

Distilling the Knowledge of BERT for Sequence-to-Sequence ASR

Hayato Futami¹, Hirofumi Inaguma¹, Sei Ueno¹, Masato Mimura¹,
Shinsuke Sakai¹, Tatsuya Kawahara¹

¹Graduate School of Informatics, Kyoto University, Sakyo-ku, Kyoto, Japan

{surname}@sap.ist.i.kyoto-u.ac.jp

Abstract

Attention-based sequence-to-sequence (seq2seq) models have achieved promising results in automatic speech recognition (ASR). However, as these models decode in a left-to-right way, they do not have access to context on the right. We leverage both left and right context by applying BERT as an external language model to seq2seq ASR through knowledge distillation. In our proposed method, BERT generates soft labels to guide the training of seq2seq ASR. Furthermore, we leverage context beyond the current utterance as input to BERT. Experimental evaluations show that our method significantly improves the ASR performance from the seq2seq baseline on the Corpus of Spontaneous Japanese (CSJ). Knowledge distillation from BERT outperforms that from a transformer LM that only looks at left context. We also show the effectiveness of leveraging context beyond the current utterance. Our method outperforms other LM application approaches such as n -best rescoring and shallow fusion, while it does not require extra inference cost.

Index Terms: speech recognition, sequence-to-sequence models, language model, BERT, knowledge distillation

1. Introduction

End-to-end models that directly map acoustic features into symbol sequences have shown promising results in automatic speech recognition (ASR). Compared to conventional DNN-HMM hybrid systems, end-to-end models have the advantages of a simplified architecture and fast decoding. There are various choices when it comes to end-to-end models: connectionist temporal classification (CTC) [1], attention-based sequence-to-sequence (seq2seq) models [2, 3], and RNN-transducer models [4, 5]. In this study, we adopt attention-based seq2seq models.

Seq2seq ASR models use paired speech and text for training. In addition, unpaired text that is more readily available can be used to improve them. An external language model (LM) is trained separately on unpaired text, and various approaches for applying the LM to ASR have been proposed. In n -best rescoring, n -best hypotheses are obtained from ASR, followed by the addition of their LM scores, and then the best-scored hypothesis among them is selected. Language model fusion approaches such as shallow fusion [6], deep fusion [7], and cold fusion [8, 9] utilize an external LM during beam-search decoding. In shallow fusion, the linearly interpolated score from both the LM and the ASR model is used in beam search during the inference stage. More recently, knowledge distillation [10]-based LM integration has been proposed [11]. In this approach, the LM (teacher model) provides soft labels to guide the seq2seq model (student model) training. The LM is used during the training stage but is not required during the inference stage.

In the above-mentioned approaches that apply LM to seq2seq ASR, n -gram, RNNLM, or transformer [12] LM is conventionally used. We call them “unidirectional” LMs, which predict each word on the basis of its left context. In this study,

we propose to apply BERT [13] as an external LM. BERT features Masked Language Modeling (MLM) in the pre-training objective, where MLM masks a word from the input and then predicts the original word. BERT can be called a “bidirectional” LM that predicts each word on the basis of both its left and right context.

Seq2seq models decode in a left-to-right way, and therefore they do not have access to the right context during training or inference. We aim to alleviate this seq2seq’s left-to-right bias, by taking advantage of BERT’s bidirectional nature. N -best rescoring with BERT was proposed in [14, 15], but the recognition result was restricted to hypotheses from left-to-right decoding. On the other hand, BERT is difficult to use in LM fusion approaches because right (future) context that has not yet been decoded cannot be accessed during inference. To solve these issues, we propose to apply BERT to ASR through knowledge distillation. BERT (teacher model) provides soft labels using both left and right contexts of a current utterance for the seq2seq model (student model) training. Furthermore, we propose to use not only right context but also context beyond utterance boundaries during the training stage. In spontaneous ASR tasks such as presentation and conversation, the speech comprises a series of utterances. In our proposed method, previous utterances, the current utterance, and future utterances are concatenated up to the fixed length of tokens and then fed into BERT. BERT provides soft labels based on context that spans across utterances, which helps achieve better seq2seq ASR training.

2. Preliminaries and related work

2.1. Sequence-to-sequence ASR

In attention-based seq2seq ASR, we model the mapping between acoustic features and symbol sequences using two distinct networks. One is an encoder network that transforms a sequence of acoustic features into a high-level representation. The other is a decoder network that predicts a sequence of symbols using the encoded representation. At each decoding step, the decoder predicts a symbol using a relevant portion of the encoded representation and previously decoded symbols. In this study, we implemented the encoder with a multi-layer bidirectional LSTM and the decoder with a unidirectional LSTM.

Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ denote a sequence of input acoustic features. Let $\mathbf{y} = (y_1, \dots, y_N)$ denote a sequence of target symbols. The target symbols are subwords in this study, and $y_i \in \{1, \dots, V\}$, where V denotes the vocabulary size. We define the seq2seq model’s output probability of subword v for the i -th target as

$$P_{ASR}^{(i,v)} = p(v | \mathbf{X}, \mathbf{y}_{<i}) \quad (1)$$

$\mathbf{y}_{<i}$ denotes the left context of y_i , that is $\mathbf{y}_{<i} = (y_1, \dots, y_{i-1})$. During the training of a seq2seq model, we minimize the fol-

lowing cross-entropy objective:

$$\mathcal{L}_{ASR} = - \sum_{i=1}^N \sum_{v=1}^V \delta(v, y_i) \log P_{ASR}^{(i,v)} \quad (2)$$

where $\delta(v, y_i)$ becomes 1 when $v = y_i$, and 0 otherwise.

2.2. BERT

BERT [13] is a mechanism for LM pre-training that consists of a multi-layer bidirectional transformer encoder [12]. BERT can be pre-trained on a large unlabeled text and then be fine-tuned on a limited labeled text. It has shown excellent results in many downstream natural language processing tasks. BERT’s success comes from learning “deep bidirectional” representations. Previous approaches to LM pre-training such as OpenAI GPT [16] (unidirectional) and ELMo [17] (shallow concatenation of left-to-right and right-to-left RNNLMs) do not perform as well as BERT because they are not “deeply bidirectional”.

BERT originally has two pre-training objectives: Masked Language Modeling (MLM) and Next Sentence Prediction (NSP). MLM randomly replaces some of the input tokens with [MASK] tokens and then predicts the original word on the basis of its both left and right context. NSP predicts whether two input sentences appear consecutively in a corpus to model sentence relationships.

2.3. Bidirectional context in seq2seq models

Seq2seq models predict each word using its left context. Due to this autoregressive property, it is difficult for seq2seq models to leverage right context during the training and inference stages. In seq2seq decoding, later predictions depend on the accuracy of previous predictions, and therefore the issue of error accumulation arises [18]. Previous studies have addressed this issue by using right context in seq2seq ASR [19] and neural machine translation (NMT) [20, 21]. In [19], a left-to-right and a right-to-left decoder generate n -best hypotheses respectively, and the two n -best hypotheses are then concatenated to make new hypotheses. In [20], a second-pass deliberation decoder that can leverage right context was proposed. Synchronous bidirectional decoding in a single model was proposed in [21].

Meanwhile, some studies have leveraged right context during the seq2seq model training by distilling the knowledge of a “bidirectional” teacher model [22, 23]. In [22], which is a succeeding work of [11], Causal cOze completeR (COR) was proposed to model both left and right context within an utterance. In COR, the output of a stack of left-to-right transformer blocks and a stack of right-to-left ones are concatenated and fed into a subsequent fusion transformer block. Compared to BERT, it only performs shallow concatenation of two directions of transformer blocks, and as such is not “deeply bidirectional”. On the other hand, we adopt BERT, which has a simpler and more general architecture. Furthermore, we use context that spans across utterances as input to BERT for better distillation, whereas the context is limited to the current utterance in [22]. In [23], a source sequence of tokens and a target sequence of tokens are fed into BERT to generate soft labels for text-to-text transduction tasks such as NMT.

2.4. Context beyond utterance boundaries in ASR

ASR is typically done at the utterance level, but context information beyond the utterance level can help improve seq2seq ASR [24, 25, 26]. A context vector generated from the previous utterance is incorporated into the decoder state in the current utterance in these studies.

In our method, context information beyond the utterance level is not incorporated into the ASR decoder but fed into BERT to predict better soft labels for ASR. During inference, BERT is not used, and therefore our method does not add any extra procedure or component to utterance-level seq2seq ASR.

3. Proposed method

3.1. Pre-training BERT

In our proposed method, BERT is used as an external LM that predicts a masked word based on its context. We need the MLM pre-training objective itself, and therefore fine-tuning for downstream tasks is not conducted. NSP is also removed from the pre-training objective. Following RoBERTa [27], BERT’s input is packed with full-length sequences sampled contiguously from the corpus.

3.2. Distilling the knowledge of BERT

In our knowledge distillation, BERT serves as a teacher model and a seq2seq ASR model serves as a student model. Pre-trained BERT provides soft labels to guide seq2seq ASR training. These soft labels encourage the seq2seq ASR model to generate more syntactically or semantically likely results.

Seq2seq ASR training with the knowledge of BERT is formulated as follows. The speech in the corpus is split into a series of utterances, and the ASR model is trained on utterance-level data. As in Section 2.1, \mathbf{X} denotes acoustic features in an utterance, and \mathbf{y} denotes a label sequence corresponding to \mathbf{X} . We utilize context beyond the current utterance as input to BERT in our method. Let $\mathbf{y}^{(L)} = (y_1^{(L)}, \dots, y_L^{(L)})$ denote a subword sequence for previous (left) utterances and $\mathbf{y}^{(R)} = (y_1^{(R)}, \dots, y_R^{(R)})$ denote one for future (right) utterances. The length of $\mathbf{y}^{(L)}$ ($= L$) and that of $\mathbf{y}^{(R)}$ ($= R$) are decided such that the sum of L , R , and N (the label length of the current utterance) is constant (e.g. $L + R + N = 256$) and that L and R are the same (i.e. $L = R$).

We define BERT’s output probability of subword v for the i -th target label as

$$P_{BERT}^{(i,v)} = p(v | [\mathbf{y}^{(L)}; \mathbf{y}_{\setminus i}; \mathbf{y}^{(R)}]) \quad (3)$$

$$= \frac{\exp(z_v/T)}{\sum_{j=1}^V \exp(z_j/T)} \quad (4)$$

where z_j is a logit before the softmax layer and T is a temperature parameter. We obtain $\mathbf{y}_{\setminus i}$ by converting the i -th token to [MASK], that is, $\mathbf{y}_{\setminus i} = (y_1, \dots, y_{i-1}, [\text{MASK}], y_{i+1}, \dots, y_N)$. $\mathbf{y}_{\setminus i}$ is concatenated with $\mathbf{y}^{(L)}$ and $\mathbf{y}^{(R)}$, then fed into BERT as $[\mathbf{y}^{(L)}; \mathbf{y}_{\setminus i}; \mathbf{y}^{(R)}]$.

Let $P_{ASR}^{(i)}$ and $P_{BERT}^{(i)}$ denote the probability distribution for the i -th target predicted by a seq2seq ASR model and by BERT, respectively. Our goal here is to distill the knowledge of BERT and transfer it to the seq2seq ASR model by making $P_{ASR}^{(i)}$ close to $P_{BERT}^{(i)}$, as illustrated in Figure 1. Thus, we minimize the Kullback-Leibler (KL) divergence between $P_{ASR}^{(i)}$ and $P_{BERT}^{(i)}$ for each i .

$$KL(P_{BERT}^{(i)} || P_{ASR}^{(i)}) = - \sum_{v=1}^V P_{BERT}^{(i,v)} \log \frac{P_{ASR}^{(i,v)}}{P_{BERT}^{(i,v)}} \quad (5)$$

$P_{BERT}^{(i)}$ is fixed during distillation, and therefore minimizing the KL divergence over the sequence is equivalent to mini-

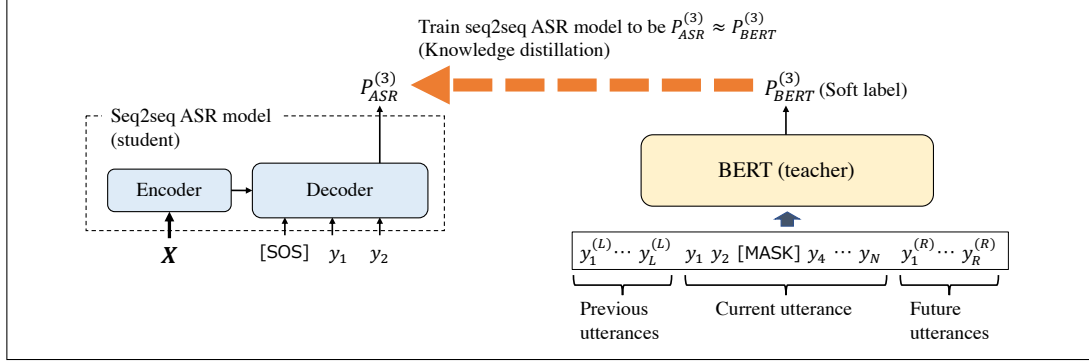


Figure 1: Illustration of our proposed method. BERT generates the soft label ($= P_{BERT}^{(3)}$) using context in which y_3 is masked and the current utterance, previous and future utterances are concatenated. The target label for $P_{ASR}^{(3)}$ is given by not only the hard label ($= y_3$) but also the soft label ($= P_{BERT}^{(3)}$).

mizing the following objective:

$$\mathcal{L}_{KD} = - \sum_{i=1}^N \sum_{v=1}^V P_{BERT}^{(i,v)} \log P_{ASR}^{(i,v)} \quad (6)$$

The final objective is linear interpolation between \mathcal{L}_{ASR} from Eq. (2) and \mathcal{L}_{KD} from Eq. (6).

$$\mathcal{L} = (1 - \alpha)\mathcal{L}_{ASR} + \alpha\mathcal{L}_{KD} \quad (0 \leq \alpha \leq 1) \quad (7)$$

This can be decomposed into a soft label based on BERT $P_{BERT}^{(i,v)}$ and an one-hot label $\delta(v, y_i)$, which can be referred as a hard label.

$$\mathcal{L} = - \sum_{i=1}^N \sum_{v=1}^V ((1 - \alpha)\delta(v, y_i) + \alpha P_{BERT}^{(i,v)}) \log P_{ASR}^{(i,v)} \quad (8)$$

We can pre-compute $P_{BERT}^{(i)}$ for all tokens in the training set. For memory efficiency, we apply top- K distillation [28]. We obtain the top- K probabilities of BERT and normalize them for distillation. BERT’s inference is generally time-consuming because it has a large set of parameters. However, this is not problematic in our method because we use BERT only for pre-computing soft labels of the training set and do not use it in the runtime.

3.3. Leveraging context beyond utterance boundaries

In our method, BERT predicts soft labels on the basis of context that spans across utterances. Tokens from previous utterances and tokens from future utterances are added to the current utterance to make up a sequence of a fixed length. We expect two benefits from looking at context beyond utterance boundaries. The first is that the ASR model can be trained with more informative soft labels. It is sometimes difficult to predict words just looking at their context within the current utterance, especially in short utterances. In this case, the top- K entries in BERT’s prediction get less syntactically or semantically relevant to corresponding hard labels. Such soft labels can have an adverse effect on seq2seq ASR training. With context beyond the current utterance, the quality of soft labels does not depend on whether the current utterance is short or long.

The other possible benefit is that we can solve the mismatch between BERT’s pre-training and distillation. While BERT is pre-trained on “full-length” sequences, the utterances are of various lengths. By adding tokens from adjacent utterances up to “full-length”, BERT is expected to perform better during distillation.

4. Experimental evaluations

4.1. Experimental conditions

We evaluated our method using the Corpus of Spontaneous Japanese (CSJ) [29] and the Balanced Corpus of Contemporary Written Japanese (BCCWJ) [30]. CSJ includes two sub-corpora, CSJ-APS and CSJ-SPS. CSJ-APS consists of about 240 hours of oral presentation speeches from academic meetings, and CSJ-SPS consists of about 280 hours of simulated presentation speeches on general topics. CSJ-eval1, which is an official test set of CSJ-APS, was used for evaluation. We also used BCCWJ-PB and BCCWJ-LB in BCCWJ as additional text for training LMs. BCCWJ-PB consists of samples extracted from published books, and BCCWJ-LB consists of samples from books registered in libraries. The text is tokenized using Byte Pair Encoding [31] of vocabulary size 7520. BCCWJ-PB and BCCWJ-LB have about 37M and 40M subword tokens, respectively. The transcriptions of CSJ-APS and CSJ-SPS have about 3.9M and 4.1M subword tokens, respectively.

In our seq2seq ASR, the encoder consists of 5 layers of bidirectional LSTMs with 320 hidden states, and the decoder consists of a single LSTM layer with 320 hidden states. We trained the seq2seq model on CSJ-APS with a batch size of 25 utterances. The average token length of utterances was about 24 (maximum: 118, minimum: 1). We used Adam [32] with the learning rate of $1e-4$ for optimizing the ASR model. SpecAugment [33] was applied to the acoustic features. We also applied label smoothing [34]. In target labels, the probability of 0.1 was distributed uniformly over all classes. In decoding, we used beam search with a beam width of 5.

We trained BERT and a unidirectional transformer LM for comparison. BERT and the transformer LM have 6 layers of transformer blocks with 512 hidden states and 8 attention heads. We trained them on BCCWJ-PB and BCCWJ-LB first, then on the transcriptions of CSJ-APS and CSJ-SPS. We sampled 150 sequences of length 256 for each pre-training step. In BERT, we randomly selected 8% of the tokens in each sequence and replaced them with [MASK] tokens. We used Adam with the learning rate of $1e-4$ with learning rate warmup over the first 10% of total steps and linear decay. K for top- K distillation was set to 8 in all our experiments. The temperature parameter T in Eq. (4) and the distillation weight α in Eq. (7) were adjusted using the development set. Our code for the proposed method is available ¹.

¹<https://github.com/hfutami/distill-bert-for-seq2seq-asr>

Table 1: The performance for ASR trained on CSJ-APS (240h) with knowledge distillation-based LM integration. “TrfLM(uni)” in the “LM” column denotes the transformer LM.

LM	Context size	WER(%)
—	—	10.31
TrfLM(uni)	utterance	9.89
TrfLM(uni)	256	10.01
BERT	utterance	9.53
BERT	256	9.19

Table 2: Ablation studies on the length of BERT’s input during pre-training and distillation.

Context size		WER(%)
Pre-training	Distillation	
64	utterance	9.91
64	64	9.69
128	utterance	9.62
128	128	9.40
256	utterance	9.53
256	64	9.28
256	128	9.28
256	256	9.19

4.2. Experimental results

We evaluated our method through ASR experiments. First, we compared the performances of the ASR models trained using BERT and the unidirectional transformer LM (TrfLM(uni)) as a teacher model. We also evaluated the effectiveness of using context beyond the current utterance. The ASR results are shown in Table 1. The result denoted as “utterance” in the “Context size” column corresponds to the ASR model guided by soft labels based on context within the current utterance. The result denoted as “256” in the “Context size” column corresponds to that guided by soft labels based on context of length 256 that spans across utterances. In TrfLM(uni), we added only previous utterances to the current utterance as context. The first line in the table denotes the baseline ASR without distillation. As shown in Table 1, knowledge distillation-based LM integration consistently improved the performance of the ASR model. We found that distillation from BERT outperformed that from the TrfLM(uni), which indicates the effectiveness of leveraging both left and right context. We also found that incorporating context beyond the current utterance was important for distillation from BERT by comparing line 4 and 5 in the table. This result improved the WER by 10.86% relatively over the baseline.

Next, we compared our method with two other LM application approaches. Shallow fusion (SF) and n -best rescoring were applied to the baseline and were compared to the ASR model trained with our method (the last line in Table 1). As shown in Figure 2, our method outperformed both shallow fusion and n -best rescoring regardless of the beam width. We also applied shallow fusion and n -best rescoring to the ASR model trained through our method and obtained some improvements, which were not as large as those applied to the baseline. This can be interpreted as the ASR model with our method had already learned the effect of applying an external LM through distillation.




Figure 2: Comparisons and combinations with other LM application approaches. “SF” denotes shallow fusion.

Table 3: The performance for ASR trained on an increased amount of data (520h, both CSJ-APS and CSJ-SPS).

LM	Context size	WER(%)
—	—	8.43
BERT	256	7.85

Next, we conducted ablation studies on context size during pre-training and distillation. The results are shown in Table 2. We found that the use of longer context in the pre-training led to better ASR performance. We also found that distillation from BERT using the same context size as pre-training performed best.

Finally, to see the effect of an increased amount of training data for ASR in our method, we trained another ASR model on both CSJ-APS and CSJ-SPS (total 520h) and evaluated the performance. As shown in Table 3, our method was still effective for this better baseline ASR model trained on an increased amount of paired data.

5. Conclusions

BERT can be pre-trained on a large unpaired text, and can also leverage not only left context but also right context that seq2seq ASR models do not have access to. In this study, we have proposed a method in which the knowledge of BERT is transferred to seq2seq ASR through a knowledge distillation framework and demonstrated its effectiveness through experiments. We found that distillation from BERT yields better ASR performance than that from the transformer LM. We also found that the knowledge of BERT based on context that spans across utterances further improved the performance of seq2seq ASR. Our proposed method outperformed other LM application approaches such as n -best rescoring and shallow fusion, including rescoring with BERT, even though our method does not require extra inference cost. As a future work, we will investigate applying other LM pre-training mechanisms such as XLNet [35] and ELECTRA [36] to ASR.

6. References

- [1] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd International Conference on Machine Learning*, 2006, pp. 369–376.
- [2] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4960–4964.
- [3] 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 28*, 2015, pp. 577–585.
- [4] A. Graves, “Sequence transduction with recurrent neural networks,” *ArXiv*, 2012.
- [5] E. Battenberg, J. Chen, R. Child, A. Coates, Y. G. Y. Li, H. Liu, S. Sathesh, A. Sriram, and Z. Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2017, pp. 206–213.
- [6] J. Chorowski and N. Jaitly, “Towards better decoding and language model integration in sequence to sequence models,” in *Proc. Interspeech*, 2017, pp. 523–527.
- [7] Çağlar Gülçehre, O. Firat, K. Xu, K. Cho, L. Barrault, H.-C. Lin, F. Bougares, H. Schwenk, and Y. Bengio, “On using monolingual corpora in neural machine translation,” *ArXiv*, 2015.
- [8] A. Sriram, H. Jun, S. Sathesh, and A. Coates, “Cold fusion: Training seq2seq models together with language models,” in *Proc. Interspeech*, 2018, pp. 387–391.
- [9] 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,” *2018 IEEE Spoken Language Technology Workshop (SLT)*, pp. 369–375, 2018.
- [10] G. E. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *ArXiv*, 2015.
- [11] Y. Bai, J. Yi, J. Tao, Z. Tian, and Z. Wen, “Learn spelling from teachers: Transferring knowledge from language models to sequence-to-sequence speech recognition,” in *Proc. Interspeech*, 2019, pp. 3795–3799.
- [12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems 30*, 2017, pp. 5998–6008.
- [13] 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)*, 2019, pp. 4171–4186.
- [14] J. Shin, Y. Lee, and K. Jung, “Effective sentence scoring method using BERT for speech recognition,” in *Proceedings of The Eleventh Asian Conference on Machine Learning*, 2019, pp. 1081–1093.
- [15] J. Salazar, D. Liang, T. Q. Nguyen, and K. Kirchhoff, “Masked language model scoring,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 2699–2712.
- [16] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, “Improving language understanding by generative pre-training,” *OpenAI technical report*, 2018.
- [17] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, “Deep contextualized word representations,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2018, pp. 2227–2237.
- [18] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer, “Scheduled sampling for sequence prediction with recurrent neural networks,” in *Advances in Neural Information Processing Systems 28*, 2015, pp. 1171–1179.
- [19] M. Mimura, S. Sakai, and T. Kawahara, “Forward-backward attention decoder,” in *Proc. Interspeech*, 2018, pp. 2232–2236.
- [20] Y. Xia, F. Tian, L. Wu, J. Lin, T. Qin, N. Yu, and T.-Y. Liu, “Deliberation networks: Sequence generation beyond one-pass decoding,” in *Advances in Neural Information Processing Systems 30*, 2017, pp. 1784–1794.
- [21] L. Zhou, J. Zhang, and C. Zong, “Synchronous bidirectional neural machine translation,” *Transactions of the Association for Computational Linguistics*, pp. 91–105, 2019.
- [22] Y. Bai, J. Yi, J. Tao, Z. Tian, Z. Wen, and S. Zhang, “Integrating whole context to sequence-to-sequence speech recognition,” *ArXiv*, 2019.
- [23] Y.-C. Chen, Z. Gan, Y. Cheng, J. Liu, and J. Liu, “Distilling knowledge learned in BERT for text generation,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 7893–7905.
- [24] S. Kim and F. Metze, “Dialog-context aware end-to-end speech recognition,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 434–440.
- [25] R. Masumura, T. Tanaka, T. Moriya, Y. Shinohara, T. Oba, and Y. Aono, “Large context end-to-end automatic speech recognition via extension of hierarchical recurrent encoder-decoder models,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 5661–5665.
- [26] S. Kim, S. Dalmia, and F. Metze, “Gated embeddings in end-to-end speech recognition for conversational-context fusion,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 1131–1141.
- [27] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “RoBERTa: A robustly optimized BERT pretraining approach,” *ArXiv*, 2019.
- [28] X. Tan, Y. Ren, D. He, T. Qin, and T.-Y. Liu, “Multilingual neural machine translation with knowledge distillation,” in *International Conference on Learning Representations*, 2019.
- [29] K. Maekawa, “Corpus of spontaneous Japanese : its design and evaluation,” *Proceedings of The ISCA & IEEE Workshop on Spontaneous Speech Processing and Recognition (SSPR)*, 2003.
- [30] K. Maekawa, M. Yamazaki, T. Ogiso, T. Maruyama, H. Ogura, W. Kashino, H. Koiso, M. Yamaguchi, M. Tanaka, and Y. Den, “Balanced corpus of contemporary written Japanese,” *Lang. Resour. Eval.*, pp. 345–371, 2014.
- [31] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016, pp. 1715–1725.
- [32] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, 2015.
- [33] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [34] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2818–2826.
- [35] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, “XLNet: Generalized autoregressive pretraining for language understanding,” in *Advances in Neural Information Processing Systems 32*, 2019, pp. 5753–5763.
- [36] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning, “ELECTRA: Pre-training text encoders as discriminators rather than generators,” in *International Conference on Learning Representations*, 2020.