



Low-Latency Sequence-to-Sequence Speech Recognition and Translation by Partial Hypothesis Selection

Danni Liu, Gerasimos Spanakis, Jan Niehues

Department of Data Science and Knowledge Engineering, Maastricht University, the Netherlands
{danni.liu, jerry.spanakis, jan.niehues}@maastrichtuniversity.nl

Abstract

Encoder-decoder models provide a generic architecture for sequence-to-sequence tasks such as speech recognition and translation. While offline systems are often evaluated on quality metrics like word error rates (WER) and BLEU scores, latency is also a crucial factor in many practical use-cases. We propose three latency reduction techniques for chunk-based incremental inference and evaluate their accuracy-latency trade-off. On the 300-hour How2 dataset, we reduce latency by 83% to 0.8 second by sacrificing 1% WER (6% rel.) compared to offline transcription. Although our experiments use the Transformer, the partial hypothesis selection strategies are applicable to other encoder-decoder models. To reduce expensive re-computation as new chunks arrive, we propose to use a unidirectionally-attending encoder. After an adaptation procedure to partial sequences, the unidirectional model performs on-par with the original model. We further show that our approach is also applicable to speech translation. On the How2 English-Portuguese speech translation dataset, we reduce latency to 0.7 second (−84% rel.) while incurring a loss of 2.4 BLEU points (5% rel.) compared to the offline system.

Index Terms: low-latency, sequence-to-sequence models, speech recognition, speech translation

1. Introduction

Encoder-decoder models with attention [1, 2] opened the possibility of training many sequence tasks in an “end-to-end” fashion. For automatic speech recognition (ASR), the roles of the acoustic and language models can be unified into one architecture [3, 4]. Likewise, for speech translation, the traditionally cascaded modules for ASR and translation can be integrated [5, 6]. In offline use-cases, promising results have been achieved, particularly with the recent adoption of self-attention mechanisms [7, 8, 9, 10]. However, latency still poses a challenge when applying encoder-decoder-attention models to online inference. Under offline training, the decoder is conditioned on full encoded sequences via a soft attention distribution. For incremental inference, when directly using a system trained on full-sequences, the train-test condition mismatch is likely to degrade performance. Several approaches have been proposed to constrain the source-target attention distribution, notably monotonic attention [11, 12, 13] and the Neural Transducer [14, 15], where the training procedure is modified accordingly to account for possible alignments between input segments and output symbols.

In this work, we take a different perspective: We start from a full-sequence model and study how to adjust its inference-time behavior for optimal latency-quality trade-off. Moreover, we aim for a single model for both offline and online inference. While the experiments use the Transformer [16], our methods are agnostic to the underlying sequence representation mecha-




Figure 1: Our incremental decoding framework, where an utterance is split into fixed-size chunks, and a subset of previous chunk-level outputs conditions the decoding of the next chunk.

nisms. Based on the observation that small input chunks contain limited acoustic context, we eliminate unstable predictions by selectively outputting chunk-level hypotheses. This partial hypothesis selection leads to an implicit look-ahead, as the model becomes more conservative with unstable predictions at chunk level. To allow re-using decoder states as input streams in, we use an encoder with no dependency on future context. To account for the different source-target attention distribution under incremental inference, we use a simple yet effective procedure to adapt to partial inputs. The experiments show the effectiveness of our approach on ASR and speech translation.

2. Related Work

Recurrent neural network transducer [17], recurrent neural aligner [18] and their self-attentional variants [19, 20] are suitable to online end-to-end ASR as they do not rely on global source-target attention. For encoder-decoder-attention models, some form of input chunking is needed. The alignment between chunks and output tokens is derived via external alignments [21, 22] or learned implicitly by connectionist temporal classification loss [23, 24]. To achieve adaptive chunk sizes, apart from monotonic chunkwise attention [12], triggered attention [25, 26] and adaptive computation time [27] are used. In simultaneous text translation, the read-write decision of the decoder is modeled by fixed schedules [28, 29], reinforcement learning agents [30, 31], or directly incorporated into the training objective [32]. Re-translation is explored in conventional ASR systems [33] and simultaneous (speech) translation [34, 35, 36].

3. Incremental Decoding

In offline systems, waiting till the end of input sequences is one of the largest factors contributing to latency. For incremental inference, we divide the input utterance into fixed-size chunks and decode every time a new chunk arrives. To avoid visual distractions from constantly modifying hypotheses, selected chunk-level predictions are *committed to* and no longer modified. Figure 1 illustrates our framework. The decoding of the next chunk is conditioned by the predictions we have *com-*

mitted to. In practice, the decoding of new chunks can either continue from a previously buffered decoder state, or start after forced decoding with the tokens we have committed to. In either approach, the source-target attention can span over all available chunks rather than only the current chunk.

3.1. Partial Hypothesis Selection

Instead of outputting all chunk-level predictions, we are selective with the partial hypotheses for the following reasons. First, the acoustic information towards chunk endpoints tends to be uncertain. Moreover, early chunks often contain very limited context. Therefore, some chunk-level predictions will be unstable, i.e. no longer part of the highest-likelihood hypothesis once new inputs appear. As the decoder is autoregressive, ingesting these predictions can lead to further outputs of low quality. In essence, by the partial hypothesis selection process, we intend to trade some latency for better output quality.

Formally, given chunk-level outputs $W^{(c)}$ from the c -th chunk, we only commit to $\text{PREFIX}(W^{(c)})$ and discard the rest. This can be seen as a form of look-ahead, since the model becomes more conservative with outputting predictions. The next chunk $c + 1$ is conditioned by $\text{PREFIX}(W^{(c)})$ instead of the full $W^{(c)}$. We describe several ways to realize the prefix function. Depending on the behavior of the prefix function, we can buffer decoder states differently to avoid re-computation.

Hold- n One of the most straightforward ways to select partial hypotheses is to *withhold* or *delete* the last n tokens in each chunk. This gives the following prefix function:

$$\text{PREFIX}(W^{(c)}) = W_{0:\max(0, |W^{(c)}| - n)}, \forall c \in \{1, \dots, C\}, \quad (1)$$

where $W^{(c)}$ is the output tokens from the c -th chunk, $|W^{(c)}|$ is the number of tokens in $W^{(c)}$, and n is the number of tokens to be deleted. In case n is greater than the number of output tokens from the current chunk, the prefix will be empty, therefore displaying nothing for this chunk. On the other hand, hold-0 means fully outputting all chunk-level predictions.

Wait- k Inspired by techniques in simultaneous translation [29], we wait for the first k chunks and subsequently outputs at a fixed rate r . The prefix function $\text{PREFIX}(W^{(c)})$ is therefore:

$$\begin{cases} \emptyset, & \text{if } c \leq k \\ W_{0:\min(|W^{(c)}|, r)}, & \text{otherwise,} \end{cases} \quad (2)$$

where $\min(|W^{(c)}|, r)$ ensures that when the output rate r exceeds the number of predicted tokens, all chunk-level outputs are shown. In this case, the effect is identical to hold-0.

Local Agreement If a partial hypothesis is unlikely to change when new inputs arrive, it can be considered promising. This idea is central to the partial track-back procedure in conventional ASR systems [37]. Here, we realize this idea by considering local agreement, and taking the agreeing prefixes of two consecutive chunks as stable hypotheses. For the first chunk, we do not display any output since there is not yet a previous chunk to compare with. From the second chunk onwards, we seek agreement with the predecessor's outputs. We take the longest common prefix of the current chunk's outputs $W^{(c)}$ and the not-yet-displayed outputs of the previous chunk. Let $\text{DISCARD}(\cdot)$ indicate the tokens that are not displayed, i.e. those in $W^{(c)}$ but not in $\text{PREFIX}(W^{(c)})$. The behavior of the local agreement strategy is defined as $\text{PREFIX}(W^{(c)}) =$

$$\begin{cases} \emptyset, & \text{if } c = 1 \\ \text{LCP}(\text{DISCARD}(W^{(c-1)}), W^{(c)}), & \text{otherwise,} \end{cases} \quad (3)$$

where $\text{LCP}(\cdot)$ indicates longest common prefix of two lists.

3.2. Unidirectional Encoder

To facilitate incremental inference, we train a Transformer with unidirectional encoder. By masking future context, we let the encoder only attend to previous and current positions. This is analogous to using unidirectional RNN encoders in recurrent models. Since each position has no future dependency, the existing decoder hidden states also do not need to be re-computed as new input chunks arrive. Note that the unidirectional model is still trained with full utterances. A later adaptation step (Sec. 4) addresses partial inputs under incremental decoding.

3.3. Latency Measurement

In the strictest sense, the latency of transcribing a word is the elapsed time between when it is said and when the corresponding transcription is generated. Therefore, for an output sequence $w_{1, \dots, T}$, the average latency is

$$\frac{1}{T} \sum_{t=1}^T (\text{outputTime}(w_t) - \text{endTime}(w_t)), \quad (4)$$

which can be rewritten to

$$\frac{1}{T} \sum_{t=1}^T \text{outputTime}(w_t) - \frac{1}{T} \sum_{t=1}^T \text{endTime}(w_t). \quad (5)$$

However, the second term in Equation (5) is usually unknown since the transcriptions are not time-aligned. Alternatively, assuming all output words correspond to the ground truth, the second term in Equation (5) remains constant and can be dropped when comparing the latency of two systems, therefore resulting in a measurement only based on the first term. Note that this metric is only valid for latency differences. When considered in isolation, it does not correspond the actual latency experienced by users. In the framework of chunk-based decoding, the output timestamps is easily derived. Since the words are displayed at the boundaries of fixed-size chunks, the output timestamp is the product of the relevant chunk index and the chunk length.

4. Adaptation to Partial Inputs

4.1. Train-Test Condition Mismatch

When training an offline system, the model is rarely incentivized to generate incomplete sentences, since most reference transcriptions are full sentences. Indeed, offline systems trained on full sequences were found to fantasize full outputs even when given partial inputs [34]. More importantly, under chunk-based incremental decoding, the decoder has to operate on partial encoder representations. During training, although the encoder has no future context at each time step, the decoder can still access the full encoded sequences. Therefore, the source-target attention can compensate for the loss of future context by placing heavier weights towards the end of input sequences. In Figure 2, we show an example from our experiments with unidirectional encoder, where one attention head aligns input frames with output tokens, while the other focuses on the tail area of the inputs, thereby compensating for the lack of future context in the encoder representation. This showcases the train-test condition mismatch that needs to be resolved.

4.2. Fine-Tuning with Partial Inputs

To bridge the mismatch between train and test conditions, we adapt the trained models using partial sequences. To avoid losing performance on offline inference after adaptation, we take




Figure 2: Example source-target attention weights at first decoder layer of our unidirectional-encoder model. The vertical axis goes top-down and corresponds to decoding the tokens sequentially. The horizontal axis indicates the input frames.

inspiration from the approach used in [34] and continue training on a 1:1 mix of both full and partial sequences. In absence of time-aligned transcriptions, we make an approximation by taking proportionally many input frames and output tokens. Specifically, given audio frames $U_{1,\dots,I}$, transcription $T_{1,\dots,J}$, and ratio p , we take the first $\lceil I \cdot p \rceil$ frames and $\lceil J \cdot p \rceil$ tokens. We expect the fine-tuning step to primarily adjust source-target attention distribution to work with both full and partial inputs.

5. Experimental Setup

5.1. Data

We use the 300-hour How2 dataset [38], which contains instruction videos from YouTube. For transcription and translation, we use the English audio with English and Portuguese subtitles respectively. We use the pre-compiled 40-dimensional filter bank, and exclude three dimensions of pitch information, the calculation of which involves future frames. To account for online inference, we do not use utterance or speaker-level normalization. An overview of the ASR dataset is shown in Table 1.

Table 1: Overall statistics of the How2 ASR dataset.

Dataset	Len (h:m)	Total utt.	Total words	Avg utt. len
train	298:12	184,949	3,304,534	5.80
dev	3:15	2,022	36,013	5.78
test	3:43	2,305	40,890	5.80

5.2. Hyperparameters

When training the Transformer model¹, we follow the reported hyperparameters from [7], including the optimizer choice, learning rate, warmup steps, dropout rate, label smoothing rate, and embedding dimension. Several hyperparameters differ. The size of the inner feed forward layer is 2048. We use 32 encoder and 12 decoder layers, and byte-pair-encoding [39] of size 10,000. The final model is an ensemble of the last 5 best checkpoints. We use a beam width of 8 when decoding.

5.3. Adaptation Procedure

For each training utterance, we choose a partial transcription from 10% to 40% of the original number of tokens. The ratio is intentionally kept low to create partial inputs lacking future information. Based on the chosen ratio, we proportionately take a subset of the input audio frame as specified in Section 4.2.

¹Implementation and experiment recipes available at: <https://github.com/dannigt/NMTGMinor.lowLatency>

Then we mix the partial sequences and their transcriptions with the original training instances. The learning rate is reduced to a quarter of before. Moreover, we use the original full-sequence dev set to avoid losing performance on offline inference.

6. Experiments and Results

In the following experiments, we use a fixed chunk size of 0.5 second. While we experimented with other chunk size (1 and 2 seconds), the key observations do not differ from those presented here. When decoding for new chunks, the partial hypotheses we have committed to are fed in via forced decoding.

6.1. Bounds for Accuracy and Latency

We first consider the upper and lower bounds for accuracy and latency. Offline transcription is expected to yield the lowest error rate but also the highest latency. On the contrary, when all chunk-level hypotheses are immediately shown (the hold-0 strategy), we expect the lowest latency at the cost of compromising accuracy. Therefore, for incremental decoding, the word error rate is bounded between that of offline and hold-0, whereas latency is bounded between that of hold-0 and offline.

6.2. Accuracy-Latency Trade-Off of Different Strategies

First, we evaluate the hypothesis selection strategies proposed in Section 3.1 in terms of accuracy-latency trade-off. Here we use the full-sequence model with the original bidirectional encoder. The results are summarized in Figure 3 for visual clarity. The horizontal axis indicates the latency difference to the hold-0 strategy, the lower bound for latency. The vertical axis indicates the WER difference to the offline system, the lower bound for accuracy. The data points on the interpolated lines correspond to different hyperparameters under the same strategy. Falling closer to the origin of the graph indicates higher efficiency in accuracy-latency trade-off.

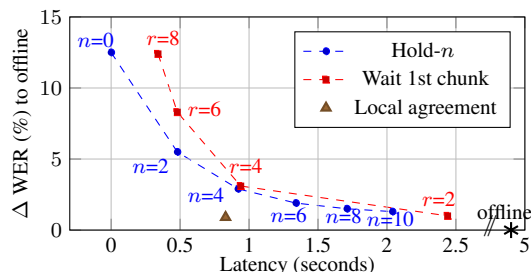


Figure 3: Accuracy-latency trade-off of different partial hypothesis selection strategies. n indicates the number of removed last tokens. r indicates the fixed output rate (tokens/second).

The first observation from Figure 3 is that the local agreement strategy outperforms other strategies by a large margin. Reducing latency by 3.8 seconds at the cost of 0.9% WER compared to offline system, it achieves the most efficient accuracy-latency trade-off. It shows that dynamically selecting partial hypotheses is preferable to hard thresholds set by hyperparameters. Comparing the other two strategies, initially surprising to us was that hold- n consistently outperforms wait- k , a proven method in simultaneous translation. A further analysis shows that this method suffers from varying utterance speed. In text translation, the input-output rate is roughly 1:1. In ASR, however, a fixed speed for all chunks is a too strong assumption.

Table 2: Incremental ASR performance on How2 300h English before and after adaptation. Latency is reported as difference to the unadapted model under the hold-0 strategy. In parentheses are changes in WER after adaptation.

Strategies	Unadapted				Adapted to Partial Inputs			
	WER (%) ↓		Δ latency (sec.) ↓		WER (%) ↓		Δ latency (sec.) ↓	
	Bidir.	Unidir.	Bidir.	Unidir.	Bidir.	Unidir.	Bidir.	Unidir.
Hold-0	37.4	40.7	0	0	28.3 (−9.1)	30.3 (−10.4)	+0.32	+0.30
Hold-2	20.4	22.0	+0.48	+0.42	18.7 (−1.7)	19.3 (−2.7)	+0.77	+0.73
Hold-4	17.8	18.8	+0.92	+0.83	17.1 (−0.7)	17.3 (−1.5)	+1.22	+1.16
Hold-6	16.8	17.6	+1.34	+1.22	16.5 (−0.3)	16.5 (−1.1)	+1.60	+1.52
Local agreement	15.8	16.8	+0.81	+0.54	15.8 (−0.0)	15.5 (−1.0)	+0.83	+0.65
Offline	14.9	14.4	+4.55	+4.47	14.4 (−0.5)	14.7 (+0.3)	+4.59	+4.45

Indeed, at a fixed output rate of $r = 2$, we get close to offline WER (15.9% vs 14.9%). However, this comes with large latency of nearly 2.5 seconds. As 2 tokens per second corresponds to a very low utterance speed, this output rate leads to most of the tokens being displayed after the utterance ends, in effect converging towards an offline system. On the other hand, increasing the output rate quickly results in more errors. Given a high output rate, the allowed number of output tokens is likely to exceed the actual number of tokens said in the chunk. In this case, we approach the other extreme, the hold-0 strategy that outputs all chunk-level tokens. Given this finding, we eliminate this option and only pursue the hold- n and the local agreement strategies in the upcoming experiments.

6.3. Low-Latency ASR

After identifying promising hypothesis selection strategies, we proceed to incremental inference. We first focus on the ASR performance of the unidirectional model. Then we study the effect of adaptation to partial sequences.

As shown in the bottom left section of Table 2, the unidirectional full-sequence model achieves similar accuracy to the bidirectional one (WER 14.4% vs 14.9%). However, the incremental inference WER with the unidirectional model is constantly around 1% higher than the bidirectional model. This suggests that the unidirectional model, when directly performing online inference, is more susceptible to the loss of future context. However, the gap closes after adaptation on partial sequences. The unidirectional model becomes on-par with its bidirectional counterpart under the more promising strategies, e.g. hold-4 and local agreement. Furthermore noteworthy is that after adaptation it still preserves the performance under offline inference, as evidenced by the WER of 14.4% and 14.7% before and after adaptation.

Contrasting the right-hand-side section of Table 2 with the left, we see the effect of adaptation. The error reduction after adaptation comes with a slight increase in latency up to 0.3 second. It is a sign that the model can better control its chunk-level outputs. Moreover, the impact is stronger with the unidirectional model. A hypothesized reason is that its encoder representation already has no dependency on future inputs, adaptation to partial sequences is therefore an easier task.

6.4. Low-Latency Speech Translation

Having seen the effectiveness of the unidirectional model in low-latency ASR, we validate the findings on speech translation, a different sequence-to-sequence task. The model here is before partial sequence adaptation.

Table 3 outlines the performance and latency of the uni-

directional model on the How2 English-Portuguese translation task. We use BLEU [40, 41] and METEOR [42] as quality metrics. The results here agree with the previous findings on ASR. In general, by selectively taking prefixes of chunk-level hypotheses, we can largely reduce latency by sacrificing some output quality. More importantly, the local agreement selection strategy remains to achieve the most efficient quality-latency trade-off. It scores similarly to the hold-4 strategy in term of the quality metrics but has 0.2 second less latency. Compared to the offline system, we see an 84% relative reduction in latency with a 5% relative loss in BLEU and METEOR.

Table 3: Incremental speech translation performance on How2 300h English-Portuguese.

Strategies	BLEU ² ↑	METEOR ↑	Δ latency (sec.) ↓
Hold-0	24.9	25.5	0
Hold-2	37.3	31.4	+0.48
Hold-4	42.2	33.6	+0.95
Hold-6	43.6	34.2	+1.38
Local agreement	42.1	33.5	+0.71
Offline	44.5	34.5	+4.36

7. Conclusion

In this paper, we explored approaches for latency reduction in sequence-to-sequence speech recognition and translation. First, we studied the accuracy-latency trade-off under chunk-based incremental inference. By selecting a subset of chunk-level outputs, we achieved large error reduction with minimal delay. Among the three partial hypothesis selection strategies, the most efficient trade-off was achieved by examining the local agreement of hypotheses. To facilitate incremental inference, we trained a unidirectional model where the encoder accesses no future context. With a simple yet effective adaptation procedure on partial inputs, the unidirectional model performed on-par with its bidirectional counterpart on ASR. On the 300-hour How2 dataset, we were able to reduce the gap to offline systems to less than 1% absolute while incurring 0.7 seconds of latency. Besides ASR, our approach was also effective on speech translation. A future direction is to incorporate adaptive chunk sizes.

8. Acknowledgment

The experiments were supported by the Google Cloud Platform research credit program.

²SacreBLEU: case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.4.3

9. References

- [1] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *Proc. ICLR*, 2015.
- [2] M. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” in *Proc. EMNLP*, 2015.
- [3] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. ICASSP*, 2016, pp. 4960–4964.
- [4] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, “End-to-end attention-based large vocabulary speech recognition,” in *Proc. ICASSP*, 2016, pp. 4945–4949.
- [5] Y. Liu, H. Xiong, J. Zhang, Z. He, H. Wu, H. Wang, and C. Zong, “End-to-End Speech Translation with Knowledge Distillation,” in *Proc. Interspeech*, 2019, pp. 1128–1132.
- [6] M. Sperber, G. Neubig, J. Niehues, and A. Waibel, “Attention-passing models for robust and data-efficient end-to-end speech translation,” *TACL*, vol. 7, pp. 313–325, 2019.
- [7] N. Pham, T. Nguyen, J. Niehues, M. Müller, and A. Waibel, “Very deep self-attention networks for end-to-end speech recognition,” in *Proc. Interspeech*, 2019, pp. 66–70.
- [8] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang, S. Watanabe, T. Yoshimura, and W. Zhang, “A comparative study on Transformer vs RNN in speech applications,” in *Proc. ASRU*, 2019.
- [9] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, “A comparison of Transformer and LSTM encoder decoder models for ASR,” in *Proc. ASRU*, 2019.
- [10] T. Nguyen, S. Stueker, J. Niehues, and A. Waibel, “Improving sequence-to-sequence speech recognition training with on-the-fly data augmentation,” in *Proc. ICASSP*, 2020, pp. 7689–7693.
- [11] C. Raffel, M. Luong, P. J. Liu, R. J. Weiss, and D. Eck, “Online and linear-time attention by enforcing monotonic alignments,” in *Proc. ICML*, 2017, pp. 2837–2846.
- [12] C. Chiu and C. Raffel, “Monotonic chunkwise attention,” in *Proc. ICLR*, 2018.
- [13] X. Ma, J. M. Pino, J. Cross, L. Puzon, and J. Gu, “Monotonic multihead attention,” in *ICLR*, 2020.
- [14] N. Jaitly, Q. V. Le, O. Vinyals, I. Sutskever, D. Sussillo, and S. Bengio, “An online sequence-to-sequence model using partial conditioning,” in *Proc. NIPS*, 2016, pp. 5067–5075.
- [15] T. N. Sainath, C.-C. Chiu, R. Prabhavalkar, A. Kannan, Y. Wu, P. Nguyen, and Z. Chen, “Improving the performance of online neural transducer models,” in *Proc. ICASSP*, 2018, pp. 5864–5868.
- [16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, u. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. NIPS*, 2017, pp. 6000–6010.
- [17] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv:1211.3711*, 2012.
- [18] H. Sak, M. Shannon, K. Rao, and F. Beaufays, “Recurrent neural aligner: An encoder-decoder neural network model for sequence to sequence mapping,” in *Proc. Interspeech*, 2017, pp. 1298–1302.
- [19] Q. Zhang, H. Lu, H. Sak, A. Tripathi, E. McDermott, S. Koo, and S. Kumar, “Transformer transducer: A streamable speech recognition model with Transformer encoders and RNN-T loss,” in *Proc. ICASSP*, 2020, pp. 7829–7833.
- [20] L. Dong, F. Wang, and B. Xu, “Self-attention aligner: A latency-control end-to-end model for ASR using self-attention network and chunk-hopping,” in *Proc. ICASSP*, 2019, pp. 5656–5660.
- [21] S. Novitasari, A. Tjandra, S. Sakti, and S. Nakamura, “Sequence-to-sequence learning via attention transfer for incremental speech recognition,” in *Proc. Interspeech*, 2019, pp. 3835–3839.
- [22] H. Inaguma, Y. Gaur, L. Lu, J. Li, and Y. Gong, “Minimum latency training strategies for streaming sequence-to-sequence ASR,” in *Proc. ICASSP*, 2020, pp. 6064–6068.
- [23] Z. Tian, J. Yi, Y. Bai, J. Tao, S. Zhang, and Z. Wen, “Synchronous transformers for end-to-end speech recognition,” in *Proc. ICASSP*, 2020, pp. 7884–7888.
- [24] H. Miao, G. Cheng, C. Gao, P. Zhang, and Y. Yan, “Transformer-based online CTC/attention end-to-end speech recognition architecture,” in *Proc. ICASSP*, 2020, pp. 6084–6088.
- [25] N. Moritz, T. Hori, and J. L. Roux, “Triggered attention for end-to-end speech recognition,” in *Proc. ICASSP*, 2019, pp. 5666–5670.
- [26] N. Moritz, T. Hori, and J. L. Roux, “Streaming automatic speech recognition with the transformer model,” in *Proc. ICASSP*, 2020, pp. 6074–6078.
- [27] M. Li, M. Liu, and H. Masanori, “End-to-end speech recognition with adaptive computation steps,” in *Proc. ICASSP*, 2019, pp. 6246–6250.
- [28] K. Cho and M. Esipova, “Can neural machine translation do simultaneous translation?” *arXiv:1606.02012*, 2016.
- [29] M. Ma, L. Huang, H. Xiong, R. Zheng, K. Liu, B. Zheng, C. Zhang, Z. He, H. Liu, X. Li, H. Wu, and H. Wang, “STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework,” in *Proc. ACL*, 2019, pp. 3025–3036.
- [30] J. Gu, G. Neubig, K. Cho, and V. O. Li, “Learning to translate in real-time with neural machine translation,” in *Proc. EACL*, Valencia, Spain, 2017, pp. 1053–1062.
- [31] F. Dalvi, N. Durrani, H. Sajjad, and S. Vogel, “Incremental decoding and training methods for simultaneous translation in neural machine translation,” in *Proc. NAACL-HLT*, Jun. 2018, pp. 493–499.
- [32] N. Arivazhagan, C. Cherry, W. Macherey, C. Chiu, S. Yavuz, R. Pang, W. Li, and C. Raffel, “Monotonic infinite lookback attention for simultaneous machine translation,” in *Proc. ACL*, Florence, Italy, 2019, pp. 1313–1323.
- [33] T. Nguyen, J. Niehues, E. Cho, T. Ha, K. Kilgour, M. Muller, M. Sperber, S. Stueker, and A. Waibel, “Low latency ASR for simultaneous speech translation,” *arXiv:2003.09891*, 2020.
- [34] J. Niehues, N. Pham, T. Ha, M. Sperber, and A. Waibel, “Low-latency neural speech translation,” in *Proc. Interspeech*, 2018, pp. 1293–1297.
- [35] N. Arivazhagan, C. Cherry, W. Macherey, and G. Foster, “Re-translation versus streaming for simultaneous translation,” *arXiv:2004.03643*, 2020.
- [36] N. Arivazhagan, C. Cherry, I. Te, W. Macherey, P. Baljekar, and G. Foster, “Re-translation strategies for long form, simultaneous, spoken language translation,” in *Proc. ICASSP*, 2020, pp. 7919–7923.
- [37] P. Brown, J. Spohrer, P. Hochschild, and J. Baker, “Partial traceback and dynamic programming,” in *Proc. ICASSP*, vol. 7, 1982, pp. 1629–1632.
- [38] R. Sanabria, O. Caglayan, S. Palaskar, D. Elliott, L. Barrault, L. Specia, and F. Metze, “How2: A large-scale dataset for multimodal language understanding,” in *Proc. ViGIL, NeurIPS*, 2018.
- [39] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proc. ACL*, 2016, pp. 1715–1725.
- [40] K. Papineni, S. Roukos, T. Ward, and W. Zhu, “BLEU: a method for automatic evaluation of machine translation,” in *Proc. ACL*, 2002.
- [41] M. Post, “A call for clarity in reporting BLEU scores,” in *Proc. WMT*, 2018, pp. 186–191.
- [42] S. Banerjee and A. Lavie, “METEOR: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proc. ACL*, Jun. 2005, pp. 65–72.