



Memory Augmented Lookup Dictionary based Language Modeling for Automatic Speech Recognition

Yukun Feng, Ming Tu, Rui Xia, Chuanzeng Huang, Yuxuan Wang

ByteDance Inc, San Jose, CA, US

{yukun.feng1, mingtu, rui.xia, huangchuanzeng, wangyuxuan.11}@bytedance.com

Abstract

Recent studies have shown that using an external Language Model (LM) benefits the end-to-end Automatic Speech Recognition (ASR). However, predicting tokens that appear less frequently in the training set is still quite challenging. The long-tail prediction problems have been widely studied in many applications, but only been addressed by a few studies for ASR and LMs. In this paper, we propose a new memory augmented lookup dictionary based Transformer architecture for LM. The newly introduced lookup dictionary incorporates rich contextual information in training set, which is vital to correctly predict long-tail tokens. With intensive experiments on Chinese and English data sets, our proposed method is proved to outperform the baseline Transformer LM by a great margin on both word/character error rate and tail tokens error rate. This is achieved without impact on the decoding efficiency. Overall, we demonstrate the effectiveness of our proposed method in boosting the ASR decoding performance, especially for long-tail tokens.

Index Terms: Automatic speech recognition, Language modeling, rare words recognition, long-tail recognition.

1. Introduction

While a lot of studies have demonstrated the superiority of end-to-end (E2E) Automatic Speech Recognition (ASR) systems [1, 2] and the effectiveness of incorporating Language Models (LM) into the E2E ASR systems[3, 4], recognition and prediction of words that appear only a few or zero times in training data are still big challenges, especially for E2E ASR systems which are optimized only on text in the training data.

Some studies have addressed this long-tail problem for E2E ASR [5, 6, 7, 8, 9, 10, 11]. The studies in [6, 9] resort to adding large corpora of textual data or adjusting the distribution of head and tail words in LM training to improve the modeling ability of tail words. In [7, 8], the authors proposed to improve the prediction of tail words with the help of large-scale pretrained LMs (BERT[12] variants) which inevitably increases the decoding computational cost. Another line of research modified the training loss or introduced extra loss terms to regularize the ASR training, and results showed improved performance on rare words [5, 11]. In [10], the authors tried to scale up the embedding capacity of an RNN LM by incorporating N-gram context embedding into the embedding layer without sacrificing decoding efficiency. However, it ignored the frequency information of words and N-grams and only reply on the input embedding layer to learn enough contextual information to predict rare words.

Since Transformer LMs have shown better performance than RNN LMs for ASR [13], in this paper we extend the Trans-

former LMs with a lookup dictionary that maps the current context to candidate tokens that have occurred during training. Inspired by [14] which focuses on effective training of BERT, we initialize a dictionary by aggregating the N-gram token IDs of the current token as keys and utilize a multi-vector array as values to enable memorization of rich context information. We now consider the dictionary’s values as the memory of the corresponding N-gram context. Specifically, the contextual memory is updated by the current token’s **subsequent** token embedding in the training based on how often the **subsequent** token occurs in the training corpus. For each key, the frequency of the **subsequent** token decides how many vectors in the corresponding multi-vector value will be updated. We then use an attention module at the last layer of the transformer blocks to map the dictionary memory to the contextualized embedding of the current token, in which the current context will query the most relevant vectors from the corresponding multi-vector memory.

We experimented on two Mandarin ASR data sets and improve 8.5% relatively of the Character Error Rate (CER) over the baseline Transformer LM. Notably, our method show 13% and 12.5% relative CER reduction on the 1-gram and 2-gram tail tokens. Also, we achieve Word Error Rate (WER) improvement on the two test sets of LibriSpeech. The results indicate the success of our method on improving not only the general ASR decoding but also the prediction of tail tokens for both Mandarin and English. We did intensive analysis to investigate the benefits of different aspects of our proposed method. Overall, this paper makes several contributions as the following:

1. We propose a new Transformer based LM for ASR equipped with a lookup dictionary consisting of multi-vector memory that builds bonding between the current context and to-be-predicted candidate tokens.
2. We incorporate the N-gram context information and the token frequency in training data into the lookup dictionary to improve the prediction of rare words.
3. Our proposed LM significantly outperforms the baseline LM in both Mandarin and English ASR **while keeping the same inference efficiency**.

2. Proposed approach

In Figure 1, we show the diagram of the Transformer LM equipped with our proposed memory augmented lookup dictionary. Each module will be introduced separately in the following subsections.

2.1. Dictionary Construction and Indexing

We initialize the dictionary as $\mathbf{D} \in \mathbb{R}^{U \times M \times d_{\text{emb}}}$, where d_{emb} is the embedding size and U is the dictionary size. Instead of

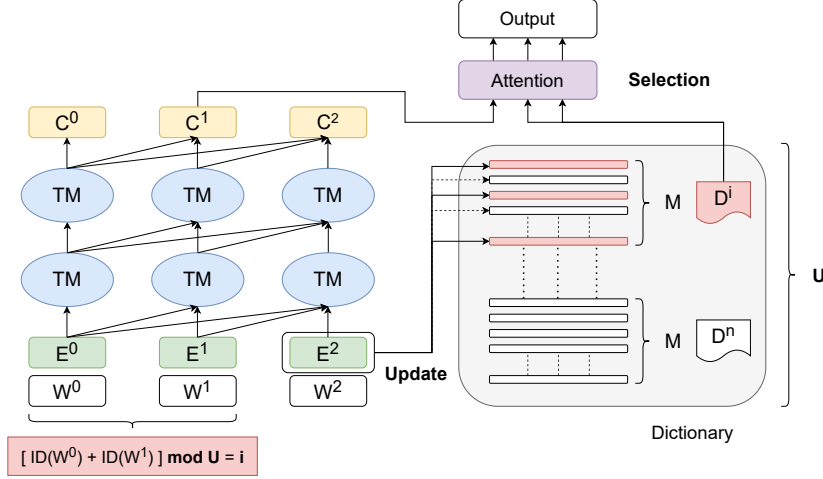


Figure 1: Overview of the proposed memory augmented lookup dictionary based Transformer LM. W^k represents the input tokens, E^k is for the token embedding, “TM” means Transformer blocks in auto-regressive manner. C^k is the contextualized embedding corresponding to the input tokens. Input and output embedding weights are shared in the LM.

using one vector as value for each key as in [14], we scale up the dictionary size by introducing an extra hyper-parameter M to form a key-value pair as (i, \mathbf{D}^i) where $\mathbf{D}^i \in \mathbb{R}^{M \times d_{\text{emb}}}$. For the k^{th} token in the input sequence, the corresponding dictionary index i is mapped through a modular hash with U , defined as:

$$i = \text{ID}(\text{Token}_k) \bmod U \quad (1)$$

where $\text{ID}()$ refers to the vocabulary id of the input token. We believe with multiple vectors stored for each entry, much richer contextual information could be memorized compared than the single vector counterpart. To consider more context in dictionary indexing, we also extend Equation 1 to N-gram case as in [10], where the dictionary index i is calculated as follows

$$i = \left(\sum_{n=k-N+1}^k \text{ID}(\text{Token}_n) \right) \bmod U \quad (2)$$

where N indicates the number of token IDs to aggregate. For example, if $N = 2$, we sum up the IDs of the current token and its previous token before the modulo operation.

To trade off information redundancy and memory capacity, collision is allowed when doing hashing. U , N and M can be adjusted appropriately. We assume this approach could utilize the dictionary memory more efficiently. We will show the influence of changing the three hyper-parameters to the performance in results part.

2.2. Dictionary Update

As the k^{th} token is mapped to the dictionary memory \mathbf{D}^i through Equation 1, each memory vector \mathbf{d}_i^m is updated by the embedding of current token’s next token \mathbf{e}_{k+1} , which can be formulated as:

$$\widetilde{\mathbf{d}}_i^m = \begin{cases} \mathbf{d}_i^m * \alpha + \mathbf{e}_{k+1} * (1 - \alpha) & \text{if } X_{k+1} = 1 \\ \mathbf{d}_i^m & \text{if } X_{k+1} = 0 \end{cases} \quad (3)$$

where α is a smoothing hyper-parameter that indicates how much information comes from \mathbf{e}_{k+1} , and we set it as 0.5 for all later experiments. We define a Bernoulli Variable

$X_{k+1} \sim \text{Bern}(P_{k+1})$, which decides how many vectors will be updated in the matrix \mathbf{D}^i . P_{k+1} indicates the update ratio, computed by the normalized occurrence of the $k + 1^{\text{th}}$ token in training data:

$$P_{k+1} = \frac{1}{\log(\text{Count of Token}_{k+1})}. \quad (4)$$

In this case, embeddings of low frequency tokens are able to contribute more to the corresponding memory compared to high frequency tokens. Token frequencies are calculated with training text and corresponding text tokenizer.

2.3. Context Selection

We use an attention module to relate the output representation of the current token to the corresponding dictionary memory. Attention performs as a mapping function for the input query (Q) and key-value (K-V) pairs, as

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_{\text{emb}}}}\right)\mathbf{V}. \quad (5)$$

As we share the input and output embedding weight in the Transformer LM, and the dictionary memory stores the candidate tokens’ embedding during training, we assume the attention module could help select useful information from the memory given the output representation of the current token. For the k^{th} token, we define the contextualized token embedding from the Transformer model as \mathbf{c}_k , and its corresponding dictionary memory is defined as \mathbf{D}^i , where i is the hashing index from Eq 1 or 2. The new output representation $\widetilde{\mathbf{c}}_k$ is computed as:

$$\widetilde{\mathbf{c}}_k = \text{Attention}(\mathbf{c}_k, \mathbf{D}^i, \mathbf{D}^i), \quad (6)$$

which will then be used to calculate the output token distribution.

2.4. Training and Inference

During training, the context selection operation was done before the dictionary update for the reason that the update information in Eq 3 for the current input sentence will not affect the

Model	Input				Search			
	Overall CER/SER	Tail-1 CER	Tail-2 CER	PPL	Overall CER/SER	Tail-1 CER	Tail-2 CER	PPL
Conformer with Language Model	6.19 / 44.91	15.87	13.82		13.51 / 45.38	22.38	21.46	
+ LM	5.55 / 41.80	13.46	11.86	36.49	9.07 / 30.11	13.86	13.78	13.92
+ LM [10]	5.42 / 40.37	12.50	11.18		8.91 / 29.79	13.50	13.44	
+ LM [14]	5.35 / 40.19	12.48	10.97		9.02 / 30.23	13.80	13.68	
+ Ours	5.09 / 38.86	11.46	10.30	31.44	8.29 / 27.60	12.34	12.13	11.22
	+8.3%	+14.9%	+13.2%		+8.6%	+11.0%	+12.0%	
+ LM_L	4.83 / 37.29	10.94	9.67	26.37	8.24 / 27.21	12.26	12.27	10.82
+ $Ours_L$	4.73 / 36.90	10.54	9.54	25.18	7.80 / 26.20	11.49	11.41	9.70
	+2.1%	+3.7%	+1.3%		+5.4%	+6.3%	+7.0%	

Table 1: Evaluation of CER, SER and language model perplexity (PPL) on two internal Chinese ASR test sets. “Input” and “Search” refer to voice input and voice search domain test sets respectively. L refers to the LM with 1024 embedding size.

current context selection. To stabilize training, we also disable the dictionary update for the first 1000 training steps to warmup the newly initialized embedding to a good distribution. During inference, the dictionary update is also disabled to avoid any information leakage for auto-regressive prediction. With the trained memory augmented Transformer LM, we apply shallow fusion to integrate the LM to ASR decoding with weight λ_{sf} . Also, Internal Language Model Estimation (ILME) [15] is adopted to suppress the internal LM of the E2E ASR and advocate the contribution of the external LM, which has been proved to be quite effective especially there is domain mismatch between textual distribution of ASR and LM training data. The weight of ILME is noted by λ_i . We also tried LM rescoring over the N-best output of beam search, and the weight of rescoring is noted by λ_{res} .

3. Experiment

3.1. Datasets

We adopt the LibriSpeech [16] dataset to evaluate the ASR performance in English. We use the standard 960 hours data for training and the “clean” and “other” test sets for evaluation. The corresponding LM is trained on PG-19 [17], an 11GB in-domain text corpus consisting of books extracted from Project Gutenberg. To match the averaged sentence length in LibriSpeech, we process the PG-19 into a sentence-level corpus. We use the unigram tokenizer [18] with vocabulary size of 5000 from ESPnet [19] for both ASR and LM training. Also, we evaluate our method on two internal Chinese video datasets. We have a 10k hours annotated audio dataset for general ASR training and two test sets: one is voice input domain (5103 utterances) and the other is voice search domain (6424 utterances), which are two different domains compared to the ASR training set. As for the LM training, we have a 60GB text corpus for the voice input domain and 2GB corpus for the voice search domain. We process the Chinese text at the character level with a vocabulary size of 11k (with both Chinese characters and English subword tokens).

Besides evaluating the overall performance on the above mentioned test sets, we also assess the ASR metrics on tail tokens. Tail tokens are defined as the tokens whose accumulated frequency in the training corpus is lower than a *threshold*, which we set as 5%, i.e. the frequency ratio of head and tail tokens is 95:5. Both 1-gram (Tail-1) and 2-gram (Tail-2) tail tokens are extracted from test sets at character-level for Chinese. For English test set, we only extracted 1-gram word-level tail

tokens.

3.2. Experimental settings

We train both Chinese and English ASR models with a LAS [2] architecture, for which we use a 12-layer Conformer[20] encoder and 6-layer Transformer decoder for Librispeech (as in ESPnet), and a 18-layer Conformer encoder and 4-layer Transformer decoder for the 10k hours Chinese dataset.

For LibriSpeech, we configure the LM as a 16-layer Transformer blocks with 1024 embedding size. It is trained on PG-19 for sentence-level language modeling with a dropout rate of 0.3 and an effective token number of 524288 in each update. Adam with betas of (0.9, 0.98), and weight decay of 0.01 is used for the optimization with 10k warmup steps. For the proposed look-up dictionary, we use 2-gram for dictionary hashing (as in Eq. 2); U is set to 5k; M is set to 64. The LM for Chinese datasets consists of 4 layer Transformer blocks with the embedding size of 384 and 1024 for small and large configuration respectively. For look-up dictionary, U is set to 10k and other hyper-parameters are the same with the Librispeech settings.

For ASR inference in this paper, we set $\lambda_{sf}=\{0.15, 0.4, 0.4\}$, $\lambda_{res}=\{0.0, 0.0, 0.1\}$, $\lambda_i=\{0.0, 0.2, 0.2\}$ for {“LibriSpeech”, “Input”, “Search”} respectively, which give the best performance. A beam size of 60 is used for the LibriSpeech and 10 for Chinese test sets. We use Word Error Rate (WER) as ASR metric for LibriSpeech test sets, and Character Error Rate (CER) and Sentence Error Rate (SER) for Chinese test sets. For both test sets, we also calculate the tail token error rate by only counting errors on tail tokens and ignoring errors on other tokens within the same testing utterances.

3.3. Results

We compare our model with the original Transformer LM, as well as two other baselines: N-gram augmented embedding for LM training in [10] and single-vector memory for BERT pre-training in [14]. In Table 1, while the original LM helps the ASR model achieve lower CER and SER, our method shows significant improvement over it and the two baseline methods. We achieve 8.3% and 8.6% CER improvement on the general “Input” and “Search” test sets over the Transformer LM, and the CER improvement of tail tokens are even higher: 13% on 1gram and 12.5% on 2gram tail tokens. As we increase the hidden size of the LM from 384 to 1024, the performance gain is not as much as the small LMs, but our method still outperforms the LM by 3.7% on overall CER and 4.6% on tail tokens CER.

Model	Clean		Other	
	Overall	Tail-1	Overall	Tail-1
Conformer	3.12%	11.92%	6.23%	24.52%
+ LM	3.08%	10.93%	5.81%	23.30%
+ Ours	3.01%	10.57%	5.73%	22.93%

Table 2: Evaluation of WER on the LibriSpeech test sets.




Figure 2: Change of the overall CER (%) on the “Search” test set with different dictionary size and different N-gram settings.

The smaller performance gain when increasing LM model size is expected given fixed amount of training data as the increased model capacity would take over part of the functionalities the memory block provides. In the meantime, the proposed LM improves the perplexity (PPL) by a great margin over the baseline model.

In Table 2, our proposed method also shows consistent improvement on the two LibriSpeech test sets. The improvement on tail word error rate is more significantly compared to the overall WER improvement as on Chinese test sets. The improvement on Librispeech test set is not that significant compared to Chinese test sets, and we believe the reason is that subword tokens are employed for English data sets. When we still uses 2-gram based hashing as for Chinese data sets, it actually incorporates less word-level contextual information (2-gram on subword level sometimes even does not correspond to a single word). We believe if we increase the N in English data sets, the performance could be further improved. We leave this for future investigation.

4. Analysis

In this section, we analyze how the different hyper-parameters, including dictionary size U , N in N-gram for hashing, memory update ratio and memory size of each entry M , affect the performance. All experiments are conducted on the Chinese “Search” test set, and the Transformer LM model with the proposed memory augmented lookup dictionary has 4 layers and 384 hidden size.

In Figure 2, we show the change of the overall CER (y axis) with the increase of dictionary size in different N-gram settings. It is clear that for each N-gram setting, increasing the dictionary size will boost the performance, and 2-gram achieves the best performance. Since the degree of collision elevates with bigger U and N , larger N means more collision; thus 4-gram performs even worse than 1 gram case when the dictionary size is not large. Considering the extra space taken by large dictionary size, we choose the 2-gram with 10k dictionary size.

Figure 3 analyzes if a large memory size M would help the selection and the overall performance. We use the Information Gain (IG) which is computed by the difference in the attention entropy (as in Eq. 5 and 6) between a randomly initialized




Figure 3: Overall CER, Gradients, and Information Gain (IG) change on the “Search” test set with the increase of memory size M

dictionary and a well-trained one. The entropy indicates how well the dictionary maps the information to the contextualized embedding of the current token \tilde{c}_k [21]. The results show the IG is highly correlated with M . Besides, we adopt the Gradient Attribution test [22, 23] to address the dictionary memory’s contribution further. It computes the normalized gradient of the model variables to reflect its contribution to the output prediction. It shows the gradients are also consistent with the previous finding that a larger memory would receive more gradients, indicating a greater contribution to the model prediction. However, considering the small relative gain and high computational cost when we increase the memory size from 64 to 128, we set the memory size as 64 in our experiments.

Finally, we want to discuss how the proposed memory augmented lookup dictionary will affect the model size and inference speed. During inference, compared to the baseline Transformer LM, the additional computation of our method is only the dictionary indexing (Eq. 2) and context selection (Eq. 6). For lookup dictionary, the indexing operation requires $O(1)$ time cost. The context selection also performs as a constant time cost as $O(M)$, where M is the memory size of the dictionary. We evaluate the Real Time Factor (RTF) on the “Search” test set on a NVIDIA A100 GPU with batch size equals to 1. The RTF is 0.124 for ASR model only, and 0.195 and 0.198 for the baseline Transformer LM and our proposed LM, respectively. We notice that such additional operations almost do not affect the decoding speed in practice though the model size increases by introducing the lookup dictionary.

5. Conclusions

In this paper, we propose a memory augmented lookup dictionary based Transformer LM to improve the language modeling in ASR, especially for long tail tokens. We have improved the baseline Transformer LM in terms of overall ASR metrics and the tail words error rate in both Chinese and English test sets. We also analyze our method under different hyper-parameter settings. Overall, the results prove the superiority of the method over the baseline Transformer LM without sacrificing inference speed. Future work includes more experiments on English data sets, especially in domain mismatch condition. We are also interested in applying the method to general language modeling tasks.

6. References

- [1] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [2] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, “Listen, attend and spell,” *arXiv preprint arXiv:1508.01211*, 2015.

- [3] S. Toshniwal, A. Kannan, C.-C. Chiu, Y. Wu, T. N. Sainath, and K. Livescu, "A comparison of techniques for language model integration in encoder-decoder speech recognition," in *2018 IEEE spoken language technology workshop (SLT)*. IEEE, 2018, pp. 369–375.
- [4] A. Kannan, Y. Wu, P. Nguyen, T. N. Sainath, Z. Chen, and R. Prabhavalkar, "An analysis of incorporating an external language model into a sequence-to-sequence model," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 1–5828.
- [5] C. Peyser, T. N. Sainath, and G. Pundak, "Improving proper noun recognition in end-to-end asr by customization of the mwer loss criterion," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7789–7793.
- [6] C. Peyser, S. Mavandadi, T. N. Sainath, J. Apfel, R. Pang, and S. Kumar, "Improving tail performance of a deliberation e2e asr model using a large text corpus," *arXiv preprint arXiv:2008.10491*, 2020.
- [7] G. I. Winata, G. Wang, C. Xiong, and S. Hoi, "Adapt-and-adjust: Overcoming the long-tail problem of multilingual speech recognition," 2021. [Online]. Available: <https://openreview.net/forum?id=34KAZ9HbJco>
- [8] K. Deng, G. Cheng, R. Yang, and Y. Yan, "Alleviating asr long-tailed problem by decoupling the learning of representation and classification," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 340–354, 2022.
- [9] W. R. Huang, C. Peyser, T. N. Sainath, R. Pang, T. Strohmman, and S. Kumar, "Sentence-select: Large-scale language model data selection for rare-word speech recognition," *ArXiv*, vol. abs/2203.05008, 2022.
- [10] W. R. Huang, T. N. Sainath, C. Peyser, S. Kumar, D. Rybach, and T. Strohmman, "Lookup-table recurrent language models for long tail speech recognition," *ArXiv*, vol. abs/2104.04552, 2021.
- [11] C.-H. H. Yang, L. Liu, A. Gandhe, Y. Gu, A. Raju, D. Filimonov, and I. Bulyko, "Multi-task language modeling for improving speech recognition of rare words," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 1087–1093.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "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)*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186. [Online]. Available: <https://aclanthology.org/N19-1423>
- [13] K. Irie, A. Zeyer, R. Schlüter, and H. Ney, "Language modeling with deep transformers," in *INTERSPEECH*, 2019.
- [14] Q. Wu, C. Xing, Y. Li, G. Ke, D. He, and T.-Y. Liu, "Taking notes on the fly helps language pre-training," in *ICLR*, 2021. [Online]. Available: https://openreview.net/forum?id=IU5Rs_wCweN
- [15] Z. Meng, N. Kanda, Y. Gaur, S. Parthasarathy, E. Sun, L. Lu, X. Chen, J. Li, and Y. Gong, "Internal language model training for domain-adaptive end-to-end speech recognition," *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7338–7342, 2021.
- [16] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [17] J. W. Rae, A. Potapenko, S. M. Jayakumar, C. Hillier, and T. P. Lillicrap, "Compressive transformers for long-range sequence modelling," *arXiv preprint*, 2019. [Online]. Available: <https://arxiv.org/abs/1911.05507>
- [18] T. Kudo, "Subword regularization: Improving neural network translation models with multiple subword candidates," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 66–75. [Online]. Available: <https://aclanthology.org/P18-1007>
- [19] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-End Speech Processing Toolkit," in *Proc. Interspeech 2018*, 2018, pp. 2207–2211.
- [20] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," 10 2020, pp. 5036–5040.
- [21] Y. Feng, F. Li, Z. Song, B. Zheng, and P. Koehn, "Learn to remember: Transformer with recurrent memory for document-level machine translation," in *Findings of the Association for Computational Linguistics: NAACL 2022*. Seattle, United States: Association for Computational Linguistics, Jul. 2022, pp. 1409–1420. [Online]. Available: <https://aclanthology.org/2022.findings-naacl.105>
- [22] M. Ancona, E. Ceolini, C. Öztireli, and M. Gross, "Towards better understanding of gradient-based attribution methods for deep neural networks," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=Sy21R9JAW>
- [23] G. Qin, Y. Feng, and B. Van Durme, "The nlp task effectiveness of long-range transformers," 2022. [Online]. Available: <https://arxiv.org/abs/2202.07856>