



HypR: A comprehensive study for ASR hypothesis revising with a reference corpus

Yi-Wei Wang¹, Ke-Han Lu², Kuan-Yu Chen¹

¹National Taiwan University of Science and Technology, Taiwan

²National Taiwan University, Taiwan

yiwei@nlp.csie.ntust.edu.tw, d12942024@ntu.edu.tw, kychen@mail.ntust.edu.tw

Abstract

With the development of deep learning, automatic speech recognition (ASR) has made significant progress. To further enhance the performance of ASR, revising recognition results is one of the lightweight but efficient manners. Various methods can be roughly classified into N -best reranking modeling and error correction modeling. The former aims to select the hypothesis with the lowest error rate from a set of candidates generated by ASR for a given input speech. The latter focuses on detecting recognition errors in a given hypothesis and correcting these errors to obtain an enhanced result. However, we observe that these studies are hardly comparable to each other, as they are usually evaluated on different corpora, paired with different ASR models, and even use different datasets to train the models. Accordingly, we first concentrate on providing an ASR hypothesis revising (HypR) ¹ dataset in this study. HypR contains several commonly used corpora (AISHELL-1, TED-LIUM 2, and LibriSpeech) and provides 50 recognition hypotheses for each speech utterance. The checkpoint models of ASR are also published. In addition, we implement and compare several classic and representative methods, showing the recent research progress in revising speech recognition results. We hope that the publicly available HypR dataset can become a reference benchmark for subsequent research and promote this field of research to an advanced level.

Index Terms: Hypothesis revising, N -best reranking, error correction, reference benchmark

1. Introduction

Automatic speech recognition (ASR) systems have significantly evolved due to the growth of deep learning research in recent years. The conventional pipeline system, which mainly consists of an acoustic model, a language model, a lexicon, and a search algorithm, has gradually been phased out [1, 2]. End-to-end-based methods, which can directly convert input speech into text, have created a new paradigm because of their simple and elegant design [3, 4]. To further boost the recognition performance, there are two main research directions. On the one hand, advanced model architectures and training objectives have been introduced [5, 6]. On the other hand, revising the ASR output to reduce the recognition errors is a lightweight but efficient strategy [7, 8, 9].

The ASR hypothesis revising methods can be roughly classified into two categories: N -best reranking methods and error correction models [10, 11]. Since the highest-scoring recognition result may not necessarily have the fewest errors, N -best reranking methods thus aim to select the best hypothesis among

a set of candidates for a given input speech. That is to say, this school of methods is designed to distinguish and select the most accurate hypothesis. Various methods can be further divided into token-level [7, 12, 13], sentence-level [8, 14], and comparison-based strategies [15, 16, 17, 18]. Alternatively, an error correction model usually consists of an error detector and an error corrector. More formally, the former aims to detect recognition errors, while the latter concentrates on correcting these errors [9, 19]. It is worthwhile to note that, along with the research trend, recently, more and more error correction models have attempted to blend these two components into an end-to-end approach [20, 21, 22, 23, 24].

2. ASR Hypothesis Revising Methods

To further enhance the recognition results, a number of ASR hypothesis revising methods are presented. N -best reranking modeling and error correction modeling are two major streams.

2.1. N -best Reranking Methods

The principal idea of N -best reranking methods is to determine an accurate score for each hypothesis, which should be highly negatively correlated with the error rate. That is, the higher the score, the lower the error rate. Therefore, according to the scores, the hypothesis with the highest score can be selected as the output from a set of recognition candidates generated by ASR for a given input speech. Depending on the granularity or manner in which scores are computed for each hypothesis, various approaches can be categorized into token-level, sentence-level, and comparison-based approaches.

Token-level methods usually employ advanced language models to calculate the score by considering the probability of each token occurring in the hypothesis. Causal language modeling (CLM) and masked language modeling (MLM) are two popular usage schemes [7]. For a given hypothesis $X = \{x_1, x_2, \dots, x_L\}$, causal language modeling decomposes X into a sequence of conditional probabilities using the chain rule and then multiplies these probabilities to obtain the final score:

$$Score_{CLM}(X) = \prod_{i=1}^L P(x_i | x_{<i}), \quad (1)$$

where $x_{<i}$ denotes the partial sequence before x_i . Even though the calculation is mathematically sound, the probability of each token only depends on what has occurred before. However, there is no doubt that all the tokens in each hypothesis are known, especially in the context of the ASR hypothesis revising task. It is certainly possible to use bidirectional information to calculate the score for each token in the hypothesis. Therefore, in contrast to CLM, masked language modeling determines the

¹<https://github.com/Alfred0622/HypR>

score as follows:

$$Score_{MLM}(X) = \prod_{i=1}^L P(x_i|x_{\setminus i}), \quad (2)$$

where $x_{\setminus i}$ represents the entire hypothesis except x_i . Compared with CLM, MLM considers more contextual information to calculate the score of each token, so it is usually expected to obtain a higher accuracy than CLM.

Typically, **sentence-level** methods first represent each candidate using a vector representation and then convert it into a score to identify the most promising hypothesis [8, 14]. During training, the ground truth for each hypothesis can be the error rate, the normalized ranking in a set of candidates, or the MLM (or CLM) score. By doing this, on the one hand, the computational cost is vastly reduced compared to the token-level methods during inference. This is because CLM and MLM methods have to perform calculations for each token in a hypothesis. In contrast, sentence-level methods first encapsulate a hypothesis into a vector, and then only one step is required for the hypothesis to obtain its own score. On the other hand, several off-the-shelf and powerful sentence embedding methods can be utilized to distill and encapsulate lexical, semantic, and even latent information from the hypothesis into a vector representation. By combining the advanced vector representation with task-oriented ground truth (e.g., error rate or normalized ranking) for training, a more reliable score can be obtained compared to the scores provided by token-level methods, which only consider token regularities.

Unlike token-level and sentence-level methods that focus on calculating a score for each candidate, **comparison-based** methods frame the N -best reranking task as a comparison problem. Specifically, the comparison-based models are trained to indicate the best among a set of candidates, so the absolute score of each candidate is not important [15]. Limited by the token numbers of model input, pairwise comparison models are most common [16, 17, 18]. Accordingly, for a set of candidates, each hypothesis should be compared with all other candidates one by one, and the hypothesis with the highest number of wins is selected as the output. Conceptually, this school of methods simplifies the reranking task from generating an accurate score for each hypothesis to determining the more correct hypothesis from a set of candidates.

2.2. Error Correction Models

Apart from the N -best reranking methods, another straightforward idea to enhance the ASR performance is to detect recognition errors and then correct these misrecognized tokens. These strategies are collectively called the error correction models, and they are usually implemented by cascading an error detection module and an error correction module in a pipeline manner [9, 19, 25, 26]. As deep learning advances, upcoming research is gradually shifting toward implementation in an end-to-end and sequence-to-sequence manner. Therefore, error correction modeling can be completed in just one step, reducing the overall model complexity and avoiding possible error propagation between components. In this context, many studies attempt to leverage powerful pre-trained language models as a cornerstone to boost performance further [22, 24], while some struggle to leverage more information extracted from ASR, such as by referring to the N -best hypotheses [21].

Several interesting observations can be made and discussed by comparing the modeling philosophies of the two streams. First, error correction modeling gives a new corrected result as

the ASR output, which may not be included in the original candidate set generated by ASR. Hence, error correction modeling seems to have the potential to exceed the performance of N -best reranking modeling because the latter has limitations imposed by the set of the N -best candidates. Second, due to its end-to-end and sequence-to-sequence manner, error correction modeling may face challenges that make training difficult, while the training process of N -best reranking modeling is usually simpler and straightforward. Finally, without the speech utterance as a reference, error correction modeling can only detect and correct misrecognized tokens by taking into account local and global contextual textual clues. As a result, unexpected errors may occur after the error correction models are processed. In contrast, N -best reranking modeling appears to be a conservative but more stable alternative.

2.3. Other Representatives and Recent Progresses

Both the mainstream methods (i.e., N -best reranking modeling and error correction modeling) have their own advantages and disadvantages, and almost all of the N -best reranking and error correction models consider acoustic information implicitly, indirectly, or even omitted. Therefore, recent research has attempted to combine acoustic and lexical information to revise the recognition results in a cascade or an end-to-end manner [10, 27]. Although these approaches further improve performance, they are limited in requiring additional speech input. Moreover, there are also studies that use a word lattice as input to increase the search space and generate better recognition results [28]. Nevertheless, encoding the topology of a word lattice as input is challenging, and these methods usually require large amounts of supervised data to train the model to leverage the lattice [29, 30].

Due to the rapid growth of generative large language models (LLMs), a popular paradigm in natural language processing research is zero-shot or one-shot strategies using these models, as they achieve remarkable results on various tasks. Following the research trend, their ability to handle the task of ASR hypothesis revising is of interest [11, 31, 32]. Additionally, combining an LLM with a pre-trained large-scale speech model through a cross-modal fusion technology for ASR hypothesis revising is also a novel research direction [33].

3. HypR: The ASR Hypothesis Revising Benchmark

In order to promote research development and examine the progress of research, this study proposes a speech recognition hypothesis revising (HypR) benchmark. HypR includes commonly used datasets, including AISHELL-1 [34], TED-LIUM 2 [35], and LibriSpeech [36]. The Transformer-based CTC-attention ASR model, built using the ESPnet toolkit, is trained for each dataset [37]. In more detail, the model architecture consists of an encoder and a decoder, constructed from 12 and 6 Transformers, respectively. CTC and cross-entropy losses are applied to train the model in a multi-task learning manner. No data augmentation methods are applied to enhance the performance of ASR. During decoding, the beam search algorithm is used to generate 50 hypotheses for each speech utterance, with or without shallow fusion with an extra language model. The language models of AISHELL-1 and TED-LIUM 2 are LSTM-based models trained from scratch, while LibriSpeech applies a Transformer-based model downloaded from the ESPnet toolkit. The detailed configurations of each dataset are summarized in

Table 1. In addition to the N -best hypotheses, checkpoints for the ASR models are also released in order to easily reproduce the contents of the benchmark, sustain development, and allow their possible use in the future.

4. Experiments

4.1. Experimental Setup

In the experiments, we compare the performance of several classic and representative N -best reranking methods and error correction models on the proposed HypR benchmark. It should be noted that without additional instructions, due to the limitations of computing resources and for a fair comparison, we used the top 10 hypotheses for each speech utterance to train, tune, and test all models in the experiments. All the hyper-parameters are selected by the development set. The model configurations, hyper-parameters, and training settings are described in detail in the repository of the benchmark. Our implementations are based on the HuggingFace library and the pre-trained language models used in this study are cloned from it [38].

4.1.1. N -best Reranking Methods

The N -best reranking methods can be divided into token-level, sentence-level, and comparison-based methods. For the token-level methods, CLM and MLM are considered in our experiment. We adopt GPT₂_{base} [39] and BERT_{base} [40] to implement the CLM and MLM methods, respectively.

RescoreBERT [8] and PBERT [14] are selected to stand for sentence-level modeling in our experiments. The fundamental difference between them is the training objective. The former aims to predict the MLM score for each given hypothesis, while the latter targets distinguishing the hypothesis with the lowest word error rate from the others. Two auxiliary discriminative training losses for RescoreBERT, namely the minimum word error rate (MWER) and matching word error distribution (MWED), are also implemented and compared in this study [8].

BERT_{sem} and BERT_{alsem} are two comparison-based methods evaluated in this study [18]. Both are implemented using BERT_{base} as the backbone model, and the settings follow previous research. BERT_{sem} concatenates a hypothesis with one of the other hypotheses and then feeds the result into the BERT model, which is fine-tuned to reveal the degree of confidence that the former candidate may have a lower error rate than the latter. Based on BERT_{sem}, BERT_{alsem} is a variant that not only stacks Bi-LSTM layers on top of BERT to obtain better representations but also considers the scores calculated by ASR to try to provide a more indicative score. Hence, the major difference between them is the statistics used in the model. BERT_{sem} only considers textual information, while BERT_{alsem} takes acoustic, linguistic, and textual information into account.

Inspired by BERT_{alsem}, which introduces LSTM layers on top of BERT to obtain better representations, we stack a Bi-LSTM layer on top of BERT for the PBERT model. Therefore, the contextualized token embeddings generated by the last layer of BERT would also go through the Bi-LSTM layer. After that, average and max pooling are performed to collect token-level clues. The two pooling vectors, the hypothesis embedding generated by BERT (i.e., the embedding of the special [CLS] token), and the ASR decoding score are concatenated to represent the target hypothesis. We denote the simple extension as PBERT+LSTM hereafter. The training objectives of PBERT and PBERT+LSTM are exactly the same.

	AISHHELL-1	TED-LIUM 2	LibriSpeech
Training Epochs	50	50	100
Batch Size	64	32	128
Learning Rate	1.0	2.0	10.0
Grad-Accum	2	2	2
Warmup Steps	25,000	25,000	25,000
Vocab Size	4,233	502	5,002
CTC Weight (Training)	0.3	0.3	0.3
CTC Weight (Decoding)	0.5	0.3	0.4
LM Weight (Decoding)	0.7	0.5	0.7

Table 1: Detailed configurations of the ASR systems.

4.1.2. Error Correction Models

BART is a widely used end-to-end and sequence-to-sequence model, so we choose it as a representative of error correction modeling [22, 24]. During training, BART aims to convert each input hypothesis into its corresponding ground truth, while only the top hypothesis is used to generate the final output for a given speech during the inference stage. Based on BART, we further implement N -best Plain and N -best Alignment methods to access more information from N -best hypotheses to try to make more accurate corrections [21]. For both methods, we use the top four hypotheses for realizing, following the setting of the original study.

4.1.3. Leveraging Large Language Models

As we lack direct access to the model parameters of ChatGPT, we employ a methodology known as prompt engineering to assess its zero-shot capabilities in this study. We interact with ChatGPT through the official API² and ask it to select or generate the most probable response from a list of 10 hypotheses. For the former, we ask ChatGPT to output a number indicating which of the hypotheses in the candidate set is selected as the output. Such a strategy limits the results to the predefined set of candidates. In contrast, for the latter, ChatGPT is asked to generate a response by considering the information provided by a set of candidates, allowing for a more open-ended result.

4.2. Performance Analysis

Table 2 summarizes the word error rate performance for various methods on the proposed HypR benchmark. Additionally, we take the AISHHELL-1 corpus as an example to report the model sizes of these methods and calculate the average computation time (in seconds) they are required to generate a final output. These statistics are listed in Table 3. Based on these results, some interesting observations and results worth exploring can be obtained. First, focusing on baseline systems, we can conclude that using language models for decoding does improve the ASR performance. Among the different datasets, the performance gap for LibriSpeech is particularly large. A possible reason for this is that the language model used for LibriSpeech is a Transformer-based model, which is more robust than the LSTM-based models used for the other datasets. By comparing the differences between the Oracle and Top1 systems, we can see that there is still a long way to go to improve recognition results, which also illustrates the value of this line of research. When comparing Oracle and Random systems, the experiments show significant differences for all datasets, demonstrating the diversity of the N -best hypotheses in HypR.

Next, we turn to examining the N -best reranking methods. At first glance, although there is not much difference in the performance of various methods, sentence-level modeling seems to be the best choice. Looking at the results in more

²gpt-3.5-turbo-0613, <https://api.openai.com>

	w/o LM								w/ LM							
	AISHELL-1		TED-LIUM2		LibriSpeech				AISHELL-1		TED-LIUM2		LibriSpeech			
	Dev	Test	Dev	Test	Dev	Other	Clean	Other	Dev	Test	Dev	Test	Clean	Other	Clean	Other
<i>Baseline</i>																
Top1	7.34	8.33	14.29	12.07	4.44	11.94	4.64	12.00	6.66	7.23	11.74	9.72	2.55	7.12	2.85	7.21
Oracle	3.98	4.74	11.49	8.55	2.28	8.17	2.43	8.14	3.70	4.14	9.05	6.28	1.15	4.48	1.35	4.45
Random	11.68	12.39	15.50	14.22	7.90	14.66	8.10	14.79	11.26	11.64	13.14	12.13	6.49	11.07	6.80	11.10
<i>N-best Reranking</i>																
<i>Token-level</i>																
CLM	6.05	6.80	13.06	10.61	3.42	9.89	3.64	9.99	5.92	6.55	11.17	8.92	2.50	6.86	2.84	7.05
MLM[7]	5.17	6.03	12.75	10.39	3.31	9.77	3.59	9.81	4.98	5.58	10.64	8.58	2.39	6.74	2.76	6.91
<i>Sentence-level</i>																
RescoreBERT[8]	5.29	6.12	13.10	10.61	3.41	9.91	3.69	9.98	4.99	5.56	11.16	8.77	2.38	6.83	2.79	6.96
+MWER	5.29	6.11	13.03	10.60	3.40	9.89	3.68	9.93	4.99	5.59	10.90	8.67	2.39	6.80	2.79	6.94
+MWED	5.30	6.10	13.02	10.66	3.37	9.91	3.66	9.97	5.06	5.62	11.06	8.83	2.38	6.81	2.79	6.95
PBERT[14]	5.08	5.89	12.81	10.35	3.36	9.83	3.63	9.97	4.76	5.35	10.49	8.51	2.35	6.69	2.74	6.81
PBERT+LSTM	5.04	5.84	12.72	10.25	3.34	9.82	3.58	9.90	4.71	5.30	10.46	8.31	2.35	6.50	2.64	6.66
<i>Comparison-based</i>																
BERT _{sem} [18]	5.27	6.05	13.08	10.53	3.46	10.18	3.83	10.20	5.00	5.63	10.80	8.85	2.50	6.84	2.85	6.96
BERT _{alsem} [18]	5.25	6.08	12.94	10.55	3.49	10.12	3.79	10.16	5.31	5.95	10.99	9.12	2.45	6.89	2.83	6.89
<i>Error Correction</i>																
BART	6.62	7.44	13.51	11.77	4.21	11.02	4.50	11.26	6.94	7.53	12.32	10.66	2.71	7.26	3.05	7.42
+4-best Plain [21]	7.10	8.10	14.14	11.81	4.24	11.60	4.44	11.69	6.38	6.99	11.96	10.04	2.69	7.21	3.00	7.37
+4-best Alignment [21]	7.37	8.35	14.89	12.99	4.58	11.97	4.81	12.09	6.76	7.31	12.44	10.50	2.85	7.40	3.18	7.64
<i>Using LLM</i>																
Selection	6.57	7.36	13.79	11.67	4.79	11.07	4.90	11.33	8.16	8.60	12.38	11.15	5.44	9.57	5.72	9.72
Generation	6.67	7.71	16.34	13.95	5.59	12.19	6.07	12.77	6.97	7.40	14.84	14.02	5.52	9.86	6.01	10.70

Table 2: Experimental results of various classic and representative methods on the HypR benchmark.

detail, for token-level modeling, MLM achieves better results than CLM, as expected. The reason for this might be that MLM considers bidirectional information to determine the score for each token in the hypothesis, while CLM only takes unidirectional clues into account. It is worth noting that although CLM and MLM have almost the same computational cost in theory, CLM is faster than MLM in practice. This is because CLM can be implemented with the attention masking strategy to reduce computation costs associated with iterative decoding, but MLM cannot do so because of information leakage [12, 41]. For that reason, CLM is more than ten times faster than MLM in our experiments. Next, we compare RescoreBERT, PBERT, and PBERT+LSTM, which are all sentence-level methods. Since MLM scores are not precisely positively correlated with obtaining the lowest error rate, there is no doubt that PBERT and PBERT+LSTM could obtain better results than RescoreBERT. There are two more notable observations. On the one side, the efficiency of RescoreBERT combined with the MWER or MWED does not appear in HypR. On the other side, it should be emphasized that our proposed simple extension PBERT+LSTM is better than other sentence-based methods and shows superior results compared to all other methods. However, due to the addition of the Bi-LSTM layer, PBERT+LSTM not only has more model parameters but is also slower than PBERT and RescoreBERT. Finally, comparison-based methods, including BERT_{sem} and BERT_{alsem}, have a high computational cost and cannot provide better results than token-level and sentence-level methods. Although BERT_{alsem} considers acoustic and language model scores while BERT_{sem} does not, the former only achieves better results on LibriSpeech, especially when paired with a language model for generating N -best hypotheses.

In the third set of experiments, we consider error correction modeling. BART and two extensions are examined in the experiments. Surprisingly, both extensions do not always deliver better results than BART. BART obtains better results than the two extensions in the "w/o LM" setting, and BART+4-best Plain outperforms BART and BART+4-best Alignment if the quality of the N -best hypotheses is at a certain level (i.e., w/ LM). The results are somewhat reasonable because these two extensions mainly focus on leveraging more hypotheses to enhance the results, so the quality of the hypotheses will definitely

	Model Parameters	Latency
<i>N-best Reranking</i>		
<i>Token-level</i>		
CLM	102.07M	0.014
MLM	102.29M	0.185
<i>Sentence-level</i>		
RescoreBERT	102.27M	0.013
PBERT	102.27M	0.014
PBERT+LSTM	140.44M	0.020
<i>Comparison-based</i>		
BERT _{sem}	102.27M	0.138
BERT _{alsem}	109.26M	0.134
<i>Error Correction</i>		
BART	140.19M	0.186
+ 4-best Plain	140.19M	0.191
+ 4-best Alignment	142.55M	0.212

Table 3: The number of model parameters and decoding time (in seconds) for various models.

affect the results.

In the last set of experiments, we are interested in the results obtained using ChatGPT. Even though ChatGPT is a generative model with a well-known advantage in performing tasks related to natural language generation, the hypothesis selection strategy produces better results than the generation strategy in most cases. We also observe that the Selection and Generation strategies perform even worse than the Top1 system in the case of decoding with a language model. The reason for this may be that the recognition candidates generated by fusing a language model are fluent, reasonable, and grammatical. Therefore, it is difficult for ChatGPT to select or generate better results. By comparing LLM-based methods with other models, the experiments show that finding better ways to leverage large language models to revise the ASR hypothesis remains urgent.

5. Conclusion

This paper presents a publicly available ASR hypothesis revising (HypR) benchmark, and a series of experiments are conducted on the corpus to perform fair analyses and comparisons. Several observations and discussions are also addressed in this study. In the future, we plan to enrich and diversify the benchmark. This will include utilizing different ASR models like RNN-T and Whisper to generate hypotheses, examining the revising performance on domain-specific corpora such as AMI and SEMEAN datasets, and more.

6. Acknowledgement

This work was supported by the National Science and Technology Council of Taiwan under Grants NSTC 112-2636-E-011-002 and NSTC 112-2628-E-011-008-MY3. We thank the National Center for High-performance Computing of the National Applied Research Laboratories (NARLabs) in Taiwan for providing computational and storage resources.

7. References

- [1] F. Jelinek, *Statistical methods for speech recognition*. MIT Press, 1998.
- [2] D. Yu and L. Deng, *Automatic speech recognition*. Springer, 2016, vol. 1.
- [3] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proceedings of ICASSP*, 2016.
- [4] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, “Hybrid CTC/Attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [5] F.-H. Yu, K.-Y. Chen, and K.-H. Lu, “Non-autoregressive ASR modeling using pre-trained language models for chinese speech recognition,” *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 30, p. 1474–1482, apr 2022.
- [6] K.-H. Lu and K.-Y. Chen, “A context-aware knowledge transferring strategy for CTC-based ASR,” in *Proceedings of SLT*, 2023.
- [7] J. Salazar, D. Liang, T. Q. Nguyen, and K. Kirchhoff, “Masked language model scoring,” in *Proceedings of ACL*, 2020.
- [8] L. Xu *et al.*, “RescoreBERT: Discriminative Speech Recognition Rescoring with BERT,” in *Proceedings of ICASSP*, 2022.
- [9] Y. Leng *et al.*, “FastCorrect: Fast Error Correction with Edit Alignment for Automatic Speech Recognition,” in *Proceedings of NIPS*, vol. 34, 2021.
- [10] C.-H. Kuo and K.-Y. Chen, “Correcting, Rescoring and Matching: An N-best list selection framework for speech recognition,” in *Proceedings of APSIPA*, 2022.
- [11] C. Chen, Y. Hu, C.-H. H. Yang, S. M. Siniscalchi, P.-Y. Chen, and E.-S. Chng, “Hyporadise: An open baseline for generative speech recognition with large language models,” in *Proceedings of NIPS*, vol. 36, 2024.
- [12] K. Song, Y. Leng, X. Tan, Y. Zou, T. Qin, and D. Li, “Transormer: Transformer for sentence scoring with sliding language modeling,” in *Proceedings of NIPS*, vol. 35, 2022.
- [13] H. Huang and F. Peng, “An empirical study of efficient ASR rescoring with Transformers,” *arXiv preprint arXiv:1910.11450*, 2019.
- [14] S.-H. Chiu and B. Chen, “Innovative BERT-based reranking language models for speech recognition,” in *Proceedings of SLT*, 2021.
- [15] M. Nie, M. Yan, and C. Gong, “Prompt-based Re-ranking Language Model for ASR,” in *Proceedings of INTERSPEECH*, 2022.
- [16] I. Illina and D. Fohr, “DNN-based semantic rescoring models for speech recognition,” in *Proceedings of TSD*, 2021.
- [17] A. Ogawa, M. Delcroix, S. Karita, and T. Nakatani, “Rescoring N-best speech recognition list based on one-on-one hypothesis comparison using encoder-classifier model,” in *Proceedings of ICASSP*, 2018.
- [18] D. Fohr and I. Illina, “BERT-based semantic model for rescoring N-best speech recognition list,” in *Proceedings of INTERSPEECH*, 2021.
- [19] H.-W. Wang, B.-C. Yan, Y.-C. Wang, and B. Chen, “Effective ASR error correction leveraging phonetic, semantic information and N-best hypotheses,” in *Proceedings of APSIPA*, 2022.
- [20] J. Guo, T. N. Sainath, and R. J. Weiss, “A spelling correction model for end-to-end speech recognition,” in *Proceedings of ICASSP*, 2019.
- [21] L. Zhu, W. Liu, L. Liu, and E. Lin, “Improving ASR error correction using N-best hypotheses,” in *Proceedings of ASRU*, 2021.
- [22] Y. Zhao, X. Yang, J. Wang, Y. Gao, C. Yan, and Y. Zhou, “BART based semantic correction for mandarin automatic speech recognition system,” in *Proceedings of INTERSPEECH*, 2021.
- [23] J. Guo, M. Wang, X. Qiao, D. Wei, H. Shang, Z. Li, Z. Yu, Y. Li, C. Su, M. Zhang, S. Tao, and H. Yang, “Ucorrect: An unsupervised framework for automatic speech recognition error correction,” in *Proceedings of ICASSP*, 2023.
- [24] S. Dutta, S. Jain, A. Maheshwari, S. Pal, G. Ramakrishnan, and P. Jyothi, “Error correction in ASR using sequence-to-sequence models,” *arXiv preprint arXiv:2202.01157*, 2022.
- [25] Y. Leng *et al.*, “FastCorrect 2: Fast error correction on multiple candidates for automatic speech recognition,” in *Proceedings of EMNLP*, 2021.
- [26] Y. Leng, X. Tan, W. Liu, K. Song, R. Wang, X.-Y. Li, T. Qin, E. Lin, and T.-Y. Liu, “Softcorrect: Error correction with soft detection for automatic speech recognition,” in *Proceedings of the AAAI*, vol. 37, no. 11, 2023.
- [27] J. Cai, M. Sunkara, X. Li, A. Bhatia, X. Pan, and S. Bodapati, “Masked audio text encoders are effective multi-modal rescorers,” in *Proceedings of ACL*, 2023.
- [28] S.-F. Tsai, S.-C. Kuo, R.-Y. Lyu, and J.-S. R. Jang, “Ensemble and re-ranking based on language models to improve ASR,” in *Proceedings of ISCSLP*, 2022.
- [29] R. Ma, H. Li, Q. Liu, L. Chen, and K. Yu, “Neural lattice search for speech recognition,” in *Proceedings of ICASSP*, 2020.
- [30] L. Dai, L. Chen, Z. Zhou, and K. Yu, “LatticeBART: Lattice-to-lattice pre-training for speech recognition,” in *Proceedings of ICASSP*, 2022.
- [31] R. Ma, M. Qian, P. Manakul, M. Gales, and K. Knill, “Can generative large language models perform ASR error correction?” *arXiv preprint arXiv:2307.04172*, 2023.
- [32] C.-H. H. Yang, Y. Gu, Y.-C. Liu, S. Ghosh, I. Bulyko, and A. Stolcke, “Generative speech recognition error correction with large language models and task-activating prompting,” in *Proceedings of ASRU*, 2023.
- [33] S. Radhakrishnan, C.-H. Yang, S. Khan, R. Kumar, N. Kiani, D. Gomez-Cabrero, and J. Tegnér, “Whispering LLaMA: A cross-modal generative error correction framework for speech recognition,” in *Proceedings of EMNLP*, 2023.
- [34] 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 *Proceedings of O-COCOSDA*, 2017.
- [35] A. Rousseau, P. Deléglise, and Y. Estève, “Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks,” in *Proceedings of LREC*, 2014.
- [36] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proceedings of ICASSP*, 2015.
- [37] 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 *Proceedings of INTERSPEECH*, 2018.
- [38] T. Wolf *et al.*, “Transformers: State-of-the-art natural language processing,” in *Proceedings of EMNLP*, 2020.
- [39] A. Radford *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, 2019.
- [40] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL*, 2019.
- [41] J. Shin, Y. Lee, S. Yoon, and K. Jung, “Fast and accurate deep bidirectional language representations for unsupervised learning,” *arXiv preprint arXiv:2004.08097*, 2020.