

EasyNLP: A Comprehensive and Easy-to-use Toolkit for Natural Language Processing

Chengyu Wang¹, Minghui Qiu^{1*}, Taolin Zhang^{1,2}, Tingting Liu^{1,2}, Lei Li^{1,2},
Jianing Wang^{1,2}, Ming Wang¹, Jun Huang¹, Wei Lin¹

¹ Platform of AI (PAI), Alibaba Group ² East China Normal University
{chengyu.wcy, minghui.qmh, huangjun.hj}@alibaba-inc.com

Abstract

Pre-Trained Models (PTMs) have reshaped the development of Natural Language Processing (NLP) and achieved significant improvement in various benchmarks. Yet, it is not easy for industrial practitioners to obtain high-performing PTM-based models without a large amount of labeled training data and deploy them online with fast inference speed. To bridge this gap, EasyNLP is designed to make it easy to build NLP applications, which supports a comprehensive suite of NLP algorithms. It further features knowledge-enhanced pre-training, knowledge distillation and few-shot learning functionalities, and provides a unified framework of model training, inference and deployment for real-world applications. EasyNLP has powered over ten business units within Alibaba Group and is seamlessly integrated to the Platform of AI (PAI) products on Alibaba Cloud. The source code of EasyNLP is released at GitHub (<https://github.com/alibaba/EasyNLP>).

1 Introduction

Pre-Trained Models (PTMs) such as BERT, GPT-3 and PaLM have achieved remarkable results in NLP. With the scale expansion of PTMs, the performance of NLP tasks has been continuously improved; thus, there is a growing trend of ultra-large-scale pre-training, pushing the scale of PTMs from millions, billions, to even trillions (Devlin et al., 2019; Brown et al., 2020; Chowdhery et al., 2022).

However, the application of large PTMs in industrial scenarios is still a non-trivial problem, with reasons as follows. i) Large PTMs are not always smarter and can make commonsense mistakes due to the lack of world knowledge (Petroni et al., 2019). Hence, it is highly necessary to make PTMs explicitly understand world facts by knowledge-enhanced pre-training, especially for supporting domain-specific applications. ii) Although large-scale PTMs have achieved good results with few

training samples, the problem of insufficient data and the huge size of models such as GPT-3 still restrict the usage of these models. Thus, few-shot fine-tuning BERT-style PTMs is more practical for online applications (Gao et al., 2021). iii) Last but not least, although large-scale PTMs have become an important part of the NLP learning pipeline, the slow training and inference speed seriously affects online applications that require higher QPS (Query Per Second) with limited computational resources.

To address these issues, we develop EasyNLP, an NLP toolkit that is designed to make the applications of large PTMs to industrial scenarios more efficiently and effectively. EasyNLP provides knowledge-enhanced pre-training functionalities to improve the knowledge understanding abilities of PTMs. Specifically, it integrates our DKPLM framework (Zhang et al., 2022) that enables the decomposition of knowledge-enhanced pre-training and task-specific learning. Hence, the resulting models can be tuned and deployed in the same way as BERT (Devlin et al., 2019). EasyNLP is equipped with a variety of popular prompt-based few-shot learning algorithms such as PET (Schick and Schütze, 2021) and P-Tuning (Liu et al., 2021b). Particularly, we propose a new few-shot learning paradigm named Contrastive Prompt Tuning (CP-Tuning) (Xu et al., 2022) that eases the manual labor of verbalizer construction based on contrastive learning. Finally, EasyNLP supports several knowledge distillation algorithms that compress large PTMs into small and efficient ones. Among them, the MetaKD algorithm (Pan et al., 2021) can significantly improve the effectiveness of the learned models with cross-domain datasets, which is particular common in industry.

Overall, our EasyNLP toolkit can provide users with large-scale and robust learning functionalities, and is seamlessly connected to the Platform of AI (PAI)¹ products. To demonstrate the useful

* Corresponding Author.

¹<https://www.alibabacloud.com/product/>

of EasyNLP, we also present the results of standard benchmarks and some real-world industrial scenarios to show how EasyNLP brings substantial improvements to these applications.

In a nutshell, the main features of the EasyNLP toolkit include the following aspects:

- **Easy-to-use and highly customizable.** In addition to providing easy-to-use commands to call cutting-edge NLP models, EasyNLP abstracts customized modules such as AppZoo and ModelZoo to make it easy to build NLP applications. It also features DataHub that provides users with a simple interface to load and process various types of NLP datasets.
- **Compatible with open-source community.** EasyNLP has rich APIs to support the training of models from other open-source libraries such as Huggingface/Transformers² with the PAI’s distributed learning framework. It is also compatible with the PTMs in EasyTransfer ModelZoo³ (Qiu et al., 2021).
- **Product-ready support.** EasyNLP is seamlessly integrated to PAI products on Alibaba Cloud to provide full model training and serving experience, including PAI-DSW for model development, PAI-DLC for cloud-native training, PAI-EAS for online serving, and PAI-Designer for zero-code model training.
- **Pre-training knowledge-enhanced PTMs.** EasyNLP also is equipped with knowledge-enhanced PTMs of various domains. Its pre-training APIs enable users to obtain customized PTMs using their own knowledge bases with just a few lines of codes.
- **Deploying large-scale PTMs.** EasyNLP provides few-shot learning capabilities based on prompts, allowing users to fine-tune large-scale PTMs with only a few training samples to achieve good results. Meanwhile, it provides knowledge distillation functionalities to help quickly distill large models to small and efficient models for online deployment.

2 Related Work

In this section, we summarize the related work on PTMs, prompt learning and knowledge distillation.

machine-learning

²<https://github.com/huggingface/transformers>

³<https://github.com/alibaba/EasyTransfer>

Pre-trained Language Models. PTMs have achieved significant improvements on various tasks by self-supervised pre-training (Qiu et al., 2020). To name a few, BERT (Devlin et al., 2019) learns bidirectional contextual representations by transformer encoders. Other transformer encoder-based PTMs include Transformer-XL (Dai et al., 2019), XLNet (Yang et al., 2019) and many others. The encoder-decoder and auto-regressive decoder architectures are used in T5 (Raffel et al., 2020) and GPT-3 (Brown et al., 2020). Knowledge-enhanced PTMs (Zhang et al., 2019; Liu et al., 2020; Sun et al., 2020) improve language understanding abilities of PTMs via injecting relational triples extracted from knowledge bases.

Prompt Learning for PTMs. Prompt learning models the probability of texts directly as the model prediction results based on language models (Liu et al., 2021a). In the literature, PET (Schick and Schütze, 2021) models NLP tasks as cloze problems and maps the results of the masked language tokens to class labels. Gao et al. (2021) generates discrete prompts from T5 (Raffel et al., 2020) to support prompt discovery. P-Tuning (Liu et al., 2021b) learns continuous prompt embeddings with differentiable parameters. Our CP-Tuning (Xu et al., 2022) optimizes the output results based on contrastive learning, without defining mappings from outputs to class labels.

Knowledge Distillation. Knowledge distillation aims at learning a smaller model from an ensemble or a larger model (Hinton et al., 2015). For large-scale PTMs, DistillBERT (Sanh et al., 2019) and PKD (Sun et al., 2019) applies the distillation loss in the pre-training and fine-tuning stages, separately. TinyBERT (Jiao et al., 2020a) further distills BERT in both stages, considering various types of signals. Due to space limitation, we do not further elaborate other approaches. Our MetaKD method (Pan et al., 2021) further improves the accuracy of the student models by exploiting cross-domain transferable knowledge, which is fully supported by EasyNLP.

3 The EasyNLP Toolkit

In this section, we introduce various aspects of our EasyNLP toolkit in detail.

3.1 Overview

We begin with an overview of EasyNLP in Figure 1. EasyNLP is built upon PyTorch and supports rich

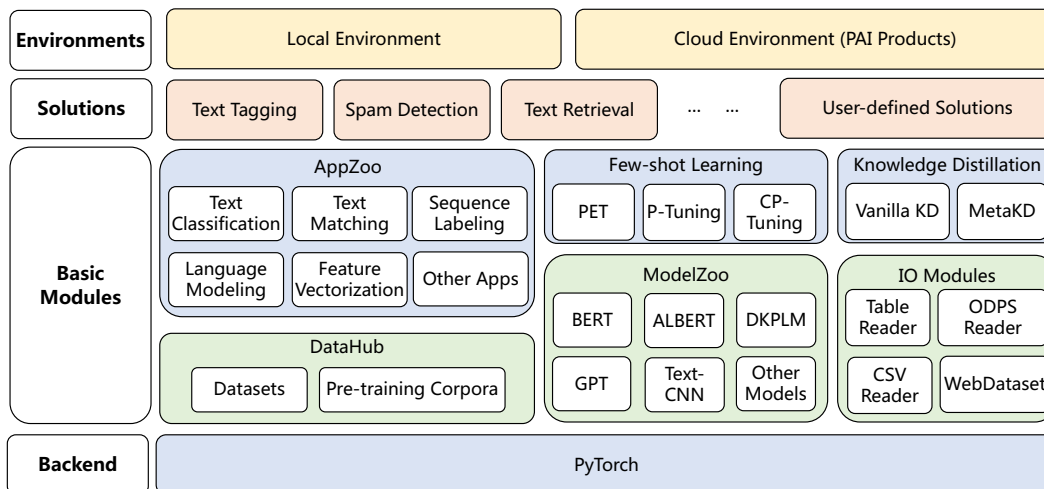


Figure 1: An overview of the EasyNLP toolkit.

data readers to process data from multiple sources. Users can load any PTMs from ModelZoo and datasets from DataHub, build their applications from AppZoo, or explore its advanced functionalities such as knowledge-enhanced pre-training, knowledge distillation and few-shot learning. The codes can run either in local environments or PAI’s products on the cloud. Users can also explore various solutions on our platform to support real-world applications. In addition, all EasyNLP’s APIs are also released to make it easy for users to customize any kinds of NLP applications.

3.2 DataHub, ModelZoo and AppZoo

DataHub. DataHub provides users with an interface to load and process various kinds of data. It is compatible with Huggingface datasets⁴ as a built-in library that supports unified interface calls and contains datasets of a variety of tasks. Some examples are listed in Table 1. Users can load the required data by specifying the dataset name through the `load_dataset` interface, and then use the `GeneralDataset` interface to automatically process the data into model input. An example of loading and pre-processing the TNEWS dataset, together with its subsequent steps, is shown in Code 1. For user-defined datasets, it is also straightforward to inherit the `GeneralDataset` class to customize the data format.

ModelZoo. PTMs such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) and T5 (Raffel et al., 2020) greatly improve the performance of NLP tasks. To facilitate user deployment of models, ModelZoo provides general pre-trained models

⁴<https://github.com/huggingface/datasets>

Task Type	Example of Datasets
Sequence Classification	TNEWS ⁵ , SogouCA ⁶
Text Generation	THUCNews ⁷ , SogouCS ⁸
Few-shot / Zero-shot Learning	BUSTM ⁹ , CHID ¹⁰
Knowledge-based NLU	OntoNotes ¹¹ , SanWen ¹²

Table 1: A partial list of datasets in EasyNLP DataHub.

```

from easynlp.dataset import load_dataset,
    GeneralDataset

# load dataset
dataset = load_dataset('clue', 'tnews')['train']
# parse data into classification model input
encoded = GeneralDataset(dataset, 'chinese-bert-base')
# load model
model = SequenceClassification('chinese-bert-base')
trainer = Trainer(model, encoded)
# start to train
trainer.train()

```

Code 1: Load the TNEWS training set and build a text classification application using EasyNLP.

as well as our own models for users to use, such as DKPLM (Zhang et al., 2022) of various domains. A few widely-used non-PTM models are also supported, such as Text-CNN (Kim, 2014).

AppZoo. To help users build NLP applications more easily with our framework, we further provide a comprehensive NLP application tool named AppZoo. It supports running applications with a few command-line arguments and provides a vari-

⁵<https://github.com/CLUEbenchmark/CLUE>

⁶<http://www.sogou.com/labs/resource/ca.php>

⁷<http://thuctc.thunlp.org/>

⁸<https://www.sogou.com/labs/resource/cs.php>

⁹<https://github.com/xiaobu-coai/BUSTM>

¹⁰<https://github.com/chujiezheng/ChID-Dataset>

¹¹<https://catalog.ldc.upenn.edu/LDC2013T19>

¹²<https://github.com/lancopku/Chinese-Literature-NER-RE-Dataset>

```

easynlp \
  --mode=train \
  --worker_gpu=1 \
  --tables=train.tsv,dev.tsv \
  --input_schema=sent:str:1,label:str:1 \
  --first_sequence=sent \
  --label_name=label \
  --label_enumerate_values=0,1 \
  --checkpoint_dir=./classification_model \
  --epoch_num=1 \
  --sequence_length=128 \
  --app_name=text_classify \
  --user_defined_parameters=
  'pretrain_model_name_or_path=bert-small-uncased'

```

Code 2: AppZoo for training a BERT-based text classifier using EasyNLP.

ety of mainstream or innovative NLP applications for users. AppZoo provides rich modules for users to build different application pipelines, including language modeling, feature vectorization, sequence classification, text matching, sequence labeling and many others. An example of training a text classifier using AppZoo is shown in Code 2.

3.3 In-house Developed Algorithms

In this section, we introduce in-house developed algorithms in EasyNLP. All these algorithms have been tested in real-world applications.

3.3.1 Knowledge-enhanced Pre-training

Knowledge-enhanced pre-training improves the performance of PTMs by injecting the relational facts from knowledge bases. Yet, a lot of existing works require additional knowledge encoders during pre-training, fine-tuning and inference (Zhang et al., 2019; Liu et al., 2020; Sun et al., 2020).

The proposed DKPLM paradigm (Zhang et al., 2022) decomposes the knowledge injection process. For DKPLM, knowledge injection is only applied during pre-training, without introducing extra parameters as knowledge encoders, alleviating the significant computational burdens for users. Meanwhile, during fine-tuning and inference stages, our model can be utilized in the same way as that of BERT (Devlin et al., 2019) and other plain PTMs, which facilitates the model fine-tuning and deployment in EasyNLP and other environments. Specifically, the DKPLM framework introduces three novel techniques for knowledge-enhanced pre-training. It recognizes long-tail entities from text corpora for knowledge injection only, avoiding learning too much redundant and irrelevant information from knowledge bases (Zhang et al., 2021). Next, the representations of entities are replaced by “pseudo token representations” derived from

knowledge bases, without introducing any extra parameters to DKPLM. Finally, a relational knowledge decoding task is introduced to force the model to understand what knowledge is injected.

In EasyNLP, we provide the entire pre-training pipeline of DKPLM for users. In addition, a collection of pre-trained DKPLMs for specific domains have been registered in ModelZoo for supporting domain-specific applications.

3.3.2 Few-shot Learning for PTMs

For low-resource scenarios, prompt-based learning leverages prompts as task guidance for effective few-shot fine-tuning. In EasyNLP, to facilitate easy few-shot learning, we integrate PET (Schick and Schütze, 2021) and P-Tuning (Liu et al., 2021b) into AppZoo that allow users call the algorithms in the similar way compared to standard fine-tuning.

It should be further noted that either PET or P-Tuning require the explicit handcraft of verbalizers, which is a tedious process and may lead to unstable results. Our CP-Tuning approach (Xu et al., 2022) enables few-shot fine-tuning PTMs without the manual engineering of task-specific prompts and verbalizers. A pair-wise cost-sensitive contrastive learning is introduced to achieve verbalizer-free class mapping by learning to distinguish different classes. Users can also explore CP-Tuning in AppZoo for any tasks that classical prompt-based methods support.

3.3.3 Knowledge Distillation for PTMs

The large model size and the long inference time hinder the deployment of large-scale PTMs to resource-constrained applications. In EasyNLP, we provide a complete learning pipeline for knowledge distillation, including data augmentation for training sets, logits extraction from teacher models and distilled training of student models.

In addition, we notice that a majority of existing approaches focus on a single domain only. The proposed MetaKD algorithm (Pan et al., 2021) explicitly leverages the cross-domain transferable knowledge to improve the accuracy of student models. It first obtain a meta-teacher model to capture transferable knowledge at both instance-level and feature-level from multiple domains. Next, a meta-distillation algorithm is employed to learn single-domain student models with selective signals from the meta-teacher. In EasyNLP, the MetaKD process is implemented as a general feature for any types of BERT-style PTMs.

PTM	AFQMC	CMNLI	CSL	IFLYTEK	OCNLI	TNEWS	WSC	Average
BERT-base	72.17	75.74	80.93	60.22	78.31	57.52	75.33	71.46
BERT-large	72.89	77.62	81.14	60.70	78.95	57.77	78.18	72.46
RoBERTa-base	73.10	80.75	80.07	60.98	80.75	57.93	86.84	74.35
RoBERTa-large	74.81	80.52	82.60	61.37	82.49	58.54	87.50	75.40
MacBERT-base	74.23	80.65	81.70	61.14	80.65	57.65	80.26	73.75
MacBERT-large	74.37	81.19	83.70	62.05	81.65	58.45	86.84	75.46

Table 2: CLUE performance of BERT, RoBERTa and MacBERT fine-tuned with EasyNLP (%).

PTM	MNLI	QNLI	QQP	RTE	SST-2	MRPC	CoLA	STSB	Average
BERT-base	84.8	91.4	91.1	68.3	92.5	88.1	55.3	89.6	82.6
BERT-large	86.6	92.4	91.2	70.8	93.4	88.2	61.1	90.1	84.2
RoBERTa-base	87.3	92.5	92.1	77.3	94.9	90.2	63.9	91.1	86.2
RoBERTa-large	90.1	94.5	92.3	87.1	96.4	91.0	67.8	92.3	88.9

Table 3: GLUE performance of BERT and RoBERTa fine-tuned with EasyNLP (%).

4 System Evaluations and Applications

In this section, we empirically examine the effectiveness and efficiency of the EasyNLP toolkit on both public datasets and industrial applications.

4.1 CLUE and GLUE Benchmarks

In order to validate the effectiveness of EasyNLP on model fine-tuning, we fine-tune PTMs on the CLUE and GLUE benchmarks (Wang et al., 2019; Xu et al., 2020). For all tasks, we use a limited hyper-parameter search space, with batch sizes in $\{8, 16, 32, 48\}$, sequence length in $\{128, 256\}$ and learning rates in $\{1e-5, 2e-5, 3e-5, 4e-5, 5e-5\}$. The underlying PTMs include BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). We also evaluate MacBERT (Cui et al., 2020) for the Chinese benchmark CLUE. We report the results over the development sets of each task in the two benchmarks, shown in Tables 2 and 3, respectively. The obtained comparable performance has shown the reliability of EasyNLP, which achieves similar performance compared to other open-source frameworks and their original implementations.

4.2 Evaluation of Knowledge-enhanced Pre-training

We report the performance of DKPLM over zero-shot knowledge probing tasks, including LAMA (Petroni et al., 2019) and LAMA-UHN (Pörner et al., 2019), with the results summarized in Table 4. Compared to strong baselines (i.e., CoLAKE (Sun et al., 2020) K-Adapter (Wang et al., 2021a) and KEPLER (Wang et al., 2021b)), we see that DKPLM achieves state-of-the-art results over three datasets (+1.57% on average). The re-

sult of DKPLM is only 0.1% lower than K-Adapter, without using any T-REx training data and larger backbones. We can see that our pre-training process based on DKPLM can effectively store and understand factual relations from knowledge bases.

Industrial Applications. Based on the proposed DKPLM framework (Zhang et al., 2022), we have pre-trained a series of domain-specific PTMs to provide model service inside Alibaba Group, such as medical and finance domains, and observed consistent improvement in downstream NLP tasks. For example, the medical-domain DKPLM improves the accuracy of a medical named entity recognition task by over 3%, compared to the standard BERT model (Devlin et al., 2019). The pre-trained model (named `pai-dkplm-medical-base-zh`) has also been released in our EasyNLP ModelZoo.

4.3 Evaluations of Few-shot Learning

We compare CP-Tuning (Xu et al., 2022) against several prompt-based fine-tuning approaches including PET (Schick and Schütze, 2021), LM-BFF (Gao et al., 2021) (in three settings where “Auto T”, “Auto L” and “Auto T+L” refer to the prompt-tuned PTM with automatically generated templates, label words and both, respectively) and P-Tuning (Liu et al., 2021b). The experiments are conducted over several text classification datasets in a 16-shot learning setting. The underlying PTM is RoBERTa (Liu et al., 2019). Readers can refer to Xu et al. (2022) for more details. From the results in Table 5, we can see that the performance gains of CP-Tuning over all the tasks are consistent, compared to state-of-the-art methods.

Industrial Applications. For business customer

Dataset	ELMo	BERT	RoBERTa	CoLAKE	K-Adapter*	KEPLER	DKPLM
Google-RE	2.2%	11.4%	5.3%	9.5%	7.0%	7.3%	10.8%
UHN-Google-RE	2.3%	5.7%	2.2%	4.9%	3.7%	4.1%	5.4%
T-REx	0.2%	32.5%	24.7%	28.8%	29.1%	24.6%	32.0%
UHN-T-REx	0.2%	23.3%	17.0%	20.4%	23.0%	17.1%	22.9%

Table 4: The performance on LAMA knowledge probing datasets. Note that K-Adapter is trained based on a large-scale model and uses a subset of T-REx as its training data.

Method	SST-2	MR	CR	MRPC	QQP	QNLI	RTE	SUBJ	Avg.
Standard Fine-tuning	78.62	76.17	72.48	64.40	63.01	62.32	52.28	86.82	69.51
PET	92.06	87.13	87.13	66.23	70.34	64.38	65.56	91.28	78.01
LM-BFF (Auto T)	90.60	87.57	90.76	66.72	65.25	68.87	65.99	91.61	78.42
LM-BFF (Auto L)	90.55	85.51	91.11	67.75	70.92	66.22	66.35	90.48	78.61
LM-BFF (Auto T+L)	91.42	86.84	90.40	66.81	61.61	61.89	66.79	90.72	77.06
P-tuning	91.42	87.41	90.90	71.23	66.77	63.42	67.15	89.10	78.43
CP-Tuning	93.35	89.43	91.57	72.60	73.56	69.22	67.22	92.27	81.24

Table 5: Comparison between CP-Tuning and baselines over the testing sets in terms of accuracy (%).

Method	Amazon	MNLI
BERT-s	87.9	81.9
BERT-mix	89.5	84.4
BERT-mtl	89.8	84.2
BERT-s \rightarrow TinyBERT	86.7	79.3
BERT-mix \rightarrow TinyBERT	87.3	79.6
BERT-mtl \rightarrow TinyBERT	87.7	79.7
MetaKD	89.4	80.4

Table 6: Evaluation of MetaKD over Amazon reviews and MNLI in terms of averaged accuracy (%).

service, it is necessary to extract the fine-grained attributes and entities from texts, which may involve a large number of classes with few training data available. By applying our algorithm in EasyNLP, the accuracy scores of entity and attribute extraction are improved by 2% and 5%. In addition, our few-shot toolkit produces the best performance on the FewCLUE benchmark (Xu et al., 2021).

4.4 Evaluations of Knowledge Distillation

We further report the performance of MetaKD (Pan et al., 2021) on Amazon reviews (Blitzer et al., 2007) and MNLI (Williams et al., 2018), where the two datasets contain four and five domain instances, respectively. In the experiments, we train the meta-teacher over multi-domain training sets, and distill the meta-teacher to each of all the domains. The teacher model is BERT-base (with 110M parameters), while the student model is BERT-tiny (with 14.5M parameters). Table 6 shows the performance of baselines and MetaKD, in terms of averaged ac-

curacy across domains. BERT-s refers to a single BERT teacher trained on each domain. BERT-mix is one BERT teacher trained on the mixture of all domain data. BERT-mtl is one teacher trained by multi-task learning over all domains. For distillation, “ \rightarrow TinyBERT” means using the KD method described in Jiao et al. (2020b) to distill the corresponding teacher model. The results show that MetaKD significantly reduces the model size while preserving a similar performance. For more details, we refer the readers to Pan et al. (2021).

Industrial Applications. Distilled PTMs have been widely used inside Alibaba Group due to the high QPS requirements of online e-commerce applications. For example, in the AliMe chatbot (Qiu et al., 2017), we distill the BERT-based query intent detection model from the base version to the tiny version, resulting in 7.2x inference speedup while the accuracy is only decreased by 1%.

5 Conclusion

In this paper, we introduced EasyNLP, a toolkit that is designed to make it easy to develop and deploy deep NLP applications based on PTMs. It supports a comprehensive suite of NLP algorithms and features knowledge-enhanced pre-training, knowledge distillation and few-shot learning functionalities for large-scale PTMs. Currently, EasyNLP has powered a number of business units inside Alibaba Cloud and provided NLP service on the cloud. The toolkit has been open-sourced to promote research and development for NLP applications.

Broader Impact

EasyNLP is a comprehensive toolkit for building various NLP applications to support industrial scenarios. It has been seamlessly integrated into the PAI products, and has been released to the open-source community. EasyNLP is also highly beneficial for academia, as it integrates state-of-the-art methods and models to make it easy for researchers to benchmark and develop their own algorithms.

Acknowledgments

The authors would like to thank Haojie Pan, Peng Li, Boyu Hou, Xiaoqing Chen, Xiaodan Wang, Xiangu Zhu and many other members of the Alibaba PAI team for their contribution and suggestions on building the EasyNLP toolkit. This work is also partially supported by Alibaba Group through Alibaba Innovative Research Program and Alibaba Research Intern Program.

References

- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *NeurIPS*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. Palm: Scaling language modeling with pathways. *CoRR*, abs/2204.02311.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. Revisiting pre-trained models for chinese natural language processing. In *EMNLP (Findings)*, pages 657–668.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. In *ACL*, pages 2978–2988.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners. In *ACL/IJCNLP*, pages 3816–3830.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020a. Tinybert: Distilling BERT for natural language understanding. In *EMNLP (Findings)*, pages 4163–4174.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020b. Tinybert: Distilling BERT for natural language understanding. In *EMNLP (Findings)*, pages 4163–4174.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP*, pages 1746–1751.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *CoRR*, abs/2107.13586.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-BERT: enabling language representation with knowledge graph. In *AAAI*, pages 2901–2908.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. GPT understands, too. *CoRR*, abs/2103.10385.
- 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.

- Haojie Pan, Chengyu Wang, Minghui Qiu, Yichang Zhang, Yaliang Li, and Jun Huang. 2021. Meta-kd: A meta knowledge distillation framework for language model compression across domains. In *ACL/IJCNLP*, pages 3026–3036.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. Language models as knowledge bases? In *EMNLP-IJCNLP*, pages 2463–2473.
- Nina Pörner, Ulli Waltinger, and Hinrich Schütze. 2019. BERT is not a knowledge base (yet): Factual knowledge vs. name-based reasoning in unsupervised QA. *CoRR*, abs/1911.03681.
- Minghui Qiu, Feng-Lin Li, Siyu Wang, Xing Gao, Yan Chen, Weipeng Zhao, Haiqing Chen, Jun Huang, and Wei Chu. 2017. Alime chat: A sequence to sequence and rerank based chatbot engine. In *ACL*, pages 498–503.
- Minghui Qiu, Peng Li, Chengyu Wang, Haojie Pan, Ang Wang, Cen Chen, Xianyan Jia, Yaliang Li, Jun Huang, Deng Cai, and Wei Lin. 2021. Easytransfer: A simple and scalable deep transfer learning platform for NLP applications. In *CIKM*, pages 4075–4084.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *CoRR*, abs/2003.08271.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108.
- Timo Schick and Hinrich Schütze. 2021. Exploiting cloze-questions for few-shot text classification and natural language inference. In *EACL*, pages 255–269.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. Patient knowledge distillation for BERT model compression. In *EMNLP-IJCNLP*, pages 4322–4331.
- Tianxiang Sun, Yunfan Shao, Xipeng Qiu, Qipeng Guo, Yaru Hu, Xuanjing Huang, and Zheng Zhang. 2020. Colake: Contextualized language and knowledge embedding. In *COLING*, pages 3660–3670.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. 2021a. K-adapter: Infusing knowledge into pre-trained models with adapters. In *ACL*, pages 1405–1418.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021b. KEPLER: A unified model for knowledge embedding and pre-trained language representation. *Trans. Assoc. Comput. Linguistics*, 9:176–194.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL-HLT*, pages 1112–1122.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. CLUE: A chinese language understanding evaluation benchmark. In *COLING*, pages 4762–4772.
- Ziyun Xu, Chengyu Wang, Peng Li, Yang Li, Ming Wang, Boyu Hou, Minghui Qiu, Chengguang Tang, and Jun Huang. 2021. When few-shot learning meets large-scale knowledge-enhanced pre-training: Alibaba at fewclue. In *NLPCC*, pages 422–433.
- Ziyun Xu, Chengyu Wang, Minghui Qiu, Fuli Luo, Runxin Xu, Songfang Huang, and Jun Huang. 2022. Making pre-trained language models end-to-end few-shot learners with contrastive prompt tuning. *CoRR*, abs/2204.00166.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, pages 5754–5764.
- Ningyu Zhang, Shumin Deng, Xu Cheng, Xi Chen, Yichi Zhang, Wei Zhang, and Huajun Chen. 2021. Drop redundant, shrink irrelevant: Selective knowledge injection for language pretraining. In *IJCAI*, pages 4007–4014.
- Taolin Zhang, Chengyu Wang, Nan Hu, Minghui Qiu, Chengguang Tang, Xiaofeng He, and Jun Huang. 2022. DKPLM: decomposable knowledge-enhanced pre-trained language model for natural language understanding. In *AAAI*.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. ERNIE: enhanced language representation with informative entities. In *ACL*, pages 1441–1451.