

Prompt-tuning in ASR systems for efficient domain-adaptation

Saket Dingliwal* Ashish Shenoy* Sravan Bodapati
Ankur Gandhe Ravi Teja Gadde Katrin Kirchhoff
{skdin, ashenoy, sravanb, aggandhe, gadderav, katrinki }@amazon.com

1 Introduction

Automatic Speech Recognition (ASR) systems form a key component of various products across industry. Many of these ASR systems rely on a complex Acoustic Model (AM) whose output is rescored by a domain-specific Language Model (LM). As we use ASR systems in new domains, the memory, maintenance and data-collection costs for these domain-specific LMs increase. Particularly, with advent of parameter-heavy Transformer based LMs (Devlin et al., 2019), maintaining multiple domain-specific LMs is practically infeasible. While on the other hand, using a generic LM for all domains falls short in performance when compared to multiple domain-specific LMs. Therefore, a need for a middle ground between performance and costs is evident.

To overcome this problem, we bring forward a methodology based on recently proposed Prompt Tuning. Lester et al. (2021) introduced this idea of learning the token embeddings of the prompt used to prime a LM to a particular task. Prompts are special tokens describing a task which when appended to the input data sample, helps the model understand and use this problem description to better solve the task. For example, to solve the machine translation task, instead of fine-tuning the Transformer model with corresponding dataset, one can achieve comparable performance by just showing text describing machine translation to the powerful Transformer-based LM. In prompt tuning, instead of providing this prompt manually to the model, one learn it from the labelled examples from the task.

To the best of our knowledge, we are the first

*equal contribution

Appearing at WeNLP (West Coast NLP) Summit 2021

one to apply prompt-tuning for domain-adaptation of ASR systems. We find that one can learn the prompt-embeddings for not only different tasks but also different domains of the same task. This means there exists a domain description, which when passed to model as prefix, helps it better score sentences from that domain. We hypothesize that this is due to the fact that self-attention in Transformer model will create a positive interaction between domain embeddings (description) and the words from sentence, resulting in better prediction of the next token. This has an important usecase in multi-domain ASR systems, where learning a small number of parameters corresponding to domain prompt embedding can save costs associated with maintaining separate copies of LMs for each domain. Particularly, with Transformer models with millions of parameters, fine-tuning can be very slow and requires a large amount of domain data to prevent overfitting. However, with prompt-tuning one can achieve similar performance with $\leq 0.1\%$ of the parameters required for fine-tuning. We provide experimental evidence by comparing the performance of prompt-tuned and fine-tuned versions of LMs to score sentences from different domains of a dialog dataset. We also provide ablations on different sizes of prompt.

2 Experiments

We design first set of experiments as a proof of concept to show that prompt-tuning works for domain adaptation of LM. For this, we used textual data from different domains of the publically available MultiDoGo dataset (Peskov et al., 2019). We split the data from each domain into train, dev and test. Then for each domain individually, we fine-tune (and prompt-tune) a pretrained GPT-2 (Radford et al., 2019) using the train set, tune the hyper-parameters using dev set and compute the perplex-




Figure 1: Comparing the performance, computation and memory costs of different methods of domain adaptation

ity numbers on the test set. We also evaluate the test set for each domain using the same out-of-the-box pretrained GPT-2. We then compare the no-adaptation, fine-tuned and prompt-tuned versions for each domain for different performance, memory and computation time metrics. Rather than fixing the size of the prompt, we repeat the above experiment with different sizes to show how we can control the trade-offs between performance and costs using this hyper-parameter. Although, we summarize the results for one such domain (fastfood) in Fig. 1, the results are consistent across different domains. All the experiments were conducted using a machine with 16 Tesla K80 GPUs and the reported times are wall clock times in seconds. As we can see in Fig. 1, with as little as 5 prompt-tokens, the perplexity improves from 130 in out-of-the-box model (green line) to 37. With 400 prompt-tokens, we achieve a very similar perplexity to fully fine-tuned model (red line) but with number of learnable parameters reducing by large orders of magnitude. As one might expect, with increase in size of prompts, the performance improves but so do the memory and computation costs.

The above experiment proves us the effectiveness of prompt-tuning in domain adaptation. With comparable perplexity as fine-tuning, it is natural to expect that prompt-tuned model can perform at par with fully tuned model in domain specific rescoring of n-best hypothesis from an AM model in any ASR system. We verify this empirically by using the above models as first pass rescorers on an internal fastfood audio dataset for a hybrid AM model and measuring the Content Word Error Rates (CWER). While the fine-tuned version brings a relative improvement of 7.3% over the no-adaptation baseline, training just 10 prompt tokens could already achieve 5.8% relative improvement. Also, with 400 prompt tokens, we could almost ex-

actly match the CWER of fine-tuned model. Since, ASR systems are used in tens of different domains, one can achieve a significant reduction in LMs that one need to maintain, by keeping a single copy of the base version of the model and saving the domain-specific prompt token embeddings.

3 Future Work

While the initial experiments gives us confidence about the potential use of prompt-tuning in LMs in ASR systems, we plan to explore a lot more as part of this ongoing work. We will be running multiple experiments to showcase the benefits in WER further in multiple domains, datasets, etc. While in our experiments, we used the complete training data to fine-tune (and prompt-tune) the models, we believe in limited domain data setting, prompt-tuned models might even outperform the corresponding fine-tuned versions as the number of learnable parameters are less. This will be especially critical for new domains in ASR where getting high-quality data is difficult. We initialized our prompt token embedding layer randomly but as suggested in (Lester et al., 2021), we plan to run an ablation on different choices of initializations. One such promising initialization can be embedding for the most frequent words in that domain. In addition to all these experiments, we also plan to compare our version of domain-adaptation with related works of AdapterHub (Pfeiffer et al., 2020), meta-learning (Bansal et al., 2020), using semantic embeddings (Shenoy et al., 2021b,a) and freezing all but last layer of Transformer model. The analysis provided in this paper will act as a proof of concept on computation and memory efficiency of prompt-tuning guiding our future directions.

References

- Trapit Bansal, Rishikesh Jha, and Andrew McCallum. 2020. [Learning to few-shot learn across diverse natural language classification tasks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Denis Peskov, Nancy Clarke, Jason Krone, Brigi Fodor, Yi Zhang, Adel Youssef, and Mona Diab. 2019. [Multi-domain goal-oriented dialogues \(MultiDoGO\): Strategies toward curating and annotating large scale dialogue data](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4526–4536, Hong Kong, China. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. [AdapterHub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Ashish Shenoy, Sravan Bodapati, and Katrin Kirchhoff. 2021a. Asr adaptation for e-commerce chatbots using cross-utterance context and multi-task language modeling. *arXiv preprint arXiv:2106.09532*.
- Ashish Shenoy, Sravan Bodapati, Monica Sunkara, Srikanth Ronanki, and Katrin Kirchhoff. 2021b. [Adapting Long Context NLM for ASR Rescoring in Conversational Agents](#). In *Proc. Interspeech 2021*, pages 3246–3250.