarizes the results across various tasks. For fine-tuning, ASA-empowered mod-

Model	Sentiment Analysis / Inference			Semantic Similarity		NER		MRC	Avg
	SST-2 (Acc)	MNLI (Acc)	QNLI (Acc)	QQP (F1)	STS-B (Spc)	WNUT-17 (F1)	DREAM (Acc)		
BERT	93.2 \pm 0.24	84.1 \pm 0.05	90.4 \pm 0.09	71.4 \pm 0.31	84.7 \pm 0.10	47.8 \pm 1.08	62.9 \pm 0.16	76.4	
BERT ^{ASA}	94.1 \star \pm 0.00	85.0 \pm 0.05	91.4 \star \pm 0.22	72.3 \pm 0.05	86.5 \star \pm 0.37	49.8 \star \pm 0.69	64.3 \star \pm 0.41	77.6 \uparrow 1.2	
BERT \dagger	93.5 \pm 0.32	84.4 \pm 0.08	90.5 \pm 0.13	71.5 \pm 0.08	85.4 \pm 0.36	49.2 \pm 0.94	61.2 \pm 0.82	76.5 \uparrow 0.1	
BERT ^{ASA} \dagger	94.0 \pm 0.05	84.7 \pm 0.06	91.5 \star \pm 0.19	72.3 \pm 0.05	86.5 \star \pm 0.23	50.3 \star \pm 0.55	63.3 \star \pm 0.28	77.5 \uparrow 1.1	
RoBERTa	95.6 \pm 0.05	87.2 \pm 0.15	92.8 \pm 0.21	72.2 \pm 0.05	88.4 \pm 0.17	54.8 \pm 0.80	67.0 \pm 0.62	79.7	
RoBERTa ^{ASA}	96.3 \pm 0.19	88.0 \pm 0.05	93.6 \pm 0.22	73.7 \star \pm 0.12	89.2 \pm 0.38	57.3 \star \pm 0.18	69.2 \star \pm 0.68	81.0 \uparrow 1.3	

Table 1: Results on different tasks (mean and variance), where \dagger refers to the longer-trained model with MLM. We run three seeds for GLUE sub-tasks (the first five, since only two test submissions are allowed each day) and five seeds for the others. For MNLI, we average the two scores of the “m” and “mm”. \star indicates the proposed approach unfolds > 1 points absolute gain.

Model	ANLI (Acc)	PAWS-QQP (Acc)	HellaSWAG (Acc)
BERT-base	48.0 \pm .68	81.7 \pm 1.24	39.7 \pm .28
BERT ^{ASA}	50.4 \star \pm .81	87.7 \star \pm 1.53	40.8 \star \pm .27
DeBERTa-large	57.6 \pm .43	95.7 \pm .38	94.3 \pm 1.02
DeBERTa ^{ASA}	58.2 \pm .94	96.0 \pm .24	95.4 \pm 1.31

Table 2: Results on robustness learning tasks when $\tau = 0.3$ over five runs. For ANLI, we put the test data of all rounds together and the model is trained with its own training data without any other data. For HellaSWAG, we report the dev.

els consistently outweigh naive BERT and RoBERTa by a large margin, lifting the average performance of BERT from **76.4** to **77.6**, and RoBERTa from **79.7** to **81.0**. ASA is supposed to perform well on small sets like STS-B (84.7 to 86.5 on BERT) and DREAM (62.9 to 64.3) with merely thousands of training samples, which tend to be more susceptible to over-fitting. However, on much larger ones like MNLI (84.1 to 85.0), QNLI (90.4 to 91.4), and QQP (72.2 to 73.7 on RoBERTa) with more than hundred-thousands of samples, it still produces powerful gain. It implies that ASA not only enhances model generalization but also language representation. For continual pre-training, ASA brings competitive performance gain when directly fine-tuning on downstream tasks.

Results on robustness To assess the impact of ASA on model robustness, we report the fine-tuning results on three challenging robustness benchmarks. These tasks contain a large number of adversarial samples in their training or test sets. From Table 2, we can see that ASA produces 2.4, 6.0, and 1.1 points of absolute gain over BERT-base on the three tasks respectively. Even on strong DeBERTa-large, ASA can still deliver considerable improvement.

	PAWS-QQP	HellaSWAG	WNUT-17
<i>Bernoulli</i>	86.1 \pm 1.2	40.5 \pm 0.2	48.2 \pm 1.0
<i>Scheduled</i>	85.4 \pm 1.5	40.2 \pm 0.2	48.7 \pm 0.9
<i>Magnitude</i>	84.6 \pm 0.9	39.9 \pm 0.3	47.3 \pm 1.1
<i>ASA</i>	87.7 \pm 1.5	40.8 \pm 0.3	49.8 \pm 0.7

Table 3: Naive masking on BERT-base over five runs.

Ablation Study VS. Naive Masking

We compare ASA with three naive masking strategies. Bernoulli distribution is a widely-used priority in network dropout (Srivastava et al. 2014). We report the best results with the masking probability selected in $\{0.05, 0.1\}$. Besides, we introduce another two potentially stronger strategies. For the first one, we dynamically schedule the masking probability for each step following the learned pattern by ASA. Different from ASA, the masked units here are Bernoulli chosen. For the second one, we always choose to mask those units with the most significant attention scores (a magnitude-based strategy). Similar to ASA, we apply the masking matrices to all self-attention layers, and the training objective corresponds to the first two terms of Eq. 9.

From Table 3, we can see that pure Bernoulli works the best among the three naive strategies, slightly better than scheduled Bernoulli. However, ASA outweighs them even by a large margin, suggesting that the worst-case masking can better facilitate model training. Besides, the magnitude-based masking turns out to be harmful. Since ASA acts as a gradient-based adversarial strategy, it may not always mask those globally most significant units in the attention matrix (this pattern can be seen in previous Figure 1).

On the other hand, it has been found in previous work that adversarial training on word embeddings can appear mediocre compared to random perturbations (Aghajanyan et al. 2021). We are positive that adversarial training benefits model training, but hypothesize that the current spatial optimization of embedding perturbations suffers from shortcomings so that it sometimes falls behind random perturbations.

	MNLI	QNLI	PAWS-QQP	H-SWAG
FreeLB	85.3 ± 0.1	91.1 ± 0.0	86.3 ± 1.3	39.6 ± 0.4
SMART	85.5 ± 0.2	91.6 ± 0.5	85.8 ± 0.8	38.2 ± 0.3
ASA	85.0 ± 0.1	91.4 ± 0.2	87.7 ± 1.5	40.8 ± 0.3

Table 4: Comparison with adversarial training on BERT-base over multiple runs (three for GLUE and five for others).

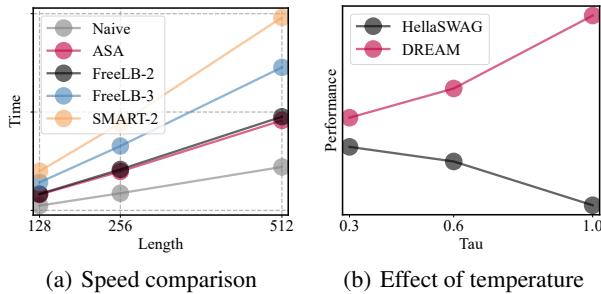


Figure 4: Speed comparison of different learning approaches across sequence lengths (128, 256, and 512). FreeLB- x means that we train the model with x inner steps.

However, the optimization of ASA is carefully designed in our paper.

VS. Adversarial Training

ASA is close to conventional adversarial training, but there are two main differences. In the text domain, adversarial training works on the input space, imposing perturbations on word embeddings, while ASA works on model structures. Besides, adversarial training normally leverages projected gradient descent (PGD) (Madry et al. 2018) to learn the adversary, while ASA is optimized through an unconstrained manner. We compare the performances on different tasks between ASA and FreeLB (Zhu et al. 2020), one of the state-of-the-art adversarial training approaches in the text domain.

From Table 4, we can see that ASA and FreeLB are competitive on MNLI and QNLI, while ASA outperforms by 1.4 and 1.2 points on PAWS-QQP and HellaSWAG. It may leave a new line for future research, where ASA can be superior to conventional adversarial training on certain tasks. Another advantage of ASA is that it only introduces one hyperparameter τ , while for FreeLB, we need to sweep through different adversarial step sizes and boundaries.

On the other hand, we find that both FreeLB and SMART can hardly induce a significant variation in the attention maps, even if the input embeddings are perturbed. The model remains focused on those pieces that are supposed to be focused on before being perturbed. This can be detrimental when one tries to explain the adversary’s behavior.

Training Speed

Figure 4 (a) summarizes the speed performances of ASA and other two state-of-the-art adversarial training algorithms in the text domain. FreeLB (Zhu et al. 2020) is currently

the fastest algorithm, which requires at least two inner steps (FreeLB-2) to complete its learning process. SMART (Jiang et al. 2020) leverages the idea of virtual adversarial training, which thus requires at least one more forward passes. We turn off FP16, fix the batch size to 16, and do the experiments under different sequence lengths. We can see that ASA is slightly faster than FreeLB-2, taking about twice the period of naive training when the sequence length is up to 512. However, training with SMART and FreeLB-3 is much more expensive, taking about three and four times that period of naive training (SMART-1 is very close to FreeLB-3, so we omit it from the figure).

Temperature Coefficient

The only hyperparameter for ASA is the temperature coefficient τ , which controls the intensity of the adversary, a lower τ corresponding to a stronger attack (higher masking proportion). In practice, τ balances the model generalization and robustness. We conduct experiments on a benign task (DREAM) and an adversarial task (HellaSWAG) respectively with τ selected in $\{0.3, 0.6, 1.0\}$. As in Figure 4 (b), the trends of two curves are opposite (we offset them vertically to make them close). A stronger adversary might lead to a decrease in generalization but benefit robustness. For example, we see the peak DREAM result when $\tau = 1.0$, while the peak HellaSWAG result when $\tau = 0.3$.

Masking Proportion

A higher masking proportion in ASA implies that the layer is more vulnerable. As in Figure 5 (a), we observe that the masking proportion almost decreases layer by layer from the bottom to the top. We attribute this to information diffusion (Goyal et al. 2020), which states that the input vectors are progressively assimilating through continuously making self-attention. Consequently, the attention scores in the bottom layers are more important so that more vulnerable, while those in the top layers become less important after their assimilation. The feed-forward layers are more contributing this time (You, Sun, and Iyyer 2020).

In Figure 5 (b), we calculate the average masking proportion of all layers and see that the situations can also be different between tasks even with the same temperature. Take sentiment analysis as an instance, the adversary merely needs to focus on specific keywords, which is enough to lead to misclassification. For NER and MRC, however, there are more sensitive words that are scattered across the sentence, and therefore a greater attack of the adversary is needed.

Why ASA Works

In addition, **ASA is effective in weakening the model reliance on keywords and encourages the model to concentrate on the broader semantics**. We show a concrete example of sentiment analysis in Figure 6. We see that the SA-trained model tends to allow more tokens to receive strong attention (potential keywords), while they are sometimes spurious features. As a result, the model gives a wrong prediction. However, the ASA-trained one locates fewer keywords but more precisely (i.e. *excitement*, *eating*, *oatmeal*)

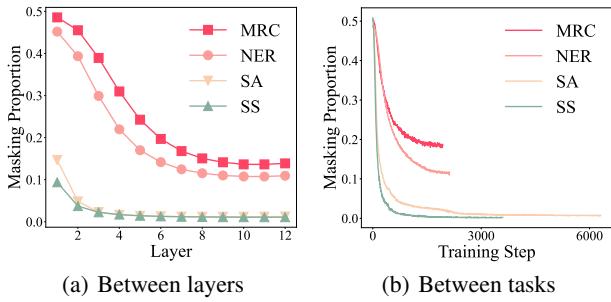


Figure 5: Masking proportion of ASA when $\tau = 0.3$ across different tasks. Note that the magnitude of adversarial perturbations is always minor but they greatly fool the model.

and thus obtains the right answer. Note that punctuation acts as a critical clue that signals the boundaries of semantics. In the top layers, they are normally given a high attention score.

Another observation is that for those examples without explicit keywords (not shown in the paper for space limitation), the SA-trained model prefers to locate many as “keywords”. Contrarily, the ASA-trained one may not locate any keywords, but rather make the predictions based on entire semantics. The above observations are well-consistent with the way ASA trains.

Related Work

Our work is closely related to Adversarial Training (AT) (Goodfellow, Shlens, and Szegedy 2015), which is a common machine learning approach to improve model robustness. In the text domain, the conventional philosophy is to impose adversarial perturbations on word embeddings (Miyato, Dai, and Goodfellow 2017). It is later found to be highly effective in enhancing model performances when applied to fine-tuning on multiple downstream tasks, e.g. FreeLB (Zhu et al. 2020), SMART (Jiang et al. 2020), InfoBERT (Wang et al. 2021), CreAT (Wu et al. 2023), while ALUM (Liu et al. 2020) provides the firsthand empirical results that adversarial training can produce a promising pre-training gain. As opposed to all these counterparts, which choose to perturb the input text or word embeddings, our work perturbs the self-attention structure. In addition, we present a new unconstrained optimization criterion to effectively learn the adversary. Our work is also related to smoothing and regularization techniques (Bishop 1995; Srivastava et al. 2014).

Adversarial training is naturally expensive. There are algorithms for acceleration, e.g. FreeAT (Shafahi et al. 2019), YOPO (Zhang et al. 2019), FreeLB (Zhu et al. 2020). This paper proposes a fast and simple implementation. Its speed performance rivals that of the current fastest adversarial training algorithm FreeLB. Another important line in adversarial training is to rationalize the behaviour of the adversary (Sato et al. 2018). In our work, we demonstrate how adversarial self-attention contributes to improving the model generalization from the perspective of feature utilization.

Our work is similar to optimizing the self-attention architecture (Vaswani et al. 2017), e.g. block-wise attention

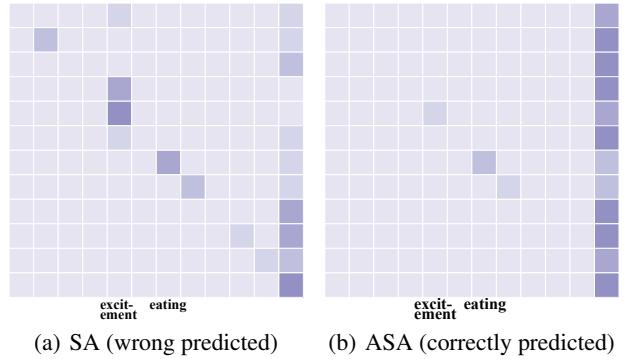


Figure 6: Comparison of attention maps between SA-train and ASA-trained models. The text is *it has all the excitement of eating oatmeal*. from SST-2). We choose the view from the 11th layer and 2nd head of BERT for instance.

(Zaheer et al. 2020), sparse attention (Shi et al. 2021), structure-induced attention (Wu, Zhao, and Zhang 2021a), Gaussian attention (You, Sun, and Iyyer 2020), synthetic attention (Tay et al. 2021), policy-based attention (Wu, Zhao, and Zhang 2021b). Most of these variants are based on a predefined priority. In comparison, our adversary derives from the data distribution itself and exploits the adversarial idea to effectively learn the self-attention structure or how to make self-attention. It is a kind of Meta-Learning (Thrun and Pratt 1998), which aims to effectively optimize the learning process, e.g. learning the update rule for few-shot learning (Ravi and Larochelle 2017), learning an optimized initialization ready for fast adaption to new tasks (Finn, Abbeel, and Levine 2017), reweighting training samples (Ren et al. 2018). Our work facilitates both fine-tuning and pre-training for pre-trained language models (PrLMs) (Devlin et al. 2019; Liu et al. 2019; Raffel et al. 2020; He et al. 2021). It is agnostic to the current pre-training paradigms, e.g. MLM (Devlin et al. 2019), RTD (Clark et al. 2020), PLM (Yang et al. 2019), and multiple objectives (Wu et al. 2022).

Conclusion

This paper presents *Adversarial Self-Attention* mechanism (ASA) to improve pre-trained language models. Our idea is to adversarially bias the Transformer attentions and facilitate model training from contaminated model structures. As it turns out, the model is encouraged to explore more on broader semantics and exploit less on keywords. Empirical experiments on a wide range of natural language processing (NLP) tasks demonstrate that our approach remarkably boosts model performances on both pre-training and fine-tuning stages. We also conduct a visual analysis to interpret how ASA works. However, the analysis in this paper is still limited.

References

Aghajanyan, A.; Shrivastava, A.; Gupta, A.; Goyal, N.; Zettlemoyer, L.; and Gupta, S. 2021. Better Fine-Tuning by Reduc-

ing Representational Collapse. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Aguilar, G.; Maharjan, S.; López-Monroy, A. P.; and Solorio, T. 2017. A Multi-task Approach for Named Entity Recognition in Social Media Data. In Derczynski, L.; Xu, W.; Ritter, A.; and Baldwin, T., eds., *Proceedings of the 3rd Workshop on Noisy User-generated Text, NUT@EMNLP 2017, Copenhagen, Denmark, September 7, 2017*, 148–153. Association for Computational Linguistics.

Bishop, C. M. 1995. Training with Noise is Equivalent to Tikhonov Regularization. *Neural Comput.*, 7(1): 108–116.

Cer, D. M.; Diab, M. T.; Agirre, E.; Lopez-Gazpio, I.; and Specia, L. 2017. SemEval-2017 Task 1. In Bethard, S.; Carpuat, M.; Apidianaki, M.; Mohammad, S. M.; Cer, D. M.; and Jurgens, D., eds., *Proceedings of the 11th International Workshop on Semantic Evaluation, SemEval@ACL 2017, Vancouver, Canada, August 3-4, 2017*, 1–14. Association for Computational Linguistics.

Clark, K.; Luong, M.; Le, Q. V.; and Manning, C. D. 2020. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, 4171–4186. Association for Computational Linguistics.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, 1126–1135. PMLR.

Ganin, Y.; and Lempitsky, V. S. 2015. Unsupervised Domain Adaptation by Backpropagation. In Bach, F. R.; and Blei, D. M., eds., *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, 1180–1189. JMLR.org.

Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Goyal, S.; Choudhury, A. R.; Raje, S.; Chakaravarthy, V. T.; Sabharwal, Y.; and Verma, A. 2020. PoWER-BERT: Accelerating BERT Inference via Progressive Word-vector Elimination. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 3690–3699. PMLR.

He, P.; Liu, X.; Gao, J.; and Chen, W. 2021. Deberta: decoding-Enhanced Bert with Disentangled Attention. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Jiang, H.; He, P.; Chen, W.; Liu, X.; Gao, J.; and Zhao, T. 2020. SMART: Robust and Efficient Fine-Tuning for Pre-trained Natural Language Models through Principled Regularized Optimization. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J. R., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 2177–2190. Association for Computational Linguistics.

Liu, X.; Cheng, H.; He, P.; Chen, W.; Wang, Y.; Poon, H.; and Gao, J. 2020. Adversarial Training for Large Neural Language Models. *CoRR*, abs/2004.08994.

Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR*, abs/1907.11692.

Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.

Miyato, T.; Dai, A. M.; and Goodfellow, I. J. 2017. Adversarial Training Methods for Semi-Supervised Text Classification. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Miyato, T.; Maeda, S.; Koyama, M.; and Ishii, S. 2019. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(8): 1979–1993.

Nie, Y.; Williams, A.; Dinan, E.; Bansal, M.; Weston, J.; and Kiela, D. 2020. Adversarial NLI: A New Benchmark for Natural Language Understanding. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J. R., eds., *Proceedings of the 58th Annual Meeting of ACL, ACL 2020, Online, July 5-10, 2020*, 4885–4901. Association for Computational Linguistics.

Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.*, 21: 140:1–140:67.

Ravi, S.; and Larochelle, H. 2017. Optimization as a Model for Few-Shot Learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Ren, M.; Zeng, W.; Yang, B.; and Urtasun, R. 2018. Learning to Reweight Examples for Robust Deep Learning. In Dy, J. G.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 4331–4340. PMLR.

Sato, M.; Suzuki, J.; Shindo, H.; and Matsumoto, Y. 2018. Interpretable Adversarial Perturbation in Input Embedding Space for Text. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 4323–4330. ijcai.org.

Shafahi, A.; Najibi, M.; Ghiasi, A.; Xu, Z.; Dickerson, J. P.; Studer, C.; Davis, L. S.; Taylor, G.; and Goldstein, T. 2019. Adversarial training for free! In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 3353–3364.

Shi, H.; Gao, J.; Ren, X.; Xu, H.; Liang, X.; Li, Z.; and Kwok, J. T. 2021. SparseBERT: Rethinking the Importance Analysis in Self-attention. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th*

International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event, volume 139 of *Proceedings of Machine Learning Research*, 9547–9557. PMLR.

Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on EMNLP, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, 1631–1642. ACL.

Srivastava, M.; Hashimoto, T. B.; and Liang, P. 2020. Robustness to Spurious Correlations via Human Annotations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 9109–9119. PMLR.

Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1): 1929–1958.

Sun, K.; Yu, D.; Chen, J.; Yu, D.; Choi, Y.; and Cardie, C. 2019. DREAM: A Challenge Dataset and Models for Dialogue-Based Reading Comprehension. *Trans. Assoc. Comput. Linguistics*, 7: 217–231.

Tay, Y.; Bahri, D.; Metzler, D.; Juan, D.; Zhao, Z.; and Zheng, C. 2021. Synthesizer: Rethinking Self-Attention for Transformer Models. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 10183–10192. PMLR.

Thrun, S.; and Pratt, L. Y. 1998. Learning to Learn: Introduction and Overview. In Thrun, S.; and Pratt, L. Y., eds., *Learning to Learn*, 3–17. Springer.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 5998–6008.

Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

Wang, B.; Wang, S.; Cheng, Y.; Gan, Z.; Jia, R.; Li, B.; and Liu, J. 2021. InfoBERT: Improving Robustness of Language Models from An Information Theoretic Perspective. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Williams, A.; Nangia, N.; and Bowman, S. R. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In Walker, M. A.; Ji, H.; and Stent, A., eds., *Proceedings of the 2018 Conference of NAACL-HLT, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, 1112–1122. Association for Computational Linguistics.

Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Scao, T. L.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. M. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45. Online: Association for Computational Linguistics.

Wu, H.; Ding, R.; Zhao, H.; Chen, B.; Xie, P.; Huang, F.; and Zhang, M. 2022. Forging Multiple Training Objectives for Pre-trained Language Models via Meta-Learning. *CoRR*, abs/2210.10293.

Wu, H.; Liu, Y.; Shi, H.; hai zhao; and Zhang, M. 2023. Toward Adversarial Training on Contextualized Language Representation. In *International Conference on Learning Representations*.

Wu, H.; Zhao, H.; and Zhang, M. 2021a. Code Summarization with Structure-induced Transformer. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, 1078–1090. Association for Computational Linguistics.

Wu, H.; Zhao, H.; and Zhang, M. 2021b. Not All Attention Is All You Need. *CoRR*, abs/2104.04692.

Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J. G.; Salakhutdinov, R.; and Le, Q. V. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 5754–5764.

You, W.; Sun, S.; and Iyyer, M. 2020. Hard-Coded Gaussian Attention for Neural Machine Translation. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J. R., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 7689–7700. Association for Computational Linguistics.

Zaheer, M.; Guruganesh, G.; Dubey, K. A.; Ainslie, J.; Alberti, C.; Ontañón, S.; Pham, P.; Ravula, A.; Wang, Q.; Yang, L.; and Ahmed, A. 2020. Big Bird: Transformers for Longer Sequences. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a Machine Really Finish Your Sentence? In Korhonen, A.; Traum, D. R.; and Márquez, L., eds., *Proceedings of the 57th Conference of ACL, ACL 2019, Florence, Italy, July 28-August 2, 2019, Volume 1: Long Papers*, 4791–4800. Association for Computational Linguistics.

Zhang, D.; Zhang, T.; Lu, Y.; Zhu, Z.; and Dong, B. 2019. You Only Propagate Once: Accelerating Adversarial Training via Maximal Principle. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 227–238.

Zhang, Y.; Baldridge, J.; and He, L. 2019. PAWS: Paraphrase Adversaries from Word Scrambling. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of NAACL-HLT, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, 1298–1308. Association for Computational Linguistics.

Zhang, Z.; Wu, Y.; Zhou, J.; Duan, S.; Zhao, H.; and Wang, R. 2020. SG-Net: Syntax-Guided Machine Reading Comprehension. In *The Thirty-Fourth AAAI, AAAI 2020, New York, NY, USA, February 7-12, 2020*, 9636–9643. AAAI Press.

Zhu, C.; Cheng, Y.; Gan, Z.; Sun, S.; Goldstein, T.; and Liu, J. 2020. FreeLB: Enhanced Adversarial Training for Natural Language Understanding. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.