

# Adversarial Training for Machine Reading Comprehension with Virtual Embeddings

Ziqing Yang<sup>1†</sup>, Yiming Cui<sup>2,1†</sup>, Chenglei Si<sup>3</sup>, Wanxiang Che<sup>2</sup>,  
Ting Liu<sup>2</sup>, Shijin Wang<sup>1</sup>, Guoping Hu<sup>1</sup>

<sup>1</sup>State Key Laboratory of Cognitive Intelligence, iFLYTEK Research, China

<sup>2</sup>Research Center for SCIR, Harbin Institute of Technology, Harbin, China

<sup>3</sup>University of Maryland, College Park MD, USA

†{zqyang5, ymcui}@iflytek.com

## Abstract

Adversarial training (AT) as a regularization method has proved its effectiveness on various tasks. Though there are successful applications of AT on some NLP tasks, the distinguishing characteristics of NLP tasks have not been exploited. In this paper, we aim to apply AT on machine reading comprehension (MRC) tasks. Furthermore, we adapt AT for MRC tasks by proposing a novel adversarial training method called PQAT that perturbs the embedding matrix instead of word vectors. To differentiate the roles of passages and questions, PQAT uses additional virtual P/Q-embedding matrices to gather the global perturbations of words from passages and questions separately. We test the method on a wide range of MRC tasks, including span-based extractive RC and multiple-choice RC. The results show that adversarial training is effective universally, and PQAT further improves the performance.

## 1 Introduction

Neural networks have achieved superior performance on many tasks, but they are vulnerable to *adversarial examples* (Szegedy et al., 2014) – examples that have been mixed with certain perturbations. *Adversarial training* (AT) (Goodfellow et al., 2015) uses both clean and adversarial examples to improve the robustness of the model for image classification.

In the field of NLP, Miyato et al. (2017) have applied adversarial training on text classification tasks and improved the model performance. From then on, many AT methods has been proposed (Wu et al., 2017; Yasunaga et al., 2018; Bekoulis et al., 2018; Zhu et al., 2020; Jiang et al., 2019; Pereira et al., 2020; Liu et al., 2020). They mostly adopt a general AT strategy, but focus less on the adaptation of AT to NLP tasks. To explore this adaptation, in this work, we aim to apply adversarial training

---

**Passage:** ... The **rock** cycle is an important concept in geology which illustrates the relationships between these three types of **rock**, and magma. When a **rock** crystallizes from melt (magma and/or lava), it is an **igneous rock**. ...

---

**Question:** An **igneous rock** is a **rock** that crystallizes from what?

---

Table 1: An example from the SQuAD dataset. We highlight two words *rock* and *igneous* for better demonstration. The words with the same color are injected with the same perturbation by PQAT. The different occurrences of the same word (for example, *rock* in passage and question) are perturbed differently depending on their roles.

on machine reading comprehension (MRC) tasks, which exhibit complex NLP characteristics.

The objective of MRC is to let a machine read the given passages and ask it to answer the related questions. There are several types of MRC tasks. In this work we focus on span-based extractive RC (Rajpurkar et al., 2016, 2018; Yang et al., 2018) and multiple-choice RC (Lai et al., 2017). To apply adversarial training on MRC tasks, we notice that there are several salient characteristics of MRC compared to other tasks such as image classification: (1) The inputs are discrete. Unlike pixels, which can take continuous values, words are discrete tokens. (2) The tokens in the input sequences are not independent. A word may occur in an input sequence several times. After the embedding layer, these occurrences are represented by the word vectors with the same value and hold the same semantic meaning (although the word may be polysemous). (3) The roles of passages and questions are different. Given a question as the query, the model needs to look up the correct answer in the passage.

People have utilized the first characteristic to apply adversarial training by perturbing input word vectors instead of tokens. However, the second and third characteristics have been largely ignored. For

example, in Table 1, which is a passage-question pair from the SQuAD dataset, the word *rock* has appeared multiple times. In the standard adversarial training, the perturbations added to each occurrence of *rock* are different, ignoring the fact that they share the same meaning. On the other hand, the multiple occurrences of the same word in the passage and question play different roles, such as the *rock* in the passage and question. It is appropriate to treat them differently.

To take the second and the third characteristics into consideration, we propose a novel adversarial training method called **PQAT**. The core of PQAT is the virtual P/Q-embeddings, which are two independent embedding spaces for passages and questions. Each time we calculate perturbations, P/Q-embeddings gather the perturbations from passages and questions for each word, then generate a global and role-aware perturbation for each word from passages and questions separately. For example, in Table 1, the perturbations on all the occurrences of *rock* in the passage and question will be gathered into two matrices separately, forming global and role-aware perturbations of *rock*. PQAT is as efficient as the standard AT with nearly no extra time cost. Also, The virtual P/Q-embeddings are only used during training. They are discarded once the training is finished. Thus PQAT does not increase the model size and inference time for predictions.

We have applied adversarial training on several MRC tasks, including span-based extractive RC and multiple-choice RC. Results show that adversarial training improves the MRC model performance universally and consistently, even over the strong pre-trained model baseline. Furthermore, the PQAT method outperforms the standard AT on both normal datasets and adversarial datasets. Lastly, our results verify the usefulness of incorporating information of task form into the design of the adversarial training method.

## 2 Standard Adversarial Training

Adversarial training first constructs adversarial examples by generating worst-case perturbations that maximize the current loss, then minimize the loss on those adversarial examples. In NLP tasks, a popular approach to generate perturbations is to perturb word vectors from the embedding layer (Miyato et al., 2017). We denote the input token sequence as  $X$  and the operation of looking up in an embedding layer  $E$  as  $emb(E, \cdot)$ . The objective




Figure 1: P/Q-embeddings collect the perturbations on each word from passages and questions separately.

of AT is

$$\min_{\theta, E} \mathbb{E}_{(X, y) \sim \mathcal{D}} \left[ \max_{\|\delta\| < \epsilon} \mathcal{L}(f_{\theta}(X_{vec} + \delta), y) \right] \quad (1)$$

where  $f_{\theta}(\cdot)$  is the model parametrized by  $\theta$  excluding word embedding layer;  $X_{vec} = emb(E, X)$  is the word vectors of input sequence.  $\mathcal{L}$  is the loss function. We perturb the word vectors with the adversarial perturbations  $\delta$ .

$\delta$  can be estimated by linearizing  $\mathcal{L}(f_{\theta}(X_{vec} + \delta), y)$  around  $X$  and perform the multiple-step projected gradient descent (PGD) (Madry et al., 2018):

$$\delta_{t+1} = \Pi_{\|\delta\| \leq \epsilon}(\delta_t + \alpha g_t / \|g_t\|) \quad (2)$$

$$g_t = \nabla_{\delta} \mathcal{L}(f_{\theta}(X_{vec} + \delta), y)|_{\delta=\delta_t} \quad (3)$$

where  $t$  is the gradient descent step,  $\Pi_{\|\delta\| \leq \epsilon}$  denotes projection  $\delta$  back onto the  $\epsilon$ -ball.  $g_t$  is the gradient of the loss with respect to perturbation  $\delta$ . The more gradient descent steps, the better approximation of  $\delta$ , but also more expensive in computation.

## 3 Adversarial Training for MRC

In the above algorithm, when generating the perturbations on  $X_{vec}$  through backward propagation, each word vector  $X_{vec}^i$  is perturbed independently, like the pixels in an image. It ignores the semantic relationship among the word vectors of a word's different occurrences. To make the perturbation on each occurrence aware of other occurrences of the same word, we adapt AT by gathering not only the perturbations on each word vector, but also the perturbations on the embedding matrix. The latter can be seen as the global perturbations, which provide context-insensitive semantic information.

The global perturbations are rather coarse-grained, since all the occurrences of the same word

receive the same global perturbation. Note that in MRC tasks, words in passages and questions play different roles. Thus, to keep this information, we distinguish the words in passages and questions by creating two *virtual* embedding matrices  $\mathbf{P}$  and  $\mathbf{Q}$ : P-embedding matrix  $\mathbf{P}$  collects the perturbations of all the words from the passages; Q-embedding matrix  $\mathbf{Q}$  for the questions. We give an illustration in Figure 1. P/Q-embedding matrices are *virtual* since they only provide perturbations, not the real word vectors. During training, perturbations from virtual embeddings and word vectors are summed up to form the adversarial input  $Z_{vec}$ . The final objective is

$$\min_{\theta, \mathbf{E}} \mathbb{E}_{(X, y) \sim \mathcal{D}} \left[ \max_{\|\delta\| < \epsilon} \mathcal{L}(f_{\theta}(Z_{vec}), y) \right] \quad (4)$$

$$Z_{vec} = [X_{vec}^P + P_{vec}; X_{vec}^Q + Q_{vec}] + \delta \quad (5)$$

$P_{vec} = emb(\mathbf{P}, X^P)$ ,  $Q_{vec} = emb(\mathbf{Q}, X^Q)$  are the perturbations from the virtual embeddings.  $X^P$  and  $X^Q$  stand for the passage and question sections in  $X$ .  $[\cdot; \cdot]$  denotes concatenation. In this way, we have generated fine-grained local perturbations  $\delta$  by standard AT, and global role-aware perturbations  $P_{vec}$  and  $Q_{vec}$  by the virtual P/Q-embeddings. We call the later process as **PQAT**, which is the main adaptation of adversarial training for MRC.

We list the overall algorithm of adversarial training for MRC in Algorithm 1. We initialize  $\mathbf{P}$  and  $\mathbf{Q}$  with the gaussian distribution. For each batch, we perform  $K$ -step gradient descent (line 9–22): we look up the original word vectors and P/Q-embedding vectors from the embedding layer  $\mathbf{E}$  and the P/Q-embedding matrices. The adversarial inputs are constructed by summing them with local perturbations  $\delta$ . Then we compute the gradients of model parameters  $\mathbf{g}_t$ , local perturbations  $\mathbf{g}_{\delta}$  and P/Q-embedding matrices  $\mathbf{g}_P$  and  $\mathbf{g}_Q$ . These gradients can be calculated in a single backward pass. Lastly, we update the virtual embeddings and local perturbations (line 18–21).

Note that P/Q-embedding matrices serve as the containers for perturbations. When the training is finished, P/Q-embedding matrices are no longer needed and can be discarded.

$\epsilon_{\delta}$ ,  $\epsilon_P$  and  $\epsilon_Q$  control the strengths of standard AT and PQAT. If  $\epsilon_{\delta} = 0$ , we have a pure P/Q-embeddings based adversarial training, i.e., PQAT; while if  $\epsilon_P = \epsilon_Q = 0$ , we recover the standard AT.

---

### Algorithm 1: Adversarial Training for Machine Reading Comprehension

---

**Notation:**  $V$  is the vocabulary size;  $D$  is the embedding dimension.

**Input:** Training samples  $\mathcal{D} = \{(X, y)\}$ ,  
P/Q-embedding matrices  $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^{V \times D}$ ,  
initialization variance  $\sigma$ , perturbation strength  
 $\{\epsilon_{\delta}, \epsilon_P, \epsilon_Q\}$ , adversarial steps  $K$ .

```

1 Initialize P/Q-embedding matrices
2  $\mathbf{P} \leftarrow \mathcal{N}(0, \sigma^2 I)$ ,  $\mathbf{Q} \leftarrow \mathcal{N}(0, \sigma^2 I)$ 
3 for batch  $B \subset \mathcal{D}$  do
4   Normalize P/Q-embedding matrices
5    $\mathbf{P} \leftarrow (\mathbf{P} - \text{mean}(\mathbf{P}) / \text{std}(\mathbf{P})) \cdot \sigma$ 
6    $\mathbf{Q} \leftarrow (\mathbf{Q} - \text{mean}(\mathbf{Q})) / \text{std}(\mathbf{Q}) \cdot \sigma$ 
7   Initialize perturbation and gradient
8    $\delta \leftarrow \frac{1}{\sqrt{D}} U(-\sigma, \sigma)$ ,  $\mathbf{g}_0 \leftarrow 0$ 
9   for  $t = 1, \dots, K$  do
10     $X_{vec} = emb(\mathbf{E}, X)$ 
11     $P_{vec} = emb(\mathbf{P}, X^P)$ 
12     $Q_{vec} = emb(\mathbf{Q}, X^Q)$ 
13     $Z_{vec} = X_{vec} + P_{vec} + Q_{vec} + \delta$ 
14     $\mathbf{g}_t = \mathbf{g}_{t-1} + \mathbb{E}[\nabla_{\theta, \mathbf{E}} \mathcal{L}(f_{\theta}(Z_{vec}), y)]$ 
15     $\mathbf{g}_{\delta} = \mathbb{E}[\nabla_{\delta} \mathcal{L}(f_{\theta}(Z_{vec}), y)]$ 
16     $\mathbf{g}_P = \mathbb{E}[\nabla_P \mathcal{L}(f_{\theta}(Z_{vec}), y)]$ 
17     $\mathbf{g}_Q = \mathbb{E}[\nabla_Q \mathcal{L}(f_{\theta}(Z_{vec}), y)]$ 
18     $\delta \leftarrow \delta + \mathbf{g}_{\delta} / \|\mathbf{g}_{\delta}\|_2 \cdot \|X_{vec}\|_2 \epsilon_{\delta}$ 
19    Update with token-wise normalization
20     $P^i \leftarrow P^i + \mathbf{g}_P^i / \|\mathbf{g}_P^i\|_2 \cdot \|X_{vec}^i\|_2 \epsilon_P$ 
21     $Q^i \leftarrow Q^i + \mathbf{g}_Q^i / \|\mathbf{g}_Q^i\|_2 \cdot \|X_{vec}^i\|_2 \epsilon_Q$ 
22  end
23   $\{\theta, \mathbf{E}\} \leftarrow \text{AdamUpdate}(\{\theta, \mathbf{E}\}, \mathbf{g}_K)$ 
24 end
```

---

## 4 Experiments Setup

**Datasets.** We perform experiments on several English MRC tasks, including span-based extractive MRC tasks – SQuAD 1.1 (Rajpurkar et al., 2016), SQuAD 2.0 (Rajpurkar et al., 2018), HotpotQA (Yang et al., 2018), and multiple-choice MRC task RACE (Lai et al., 2017). We also test model robustness on the adversarial datasets AddSent and AdOneSent (Jia and Liang, 2017).

**Model Settings.** We build the MRC model with RoBERTa (Liu et al., 2019), following the standard model structure for SQuAD and RACE (Devlin et al., 2018). For HotpotQA, we follow the model in Shao et al. (2020). It uses RoBERTa as the encoder followed by a multi-task prediction layer. We denote the passage as  $P$  and the question as  $Q$ . To construct the inputs, for span-based extractive RC, we concatenate each  $P$  and  $Q$  with model-dependent special tokens; for multiple-choice RC with  $m$  options for each example, we append each option to the concatenation of  $P$  and  $Q$ , and construct  $m$  input sequences from each example.

When applying AT or PQAT, we only perturb

Model	SQuAD 1.1		SQuAD 2.0		HotpotQA		RACE
	EM	F1	EM	F1	joint EM	joint F1	Acc
<i>BASE setting</i>							
RoBERTa	84.72	91.54	79.77	83.18	41.70	69.30	74.75
PQAT	<b>85.87</b>	<b>92.33</b>	<b>81.66</b>	<b>84.79</b>	<b>43.03</b>	<b>70.40</b>	<b>76.32</b>
<i>LARGE setting</i>							
RoBERTa	87.76	93.90	85.67	88.86	45.91 <sup>†</sup>	73.93 <sup>†</sup>	84.66
PQAT	<b>88.32</b>	<b>94.34</b>	<b>86.35</b>	<b>89.49</b>	<b>46.79</b>	<b>74.63</b>	<b>86.02</b>

Table 2: Results on the development sets of SQuAD 1.1, SQuAD 2.0 and HotpotQA, and results on the test set of RACE. <sup>†</sup>: the results are taken from Shao et al. (2020).

Model	SQuAD 1.1 EM	SQuAD 2.0 EM	RACE Acc
<i>BASE setting</i>			
<u>PQAT</u>	<u>85.87</u>	<b>81.66</b>	<u>76.32</u>
	(0.08)	(0.21)	(0.32)
PQAT + AT	<b>85.96</b>	81.11 ↓	<b>76.50</b>
	(0.10)	(0.14)	(0.35)
AT	85.64 ↓	81.23 ↓	75.94 ↓
	(0.15)	(0.30)	(0.37)

Table 3: Comparison of PQAT, standard AT and the combination. AT is short for Standard AT. Arrows indicate the drops relative to the PQAT. Numbers in the parentheses are the standard deviations.

the word embeddings and leave the position embeddings unchanged. For PQAT on RACE, we let the Q-embedding matrix collect perturbations from both questions and options.

**Training Settings and Hyperparameters.** All the models are implemented with Transformers (Wolf et al., 2019) and trained on a single Nvidia V100 GPU. To improve the stability and reduce the uncertainty of the results, we run each experiment four times with different seeds and report the mean value of performance. We use AdamW as our optimizer with batch size 24 and learning rate  $3e-5$  for RoBERTa<sub>BASE</sub> and  $2e-5$  or  $1e-5$  for RoBERTa<sub>LARGE</sub>. The maximum number of epochs is set to 3 for SQuAD and 5 for RACE and HotpotQA. A linear learning rate decay schedule with warmup ratio 0.1 was used. For PQAT,  $\epsilon_\delta$  is set to 0,  $\epsilon_P$  and  $\epsilon_Q$  is set to  $4e-2$  for RACE and  $2e-2$  for other tasks. The variance  $\sigma$  is  $1e-2$ . We set the number of gradient descent steps  $K = 2$  to balance speed and performance.

## 5 Results

### 5.1 Overall Results

The overall results are summarized in Table 2, where we compare PQAT with the baseline. PQAT

Model	AddSent	AddOneSent	Dev
R.M-Reader <sup>†</sup>	58.5	67.0	86.6
KAR <sup>‡</sup>	60.1	72.3	83.5
ALUM <sub>BERT-BASE</sub> <sup>§</sup>	60.4	69.8	90.8
RoBERTa <sub>BASE</sub>	59.7	68.8	91.5
PQAT	<b>64.7</b>	<b>73.6</b>	<b>92.3</b>
AT	63.2	72.6	92.1

Table 4: Model performance (F1) on AddSent, AddOneSent and SQuAD 1.1 dev set. AT is short for Standard AT. <sup>†</sup>: Wang and Jiang (2018). <sup>‡</sup>: Hu et al. (2018), <sup>§</sup>: Liu et al. (2020).

is able to boost model performance across all MRC tasks and outperforms the RoBERTa baseline significantly. On HotpotQA, which is a complicated MRC task that features multi-hop questions and asks for multiple kinds of predictions, PQAT still outperforms the baseline by 1.3/1.1 on Joint EM/Joint F1. On RACE, PQAT improves the performance significantly by 1.5% in accuracy. The universal improvements on various kinds of MRC tasks prove the wide applicability of PQAT.

### 5.2 Comparison

We compare different adversarial training methods and their combinations by tuning the strengths of perturbations  $\{\epsilon_\delta, \epsilon_P, \epsilon_Q\}$ . The results are in Table 3. The underlined scores are the ones reported in Table 2. Firstly, to test the effectiveness of standard AT, we disable PQAT with  $\epsilon_P = \epsilon_Q = 0$  and enable standard AT with  $\epsilon_\delta = 2e-3$  for RACE and  $1e-2$  for other tasks<sup>1</sup>. Other settings are unchanged, and we still follow Algorithm 1. PQAT consistently outperforms standard AT on the three tasks. Then we enable both PQAT and standard AT by setting all the strengths  $\{\epsilon_\delta, \epsilon_P, \epsilon_Q\}$  to non-zero values. The performance gets slightly better on SQuAD 1.1 and RACE, but gets worse on SQuAD 2.0.

<sup>1</sup>We have searched from  $1e-3$  to  $1e-1$  and taken the best value.

Compared with the standard AT, PQAT achieves higher performance by itself. Therefore PQAT could be a better alternative to applying adversarial training on MRC tasks.

### 5.3 Robustness on Adversarial Datasets

We assess the robustness of MRC models with AddSent and AddOneSent. AddSent and AddOneSent are two adversarial datasets built on SQuAD 1.1. In both datasets, passages are appended with distracting sentences. MRC models that heavily rely on text matching may be easily fooled to predict wrong answers from the distracting sentences.

The results are shown in Table 4. With the standard adversarial training (AT), the MRC model improves its robustness by about 5% over RoBERTa<sub>BASE</sub> in F1. PQAT further improves the performance over AT by about 1% on both AddSent and AddOneSent.

## 6 Conclusion

We have applied adversarial training on a wide range of MRC tasks, including span-based extractive RC and multiple-choice RC. Especially, we have proposed a novel adversarial training method PQAT, which uses virtual P/Q-embedding matrices to generate global and role-aware perturbations that consider the characteristics of MRC tasks. Our experiments demonstrate that adversarial training improves the MRC model performance universally and consistently, even over the strong pre-trained model baseline. The PQAT method further improves the model performance over the standard AT on both normal datasets and adversarial datasets.

## Acknowledgments

We would like to thank all anonymous reviewers for their valuable comments on our work. This work is funded by National Key R&D Program of China (No.2018YFC0831601).

## References

Giannis Bekoulis, Johannes Deleu, Thomas Demeester, and Chris Develder. 2018. [Adversarial training for multi-context joint entity and relation extraction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2830–2836, Brussels, Belgium. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep

bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *International Conference on Learning Representations*.

Minghao Hu, Yuxing Peng, Zhen Huang, Xipeng Qiu, Furu Wei, and Ming Zhou. 2018. [Reinforced mnemonic reader for machine reading comprehension](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 4099–4106.

Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.

Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2019. [SMART: robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization](#). *CoRR*, abs/1911.03437.

Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale Reading comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.

Xiaodong Liu, Hao Cheng, Pengcheng He, Weizhu Chen, Yu Wang, Hoifung Poon, and Jianfeng Gao. 2020. [Adversarial training for large neural language models](#). *CoRR*, abs/2004.08994.

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](#). *CoRR*, abs/1907.11692.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 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.

Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. 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*.

Lis Pereira, Xiaodong Liu, Fei Cheng, Masayuki Asahara, and Ichiro Kobayashi. 2020. [Adversarial training for commonsense inference](#). *CoRR*, abs/2005.08156.

- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Nan Shao, Yiming Cui, Ting Liu, Shijin Wang, and Guoping Hu. 2020. Is Graph Structure Necessary for Multi-hop Question Answering? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7187–7192, Online. Association for Computational Linguistics.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Chao Wang and Hui Jiang. 2018. Exploring machine reading comprehension with explicit knowledge. *CoRR*, abs/1809.03449.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Yi Wu, David Bamman, and Stuart Russell. 2017. Adversarial training for relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1778–1783, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380.
- Michihiro Yasunaga, Jungo Kasai, and Dragomir Radev. 2018. Robust multilingual part-of-speech tagging via adversarial training. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 976–986, New Orleans, Louisiana. Association for Computational Linguistics.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 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.