Building a QA System (Robust QA track)

Ammaar Adam, Coleman Smith, Takao Yatagai

CS224N, Computer Science, Stanford University

Stanford
Computer Science

Problem

- QA problem in the NLP community has gained much attention in recent years.
- An effective QA system would be extremely useful in our daily lives.
- Examples are smart assistants and for search engines.
- Difficulties of the QA system is that the question and answers can span so many different domains.
- This variability in the domains makes the problem quite difficult to tackle.
- The RobustQA track emulates this challenge by having to perform well on out-of-domain datasets, given many in-domain samples, and a very small set of out-of-domain samples.

Background

- We want to investigate the effects that meta-learning has on improving model accuracy in the context of a question answering system.
- Work based on findings from [1] and [2].
- DistilBERT baseline model is not able to adapt to out-of-domain texts.
- We try to use meta-learning to get the model to better adapt to new domains.
- We were able to contribute a codebase and results.
- This allowed us to test our meta-learning methods - using DistilBERT and two classification head - on its ability to adapt to out-of-domain questions.

Stanford
University

Results

Method	Dev F1	Dev EM	Test F1	Test EM
In Domain Datasets				
Baseline	70.18	54.30	N/A	N/A
LinearLearner	65.01	48.07	N/A	N/A
MLPLearner	68.57	52.15	N/A	N/A
Out of Domain Datasets				
Baseline	47.57	31.41	N/A	N/A
LinearLearner	42.55	24.87	55.063	35.872
MLPLearner	44.77	29.32	58.331	39.358

Analysis

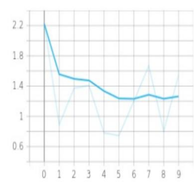

Figure 2: Loss curve for Meta-Learning Distilbert Model + LinearLayer [Loss over epochs]

Figure 3: Loss curve for Meta-Learning Distilbert Model + MLP [Loss over epochs]

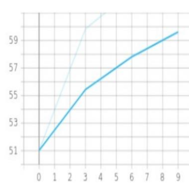

Figure 4: In-Domain F1 Score for Meta-Learning Distilbert Model + LinearLearner [F1 score over epochs]

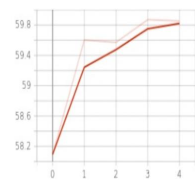

Figure 5: In-Domain F1 Score for Meta-Learning Distilbert Model + MLP [F1 score over epochs]

Methods

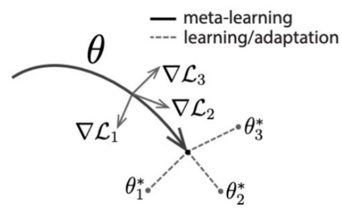

Figure 1: Diagram of the model-agnostic meta-learning algorithm (MAML), which optimizes for a representation θ that can quickly adapt to new tasks. Taken from [2].

- MAML mainly consists of two steps:
  1. We use the support set as examples from the new task, to adapt our model to this new domain. The adaptation is shown as θi in the diagram above.
  2. Once our model has adapted to the new task, we then see how well the adapted model performs on the query set, another set of samples from the new task.
- We perform forward propagation on the query set
- Then we calculate the loss we get from the query set of this new task.
- Key point behind meta-learning is that we aggregate the losses we get from adapting and evaluating on different domains and use the aggregated losses for back-propagation.

Conclusions

- A meta-learning approach to the generalized question-answer task clearly shows potential for improving upon more well-established techniques in the field.
- Meta-learning during training could be a viable way for the model to adapt to different domains better.
- The main bottleneck of the project has been the compute resources and time.

References

1. Arzoo Katiyar Kilian Q. Weinberger Tianyi Zhang, Felix Wu and Yoav Artzi. Revisiting few- sample bert fine-tuning. In The International Conference on Learning Representations (ICLR), 2021.
2. Chelsea Finn, Pieter Abbeel, and Sergey Levine. Modeal-agnostic meta-learning for fast adaptation of deep networks. In arXiv preprint arXiv:1703.03400, 2017