

QADiver: Interactive Framework for Diagnosing QA Models

Gyeongbok Lee*
Yonsei University and
Clova AI Research, NAVER Corp.
alias_n@yonsei.ac.kr

Sungdong Kim
Clova AI Research, NAVER Corp.
sungdong.kim@navercorp.com

Seung-won Hwang
Yonsei University
seungwonh@yonsei.ac.kr

Abstract

Question answering (QA) extracting answers from text to the given question in natural language, has been actively studied and existing models have shown a promise of outperforming human performance when trained and evaluated with SQuAD dataset. However, such performance may not be replicated in the actual setting, for which we need to diagnose the cause, which is non-trivial due to the complexity of model. We thus propose a web-based UI that provides how each model contributes to QA performances, by integrating visualization and analysis tools for model explanation. We expect this framework can help QA model researchers to refine and improve their models.

Introduction

With the help of large-scale question answering datasets and deep learning frameworks released to the public, question answering (QA) models have been improved rapidly as community efforts. Recently, the reported accuracy of state-of-the-art models has exceeded human performance in SQuAD task, where plausible answers are extracted for the given question and context document pair. However, it is still non-trivial to reproduce such accuracy reported in the paper in production settings.

This paper proposes a diagnosing tool, for troubleshooting the performance gap. For example, is the performance lower than expected because of the biased training to certain types of questions or texts? Is the model attending to wrong words for question or text understanding? Is the embedding suitable for the given task? At the same time, there is a concerning observation that these models can be easily perturbed by a simple adversarial example added (Jia and Liang 2017). A desirable tool may support developers to easily perturb training to identify such vulnerability.

We will demonstrate **QADiver**, a data-centric diagnosing framework for QA model, with diverse interactive visualization and analysis tools for a full pipeline of the attention-based QA model. The framework is connected with the target model to retrieve answer span prediction, inner-model

*This work was done during an internship at Clova AI Research, NAVER.
Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

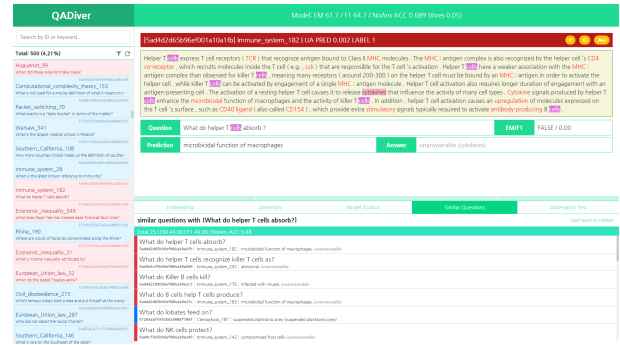


Figure 1: Snapshot of QADiver Framework

values (such as attention), and no answer probability for given context and question from the model.

Our framework targets SQuAD 2.0 dataset (Rajpurkar, Jia, and Liang 2018), a machine reading comprehension benchmark that contains both answerable and unanswerable question for given context. In this task, not only finding plausible answer span in context but predicting question is answerable or not is also crucial. By exploring large question-answer instances, we expect developers can make a better diagnosis and find insights, than those made from qualitative observations of a few instances. The demonstration video for the framework is available in <https://youtu.be/V6c8nls6Qcc>.

Key Features

Figure 1 shows a snapshot of our system, where the cause of low accuracy can be diagnosed as one or more of the following problems:

Dataset Bias

Left-sidebar in Figure 1 shows data instances form SQuAD 2.0 development set instances. To support a quick exploration of how such instances are distributed for answers predicted right or wrong, we color each to blue and red respectively. This color code is used for the entire system. The top right area in Figure 1 shows more detailed statistics about the instance selected: a question, corresponding context, gold answers and prediction result from the model

including unanswerable probability and EM/F1 scores. Answer span and out-of-vocabulary words are highlighted. The user can also switch between the original text and preprocessed text by the model for both context and question to identify cases where the underlying NLP models, such as tokenizer, are not working as intended.

Embedding Analysis

QA model performance also depends on the effectiveness of word embedding. To visualize such effectiveness, we allow developers to visualize word vector (and nearby words in the space by clicking a word. The user can restrict words in the context, or use the whole vocabulary set for the dataset. For a fast similarity search of dense vectors, we use FAISS (Johnson, Douze, and Jégou 2017) library.

Neural Model Internals

Visual representation for internal states of the neural model shows whether right words are highlighted for understanding questions or finding answers. The most common example is visualizing attention and output layer in the model. For attention layer, we visualize a context-question attention matrix heatmap for the given instance. Similarly, we provide an interpreted version of the model output used for answer span prediction and answerability decision by listing the top-k words with the highest weights. Users can also see the list of answer span candidates (including “unanswerable” case) and its certainty as a colored heatmap.

Question Bias

A model can be biased to answer a certain type of questions particularly well. To diagnose such a case, we identify similar questions to the given instance, labeled with prediction result and evaluation metrics (EM/F1). For a desirable question embedding, projecting similar question types close in the embedding space, we use well-studied features from the question and gold answer such as: answer length, the existence of number and entity, and 2-word question prefix like *What is* and *How many*. Feature values of questions in the same class are mean-aggregated to generate a global statistic vector. To represent the characteristic of each question, local features like word match ratio between context and question and one-hot vector for frequent words are used so that similar types of questions have high similarity. As each question is vectorized, top similar questions for the instance can be retrieved from the whole dataset by the similarity search.

Adversarial Test

As we overviewed, many existing QA models are reportedly lacking the robustness over adversarial examples. Using our tool, the user can easily perform the adversarial test for each instance in two ways: manual modification and rule-based test. First, the user can modify some words in context document and question from the data viewer by double-clicking target word and replacing to another. After this edit, we show an updated prediction and EM/F1 score.

Instead of perturbing each word, which may be costly, the user may create reusable adversarial rules, about word and

its part-of-speech (POS) tag. We use NLTK toolkit (Bird, Klein, and Loper 2009) for POS tagging and word tokenization used in rule matching. We also provide pre-defined adversarial rules from SEAR (Ribeiro, Singh, and Guestrin 2018) for those wanting to check the robustness for common cases.

Related Work

Due to the growing demand for model interpretability, many visualization tools for QA models were proposed: AllenNLP machine comprehension demo¹ provides answer span highlights and attention matrix visualization. (Rücklé and Gurevych 2017) shows attention visualization on context and question text, and provides a comparison between the two model. (Liu et al. 2018) proposes bipartite graph attention representation and hierarchical visual for highly asymmetric attention. This tool also supports word- and attention-level perturbation by user edit.

Acknowledgements

This work was supported by the ICT R&D program of MSIT/IITP. [No.2017-0-01778, Development of Explainable Human-level Deep Machine Learning Inference Framework]

References

- Bird, S.; Klein, E.; and Loper, E. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Jia, R., and Liang, P. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on EMNLP*.
- Johnson, J.; Douze, M.; and Jégou, H. 2017. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*.
- Liu, S.; Li, T.; Li, Z.; Srikumar, V.; Pascucci, V.; and Bremer, P. T. 2018. Visual interrogation of attention-based models for natural language inference and machine comprehension. In *Proceedings of the 2018 Conference on EMNLP: System Demonstrations*.
- Rajpurkar, P.; Jia, R.; and Liang, P. 2018. Know what you don’t know: Unanswerable questions for squad. *Proceedings of the 56th Annual Meeting of the ACL*.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2018. Semantically equivalent adversarial rules for debugging nlp models. In *Proceedings of the 56th Annual Meeting of the ACL*, volume 1, 856–865.
- Rücklé, A., and Gurevych, I. 2017. End-to-end non-factoid question answering with an interactive visualization of neural attention weights. In *Proceedings of the 55th Annual Meeting of the ACL-System Demonstrations (ACL 2017)*, 19–24. Association for Computational Linguistics.

¹<http://demo.allennlp.org/machine-comprehension>