



Factual Consistency Oriented Speech Recognition

Naoyuki Kanda, Takuya Yoshioka, Yang Liu

Microsoft, USA

{nakanda, tayoshio, yaliu10}@microsoft.com

Abstract

This paper presents a novel optimization framework for automatic speech recognition (ASR) with the aim of reducing hallucinations produced by an ASR model. The proposed framework optimizes the ASR model to maximize an expected factual consistency score between ASR hypotheses and ground-truth transcriptions, where the factual consistency score is computed by a separately trained estimator. Experimental results using the AMI meeting corpus and the VoxPopuli corpus show that the ASR model trained with the proposed framework generates ASR hypotheses that have significantly higher consistency scores with ground-truth transcriptions while maintaining the word error rates close to those of cross entropy-trained ASR models. Furthermore, it is shown that training the ASR models with the proposed framework improves the speech summarization quality as measured by the factual consistency of meeting conversation summaries generated by a large language model.

Index Terms: speech recognition, speech summarization, hallucination errors

1. Introduction

Thanks to substantial progress in deep learning techniques, such as advanced model architectures [1, 2] and training criteria [3, 4, 5, 6] as well as the utilization of large-scale training data [7, 8, 9, 10], the automatic speech recognition (ASR) accuracy has significantly improved. For example, latest ASR systems achieved lower word error rates (WERs) than human transcribers for many public test sets, such as the LibriSpeech [11, 12] and Switchboard benchmarks [13, 14]. It was also reported that a large ASR model trained on 68K hours of transcribed data curated from the Web achieved domain robustness and accuracy closer to that of the human transcribers [10].

However, the ASR technology is still prone to recognition errors under different conditions. To complicate matters, an ASR model with a strong language modeling (LM) capability can produce fluent but factually ungrounded, or hallucinated, ASR errors, regardless of whether the LM is part of an end-to-end ASR model [15, 16, 17] or explicitly provided as with hybrid ASR [18]. These errors often make the transcriptions difficult or confusing for humans and downstream natural language processing (NLP) modules, such as an automatic summarization module, to comprehend and thus have a detrimental effect.

Many approaches were explored to alleviate the negative effect that the ASR errors have on the downstream tasks. A conventional approach is using N-best hypotheses or lattices in the subsequent NLP modules [19, 20]. However, feeding the N-best/lattice representations requires significant changes to the NLP modules, and how to do so is not obvious for

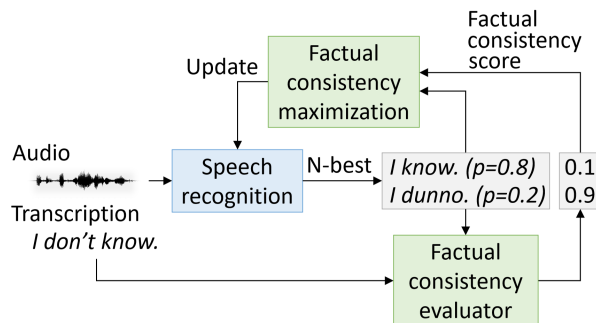


Figure 1: Factual consistency maximization of ASR.

many state-of-the-art NLP models pre-trained on regular texts, such as GPT-3 [21]. Another approach is joint optimization of the ASR and NLP modules either by end-to-end modeling [22, 23, 24, 25] or by using an NLP task-oriented loss [26]. While this approach has shown promising improvements on the downstream NLP tasks, using a task-specific architecture or loss may make the ASR model suitable only for the specific task it was trained for. In addition, joint modeling approach is difficult to apply to a task that handles very long sequences, such as a speech summarization task. For example, the current text summarization model is often trained to process more than 8,000 tokens, which roughly corresponds to 30–40 audio minutes. It is not trivial to train joint models for such long audio inputs, and therefore existing joint ASR and summarization models are limited to handling short durations such as 100 seconds [24] or 300 seconds [25].

In this paper, we propose an ASR optimization scheme for factual consistency, with the aim of reducing hallucinations of the ASR model. Figure 1 shows an overview of our proposed method along with illustrative transcriptions and scores. In this example, the ground-truth transcription of the input audio is “I don’t know.” Meanwhile, the ASR system has produced “I know.” and “I dunno.” with posterior probabilities of 0.8 and 0.2, respectively. In terms of the WER ignoring the punctuations, the second hypothesis would be worse than the first one (67% vs. 33%). However, the second hypothesis is factually or semantically more consistent with the ground-truth transcription. Our proposed optimization method encourages the ASR model to produce the second hypothesis over the first one by maximizing an expected factual consistency score.

Section 2 reviews related works. Section 3 describes the proposed factual consistency maximization framework. Section 4 reports experimental results using the AMI meeting corpus [27] and the VoxPopuli corpus [28], where the effectiveness of our method is demonstrated for both utterance-wise evaluation and speech summarization. Section 5 concludes the paper.

2. Related works

2.1. Quality measurement for generated texts

A lot of effort has been made for measuring the quality of texts generated by ASR or NLP models. While the ASR models are typically evaluated with the WER or character error rate depending on the language, several task-specific metrics were also utilized, such as argument WER [22] for spoken language understanding tasks. In NLP, various quality metrics were proposed for different tasks, including BLEU [29], ROUGE [30], and METEOR [31]. Recent work showed metrics based on pre-trained LMs are better correlated with human evaluation scores. BERTScore [32], MoverScore [33], and several other methods were proposed to estimate a general semantic similarity between two sentences based on pre-trained LMs.

Factual consistency between two texts [34, 35, 36], which is commonly used for evaluating natural language generation systems, is one key approach to taking account of hallucinations in the text quality assessment. In this work, we mainly use UniEval [37] to compute the factual consistency score because of its superior correlation with human ratings. UniEval was trained by fine-tuning T5-large [38] based on artificially manipulated documents to provide multiple explainable dimensions of generated texts such as consistency and fluency, resulting in multiple models with different flavors. UniEval-sum was specially trained for summarization to evaluate multiple aspects related to summarization (e.g., consistency, fluency, coherence) while UniEval-fact was trained only for factual consistency scoring. Both models produce a factual consistency score between 0 (low consistency) and 1 (high consistency) given two texts. In our work, we used UniEval-fact for ASR optimization and utterance-wise evaluation while UniEval-sum for speech summarization evaluation. We also used FactCC [34], which is another factual consistency evaluator based on BERT [39], in the utterance-wise evaluation to further validate the effectiveness of the proposed method.

2.2. ASR optimization

An ASR model is usually trained with a cross-entropy (CE) loss between the model’s output distribution and the ground-truth labels. Since the CE loss does not directly minimize the WER, alternative training schemes such as minimum WER training [3, 5, 6] are also widely adopted. Our work can be considered as a variant of minimum WER training, where we optimize the ASR model based on the UniEval-based factual consistency score instead of the WER.

Joint optimization of ASR and NLP models, such as end-to-end speech summarization [24, 25] or end-to-end spoken natural language understanding [22, 23], is also related to our work. Notably, Rao et al. [26] proposed to integrate evaluation metrics for intent classification and slot filling into the ASR training criterion. While Rao’s work shares the same backbone in the optimization scheme with ours, we attempt to avoid using a task-specific metric and conduct evaluation with a task that the model is not directly trained for.

3. Factual consistency maximization training of ASR

3.1. Training objective function

Suppose we have ASR training dataset $\mathcal{D} = \{X_r, Y_r\}_{r \in \Omega}$, where X_r is the r -th training audio sample, Y_r is the corre-

sponding ground-truth transcription, and Ω is the sample index set. We propose to maximize the following training criterion:

$$\mathcal{F}^{\text{FC}} = \sum_{r \in \Omega} \bar{C}_r, \quad (1)$$

where \bar{C}_r is the expected consistency score for the r -th training sample, which is estimated as follows:

$$\bar{C}_r = \sum_{Y \in \mathcal{B}(X_r)} \hat{P}(Y|X_r) \{|Y_r| \cdot \text{Consistency}(Y; Y_r)\}. \quad (2)$$

Here, $\mathcal{B}()$ is a beam search function which returns the N -best hypotheses for X_r based on the current ASR model parameters. Function $\text{Consistency}(Y; Y_r)$ calculates a consistency score of Y measured against the ground-truth transcription Y_r , which is obtained by a factual consistency evaluator. $|Y_r|$ is the number of words in Y_r . Finally, $\hat{P}(Y|X_r)$ is the normalized posterior probability of Y that is defined as

$$\hat{P}(Y|X_r) = \frac{P(Y|X_r)}{\sum_{Y' \in \mathcal{B}(X_r)} P(Y'|X_r)}, \quad (3)$$

where $P(Y|X_r)$ is the raw posterior probability of Y given X_r .

3.2. Training procedure

Now, we describe the training procedure for attention encoder-decoder-based ASR models [17]. The training algorithms for other ASR model types, such as connectionist temporal classification [15] or recurrent neural network transducer [16], can be derived based on similar formulations.

With attention encoder-decoder-based ASR, the decoder module iteratively generates an output distribution over recognition tokens (i.e., characters, subwords, or words). Given the output $o_{n,i}$, where n is the decoder iteration index and i is the token index, the derivative of the objective function with respect to $\log(o_{n,i})$ is calculated for each N -best hypothesis $Y \in \mathcal{B}(X_r)$ if and only if the token index at the n -th position of Y is equal to i .

$$\frac{\partial \mathcal{F}^{\text{FC}}}{\partial \log(o_{n,i})} = \hat{P}(Y|X_r) \{|Y_r| \cdot \text{Consistency}(Y; Y_r) - \bar{C}_r\}. \quad (4)$$

Otherwise, $\log(o_{n,i})$ is zero.

The overall training procedure is as follows. For each training sample $\{X_r, Y_r\}$, we first decode X_r with beam search to generate the N -best hypotheses Y based on the current model parameters. We then compute the normalized posterior probability of each hypothesis based on Eq. (3). Next, we compute the empirical expected consistency score \bar{C}_r per Eq. (2). We then compute the training objective derivative with respect to $\log(o_{n,i})$ by using Eq. (4). Finally, we perform back-propagation to compute the gradient, and update the model parameter values based on the gradient ascent algorithm.

Note that maximizing the factual consistency does not necessarily improve the WER as illustrated in Fig. 1. Also, the factual consistency maximization tends to encourage a hypothesis with fewer output tokens to avoid hallucination errors. Therefore, it is important to have some safeguard to avoid severe deletion errors, such as interpolating a conventional CE loss. In this work, we adopted a simple early-stopping approach where we applied the factual consistency maximization to a well-trained model based on the CE loss, and limited the number of updates to a relatively small number.

4. Experiments

We evaluated the proposed ASR training framework in two scenarios. In the first experiment, both the training and evaluation of ASR models used short utterances (Section 4.1) consistently. This experiment was carried out to validate the correctness and effectiveness of our formulation. In the second experiment, we conducted a speech summarization experiment (Section 4.2). We applied speech summarization to the meeting transcriptions generated by the ASR models and measured the impacts of the different ASR models on the summarization quality. This experiment was performed to investigate whether optimizing the ASR models for factual consistency helps improve a downstream task by using the speech summarization as an example. Note that improving the summarization quality is not straightforward since the ASR models are optimized at the utterance level despite the evaluation being performed with much longer segments. We also used a state-of-the-art large language model-based (LLM) summarization model to obtain relevant conclusions, unlike previous speech summarization works.

4.1. Utterance-wise evaluation

4.1.1. Experiment settings

In this experiment, we evaluated ASR models by using both the utterance-level factual consistency score and WER to examine the impact of the proposed training method. We used the AMI meeting corpus [27] and the VoxPopuli corpus [28].

The AMI meeting corpus contains approximately 100 hours of meeting recordings captured by both independent headset microphones (IHM) and multiple distant microphones. In our experiment, we used the IHM audio with ground-truth utterance boundaries. We followed the Kaldi recipe [40] for partitioning the data into training, development, and evaluation sets, resulting in 80.2 hours, 9.7 hours, and 9.1 hours of recordings for these splits, respectively. Unlike a conventional ASR experiment using case-normalized and non-punctuated transcriptions for training and testing, we used the case-sensitive and punctuated-transcriptions based on the official transcription¹ throughout our experiments. We chose to do so because the factual consistency evaluator (in our case, UniEval [37]) is case- and punctuation-sensitive. Only for the WER scoring, we applied the text normalizer provided by [10].

The VoxPopuli corpus consists of 100K hours of unlabelled speech data in 23 languages, as well as 1.8K hours of transcribed speech in 16 languages. In our experiment, we used the transcribed English subset. The official dataset is split into training, development, and testing sets, which contains 522.6 hours, 5.0 hours, and 4.9 hours of speech, respectively. As with the AMI meeting corpus, we used the case-sensitive, punctuated transcriptions. Note that we excluded roughly 3% of the development and testing sets (58 out of 1753 sentences in the development set and 52 out of 1842 sentences in the testing set) that did not have case-sensitive, punctuated transcriptions.

We used the Whisper Base model [10] as our initial ASR model. In the experiment with the AMI meeting corpus, we first fine-tuned the Whisper Base on the AMI training set with the CE loss. We conducted 2,500 training iterations with 8 GPUs, each consuming mini-batches of 15,000 frames. We used a linear decay learning rate schedule with an initial learning rate of 1e-5. After the CE-loss-based fine-tuning, we further updated

¹We used the words and punctuations of the official annotation without modifications. The only exception was that we removed underscores from a word (e.g., X.M.L. was converted to XML).

Table 1: WER (%), UniEval-fact consistency score (UE) and FactCC consistency score (FCC) for AMI-IHM development and evaluation sets. FCM: factual consistency maximization.

| ASR model | AMI-IHM dev | | | AMI-IHM eval | | |
|--------------|-------------|--------------|--------------|--------------|--------------|--------------|
| | WER (↓) | UE (↑) | FCC (↑) | WER (↓) | UE (↑) | FCC (↑) |
| Whisper Base | 19.7 | 0.727 | 0.925 | 20.3 | 0.719 | 0.913 |
| ↪ CE-loss | 11.5 | 0.785 | 0.930 | 12.6 | 0.787 | 0.932 |
| ↪ FCM | 11.4 | 0.799 | 0.942 | 12.5 | 0.801 | 0.940 |

Table 2: WER (%), UniEval-fact consistency score (UE) and FactCC consistency score (FCC) for VoxPopuli development and test sets. FCM: factual consistency maximization.

| ASR model | VoxPopuli dev* | | | VoxPopuli test* | | |
|--------------|----------------|--------------|--------------|-----------------|--------------|--------------|
| | WER (↓) | UE (↑) | FCC (↑) | WER (↓) | UE (↑) | FCC (↑) |
| Whisper Base | 9.4 | 0.755 | 0.896 | 9.4 | 0.747 | 0.887 |
| ↪ CE-loss | 8.2 | 0.766 | 0.901 | 8.3 | 0.757 | 0.886 |
| ↪ FCM | 8.4 | 0.795 | 0.920 | 8.5 | 0.791 | 0.911 |

* Utterances without case and punctuated transcriptions were excluded.

the model based on the proposed factual consistency maximization by using the same AMI training data. We used UniEval-fact to compute the consistency score and used an N-best size of 4 to compute the gradient. We conducted 2,500 training iterations with 8 GPUs, each consuming 1 sample for one training iteration. A linear decay learning rate schedule was used, starting with a learning rate of 1e-6. The VoxPopuli experiment was conducted with almost the same setting but using 6,250 training iterations with 8 GPUs (i.e., a total of 50,000-minibatch consumption) for both the CE loss-based training and factual consistency maximization.

In the evaluation, we performed beam search decoding with a beam size of 4. We set the language tag to English to enforce the model to output only English transcriptions [10]. We evaluated the WER by applying the text normalizer provided in [10]. We also computed the factual consistency score of the ASR hypothesis given the ground-truth transcription using two different factual consistency evaluators: UniEval-fact [37] and FactCC [34]. For UniEval-fact, we report the average factual consistency score over the test utterances. For FactCC, we report the ratio of the hypotheses that were judged to be consistent with the ground-truth transcriptions as FactCC makes a binary decision as to the factual consistency. For both consistency metrics, a higher score means the ASR hypothesis is more factually consistent with the ground-truth transcription.

4.1.2. Results

The results for AMI-IHM and VoxPopuli are shown in Tables 1 and 2, respectively. We can see that the CE-loss-based fine-tuning achieved significant gains in both WER and consistency score for most evaluation sets except for the FCC score on VoxPopuli test set. Furthermore, we can see that the proposed factual consistency maximization further improved all the consistency metrics while keeping the WER almost intact. For VoxPopuli, as one may expect, marginal WER degradation was observed because our training objective was not aimed at minimizing the label prediction errors as illustrated in Fig. 1.

4.2. Speech summarization evaluation

4.2.1. Experiment settings

We also evaluated the proposed ASR training method by performing automatic summarization using the ASR-generated transcriptions. In this experiment, we used the AMI meeting

corpus.

We opted to create our own experiment setup due to the challenges that we faced when we attempted to use existing speech summarization corpora. We initially tried to use the ROUGE score [30] based on the human-crafted abstractive summaries provided in the AMI official annotations. However, a preliminary experiment revealed that the number of sessions in the AMI development and evaluation sets was too small (18 and 16 sessions, respectively) to discuss the statistical differences in the summarization quality between different systems. We also considered other speech summarization datasets, namely the How2 dataset [41] and the SLUE-TED corpus [25]. However, the How2 dataset provides only feature files that are not compatible with our ASR models, and the SLUE-TED corpus has yet to be released at the time of writing. In addition to the lack of the evaluation data, recent studies also reported that the correlation between the ROUGE and human scores is much weaker than LM-based metrics (e.g., [37]).

With these considerations, we designed our speech summarization experiment as follows.

1. Split the AMI development and evaluation set recordings into non-overlapping chunks of 60 seconds.
2. Conduct the following steps for each 60-s chunk.
 - (a) Apply ASR. We used the IHM audio with ground-truth segmentation.
 - (b) Convert the ASR hypotheses to a speaker-attributed transcription format with a template of “Speaker [mic-index]: [ASR-hypothesis]\n”. We used the headset microphone index as the speaker index and sorted the hypotheses based on the start time of each utterance. For example, three ASR hypotheses consisting of “Hello. (mic-index=1, start=1s)”, “Hi, how are you? (mic-index=2, start=3s)”, and “I’m fine. (mic-index=1, start=5s)” would be converted to “Speaker 1: Hello.\nSpeaker 2: Hi, how are you?\nSpeaker 1: I’m fine.\n”.
 - (c) Generate a summary text for the speaker-attributed transcription obtained in the previous step. We used InstructGPT (text-davinci-002) [42] to generate the summary by prompting the model as “[speaker-attributed transcription]\nSummarize the conversation above.\n”. We employed the least randomness configuration by setting temperature, top_p, and max_tokens at 0.0, 1.0, and 200, respectively.
 - (d) Compute the factual consistency score of the generated summary by using the ground-truth speaker-attributed transcription as the reference, which was formatted in the same way as 2 (b). We used UniEval-sum [37] for the consistency scoring.
3. Compute the average of the factual consistency scores over all the 60-s chunks.

As a result, we obtained 574 and 542 speech summarization testing materials (i.e., 60-s chunks) for the development and evaluation sets, respectively, which enabled us to draw statistically relevant conclusions regarding the summarization quality. Note that UniEval-sum can also produce other metrics, such as the fluency and coherence of the generated summaries. As these metrics are irrelevant to the ASR errors (that is, the fluency and coherence scores are determined primarily by the summarization model quality rather than the quality of the input text produced by the ASR), we only report the UniEval-sum factual consistency result, which would best reflect the impact of the

Table 3: Summarization consistency score based on UniEval-sum for AMI development and evaluation sets. FC: factual consistency.

| ASR model | Summarization consistency (\uparrow) | |
|---------------------------------------|--|--------------|
| | AMI dev | AMI eval |
| Whisper Base | 0.697 | 0.739 |
| \hookrightarrow CE-loss fine-tuning | 0.698 | 0.745 |
| \hookrightarrow FC maximization | 0.706 | 0.748 |
| Ground-truth transcription | 0.709 | 0.753 |

ASR performance on the summarization.

4.2.2. Results

Table 3 shows the speech summarization experiment result. The ASR models trained in the utterance-wise evaluation experiment (Section 4.1) were used. We also measured the consistency score for the summaries generated by using the ground-truth transcriptions as input. This can be considered as the upper bound of the consistency score that can be achieved with the employed summarization model. First, we can observe that the summarization consistency score obtained with Whisper Base was significantly lower than that of the ground-truth transcription. This shows the negative impact of the ASR errors on the summarization quality. It can also be seen that the CE fine-tuning slightly improved the consistency score. The consistency score for the development set was only improved by 0.1% despite the significant WER improvement from 19.7% to 11.5% (Table 1). On the other hand, the proposed factual consistency maximization further improved the consistency scores for both the development and evaluation sets, coming close to the ground-truth-based upper bound scores.

We further conducted paired t-tests for the results of Table 3. The consistency score difference between the summary texts obtained with Whisper Base (0.739) and those with the CE fine-tuned model (0.745) was judged to be not significant with 95% confidence. On the other hand, the difference between the summarization with Whisper Base (0.739) and that with the proposed factual consistency-maximized model (0.748) was judged to be significant with 95% confidence. Finally, the difference between the summarization with the CE fine-tuned model and that with the factual consistency-maximized model was also judged to be significant with 95% confidence for the development set (0.698 vs 0.706) while it was not the case for the evaluation set (0.745 vs 0.748), which we would attribute to the yet limited number of evaluation samples. Overall, we observed a clear trend showing the effectiveness of the proposed factual consistency maximization method even when using the powerful LLM for the speech summarization.

5. Conclusions

In this paper, we proposed a novel ASR training framework, which optimizes the ASR model for maximizing the factual consistency between an ASR hypothesis and a ground-truth transcription. In our experiments using the AMI meeting corpus and the VoxPopuli corpus, we showed that the ASR model optimized by the proposed framework produced hypotheses which had significantly higher consistency scores with the ground-truth transcriptions while keeping the WER almost intact. We also showed that using the ASR models trained with the proposed framework improved the speech summarization quality for meeting conversations, demonstrating the usefulness of the proposed method in a challenging downstream NLP task.

6. References

- [1] A. Vaswani, N. Shazeer, N. Parmar *et al.*, “Attention is all you need,” in *Proc. NIPS*, 2017, pp. 6000–6010.
- [2] A. Gulati, J. Qin, C.-C. Chiu *et al.*, “Conformer: Convolution-augmented Transformer for speech recognition,” *Proc. Interspeech*, pp. 5036–5040, 2020.
- [3] A. Graves and N. Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *Proc. ICML*. PMLR, 2014, pp. 1764–1772.
- [4] D. Povey, V. Peddinti, D. Galvez *et al.*, “Purely sequence-trained neural networks for ASR based on lattice-free MMI,” in *Proc. Interspeech*, 2016, pp. 2751–2755.
- [5] M. Shannon, “Optimizing expected word error rate via sampling for speech recognition,” *Proc. Interspeech*, pp. 3537–3541, 2017.
- [6] R. Prabhavalkar, T. N. Sainath, Y. Wu *et al.*, “Minimum word error rate training for attention-based sequence-to-sequence models,” in *Proc. ICASSP*, 2018, pp. 4839–4843.
- [7] S. H. K. Parthasarathi and N. Strom, “Lessons from building acoustic models with a million hours of speech,” in *Proc. ICASSP*, 2019, pp. 6670–6674.
- [8] N. Kanda, G. Ye, Y. Wu *et al.*, “Large-scale pre-training of end-to-end multi-talker ASR for meeting transcription with single distant microphone,” in *Proc. Interspeech*, 2021, pp. 3430–3434.
- [9] Y. Zhang, D. S. Park, W. Han *et al.*, “BigSSL: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [10] A. Radford, J. W. Kim, T. Xu *et al.*, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.04356*, 2022.
- [11] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [12] D. Amodei, S. Ananthanarayanan, R. Anubhai *et al.*, “Deep Speech 2: End-to-end speech recognition in English and Mandarin,” in *Proc. ICML*, 2016, pp. 173–182.
- [13] J. J. Godfrey, E. C. Holliman, and J. McDaniel, “SWITCHBOARD: Telephone speech corpus for research and development,” in *Proc. ICASSP*, vol. 1, 1992, pp. 517–520.
- [14] W. Xiong, J. Droppo, X. Huang *et al.*, “Achieving human parity in conversational speech recognition,” *arXiv preprint arXiv:1610.05256*, 2016.
- [15] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006, pp. 369–376.
- [16] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [17] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, “End-to-end continuous speech recognition using attention-based recurrent NN: First results,” in *NIPS Workshop on Deep Learning*, 2014.
- [18] F. Seide, G. Li, and D. Yu, “Conversational speech transcription using context-dependent deep neural networks,” in *Proc. Interspeech*, 2011, pp. 437–440.
- [19] D. Hakkani-Tür, F. Béchet, G. Riccardi, and G. Tur, “Beyond ASR 1-best: Using word confusion networks in spoken language understanding,” *Computer Speech & Language*, vol. 20, no. 4, pp. 495–514, 2006.
- [20] F. Morbini, K. Audhkhasi, R. Artstein *et al.*, “A reranking approach for recognition and classification of speech input in conversational dialogue systems,” in *Proc. SLT*, 2012, pp. 49–54.
- [21] T. Brown, B. Mann, N. Ryder *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [22] P. Haghani, A. Narayanan, M. Bacchiani *et al.*, “From audio to semantics: Approaches to end-to-end spoken language understanding,” in *Proc. SLT*, 2018, pp. 720–726.
- [23] M. Rao, A. Raju, P. Dheram *et al.*, “Speech to semantics: Improve ASR and NLU jointly via all-neural interfaces,” in *Proc. Interspeech*, 2020, pp. 876–880.
- [24] R. Sharma, S. Palaskar, A. W. Black, and F. Metze, “End-to-end speech summarization using restricted self-attention,” in *Proc. ICASSP*, 2022, pp. 8072–8076.
- [25] S. Shon, S. Arora, C.-J. Lin *et al.*, “SLUE phase-2: A benchmark suite of diverse spoken language understanding tasks,” *arXiv preprint arXiv:2212.10525*, 2022.
- [26] M. Rao, P. Dheram, G. Tiwari *et al.*, “Do as i mean, not as i say: Sequence loss training for spoken language understanding,” in *Proc. ICASSP*, 2021, pp. 7473–7477.
- [27] J. Carletta, S. Ashby, S. Bourban *et al.*, “The AMI meeting corpus: A pre-announcement,” in *International workshop on machine learning for multimodal interaction*, 2006, pp. 28–39.
- [28] C. Wang, M. Riviere, A. Lee *et al.*, “VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” in *Proc. ACL-IJCNLP*, 2021, pp. 993–1003.
- [29] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: a method for automatic evaluation of machine translation,” in *Proc. ACL*, 2002, pp. 311–318.
- [30] C.-Y. Lin, “ROUGE: A package for automatic evaluation of summaries,” in *Text Summarization Branches Out*, 2004, pp. 74–81.
- [31] S. Banerjee and A. Lavie, “METEOR: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proc. of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.
- [32] T. Zhang, V. Kishore, F. Wu *et al.*, “BERTScore: Evaluating text generation with BERT,” *arXiv preprint arXiv:1904.09675*, 2019.
- [33] W. Zhao, M. Peyrard, F. Liu *et al.*, “MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance,” in *Proc. EMNLP-IJCNLP*, 2019, pp. 563–578.
- [34] W. Kryściński, B. McCann, C. Xiong, and R. Socher, “Evaluating the factual consistency of abstractive text summarization,” in *Proc. EMNLP*, 2020, pp. 9332–9346.
- [35] A. Wang, K. Cho, and M. Lewis, “Asking and answering questions to evaluate the factual consistency of summaries,” *arXiv preprint arXiv:2004.04228*, 2020.
- [36] M. Cao, Y. Dong, J. Wu, and J. C. K. Cheung, “Factual error correction for abstractive summarization models,” in *Proc. EMNLP*, 2020, pp. 6251–6258.
- [37] M. Zhong, Y. Liu, D. Yin *et al.*, “Towards a unified multi-dimensional evaluator for text generation,” in *Proc. EMNLP*, 2022, pp. 2023–2038.
- [38] C. Raffel, N. Shazeer, A. Roberts *et al.*, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [39] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
- [40] D. Povey, A. Ghoshal, G. Boulianne *et al.*, “The Kaldi speech recognition toolkit,” in *Proc. ASRU*, 2011.
- [41] R. Sanabria, O. Caglayan, S. Palaskar *et al.*, “How2: a large-scale dataset for multimodal language understanding,” in *Proceedings of the Workshop on Visually Grounded Interaction and Language*. NeurIPS, 2018.
- [42] L. Ouyang, J. Wu, X. Jiang *et al.*, “Training language models to follow instructions with human feedback,” *arXiv preprint arXiv:2203.02155*, 2022.