MatchZoo: A Toolkit for Deep Text Matching

Yixing Fan, Liang Pang, JianPeng Hou, Jiafeng Guo, Yanyan Lan, Xueqi Cheng

CAS Key Lab of Network Data Science and Technology, Institute of Computing Technology,

Chinese Academy of Sciences

Beijing, China

{fanyixing,pangliang,houjianpeng}@software.ict.ac.cn,{guojiafeng,lanyanyan,cxq}@ict.ac.cn

ABSTRACT

In recent years, deep neural models have been widely adopted for text matching tasks, such as question answering and information retrieval, showing improved performance as compared with previous methods. In this paper, we introduce the MatchZoo toolkit that aims to facilitate the designing, comparing and sharing of deep text matching models. Specifically, the toolkit provides a unified data preparation module for different text matching problems, a flexible layer-based model construction process, and a variety of training objectives and evaluation metrics. In addition, the toolkit has implemented two schools of representative deep text matching models, namely representation-focused models and interaction-focused models. Finally, users can easily modify existing models, create and share their own models for text matching in MatchZoo.

CCS CONCEPTS

•**Information systems → Open source software;**

KEYWORDS

matching; deep learning; toolkit

1 INTRODUCTION

Many natural language processing (NLP) tasks can be formulated as a matching problem between two texts, such as paraphrase identification, question answering, and information retrieval (IR). In recent years, a number of deep neural models have been designed and applied to such kind of tasks, e.g., DSSM [5], MatchPyramid [6] and DRMM [3], leading to state-of-the-art performances as compared with previous methods. However, these deep text matching models are often designed with special layers, implemented in different environments, and targeted to different matching tasks, making it difficult for reproduction and comparison.
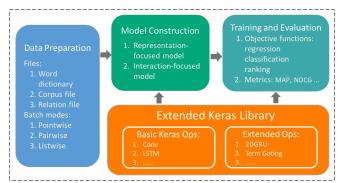
Figure 1: The Architecture of the MatchZoo toolkit.

In this paper, we introduce the MatchZoo toolkit that aims to facilitate the designing, comparing and sharing of deep text matching models. Firstly, the toolkit provides a data preparation module to convert dataset of different text matching tasks into a unified format so that they can be used as the input of different deep matching models. Secondly, the toolkit allows users to easily construct a deep matching model layer by layer based on the Keras [2] libarary. We have extended the Keras libarary to include layer interfaces that are specifically designed for text matching problems. Moreover, the toolkit has implemented two schools of representative deep text matching models, namely representation-focused models and interaction-focused models, making it an off-the-shelf model repository. Users can directly apply these models or modify them by simple configuration. Finally, the toolkit provides a variety of training objectives for optimization, as well as a set of evaluation metrics for comparison between different deep matching models. The MatchZoo toolkit is licensed under the permissive Apache license 2.0[1].

2 THE MATCHZOO TOOLKIT

The architecture of the MatchZoo toolkit is depicted in Figure 1. There are three major modules in the toolkit, namely data preparation, model construction, training and evaluation, respectively. These three modules are actually organized as a pipeline of data flow. The toolkit is built upon the Keras library, so that it can run over TensorFlow [1], CNTK[11] and Theano[2], and work seamlessly with CPUs and GPUs. We will introduce each of the three modules in more details in the following.

---

[1]The toolkit is licensed under the permissive Apache open-source license and can be found at https://github.com/faneshion/MatchZoo

[2]http://deeplearning.net/software/theano/

2.1 Data Preparation

The data preparation module aims to convert dataset of different text matching tasks into a unified format as the input of deep matching models. Users provide datasets which contain pairs of texts along with their labels, and the module produces the following files.

- **Word Dictionary** records the mapping from each word to a unique identifier called *wid*. Words that are too frequent (e.g. stopwords), too rare or noisy (e.g. fax numbers) can be filtered out by predefined rules.
- **Corpus File** records the mapping from each text to a unique identifier called *tid*, along with a sequence of word identifiers contained in that text. Note here each text is truncated or padded to a fixed length customized by users.
- **Relation File** is used to store the relationship between two texts, each line containing a pair of *tid*s and the corresponding label.

Beyond these basic files, users are also allowed to provide additional feature files, e.g., word embedding files or human-crafted feature files, as model input. After converting the raw dataset to the above unified format, the module provides three types of data batch modes, i.e., generating a batch of data in pointwise, pairwise or listwise manner.

2.2 Model Construction

In the model construction module, we employ Keras libarary to help users build the deep matching model layer by layer conveniently. The Keras libarary provides a set of common layers widely used in neural models, such as convolutional layer, pooling layer, dense layer and so on. To further facilitate the construction of deep text matching models, we extend the Keras libarary to provide some layer interfaces specifically designed for text matching. We list a few examples here.

- **Matching_Matrix** layer builds a word-by-word matching matrix based on dot product, cosine similarity or indicator function [6].
- **Term_Gating** layer builds a term gating layer based on word embeddings from texts [3].
- **2D-GRU** layer builds a two dimensional recurrent neural network using the Gated Recurrent Unit (GRU) [10] .

Moreover, the toolkit has implemented two schools of representative deep text matching models, namely representation-focused models and interaction-focused models [3].

- **Representation-based models** include ARC-I [4], DSSM [5], CDSSM [8], MV-LSTM [9], and CNTN [7];
- **Interaction-based models** include DRMM [3], ARC-II [4], MatchPyramid[6], and Match-SRNN [10];

Users can apply these models out-of-the-box or modify them by simple configuration. Moreover, as a shared model repository, users can contribute their own matching models into this toolkit easily.

2.3 Training and Evaluation

For learning the deep matching models, the toolkit provides a variety of objective functions for regression, classification and ranking. For example, the ranking-related objective functions include several well-known pointwise, pairwise and listwise losses. It is flexible for users to pick up different objective functions in the training phase for optimization. Once a model has been trained, the toolkit could be used to produce a matching score, predict a matching label, or rank target texts (e.g., a document) against an input text.

For evaluation, the toolkit provides several widely adopted evaluation metrics, such as Precision, MAP, NDCG and so on. For IR tasks, the toolkit can also output the ranking results as a TREC-compatible file that can be used as the input to the trec_eval[3] evaluation utitlity.

3 CONCLUSIONS

In this paper, we introduce the MatchZoo toolkit which helps the design, comparison and share of deep text matching models. The toolkit contains three major modules, namely data preparation, model construction, training and evaluation. Meanwhile, the toolkit provides implementations of state-of-the-art deep matching models on NLP and IR related tasks. Users can easily modify these off-the-shelf models or implement their own models by using the layer interfaces provided by the toolkit.

For the future work, we would like to continuously extend the model repository and layer interfaces in the MatchZoo toolkit for better usage. It would also be interesting to accommodate other types of data beyond text for new matching scenarios.

REFERENCES

[1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. (2015). http://tensorflow.org/ Software available from tensorflow.org.
[2] François Chollet and others. 2015. Keras. https://github.com/fchollet/keras. (2015).
[3] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. 2016. A deep relevance matching model for ad-hoc retrieval. In *CIKM*. ACM, 55–64.
[4] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *NIPS*. 2042–2050.
[5] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *CIKM*. ACM, 2333–2338.
[6] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text Matching as Image Recognition. In *AAAI*.
[7] Xipeng Qiu and Xuanjing Huang. 2015. Convolutional Neural Tensor Network Architecture for Community-Based Question Answering.. In *IJCAI*. 1305–1311.
[8] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. Learning semantic representations using convolutional neural networks for web search. In *WWW*. 373–374.
[9] Shengxian Wan, Yanyan Lan, Jiafeng Guo, Jun Xu, Liang Pang, and Xueqi Cheng. 2016. A deep architecture for semantic matching with multiple positional sentence representations. In *Thirtieth AAAI Conference on Artificial Intelligence*.
[10] Shengxian Wan, Yanyan Lan, Jun Xu, Jiafeng Guo, Liang Pang, and Xueqi Cheng. 2016. Match-SRNN: Modeling the Recursive Matching Structure with Spatial RNN. *arXiv preprint arXiv:1604.04378* (2016).
[11] Dong Yu, Adam Eversole, Mike Seltzer, Kaisheng Yao, Oleksii Kuchaiev, Yu Zhang, Frank Seide, Zhiheng Huang, Brian Guenter, Huaming Wang, Jasha Droppo, Geoffrey Zweig, Chris Rossbach, Jie Gao, Andreas Stolcke, Jon Currey, Malcolm Slaney, Guoguo Chen, Amit Agarwal, Chris Basoglu, Marko Padmilac, Alexey Kamenev, Vladimir Ivanov, Scott Cypher, Hari Parthasarathi, Bhaskar Mitra, Baolin Peng, and Xuedong Huang. 2014. *An Introduction to Computational Networks and the Computational Network Toolkit*. Technical Report. https://www.microsoft.com/en-us/research/publication/an-introduction-to-computational-networks-and-the-computational-network-toolkit/

---

[3]https://github.com/usnistgov/trec_eval