

Learn Spelling from Teachers: Transferring Knowledge from Language Models to Sequence-to-Sequence Speech Recognition

Ye Bai^{1,2}, Jiangyan Yi¹, Jianhua Tao^{1,2,3}, Zhengkun Tian^{1,2}, Zhengqi Wen¹

¹NLPR, Institute of Automation, Chinese Academy of Sciences, China

²School of Artificial Intelligence, University of Chinese Academy of Sciences, China

³CAS Center for Excellence in Brain Science and Intelligence Technology

{ye.bai, jiangyan.yi, jhtao, zhengkun.tian, zqwen}@nlpr.ia.ac.cn

Abstract

Integrating an external language model into a sequence-to-sequence speech recognition system is non-trivial. Previous works utilize linear interpolation or a fusion network to integrate external language models. However, these approaches introduce external components, and increase decoding computation. In this paper, we instead propose a knowledge distillation based training approach to integrating external language models into a sequence-to-sequence model. A recurrent neural network language model, which is trained on large scale external text, generates soft labels to guide the sequence-to-sequence model training. Thus, the language model plays the role of the teacher. This approach does not add any external component to the sequence-to-sequence model during testing. And this approach is flexible to be combined with shallow fusion technique together for decoding. The experiments are conducted on public Chinese datasets AISHELL-1 and CLMAD. Our approach achieves a character error rate of 9.3%, which is relatively reduced by 18.42% compared with the vanilla sequence-to-sequence model.

Index Terms: knowledge distillation, external language models, end-to-end, sequence-to-sequence models

1. Introduction

Attention based sequence-to-sequence (Seq2Seq) models have achieved promising performance in automatic speech recognition (ASR) [1, 2, 3, 4]. A Seq2Seq model consists of two components: an encoder encodes the acoustic feature sequence into a high level representation, and a decoder generates the corresponding word sequence. The encoder leverages attention mechanism to fuse extracted features into a fixed-dimensional vector for capturing global semantic information of a speech signal. The decoder is a conditional language model (LM) to capture linguistic information of transcriptions. During decoding stage, the decoder predicts the current word in terms of the acoustic encoding of the encoder, history context, and the previous word at each step. This architecture is also referred to as Listen, Attend, and Spell [2].

Compared with speech transcriptions, abundant unsupervised text corpora, which have rich linguistic information, are easier to obtain. Large scale external text data is commonly used to train language models (LMs) to improve ASR performance in conventional hidden Markov model (HMM) or connectionist temporal classification (CTC) based ASR pipelines. However, because the encoder and the decoder are optimized jointly, it is non-trivial to integrate an external LM into a Seq2Seq model.

Shallow fusion and deep fusion are two approaches to

integrating an LM into a Seq2Seq model [5]. Shallow fusion performs log-linear interpolation between the decoder of a Seq2Seq model and an external LM during beam search. The external LM can be n -gram LM or neural network language models (NNLMs). It has achieved success in ASR tasks [1, 6]. Various deep fusion approaches leverage a neural network to fuse hidden representations of the Seq2Seq decoder and the external neural network based LM [5]. Cold fusion and component fusion utilize a pre-trained recurrent neural network language model (RNNLM) and gating mechanism to improve ASR performances [7, 8]. These fusion approaches have shown promising performance. However, the neural network of the external LM increases complexity of the Seq2Seq model. Specifically, the fusion network introduces external parameters into the Seq2Seq model for deep fusion. Both shallow fusion and deep fusion need the external LM during test stage. It introduces external complexity into the ASR system.

We propose a knowledge distillation (KD) [9] based training approach to integrating an external LM into a Seq2Seq model. First, an RNNLM is trained on large scale text data. Then, the RNNLM is used to generate soft labels of speech transcriptions to train the Seq2Seq model. This training approach is also named as Teacher/Student model: the teacher (RNNLM) provides soft labels as prior knowledge to “teach” the student (Seq2Seq decoder). Thus, we refer to the proposed training approach as “Learn Spelling from Teachers” (LST). LST is simple to implement: it does not modify the model structure, and only needs to train an RNNLM to generate soft labels. With LST, the external LM is only needed during training, so it does not increase complexity of the model for testing. Furthermore, LST and shallow fusion can be used together to achieve better performance. We conduct experiments on publicly available AISHELL-1¹ dataset [10] and CLMAD² text dataset [11] to show the effectiveness of the proposed LST. We use Speech-Transformer [3] as the backbone network. Our proposed approach reduced character error rate (CER) from 11.4% to 9.3%. we further utilized shallow fusion for the model trained with LST, and achieved CER of 8.3%.

The rest of this paper is organized as follows. Section 2 introduces the background. Section 3 introduces the proposed LST. Section 4 introduces the related work. Section 5 describes the experimental results. Section 6 summarizes the paper.

2. Background: Seq2Seq models for ASR

A basic Seq2Seq model is shown in Fig. 1(a). First, a speech signal is processed into an acoustic feature sequence. Then, an

¹<http://openslr.org/33/>

²<http://openslr.org/55/>




Figure 1: (a) illustrates a basic encoder-decoder architecture for ASR. x_1, \dots, x_t represent acoustic features, c_{t-1} denotes the context of $t - 1$ step, and y_{t-1} denotes the previous token. The decoder predicts the current token in terms of the context c_{t-1} , the previous token y_{t-1} , and the acoustic vector generated by the encoder. The loss is computed with the softmax function of the decoder and the current ground truth token y_t . (b) illustrates the proposed “Learn Spelling from Teachers” (LST) approach. The RNNLM generates soft labels to train the Seq2Seq model, and it is removed during testing.

encoder network encodes the sequence into a high level acoustic representation. The encoder can be a recurrent neural network [2, 12] or a transformer [3]. The decoder is a conditional LM: given the high level acoustic representation, the previous token, and history context, it predicts the current token. The probability distribution on the vocabulary is computed by a softmax function.

The attention is an important mechanism to capture the relationship between the acoustic representations and the current state of the decoder. The attention scores are computed in terms of the current state of the decoder and the high level acoustic representations, and then the acoustic information and the decoder state are fused.

The encoder and decoder are trained jointly. The training criterion is cross entropy:

$$L_{CE}(\theta) = - \sum_{k=1}^K \delta(k, y_t) \log P_{S2S}(k|y_{t-1}, c_{t-1}, \mathbf{x}; \theta), \quad (1)$$

where k is the index of each token, K is the vocabulary size, y_t is the index of the corresponding ground truth token at step t , y_{t-1} is the previous token, c_{t-1} is the history context, \mathbf{x} is the acoustic features, P_{S2S} represents probability, and θ stands the parameters of the whole network. $\delta(\cdot, \cdot)$ is 1 if the two variables are equal, and 0 otherwise.

3. Distilling knowledge from external LMs

The basic idea of “Learn Spelling from Teachers” (LST) is: first, train an RNNLM on an external large scale text corpus, and then use this RNNLM to guide Seq2Seq model training. Besides 1-of-K hard labels provided by the transcriptions, the RNNLM provides soft labels, which carries the knowledge of the text corpus. The soft labels are probabilities estimated by the RNNLM. Fig. 2 shows the hard labels and soft labels of tokens in the vocabulary at one time step in a sequence. The soft labels contain more information than hard labels, e.g. some tokens have relatively large probabilities, and some tokens have very small probabilities.

Given the context and the previous token, the probability of k -th token in vocabulary estimated by the RNNLM is

$$P_{LM}(k|y_{t-1}, h_{t-1}) = \frac{\exp(z_k/T)}{\sum_{i=1}^K \exp(z_i/T)}, \quad (2)$$

where z_i is i -th node of latent variable before the softmax function, K is the vocabulary size, y_{t-1} is the previous token, h_{t-1}

0	2.50e-05	璃 檬 : 的 们
0	2.68e-05	
...	...	
0	0.0062	
1	0.078	
Hard Labels	Soft Labels	Corresponding Tokens

Figure 2: Hard labels and soft labels at one time step of a sequence for training. The values of the soft labels reflect knowledge of the external LM.

is the history context, and T is a parameter called temperature to smooth the outputs.

To make the Seq2Seq model learn the knowledge from the RNNLM, we minimize the Kullback-Leibler divergence (KLD) between estimated probability of the RNNLM P_{LM} and the estimated probability of the Seq2Seq model P_{S2S} . Let $P_{S2S}^k = P_{S2S}(k|y_{t-1}, c_{t-1}, \mathbf{x}; \theta)$, and $P_{LM}^k = P_{LM}(k|y_{t-1}, h_{t-1})$, the KLD is

$$D_{KL}(P_{LM}||P_{S2S}) = - \sum_{k=1}^K P_{LM}^k \log \frac{P_{S2S}^k}{P_{LM}^k}. \quad (3)$$

Because P_{LM} is fixed during training the Seq2Seq model, the loss function is equivalent to the cross entropy form:

$$L_{LST}(\theta) = - \sum_{i=1}^K P_{LM}^k \log P_{S2S}^k. \quad (4)$$

We refer to the above loss as LST loss.

The cross entropy loss in Eq. (1) and the LST loss in Eq. (4) are weighted with a coefficient $\lambda \in [0, 1]$, then the final loss is

$$L(\theta) = \lambda L_{CE}(\theta) + (1 - \lambda) L_{LST}(\theta). \quad (5)$$

We can simplify the above equation as the label interpolation form:

$$L(\theta) = - \sum_{k=1}^K (\lambda \delta(k, y_t) + (1 - \lambda) P_{LM}^k) \log P_{S2S}^k. \quad (6)$$

Thus, compared with the vanilla Seq2Seq model, we just modify the labels rather than the loss function during training

stage. $L(\theta)$ combines the knowledge from transcriptions and the knowledge from the LM. The LST is illustrated in Fig. 1(b).

Comparing Fig. 1(b) with Fig. 1(a), we can see that LST is only used for training, and the external RNNLM is removed during testing. So the computation is the same as the original Seq2Seq for testing. In order to achieve better performance, shallow fusion can be further used with LST during decoding. In addition, besides ASR, our proposed LST can be generally used for Seq2Seq models.

4. Related work

Knowledge distillation. KD was proposed for model compression [9]. It is also referred to as teacher-student learning. Yoon et al. proposed to use KD to reduce the size of a Seq2Seq model for machine translation [13]. It has also been used for domain adaptation for acoustic models [14] and language models [15]. Different from these work, our work focuses on integrating external language models for Seq2Seq ASR systems.

Label smoothing. Label smoothing have been used to prevent the Seq2Seq ASR model making overconfident predictions [4, 3, 16]. It can be seen as a special case of KLD regularization when assuming the prior label distribution is uniform [17]. Unlike label smoothing, LST leverages an RNNLM to provide a context-dependent prior distribution rather than a simple uniform distribution. Instead of assumption, the prior distribution is estimated with a data-driven method. Besides solving the overconfident problem, LST introduces knowledge from an external large scale text data corpus.

5. Experiments

5.1. Datasets

We use a Chinese corpus AISHELL-1 to evaluate our proposed approach [10]. The training set contains 150 hours of speech (120,098 utterances) recorded by 340 speakers. The development set contains 20 hours of speech (14,326 utterances) recorded by 40 speakers. And the test set contains 10 hours of speech (7,176 utterances) recorded by 20 speakers. The speakers of the training set, development set, and test set are not overlapped. All the recordings of the corpus are in 16 kHz WAV format. The content of the speech is news with different topics.

A subset of CLMAD [11, 18] text dataset is used as external text dataset³. We use an open source tool XenC to extract the subset of CLMAD which is topic matched with AISHELL-1 [19]. The preprocessing steps are as follows:

- (1) Select 3 million sentences which have small cross entropy with AISHELL-1 training transcriptions [20];
- (2) Remove the sentences whose lengths are longer than 50;
- (3) Mix the remained sentences with training transcriptions (which are duplicated 10 times to improve proportion);
- (4) Re-segment the word sequences into characters.

The information of the text data is shown in Table 1.

5.2. Experimental setup

In this paper, we employ Speech-Transformer [3, 21], a non-recurrent Seq2Seq model for speech recognition, as the backbone network. Instead of hidden states and recurrent structures of RNNs, the transformer models the context by computing attention directly. Please see [3, 12, 21] for details of the transformer.

³This subset of the external text has been shared with OneDrive: <https://1drv.ms/u/s!An08U7hvUohBb234-V-Z0Qb.Zcc>

Table 1: *The description of the text.*

	#Sentences	#Characters	Size
Training Trans.	120,098	1,730,113	6.6MB
Test Trans.	7,176	104,765	0.4MB
External Text	3,703,982	75,893,998	290MB

The acoustic features are 80-dimension Mel-filter bank features (FBANK), which are extracted every 10ms with 25ms of frame length. Each frame is spliced with three left frames. So, the input of the network is 320-dimensional. The sequence is subsampled every three frames. The Speech-Transformer consists of 6 blocks of an encoder and 6 blocks of a decoder. The dimensionality of the model is 512, and the number of inner nodes of the fully connected feed-forward network is 2048. The number of heads is 8. The modeling units of the decoder are 4232 characters, including three special symbols “<unk>”, “<sos>”, “<eos>”, which represent unknown character, start of a sentence, end of a sentence, respectively. The character embeddings is shared with the output weights of the decoder [22]. Following [3], we use Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$. The learning rate α is updated as follows:

$$\alpha = k \cdot d_{\text{model}}^{-0.5} \cdot \min(n^{-0.5}, n \cdot \text{warmup}^{-1.5}), \quad (7)$$

where d_{model} is the dimensionality of the model, n is the step number, k is a tunable parameter, and the learning rate increases linearly for *warmup* steps. We set $k = 0.5$, *warmup* = 8000. The model is trained for 50 epochs. There are utterances containing about 20K frames in one batch. The development set is used for validation. Only the model which achieves the lowest cross entropy on development set is stored as the final model.

The external RNNLM is a two layers of long short-term memory (LSTM) network. The modeling units are the same as the Seq2Seq model. The RNNLM is trained on the external text. The embedding size of the RNNLM is 300, and the number of LSTM cells of each layer is 1024. The RNNLM is trained on external text. The stochastic gradient decent (SGD) with momentum as the optimizer for training the RNNLM. The momentum is set to 0.9, and the learning rate is set to 1.0. The RNNLM is trained for 5 epochs.

For decoding, we set beam width to 5 for beam search, and maximum decoding length to 60.

5.3. Results and analysis

5.3.1. The effectiveness of external text

Firstly, we demonstrate the effectiveness of the external text data and the RNNLM. We compute the perplexities on AISHELL-1 test transcriptions, which are shown in Table 2. Note that the data is in character level, so the perplexities are relatively smaller. We can see that compared with the 3-gram with Kneser-Ney smoothing trained on training transcriptions, the 3-gram trained on external text achieves a significant reduction of perplexity. Moreover, the RNNLM achieves about a 28% relative reduction over the 3-gram trained on the external text.

5.3.2. The impact of hyper-parameters

Table 3(a) shows the the character error rates (CERs) on development set with different temperature T in Eq. (2) when λ is fixed at 0.9. The temperature controls the smoothness of the soft labels generated by the RNNLM. When it is too small, the soft labels are too sharp, and the Seq2Seq training is perturbed

Table 2: The perplexities on transcriptions of AISHELL-1 development set.

LM	PPL
3-gram (Training Trans.)	70
3-gram (Ext. Text)	47
RNNLM (Ext. Text)	34

Table 3: Comparisons of different hyper-parameters on the development set.

(a) Varying T for RNNLM softmax ($\lambda = 0.9$).

Temperature	CER%
$T = 1.0$	10.5
$T = 3.0$	8.3
$T = 5.0$	8.0
$T = 7.0$	8.2
$T = 10.0$	8.2
$T = 15.0$	9.0

(b) Varying λ for interpolation ($T = 5.0$).

Weight	CER%
$\lambda = 0.3$	9.4
$\lambda = 0.5$	9.2
$\lambda = 0.8$	9.2
$\lambda = 0.85$	8.9
$\lambda = 0.9$	8.0
$\lambda = 0.95$	9.3

heavily. When it is too large, the soft labels are too smoothed to affect the training. We can see that when the temperature is set to 5.0, the model achieves the best performance.

Then we fix T at 5.0 and evaluate the influence of λ . The parameter λ controls the proportion of the ground truth hard labels to the soft labels of the RNNLM. The results are shown in Table 3(b). We can see that when $\lambda = 0.9$, the model achieves the best performance on the development set. According to the above results in Table 3, we select $\lambda = 0.9$ and $T = 5.0$ as the final hyper-parameters. We refer to the model trained with $\lambda = 0.9$ and $T = 5.0$ as “Seq2Seq+LST” in the rest.

5.3.3. The effectiveness of the proposed approach

Table 4 gives the results on the test set of each model. “Seq2Seq” is the plain Seq2Seq without regularization. Compared to “Seq2Seq”, “Seq2Seq+LST” achieves an 18.42% relative reduction in character error rate.

We report results of two KLD based regularization approaches, namely label smoothing and unigram smoothing. For label smoothing, the prior label distribution is assumed to be a uniform distribution. The label smoothing achieves a 4.4% relative reduction in character error rate. For unigram smoothing, the prior label distribution is assumed to be the frequency of each label. The frequency is estimated on the external text. Because the unigram is too sharp, it introduces noises and affects training. We add 0.1 to the frequencies and re-normalize them to smooth the unigram. We can see that the original unigram hurts the performance, and the smoothed unigram improves the performance. Both label smoothing and unigram smoothing are effective for regularizing the model. The unigram should be smoothed for training to reduce the sharp problem.

From Table 4, we can see that “Seq2Seq+LST” outperforms both label smoothing and unigram smoothing (without shallow fusion). We analyze that the assumptions of label smoothing (uniform distribution) and unigram smoothing (unigram frequency) do not match real situations. However, LST, which is a data-driven approach, does not assume prior distributions.

We further leverage shallow fusion with the RNNLM for each model. The weight of LM is 0.1. The RNNLM is the same one which is used for LST. We can see that shallow fusion improves performances for all models. “Seq2Seq + LST” model outperforms “Seq2Seq + Label Smoothing + SF” model (CER 9.9%), which demonstrates that LST is an effective way

Table 4: Comparisons on the test set. LST represents our proposed “Learn Spelling from Teachers” approach. SF means using shallow fusion during decoding.

Model	CER%
Seq2Seq (baseline)	11.4
Seq2Seq + SF	10.5
Seq2Seq + Label Smoothing	10.9
Seq2Seq + Label Smoothing + SF	9.9
Seq2Seq + Original Unigram Smoothing	15.7
Seq2Seq + Smoothed Unigram Smoothing	10.0
Seq2Seq + Smoothed Unigram Smoothing + SF	8.7
Seq2Seq + Proposed LST	9.3
Seq2Seq + Proposed LST + SF	8.3




Figure 3: The loss curves of Seq2Seq model (left) and Seq2Seq model with LST (right). For Seq2Seq model, the training loss is lower than validation loss. However, with LST, the training loss is higher than the validation loss. Moreover, the validation loss in the right figure is a little smaller than the left one.

to improve the performance of Seq2Seq models. Moreover, the model which uses LST and shallow fusion together, i.e. “Seq2Seq + LST + SF”, achieves the best CER of 8.3%.

To further show the effect of our proposed approach, we draw the loss curves with baseline “Seq2Seq” and the proposed “Seq2Seq+LST” in Fig. 3. For “Seq2Seq” model, the training loss is lower than validation loss. However, for “Seq2Seq+LST”, the training loss is higher than validation loss. The final validation loss of “Seq2Seq+LST” is a little bit smaller than “Seq2Seq”. This result shows regularization effect of LST.

6. Conclusions

In this paper, we propose LST training approach to integrating an external RNNLM into a Seq2Seq model. An RNNLM is first trained on large scale external text data. Then, the RNNLM provides soft labels of training transcriptions to train the Seq2Seq model. We used transformer based Seq2Seq as backbone, and conducted experiments on public available Chinese datasets AISHELL-1 (speech) and CLMAD (external text). The experiments demonstrate the effectiveness of our proposed approach. We will try integrating more powerful language models into Seq2Seq systems in the future.

7. Acknowledgements

This work is supported by the National Key Research & Development Plan of China (No.2017YFB1002801), the National Natural Science Foundation of China (NSFC) (No.61425017, No.61831022, No.61773379, No.61603390), the Strategic Priority Research Program of Chinese Academy of Sciences (No.XDC02050100), and Inria-CAS Joint Research Project (No.173211KYSB20170061).

8. References

- [1] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," *international conference on acoustics, speech, and signal processing*, pp. 4945–4949, 2016.
- [2] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [3] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.
- [4] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, E. Gonina *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4774–4778.
- [5] C. Gulcehre, O. Firat, K. Xu, K. Cho, L. Barrault, H. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On using monolingual corpora in neural machine translation," *arXiv: Computation and Language*, 2015.
- [6] 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," pp. 5824–5828, 2018.
- [7] A. Sriram, H. Jun, S. Satheesh, and A. Coates, "Cold fusion: Training seq2seq models together with language models." pp. 387–391, 2018.
- [8] S. Changhao, C. Wen, G. Wang, D. Su, M. Luo, and D. Yu, "Component fusion: Learning replaceable language model component for end-to-end speech recognition system," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019.
- [9] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [10] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [11] Y. Bai, J. Tao, J. Yi, Z. Wen, and C. Fan, "CLMAD: A chinese language model adaptation dataset," in *The Eleventh International Symposium on Chinese Spoken Language Processing (ISCSLP 2018)*, 2018.
- [12] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Advances in neural information processing systems*, 2015, pp. 577–585.
- [13] Y. Kim and A. M. Rush, "Sequence-level knowledge distillation," *arXiv preprint arXiv:1606.07947*, 2016.
- [14] J. Li, M. L. Seltzer, X. Wang, R. Zhao, and Y. Gong, "Large-scale domain adaptation via teacher-student learning," *arXiv preprint arXiv:1708.05466*, 2017.
- [15] J. Andrés-Ferrer, N. Bodestab, and P. Vozila, "Efficient language model adaptation with noise contrastive estimation and kullback-leibler regularization," *Proc. Interspeech 2018*, pp. 3368–3372, 2018.
- [16] J. Chorowski and N. Jaitly, "Towards better decoding and language model integration in sequence to sequence models," *Proc. Interspeech 2017*, pp. 523–527, 2017.
- [17] G. Pereyra, G. Tucker, J. Chorowski, Ł. Kaiser, and G. Hinton, "Regularizing neural networks by penalizing confident output distributions," *arXiv preprint arXiv:1701.06548*, 2017.
- [18] J. Li and M. Sun, "Scalable term selection for text categorization," in *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, 2007.
- [19] A. Rousseau, "Xenc: An open-source tool for data selection in natural language processing," *The Prague Bulletin of Mathematical Linguistics*, vol. 100, pp. 73–82, 2013.
- [20] R. C. Moore and W. Lewis, "Intelligent selection of language model training data," in *Proceedings of the ACL 2010 conference short papers*. Association for Computational Linguistics, 2010, pp. 220–224.
- [21] S. Zhou, L. Dong, S. Xu, and B. Xu, "Syllable-based sequence-to-sequence speech recognition with the transformer in mandarin chinese," *arXiv preprint arXiv:1804.10752*, 2018.
- [22] O. Press and L. Wolf, "Using the output embedding to improve language models," *arXiv preprint arXiv:1608.05859*, 2016.