

Rethinking Denoised Auto-Encoding in Language Pre-Training

Fuli Luo^{1,*}, Pengcheng Yang^{1,2,*}, Shicheng Li², Xuancheng Ren²,
Xu Sun², Songfang Huang¹, Fei Huang¹

¹Alibaba Group

²Key Laboratory of Computational Linguistics, Peking University

{lfl1259702, songfang.hsf, f.huang}@alibaba-inc.com

{yang_pc, lisc99, renxc, xusun}@pku.edu.cn

Abstract

Pre-trained self-supervised models such as BERT have achieved striking success in learning sequence representations, especially for natural language processing. These models typically corrupt the given sequences with certain types of noise, such as masking, shuffling, or substitution, and then try to recover the original input. However, such pre-training approaches are prone to learning representations that are covariant with the noise, leading to the discrepancy between the pre-training and fine-tuning stage. To remedy this, we present *ContrAstive Pre-Training (CAPT)* to learn noise invariant sequence representations. The proposed CAPT encourages the consistency between representations of the original sequence and its corrupted version via unsupervised instance-wise training signals. In this way, it not only alleviates the pretrain-finetune discrepancy induced by the noise of pre-training, but also aids the pre-trained model in better capturing global semantics of the input via more effective sentence-level supervision. Different from most prior work that focuses on a particular modality, comprehensive empirical evidence on 11 natural language understanding and cross-modal tasks illustrates that CAPT is applicable for both language and vision-language tasks, and obtains surprisingly consistent improvement, including 0.6% absolute gain on GLUE benchmarks and 0.8% absolute increment on NLVR².

1 Introduction

Recently, pre-trained self-supervised models such as BERT (Devlin et al., 2019) have attracted an increasing amount of attention in natural language processing and vision-language processing. Benefiting from common knowledge contained in massive unlabeled data (Liu et al., 2019), the pretraining-finetuning framework has become

*Equal Contribution.

Models	Noise types
BERT (Devlin et al., 2019)	Mask tokens
SpanBERT (Joshi et al., 2019)	Mask spans
RoBERTa (Liu et al., 2019)	Mask token
XLNet (Yang et al., 2019)	Shuffle token
ELECTRA (Clark et al., 2019)	Replace tokens
StructBERT (Wang et al., 2019b)	Mask + Shuffle tokens
BART (Lewis et al., 2019)	Mask + Shuffle + Replace.
UNITER (Chen et al., 2019)	Mask tokens/regions
LXMERT (Tan and Bansal, 2019)	Mask tokens/regions

Table 1: The type of noise used in the current natural language (upper) and vision-language (lower) sequence representation models.

a representative paradigm for advancing various language-related downstream tasks.

Most endeavors on pre-trained representation models rely on elaborately designed self-supervised tasks, which typically corrupt the given sequence with certain types of noise (e.g., masking in BERT in Table 1), and then train the model to recover the original sequence. As a consequence, the learned representations tend to be covariant with the input noise of pre-training in this paradigm. However, when transferred to downstream tasks, the pre-trained model is responsible for encoding the original sequence without noise, and is expected to obtain noise invariant representations. Such pretrain-finetune discrepancy not only impedes fast fine-tuning, but also may result in suboptimal sequence representations, thus affecting the performance in downstream tasks.

To remedy this, we present *ContrAstive Pre-Training (CAPT)* to learn noise invariant (or denoised) sequence representations. The core idea of CAPT is to enhance the consistency between semantic representations of the original sequence and that of corresponding corrupted version (e.g. the masked sequence) via unsupervised instance-wise training signals. As shown in Figure 1, our approach strives to pull the representation of the corrupted sequence towards that of the original

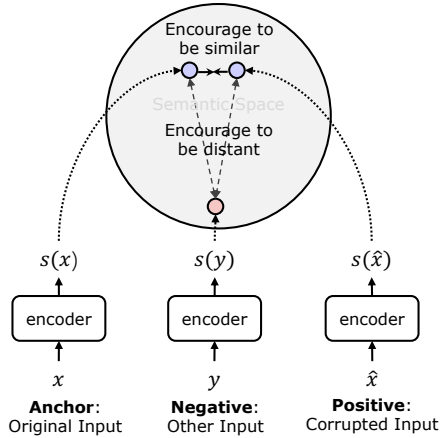


Figure 1: ContrAstive Pre-Training (CAPT) encourages the representations of inputs sharing semantics like (x, \hat{x}) to be similar, while penalizing the representations of inputs expressing different semantics like (x, y) and (\hat{x}, y) to be distant.

instance in the semantic space, while pushing it away from representations of other instances. Such training objectives are formulated as a multi-class classification task, which aims at classifying the original sequence to the class of its corrupted version and vice versa, while classifying different instances into different classes. Moreover, in order to enable the model to learn from more “difficult” and “diverse” instances, two effective methods are proposed to further enhance the capability of the model to extract noise-concentrated and instance-diffused features. With such training objective, the pre-trained model is encouraged to learn noise invariant representations, thereby alleviating the pretrain-finetune discrepancy to some extent.

As an additional benefit, our approach also assists the pre-trained model to more effectively capture the global semantics of the input. Most prior work only focuses on token-level pre-training tasks (e.g. masked language modeling), which lacks the modeling of global semantics of the input. Some other efforts alleviate this problem by introducing sentence-level pre-training tasks (e.g. next sentence prediction) that rely on the relative position of segments in the document. However, the semantic connection between these segments tends to be excessively loose, which may result in confusing gradient signals (Liu et al., 2019). By contrast, our CAPT offers incentives for representations of inputs sharing the same semantics (the original instance and its corrupted version) to be similar, while the representations of inputs expressing dif-

ferent semantics (different instances) are penalized to be distinguished from each other. Such more reasonable sentence-level supervision enables our approach to look beyond the local structures of input sequences and become more aware of the global semantics.

We perform the evaluation on a comprehensive suite of benchmark, covering 8 natural language understanding and 3 cross-modal tasks. Extensive empirical evidence demonstrates that our approach can achieve consistent improvements over the baselines in both language and vision-language domains. To be more specific, our CAPT raises the performance of RoBERTa (Liu et al., 2019) from 88.9% to 89.5% on the GLUE dev set, and also surpasses LXMERT (Tan and Bansal, 2019) by 0.5%, 0.6% and 0.8% on VQA, GQA and NLVR², respectively.

2 Methodology

2.1 Contrastive Pre-training

The proposed CAPT has excellent versatility, which can be built on various pre-trained models in either language or vision-language domains. Therefore, we use the symbol \mathcal{E} to represent a series of generalized pre-trained models. Starting from the property of the semantic representation that inputs sharing semantics should exhibit similar representations, our CAPT strives to capture global semantics of the input more effectively. Different from prior work that tends to learn representations covariant with the noise of pre-training, our CAPT aims at aiding \mathcal{E} in learning noise invariant sequence representations by enhancing the consistency between representations of the original sequence and its corrupted version.

Specifically, for a pre-trained model \mathcal{E} and an input sequence x , the model-specific noise (e.g. masking in BERT) can be added to x to construct its corrupted version \hat{x} . Then, the pre-trained model \mathcal{E} encodes x or \hat{x} with self-attention mechanism (Vaswani et al., 2017) to obtain hidden representations $h(x) = \mathcal{E}(x)$ or $h(\hat{x}) = \mathcal{E}(\hat{x})$. Both $h(x)$ and $h(\hat{x})$ belong to the representation space $\mathbb{R}^{m \times d}$, where m denotes the length of the input sequence and d is the dimension of hidden representation.

Different from prior work (Devlin et al., 2019), we apply an extra aggregation layer \mathcal{A} to obtain the global semantic representation of the input. Here \mathcal{A} can be implemented as a multi-layer perceptron

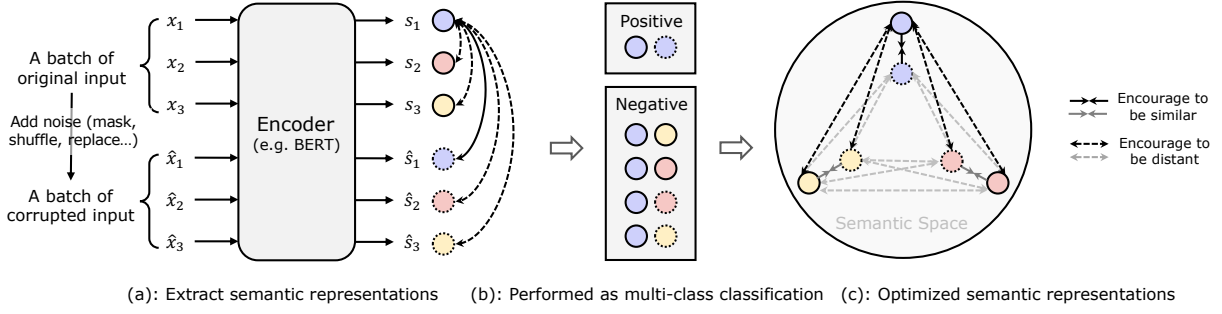


Figure 2: ContrActive Pre-Training (CAPT) performed in a training batch. (a): A batch of original sequences and corresponding corrupted versions are fed into the encoder to extract global representations. (b, c): CAPT aims at classifying the original sequence (e.g. x_1) to the class of its corrupted version (e.g. \hat{x}_1) and vice versa to encourage them to be similar in the semantic space, while classifying different instances into other classes to encourage them to be distant (see Eq. (2)).

with the representation of special classification token or the mean-pooling of all token representations as input. The final global semantic representations $s(\mathbf{x}) \in \mathbb{R}^d$ and $s(\hat{\mathbf{x}}) \in \mathbb{R}^d$ of \mathbf{x} and $\hat{\mathbf{x}}$ are computed as:

$$\begin{aligned} s(\mathbf{x}) &= \ell_2^{\text{norm}} \circ \mathcal{A}(\mathbf{h}(\mathbf{x})) \\ s(\hat{\mathbf{x}}) &= \ell_2^{\text{norm}} \circ \mathcal{A}(\mathbf{h}(\hat{\mathbf{x}})) \end{aligned} \quad (1)$$

where $\ell_2^{\text{norm}}(\cdot)$ represents ℓ_2 -normalization and \circ denotes the composition of operations. In order to obtain noise invariant sequence representations, we expect $s(\mathbf{x})$ and $s(\hat{\mathbf{x}})$ to be as similar as possible, which can also be derived from the characteristic that \mathbf{x} and $\hat{\mathbf{x}}$ share semantics. At the same time, the global semantic representations of different instances should be distinguished from each other to extract the high-level specific signals of the input. Motivated by this, we employ contrastive loss (Hadsell et al., 2006) to model such training objectives, which can be formalized as a multi-class classification task. We represent a training batch of the original sequences as $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ where n is the batch size, and its corresponding corrupted data is denoted as $\{\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n\}$. Intuitively, the loss should be low when s_i is similar to its corrupted version \hat{s}_i (positive example) and dissimilar to all other inputs (negative examples). Thus, the training loss for the original sequence \mathbf{x}_i is defined as:

$$\mathcal{L}(\mathbf{x}_i) = -\log \frac{\exp(\mathbf{s}_i \cdot \hat{\mathbf{s}}_i / \tau)}{\sum_j \exp(\mathbf{s}_i \cdot \hat{\mathbf{s}}_j / \tau) + \sum_{j \neq i} \exp(\mathbf{s}_i \cdot \mathbf{s}_j / \tau)} \quad (2)$$

where $\mathbf{s}_i = s(\mathbf{x}_i)$, $\hat{\mathbf{s}}_i = s(\hat{\mathbf{x}}_i)$, and τ is the temperature presented in Section 2.2. Similarly, the training loss for the corrupted sequence $\hat{\mathbf{x}}_i$ can be

defined as:

$$\mathcal{L}(\hat{\mathbf{x}}_i) = -\log \frac{\exp(\hat{\mathbf{s}}_i \cdot \mathbf{s}_i / \tau)}{\sum_j \exp(\hat{\mathbf{s}}_i \cdot \mathbf{s}_j / \tau) + \sum_{j \neq i} \exp(\hat{\mathbf{s}}_i \cdot \hat{\mathbf{s}}_j / \tau)} \quad (3)$$

Eq. (2) and Eq. (3) essentially correspond to the log loss of a softmax-based classifier measuring semantic similarity by dot product. The classifier treats each instance as a distinct class, and aims to classify \mathbf{x}_i to the class of $\hat{\mathbf{x}}_i$ and vice versa. More vividly, as shown in Figure 2, Eq. (2) and Eq. (3) strive to pull the original representation s_i towards the representation \hat{s}_i of the corrupted sequence $\hat{\mathbf{x}}_i$, and push it away from global semantic representations of other sequences. By maximizing the semantic similarity of global representations of \mathbf{x}_i and $\hat{\mathbf{x}}_i$, the model is encouraged to learn noise-invariant and instance-diffused representations. On this account, the self-supervised representation model is pre-trained in a manner that is more applicable for noise-free data distribution. This alleviates the pretrain-finetune discrepancy induced by the noise of pre-training to some extent, leading to improved performance on downstream scenarios. Besides, by introducing more reasonable sentence-level supervision, our approach can also capture global semantics of the input more effectively.

For the original training batch $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ and the constructed corrupted inputs $\{\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n\}$, the final contrastive loss is the total sum of losses of all instances, which can be formulated as:

$$\mathcal{L}_{\text{capt}} = \sum_{i=1}^n \{\mathcal{L}(\mathbf{x}_i) + \mathcal{L}(\hat{\mathbf{x}}_i)\} \quad (4)$$

2.2 Model Extensions

We improve the proposed CAPT methods by proposing two extensions to the model: adaptive temperature and memory queue.

Adaptive Temperature. Prior work (Chen et al., 2020) has illustrated that the temperature τ controlling the concentration level of the sample distribution in Eq. (2) and Eq. (3) exhibits significant impact on model performance. A suitable temperature can help the model learn from hard negatives via the gradient (Chen et al., 2020). Thus, it needs to be tuned elaborately to obtain a satisfactory fixed value. However, the optimal τ is constantly evolving as training proceeds. In fact, the gradient of the loss function Eq. (2) with respect to the representation s_i can be derived as:

$$\frac{1}{\tau Z} \left[\sum_{j \neq i} \left(\exp(s_i \cdot s_j / \tau) s_j + \exp(s_i \cdot \hat{s}_j / \tau) \hat{s}_j \right) + \left(\exp(s_i \cdot \hat{s}_i / \tau) - Z \right) \hat{s}_i \right] \quad (5)$$

where $Z = \sum_j \exp(s_i \cdot \hat{s}_j / \tau) + \sum_{j \neq i} \exp(s_i \cdot s_j / \tau)$ is the normalization factor.

We found that the norm of the gradient in Eq. (5) tends to be inversely proportional to τ . Ideally, the model update should be carefully controlled not both in the early stage of contrastive pre-training to stabilize the training, but also in the later stage to avoid ending up bouncing around the minimum or getting stuck in local optima. Therefore, we implement the temperature τ as the following inverted triangle schedule instead of a predetermined fixed value:

$$\tau(t) = \frac{1}{T} \left| t - \frac{T}{2} \right| + 0.05 \quad (6)$$

where t denotes the learning step and T refers to the preset total number of updates.

Memory Queue. As depicted in Figure 2, we treat different instances from the same batch as negative samples. Several recent studies (Chen et al., 2020; Dai and Lin, 2017) have illustrated that contrastive learning can benefit from larger negative representations. However, due to massive model parameters and limited machine memory, implementation with large batch size tends to be infeasible in many circumstances. To remedy this, we employ a dynamic memory queue \mathcal{Q} to store the desired negative representations (He et al., 2019). At each learning step, the aggregated representations of the current batch of original inputs and its corresponding corrupted inputs are enqueued into \mathcal{Q} . Once reaching the preset capacity of \mathcal{Q} , the oldest redundant representations are dequeued. Different from (Wu et al., 2018), the negative representations

stored in \mathcal{Q} are updated along with the training process to provide competitive confounders and informative signals for positive instances.

3 Experiments on Language Tasks

This section presents experiments on language tasks, including specific implementation and detailed results.

3.1 Implementation

For learning language representations, the noise corrupting the input sentence x can be implemented as masking like BERT or shuffling like BART. In our implementation, the main experiments follow the same corruption as BERT. That is, we randomly mask 15% tokens of x to construct its corrupted version \hat{x} . Then, both x and \hat{x} in the training batch are fed into the encoder to compute the CAPT loss and mask language model (MLM) loss simultaneously. The final training loss is the sum of the above two.¹ More analysis of the influence of other corruption approaches can be found in Section 5.2.

During pre-training, we train a small model (Section 5.1) to validate the influence of key components of CAPT, and a large model (Section 3.3 and Section 5.2) to demonstrate the effectiveness of CAPT for learning denoised text representation at a large scale. The small model, which is designed as a 6-layer Transformer with 256 hidden size and 4 attention heads, is trained on BookCorpus and English Wikipedia datasets. For the large CAPT model, we adopt RoBERT-Large model and training settings, which is a 24-layer Transformer with 1024 hidden size, 16 attention heads, and is trained on larger datasets. Readers can refer to (Liu et al., 2019) for the statistics of the dataset and processing details. The aggregation layer \mathcal{A} that takes the representation of special classification token as input is implemented as a nonlinear projection with one hidden layer. The inner hidden size of \mathcal{A} is set to the same as FFN inner hidden size and the output hidden size is set to the same as Transformer hidden size. The queue size is set to 8192 and we use Adam optimizer. The peak learning rate with linear warmup and decay is set to $5e-4$ and $6e-4$ for small and large models, respectively².

¹The reason is that the MLM excels at learning representations conditioned on masking noise and CAPT can help to learning a noise invariant representation.

²More details about pre-training hyper-parameters can be found in Appendix A.

Model	CoLA	SST-2	MRPC	STS-B	QQP	MNLI	QNLI	RTE	Avg
<i>Single-model test results via a leaderboard submission</i>									
BERT (Devlin et al., 2019)	60.5	94.9	85.4	87.6	89.3	86.7	92.7	70.1	83.4
SpanBERT (Joshi et al., 2019)	64.3	94.8	87.9	89.9	89.5	88.1	94.3	79.0	86.0
RoBERTa (Liu et al., 2019)	63.8	96.3	88.1	91.9	90.0	89.8	94.8	86.0	87.6
CAPT (Ours)	64.7	96.8	88.9	91.5	89.7	90.0	94.7	86.9	87.9
<i>Averaged dev results over 5 random seeds</i>									
BERT (Devlin et al., 2019)	60.6	93.2	88.0	90.0	91.3	86.6	92.3	70.4	84.0
XLNet (Yang et al., 2019)	63.6	95.6	89.2	91.8	91.8	89.8	93.9	83.8	87.4
RoBERTa (Liu et al., 2019)	68.0	96.4	90.9	92.4	92.2	90.2	94.7	86.6	88.9
ELECTRA (Clark et al., 2019)	69.3	96.0	90.6	92.1	92.4	90.5	94.5	86.8	89.0
CAPT (Ours)	69.2	96.5	92.1	92.5	92.3	90.7	95.0	88.0	89.5

Table 2: GLUE test and dev results of large models (24-layer transformer). We list the results on each set that are available in the published papers. ‘‘Avg’’ denotes the average score in terms of the reported metrics, which is slightly different from that in the GLUE leaderboard.

3.2 Evaluation Tasks

We perform the evaluation on the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2019a). Following previous work (Devlin et al., 2019; Liu et al., 2019), we experiment on 8 natural language understanding tasks, including linguistic acceptability (CoLA), sentiment analysis (SST), text paraphrase (MRPC and QQP), sentence similarity (STS-B), and natural language inference (MNLI, QNLI, and RTE).

For fine-tuning, we adopt the same settings and hyper-parameters as those of RoBERTa. All GLUE tasks are framed as single-sentence or sentence-pair classification tasks, except for STS-B which is a regression task. Extra multi-layer perceptrons (MLP) are added to perform classification or regression with the representation of special classification token as the input. The evaluation metrics include Matthews correlation for CoLA, Pearson correlation for STS-B, and accuracy for other tasks.

3.3 Results

Following prior work (Liu et al., 2019; Yang et al., 2019), we report results on both dev and test data sets. For the dev set, we report the median of multiple random fine-tuning runs like RoBERTa to show reliable results. For the test set, since the ground-truth labels are not obtainable, we only made a single-model submission to the GLUE evaluation server. Note that most systems on the GLUE leaderboard adopt different ensemble metrics and task-specific fine-tuning methods (e.g. formulating QNLI as a ranking task or using multi-task fine-tuning), increasing the difficulty for a direct

comparison. Thus, following (Devlin et al., 2019), we only report non-ensemble single-task results.

Table 2 presents the performance of representative models on the GLUE benchmark. We can see that the proposed CAPT obtains best results on most of datasets. In more detail, CAPT outperforms RoBERTa that is a very strong baseline model on most language understanding datasets, by 0.3% and 0.6% improvements of average test and dev score. Note that CAPT and RoBERTa are nearly identical in terms of the model architecture and fine-tuning hyperparameters, the only difference being the incorporation of contrastive pre-training in CAPT. Thus we can attribute the improvements of performance on downstream tasks to the role contrastive pre-training plays in learning noise invariant sequence representations. We believe the reason behind its success is the capability of CAPT to alleviate the pretrain-finetune discrepancy induced by the noise of pre-training.

In particular, we find that CAPT performs extremely well on natural language inference (RTE, MNLI) which requires a deep understanding of sentence semantics, with 1.0% absolute improvement of average accuracy over RoBERTa on the dev set. This phenomenon can be possibly explained by the fact that our CAPT can better capture the global semantics of the input sequence due to the more effective sentence-level supervision provided by the contrastive training via negative sampling from a memory queue, resulting in superior model performance therein. More analysis about the influence of memory queue can be found in Section 5.1.

Model	VQA		GQA		NLVR ²	
	test-dev	test-std	test-dev	test-std	dev	test-p
SOTA (No pre-training)	70.63	70.90	55.8	56.1	54.80	53.50
ViLBERT (Lu et al., 2019)	70.55	70.92	–	–	–	–
VisualBERT (Li et al., 2019b)	70.80	71.00	–	–	67.40	67.00
VL-BERT (Su et al., 2019)	71.79	72.22	–	–	–	–
LXMERT (Tan and Bansal, 2019)	72.42	72.54	59.95	60.33	74.82	74.41
CAPT (Ours)	72.78	73.03	60.48	60.93	75.12	75.13

Table 3: Comparison to the state-of-the-art systems with the single model on VQA, GQA and NLVR². The results of both VQA and GQA are reported on the “test-dev” split (used for validation on the official server) and the “test-std” split (used for maintaining the public leaderboard). The NLVR² results are reported on the local dev set (“dev”) and the public test set (“test-p”). The results of baselines except LXMERT are obtained from prior work.

4 Experiments on Vision-language Tasks

4.1 Implementation

Different from language tasks, the input sequence in the vision-language domain consists of visual region features paired with textual words. In this scenario, we build CAPT on LXMERT (Tan and Bansal, 2019), a representative cross-modal representation model that separately encodes visual and textual features and then introduces a cross-modal layer to integrate them. Same as Section 3.1, we construct the corrupted version \hat{x} of the original input x by masking part of visual features or textual words. In addition to the proposed CAPT which can learn sequence-level representations, following (Tan and Bansal, 2019), we also adopt three other pre-training tasks to learn more fine-grained word/region-level representations. These tasks include: masked language modeling (MLM) that predicts the masked words based on the corrupted input \hat{x} , masked region modeling (MRM) that predicts the masked visual region objects based on the corrupted input \hat{x} , and image-text matching (ITM) that predicts whether the input word sequence is semantically matched with the visual features. Due to space limitations, we do not elaborate on the detailed model architecture and these pre-training tasks here. We strongly recommend readers to refer to (Tan and Bansal, 2019) for the details. We also set the queue size to 8192 and the final training loss for visual-linguistic CAPT is defined as the sum of all the above training objectives.

We use the preprocessed data provided by (Tan and Bansal, 2019), which mainly includes CoCo (Lin et al., 2014) and Visual Genome (Krishna et al., 2017). Only 36 objects detected by Faster-RCNN are kept for each image. The model

architecture is the same as (Tan and Bansal, 2019), which consists of 9 language layers, 5 vision layers, and 5 cross-attention layers, with 768 hidden size. We employ Adam optimizer with the peak learning rate $1e - 4$ paired with linear warmup and decay. The batch size and dropout are set to 512 and 0.1, respectively.

4.2 Evaluation Tasks

We perform evaluation on three benchmark tasks: VQA, GQA, and NLVR². VQA (Goyal et al., 2017) aims to select the correct answer based on both the question and its paired image, while GQA (Hudson and Manning, 2019) shares the same task setting but require more reasoning. The goal of NLVR² (Suhr et al., 2019) is to predict whether the statement correctly describes the two images. All three tasks use accuracy (Acc) as the evaluation metric.

For VQA and GQA, we add extra multi-layer perceptrons (MLP) that take the representation of [CLS] as the input to perform classification. Since each instance in NLVR² is composed of two images (v_1, v_2) and a sentence s , we use the representation model to encode (v_1, s) and (v_2, s) , respectively. Then, a similar MLP takes the concatenation of the [CLS] representations of both (v_1, s) and (v_2, s) as the input to perform classification. We adopt the same hyper-parameters with LXMERT for fine-tuning to make a fair comparison.

4.3 Results

Table 3 presents the comparison between our approach and several representative systems on the vision-language tasks. We observe the consistent performance boost for our CAPT on all three tasks. For instance, it yields a 0.6% gain over the base




Figure 3: Left: The influence of adaptive temperature and the size of negative samples on the GLUE dev set of small CAPT models. Center: The validation curves regarding the accuracy of both CAPT and the base architecture LXMERT on the GQA dev set. Right: The GLUE dev results of large CAPT models that perform two different ways to construct the corrupted sequence.

architecture LXMERT on GQA and also surpasses various baselines on two other tasks. Such improvements indicate the enhanced capability of CAPT to learn noise invariant representations as well as capture the joint semantic representation of the input image-text pair. It is also worth noting that the performance gain of our approach on GQA is greater than that on VQA. The reason may be that GQA pays more attention to visual reasoning, which imposes higher demands on the modeling of joint semantics of visual-textual information. Correspondingly, our approach displays competence in the modeling of the global semantics of the input by introducing more effective sentence-level supervision, thereby attaining better results. Different from VQA and GQA, the goal of NLVR² is to determine whether a language caption is true about a pair of images. The increased accuracy on this task demonstrates the universal efficacy of our CAPT under a variety of task settings.

5 Further Analysis

5.1 Ablation Study

We conduct an ablation study to verify the effectiveness of adaptive temperature and memory queue proposed in Section 2.2. As illustrated in Figure 3 (Left), fixing the size of negative samples, adaptive temperature exhibits consistent superiority over manually tuned constant temperature. Our elaborately designed inverted triangle schedule regarding temperature allows the self-adjustment of the gradients at different stages of contrastive pre-training, leading to a significant gain in model performance. Figure 3 (Left) also demonstrates that the GLUE score of the model consistently increases with the size of stored negative representations, and the ab-

sence of memory queue (corresponding to the size of 128 in the figure) could result in a considerable degradation in performance. As depicted in prior work (Chen et al., 2020), large-scale negative samples can assist the model to capture distinctive high-level information of the inputs, therefore enhancing its capacity of feature extraction.

5.2 Analysis of Corruption noise

We conduct deeper analysis by constructing \hat{x} by means of the shuffling noise to gain further insight into the influence of different corruption methods. For implementation, we apply a random permutation to the input sequence x within a fixed window k to construct its corrupted version \hat{x} . Following previous work (Lample et al.; Wang et al., 2019b), we set the window k to 3.

Figure 3 (Right) presents the GLUE average score on the dev set when the corrupted sequence \hat{x} is constructed by shuffling or masking part of tokens. It suggests that both masking and shuffling contribute to the improvement of performance over the baseline which does not perform any corruption noise to compute contrastive loss. The results also reveal the superiority of the masking operation over the shuffling operation. We speculate that the reason behind this phenomenon may be twofold:

In the first place, since the baseline model is pre-trained with masked language modeling alone, the learned representations are covariant only with the masking noise. Thus, using the masking noise when applying CAPT to the baseline model makes more sense than using the shuffling noise.

The other reason may also lie in the observation that masking endows the model with the ability to “associate” words, while shuffling only allows the model to learn to “reorder” words. In other words,

masking is likely to bring more informative supervision signals than shuffling when learning language representations. Due to the analysis above, we opt for masking as the noise when implementing our final version of CAPT.

5.3 Alleviation of Pretrain-finetune Discrepancy

In order to verify that the proposed CAPT can effectively alleviate the pretrain-finetune discrepancy induced by the noise of pre-training, we plot validation curves regarding the accuracy of both CAPT and the base architecture LXMERT on GQA. Figure 3 (Center) presents the corresponding results, illustrating that the self-supervised model pre-trained by CAPT not only exhibits increased fine-tuning speed but also obtains better performance. As shown in Figure 3 (Center), in the early stage of fine-tuning (the first 1K steps), the proposed CAPT managed to maintain an absolute lead of 5%-10% over LXMERT consistently. This demonstrates that representations learned by CAPT are more applicable for the data distribution in downstream tasks, rendering model transfer more effective. By narrowing the differences between representations of the original sequence and its corrupted version, the model is encouraged to learn noise invariant sequence representations. In this way, the pre-trained representation model is congruous with the task setting of noiseless inputs in downstream scenarios, leading to better model performance.

6 Related Work

Pre-trained Language Representations. This task strives to build linguistic representations benefiting various downstream tasks. One line of research focuses on autoregressive (AR) pre-training, while the other centers on denoising autoencoding (DAE). Representative work of AR pre-training includes ELMo (Peters et al., 2018) and GPT (Radford, 2018), which aim to predict the next word based on previous tokens but lack the modeling of bidirectional context. The other research line is built upon DAE, which strives to reconstruct the original sequence based on the corrupted input by jointly attending to both the left and right context. Main efforts focus on *token-level* pre-training tasks. However, DAE introduces the noise discarded on downstream tasks during pre-training, which is prone to learn representations covariant with the input noise, leading to the pretrain-finetune

discrepancy. XLNet takes one step forward to solve this problem via permutation language modeling (PLM), but leave its own limitation: Each token does not know the “position” of future tokens in the permuted sentence (Song et al., 2020). Therefore, XLNet also brings a discrepancy between pre-training and fine-tuning. Since CAPT can be adapted to model with arbitrary noise transformations to the input, it can also be built based on the shuffling noise of PLM.

Besides, most AR and DAE based pre-training tasks neglect the modeling of global semantics of the input. Some DAE based approaches address this problem by incorporating supervisions regarding the entire segment through *sentence-level* tasks (e.g. next or adjacent sentence prediction (Devlin et al., 2019; Wang et al., 2019b)). However, such training relies heavily on the relative position of segments, which suffers from excessively loose semantic connections. Thus, it tends to result in confusing gradient signals. In comparison, CAPT encourages the semantic consistency of the original sequence and its corrupted version via unsupervised contrastive loss. This not only alleviates the pretrain-finetune discrepancy, but also better captures the global semantics of the input.

Pre-trained Vision-language Representations.

This direction attempts to build generic representation models for vision-language tasks. In terms of model architecture, one research line focuses on *one-stream* BERT-based architecture, which strives to learn generic image-text representations with a unified model. The corresponding representative work includes VideoBERT (Sun et al., 2019), VisualBERT (Li et al., 2019b), UNITER (Chen et al., 2019), Unicoder-VL (Li et al., 2019a), etc. In contrast, the other line such as ViLBERT (Lu et al., 2019) and LXMERT (Tan and Bansal, 2019) focuses on the *two-stream* architecture. They first separately encode visual and textual features and then interact with each other in the co-attention layers. As for pre-training tasks, different work exhibits commonalities, all focusing on MRM, MLM, and several specific tasks (e.g. ITM). However, most of these tasks are prone to learning noise covariant representations in the pre-training stage. Compared with these endeavors, our CAPT benefits the pre-trained model to learn noise invariant vision-language representations via elaborately-designed semantic contrastive loss, thereby bringing better model performance.

Contrastive Learning. It serves as an unsupervised objective, with the main idea is to construct or collect pairs of related (similar) data as positive samples and pairs of unrelated data as negative samples, and then learn to classify them via the contrastive loss. For example, the positive samples can be nodes connected by the same edge in graph representation learning (Bordes et al., 2013; Grover and Leskovec, 2016; Velickovic et al., 2019), images processed by pretext tasks in image representation learning (Wu et al., 2018; Ye et al., 2019), and bilingual sentence pairs in cross-lingual pre-training (Chi et al., 2020), etc. The contrastive loss can come in several forms, including noise contrastive estimation (Gutmann and Hyvärinen, 2010; Oord et al., 2018; Hjelm et al., 2019), instance-wise classification (Wu et al., 2018), and etc. Inspired by these works, we adapt contrastive learning to the natural language and vision-language domains to learn noise invariant sequence representations, demonstrating its effectiveness in improving various pre-trained models.

7 Conclusion

This work presents contrastive pre-training for learning denoised sequence representations in a self-supervised manner. By enhancing the consistency between representations of the original sequence and the corresponding corrupted version, the pre-trained model is encouraged to learn noise invariant sequence representations. On this account, the proposed approach not only alleviates the pretrain-finetune discrepancy induced by the noise of pre-training, but also better captures the global semantics of the input via more effective sentence-level supervision. Extensive experiments demonstrate the effectiveness and versatility of our approach, which can achieve consistent improvements over baselines in both language and vision-language domains.

Broader Impact

This section highlights the potential impact of this work, detailed as follows.

Beneficiaries of this work and detailed benefits.

Both academic researchers and industrial engineers engaged in language-related tasks can benefit from our work. Pre-trained self-supervised representation models have become a research mainstream in the field of language. Plenty of efforts has been

paid to promote relevant research. However, existing systems are prone to learn noise-covariant representations as well as cannot effectively capture the global semantic representation of the input, which results in the sub-optimal performance in downstream tasks. Our proposed CAPT not only alleviates the pretrain-finetune discrepancy induced by the noise of pre-training, but also aids the pre-trained model in better capturing global semantics of the input via more effective sentence-level supervision. This not only provides helpful insight for advancing related research in this field, but also offers an effective means for engineers to improve industrial systems and user experience.

Consequences of system failure. Existing deep learning systems, especially pre-trained representation models, require large-scale computational overhead. We will spare no effort to provide the community with detailed instructions on this work. However, the pre-training may suffer from the risk of divergence due to some unexpected reasons such as the wrong setting of hyper-parameters. In such extremely unfortunate situations, the wasted computing resources may result in unnecessary pollution to the environment.

Potential adverse effects. As an academic research paper, our work aims at providing the community with more effective tools and helpful insights. However, everything has its double-sidedness. The key factor determining the outcome lies in the individual who owns the tool, not the tool itself. As a representative deep learning paradigm, our approach may also adhere to the risks suffered from all deep learning models. When employed improperly by others, it may cause potential adverse effects on society. For instance, almost all pre-trained models (also including our approach) can be used to make fake texts. Nevertheless, we still want to emphasize that technology is innocent.

References

- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, pages 2787–2795.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A simple framework for contrastive learning of visual representations. volume abs/2002.05709.

- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Uniter: Learning universal image-text representations.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2020. InfoXLM: An information-theoretic framework for cross-lingual language model pre-training.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2019. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- Bo Dai and Dahua Lin. 2017. Contrastive learning for image captioning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 898–907.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 6325–6334.
- Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010*, pages 297–304.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), 17-22 June 2006, New York, NY, USA*, pages 1735–1742.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2019. Momentum contrast for unsupervised visual representation learning.
- R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Drew A. Hudson and Christopher D. Manning. 2019. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6700–6709.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2019. SpanBERT: Improving pre-training by representing and predicting spans.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. volume 123, pages 32–73.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. Unsupervised machine translation using monolingual corpora only. In *6th International Conference on Learning Representations, ICLR 2018*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension.
- Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and Ming Zhou. 2019a. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019b. VisualBERT: A simple and performant baseline for vision and language.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: common objects in context. In *Computer Vision - ECCV 2014*, pages 740–755.
- 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.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237.

Alec Radford. 2018. Improving language understanding by generative pre-training.

Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. Mpnnet: Masked and permuted pre-training for language understanding.

Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. VL-BERT: Pre-training of generic visual-linguistic representations.

Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. A corpus for reasoning about natural language grounded in photographs. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 6418–6428.

Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. VideoBERT: A joint model for video and language representation learning.

Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 5099–5110.

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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pages 5998–6008.

Petar Velickovic, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R. Devon Hjelm. 2019. Deep graph infomax. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. 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*.

Wei Wang, Bin Bi, Ming Yan, Chen Wu, Zuyi Bao, Liwei Peng, and Luo Si. 2019b. StructBERT: Incorporating language structures into pre-training for deep language understanding.

Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019*, pages 5754–5764.

Mang Ye, Xu Zhang, Pong C. Yuen, and Shih-Fu Chang. 2019. Unsupervised embedding learning via invariant and spreading instance feature. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6210–6219.

A Linguistic CAPT Pre-training Details

Table 4 describes the hyperparameters for pretraining of the base and large linguistic CAPT models.

Hyperparam	CAPT-Small	CAPT-Large
Number of Layers	6	24
Attention heads	4	16
Hidden size	256	1024
FFN inner hidden size	1024	4096
Aggregation layer inner hidden size	1024	4096
Aggregation layer output hidden size	256	1024
Dropout	0.1	0.1
Batch Size	2k	8k
Max Steps	200k	500k
Warmup Steps	10k	30k
Peak Learning Rate	5e-4	6e-4
Learning Rate Decay	Linear	Linear
Weight Decay	0.01	0.01
Adam ϵ	1e-6	1e-6
Adam β_1	0.9	0.9
Adam β_2	0.98	0.98

Table 4: Hyperparameters for pretraining small and large Linguistic CAPT models.

B Summary of GLUE Tasks

Table 5 shows the details of the each tasks in GLUE benchmark.

	CoLA	SST	MRPC	STS	QQP	MNLI	QNLI	RTE
#Train	8.5k	67k	3.5k	5.7k	363k	392k	108k	2.5k
#Dev	1k	872	3.7k	1.5k	40k	20k	5.7k	276
#Test	1k	1.8k	408	1.4k	391k	20k	5.7k	3k
Task	SSC	SSC	SPC	SPR	SPC	SPC	SPC	SPC
#Classes	2	2	2	-	2	3	2	2
Metrics	Mcc	Acc	Acc	Pcc	Acc	Acc	Acc	Acc

Table 5: Summary of the GLUE benchmark. All the tasks are framed as Single Sentence Classification (SSC), Sentence Pair Classification (SPC), and Sentence Pair Regression (SPR).