

# Streaming Models for Joint Speech Recognition and Translation

Orion Weller<sup>1\*</sup>, Matthias Sperber<sup>2</sup>, Christian Gollan<sup>2</sup>, Joris Kluivers<sup>2</sup>

<sup>1</sup>Brigham Young University

<sup>2</sup>Apple

oweller@byu.edu, {m\_sperber, cgollan, jkluivers}@apple.com

## Abstract

Using end-to-end models for speech translation (ST) has increasingly been the focus of the ST community. These models condense the previously cascaded systems by directly converting sound waves into translated text. However, cascaded models have the advantage of including automatic speech recognition output, useful for a variety of practical ST systems that often display transcripts to the user alongside the translations. To bridge this gap, recent work has shown initial progress into the feasibility for end-to-end models to produce both of these outputs. However, all previous work has only looked at this problem from the consecutive perspective, leaving uncertainty on whether these approaches are effective in the more challenging streaming setting. We develop an end-to-end streaming ST model based on a re-translation approach and compare against standard cascading approaches. We also introduce a novel inference method for the joint case, interleaving both transcript and translation in generation and removing the need to use separate decoders. Our evaluation across a range of metrics capturing accuracy, latency, and consistency shows that our end-to-end models are statistically similar to cascading models, while having half the number of parameters. We also find that both systems provide strong translation quality at low latency, keeping 99% of consecutive quality at a lag of just under a second.

## 1 Introduction

Speech translation (ST) is the process of translating acoustic sound waves into text in a different language than was originally spoken in.

This paper focuses on ST in a particular setting, as described by two characteristics: (1) We desire models that translate in a streaming fashion, where

users desire the translation before the speaker has finished. This setting poses additional difficulties compared to consecutive translation, forcing systems to translate without knowing what the speaker will say in the future. (2) Furthermore, the speaker may want to verify that their speech is being processed correctly, intuitively seeing a streaming transcript while they speak (Fügen, 2008; Hsiao et al., 2006). For this reason, we consider models that produce both transcripts and translation jointly.<sup>1</sup>

Previous approaches to streaming ST have typically utilized a cascaded system that pipelines the output of an automatic speech recognition (ASR) system through a machine translation (MT) model for the final result. These systems have been the preeminent strategy, taking the top place in recent streaming ST competitions (Pham et al., 2019; Jan et al., 2019; Elbayad et al., 2020; Ansari et al., 2020). Despite the strong performance of these cascaded systems, there are also some problems: error propagation from ASR output to MT input (Ruiz and Federico, 2014); ASR/MT training data mismatch and loss of access to prosodic/paralinguistic speech information at the translation stage (Sperber and Paulik, 2020); and potentially sub-optimal latencies in the streaming context. End-to-end (E2E) models for ST have been proposed to remedy these problems, leveraging the simplicity of a single model to sidestep these issues. E2E models are also appealing from computational and engineering standpoints, reducing model complexity and decreasing parameter count.

Although initial research has explored E2E models for joint speech recognition and translation, no previous works have examined them in the streaming case, a crucial step in using them for many real-world applications. To understand this area more fully, we develop an E2E model to compare

<sup>1</sup>This corresponds to the *mandatory transcript* case in the proposed categorization by Sperber and Paulik (2020).

\*Work done during an internship with Apple

with its cascading counterpart in this simultaneous joint task. We build off the models proposed by Sperber et al. (2020) in the consecutive case, extending them for use in the streaming setting. We also use the re-translation technique introduced by Niehues et al. (2018) to maintain simplicity while streaming. To reduce model size, we introduce a new method for E2E inference, producing both transcript and translation in an interleaved fashion with one decoder.

As this task requires a multi-faceted evaluation along several axes, we provide a suite of evaluations to highlight the differences of these major design decisions. This suite includes assessing translation quality, transcription quality, lag of the streaming process, output flicker, and consistency between the transcription and translation. We find that our E2E model performs similarly to the cascaded model, indicating that E2E networks are a feasible and promising direction for streaming ST.

## 2 Proposed Method

**Network Architecture** In the ST survey provided by Sperber et al. (2020), they introduce several E2E models that could be used for the joint setting. As our work focuses on providing a simple but effective approach to streaming ST, we focus on the *CONCAT* model, which generates both the transcript and translation in a concatenated fashion. We compare this E2E model against the standard cascading approach, following the architecture and hyperparameter choices used in Sperber et al. (2020). All audio input models use the same multi-layer bidirectional LSTM architecture, stacking and downsampling the audio by a factor of three before processing. We note that although bidirectional encoders are unusual with standard ASR architectures, re-translation makes them possible. The cascaded model’s textual encoder follows the architecture described in Vaswani et al. (2017) but replaces self-attention blocks with LSTMs. Decoder networks are similar, but use unidirectional LSTMs. More implementation details can be found in Appendix A.

In order to reduce model size and inference time for E2E networks, we introduce a novel method for interleaving both transcript and translation in generation, removing the need to use separate decoders. This method extends the *CONCAT* model proposed by Sperber et al. (2020) to jointly decode according to the ratio given by the parameter  $\gamma$  (Fig-

Interleaving at $\gamma = 0.0$	
Outputs	Do you want to go EOS1 Wollen Sie gehen EOS2
Language tokens	src src src src src src trg trg trg trg
Target Inputs	<s> Do you want to go <s> Wollen Sie gehen
Interleaving at $\gamma = 0.5$	
Outputs	Do Wollen you Sie want gehen to EOS2 go EOS1
Language tokens	src trg src trg src trg src trg src src
Target Inputs	<s> <s> Do Wollen you Sie want gehen to go
Interleaving at $\gamma = 1.0$	
Outputs	Wollen Sie gehen EOS2 Do you want to go EOS1
Language tokens	trg trg trg trg src src src src src
Target Inputs	<s> Wollen Sie gehen <s> Do you want to go

Figure 1: Example token representations (En→De) for three different interleaving parameters (Section 2). Language tokens indicate whether the data corresponds to the source transcript or the target translation and are used with a learned embedding that is summed with the word embeddings, as described in Sperber et al. (2020).

ure 1). When  $\gamma = 0.0$ , we generate the transcript tokens until completion, followed by the translation tokens (vice versa for  $\gamma = 1.0$ ). At  $\gamma = 0.0$ , our model is equivalent to the previously proposed model. Defining  $\text{count}_i$  as the count of  $i$  tokens previously generated, transcription tokens as  $st$  and translation tokens as  $tt$ , we generate the next token as a transcription token if:

$$(1.0 - \gamma) * (1 + \text{count}_{tt}) > \gamma * (1 + \text{count}_{st})$$

This approach enables us to produce tokens in an interleaving fashion, given the hyperparameter  $\gamma$ .

**Re-translation** We use the re-translation method (Niehues et al., 2018; Arivazhagan et al., 2020a,b) as it provides a simple way to handle the streaming case. This method works by simply re-translating the utterance as new data arrives, updating its former prediction. As we are generating both transcript and translation, this avoids the challenging issue of combining the requirements for both components: streaming speech models need to manage the audio signal variability across time while streaming translation models need to overcome issues with reordering and lack of future context.

Alternative strategies to the re-translation approach include the chunk-based strategy explored by Liu et al. (2020), which commits to all previous output chunks and Ren et al. (2020) who utilize an additional segmenter model trained via CTC




Figure 2: Left: average lag in seconds vs BLEU score. Right: average lag in seconds vs WER score. All points are the mean of each configuration’s score across the eight target languages. Configurations are the cross product of the values for  $K$  and  $F$ , see Section 2: Inference. Note that points near 1.0 AL have appx. 99% of the unconstrained BLEU score. Results for the E2E model use  $\gamma = 0.5$ .

Metric	Params	Model	De	Es	Fr	It	Nl	Pt	Ro	Ru	Mean
BLEU $\uparrow$	217M	Cascade	<b>18.8</b>	<b>22.7</b>	<b>27.0</b>	<b>18.9</b>	<b>22.5</b>	<b>21.9</b>	<b>17.9</b>	<b>13.0</b>	<b>20.3</b>
	107M	E2E $\gamma=0.0$	18.1	<b>23.1</b>	<b>27.0</b>	<b>18.7</b>	<b>22.3</b>	<b>22.2</b>	<b>17.6</b>	12.2	20.2
	107M	E2E $\gamma=0.3$	17.7	22.6	26.3	18.0	21.5	21.5	17.0	12.1	19.6
	107M	E2E $\gamma=0.5$	18.2	<b>22.8</b>	<b>27.0</b>	<b>18.6</b>	21.9	21.9	17.1	12.0	19.9
	107M	E2E $\gamma=1.0$	18.2	<b>22.8</b>	<b>27.1</b>	<b>18.9</b>	<b>22.2</b>	<b>22.3</b>	<b>17.6</b>	<b>12.7</b>	<b>20.2</b>
WER $\downarrow$	217M	Cascade	25.9	24.0	23.1	25.6	28.5	26.4	24.4	<b>23.1</b>	25.1
	107M	E2E $\gamma=0.0$	24.2	23.5	23.3	<b>23.0</b>	23.4	25.3	24.1	23.6	23.8
	107M	E2E $\gamma=0.3$	24.1	23.6	22.9	23.8	23.4	25.7	24.1	24.1	24.0
	107M	E2E $\gamma=0.5$	24.5	23.9	22.9	23.8	23.4	25.7	24.3	23.6	24.0
	107M	E2E $\gamma=1.0$	<b>23.6</b>	<b>22.9</b>	<b>22.3</b>	<b>23.0</b>	<b>22.4</b>	<b>24.7</b>	<b>23.4</b>	<b>22.7</b>	<b>23.1</b>

Table 1: BLEU and WER scores for models trained on different target languages. Bold scores indicate results that are statistically similar to the best score using a bootstrap permutation test with  $\alpha = 0.05$ .

(Graves et al., 2006) to create segments that are translated via wait- $k$  (Ma et al., 2019). Although these approaches show effective results, they add additional complexity without addressing issues particular to streaming transcription.

**Inference** In order to generate quality-latency curves, we use several techniques to reduce latency and flicker at the cost of quality. The first is the mask- $k$  method proposed by Arivazhagan et al. (2020b), masking the last  $K$  output tokens. The second method is a form of constrained decoding: we define a hyperparameter  $F$  that sets the number of free tokens allowed to change in the next re-translation. Thus, we constrain future output to match the first  $len(tokens) - F$  tokens of the current output. All models use values  $\{0, 1, 2, 3, 4, 5, 7, 10, 100\}$  for  $K$  and

$\{0, 1, 2, 3, 4, 5, 7, 10, 15, 20, 25, 100\}$  for  $F$ . For interleaving models, we set  $K$  and  $F$  on both transcript and translation tokens.

### 3 Experimental Settings

**Data** We use the MuST-C corpus (di Gangi et al., 2019) since it is the largest publicly available ST corpus, consisting of TED talks with their English transcripts and translations into eight other language pairs. The dataset consists of at least 385 hours of audio for each target language.

We utilize the log Mel filterbank speech features provided with the corpus as input for the ASR and E2E models. To prepare the textual data, we remove non-speech artifacts (e.g. “(laughter)” and speaker identification) and perform subword tokenization using SentencePiece (Kudo and Richard-

Model	En→De Incr.	En→De Full	En→Es Incr.	En→Es Full	Mean Incr.	Mean Full
Cascade	<b>13.8</b>	<b>13.2</b>	<b>12.2</b>	<b>11.6</b>	<b>14.1</b>	<b>13.4</b>
E2E $\gamma=0.0$	17.6	16.7	14.9	13.8	17.0	16.0
E2E $\gamma=0.3$	17.2	16.6	14.3	13.7	16.6	15.8
E2E $\gamma=0.5$	17.8	16.5	14.8	13.3	17.3	15.7
E2E $\gamma=1.0$	17.3	16.8	14.9	13.7	16.9	15.8

Table 2: Consistency scores for En→De, En→Es, and average results over all languages; lower is better (see Sperber et al. (2020)). *Incr.* stands for the incremental consistency score, or the average consistency throughout re-translation. Bold scores indicate results that are statistically similar to the best score using a bootstrap permutation test with  $\alpha = 0.05$ .

son, 2018) on the unigram setting. Following previous work for E2E ST models, we use a relatively small vocabulary and share transcription and translation vocabularies. We use MuST-C dev for validation and report results on tst-COMMON, utilizing the segments provided (Appendix D).

**Prefix Sampling** We implement techniques developed by Niehues et al. (2018); Arivazhagan et al. (2020b) for improving streaming ST, sampling a random proportion of each training instance as additional data to teach our models to work with partial input. See Appendix C for implementation details.

**Metrics** We evaluate these models on a comprehensive suite of metrics: sacrebleu (*BLEU*, Post (2019)) for translation quality, word error rate (*WER*, (Fiscus, 1997)) for transcription quality, average lag (*AL*, Ma et al. (2019)) for the lag between model input and output, and normalized erasure (*NE*, Arivazhagan et al. (2020a)) for output flicker. Measuring consistency is a nascent area of research; we use the robust and simple lexical consistency metric defined by Sperber et al. (2020), which uses word-level translation probabilities. To show how consistent these results are while streaming, we compute an incremental consistency score, averaging the consistency of each re-translation.

## 4 Results

Results for the quality-latency curves created by the use of constrained decoding and mask- $k$  (Section 3) are shown in Figure 2. Unconstrained settings are used for all results in table form. For convenience, bold scores indicate the highest performing models in each metric according to a bootstrap permutation test.

**Translation Quality** We see in Table 1 that the cascaded model slightly outperforms some E2E

models, while achieving statistically similar performance to the  $\gamma = 1.0$  model. We note however, that the cascaded model has nearly twice as many parameters as the E2E models (217M vs 107M). When we examine these models under a variety of different inference conditions (using constrained decoding and mask- $k$  as in Arivazhagan et al. (2020a)), we further see this trend illustrated through the quality vs latency trade-off (left of Figure 2), with both models retaining 99% of their BLEU at less than 1.0 AL.

**Transcription Quality** Conversely, Table 1 and the right of Figure 2 show that the  $\gamma = 1.0$  E2E model performs similarly or slightly better than the cascaded model across all inference parameters and all target languages. With an AL of 1.5, the E2E model loses only 3% of its performance.

**Consistency** The E2E models perform worse than the cascaded on consistency, with the best models being approximately 18% less consistent (Table 2). The cascaded model also maintains better scores through each re-translation (*Incr.*)<sup>2</sup>

**Flicker** We note that the flicker scores for cascade and E2E models are similar, with both having normalized erasure scores of less than 1 and the majority of inference settings having less than the “few-revision” threshold of 0.2 (proposed by Arivazhagan et al. (2020a)). More NE details are found in Appendix B.

**Interleaving Rate** Table 1 also shows us the overall results for different interleaving rates. We see that interleaving at a rate of 1.0 has the best

<sup>2</sup>Initial experiments indicate that the triangle E2E architecture (Sperber et al., 2020) model may perform better on consistency in our streaming setting, but due to time constraints we were not able to explore this further. Future work exploring alternative architectures or decoding techniques (Le et al., 2020) may provide fruitful avenues of research.

quality scores (0.7 less WER than the next best rate, the base  $\gamma = 0.0$  model) but the worst consistency (Table 2). Conversely,  $\gamma = 0.3$  has the worst quality scores but the best consistency.

## 5 Conclusion

We focus on the task of streaming speech translation, producing both a target translation and a source transcript from an audio source. We develop an end-to-end model to avoid problems that arise from the use of cascaded models for streaming ST. We further introduce a new method for joint inference for end-to-end models, generating both translation and transcription tokens concurrently. We show that our novel end-to-end model, with only half the number of parameters, is comparable to standard cascaded models across a variety of evaluation categories: transcript and translation quality, lag of streaming, consistency between transcript and translation, and re-translation flicker. We hope that this will spur increased interest in using end-to-end models for practical applications of streaming speech translation.

## References

- Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, Alexander Waibel, and Changhan Wang. 2020. **FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN**. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34, Online. Association for Computational Linguistics.
- Naveen Arivazhagan, Colin Cherry, Wolfgang Macherey, and George Foster. 2020a. Re-translation versus streaming for simultaneous translation. *arXiv preprint arXiv:2004.03643*.
- Naveen Arivazhagan, Colin Cherry, Isabelle Te, Wolfgang Macherey, Pallavi Baljekar, and George Foster. 2020b. Re-translation strategies for long form, simultaneous, spoken language translation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7919–7923. IEEE.
- Maha Elbayad, Ha Nguyen, Fethi Bougares, Natalia Tomashenko, Antoine Caubrière, Benjamin Lecouteux, Yannick Estève, and Laurent Besacier. 2020. On-trac consortium for end-to-end and simultaneous speech translation challenge tasks at iwslt 2020. *arXiv preprint arXiv:2005.11861*.
- Jonathan G Fiscus. 1997. A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (rover). In *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, pages 347–354. IEEE.
- Christian Fügen. 2008. *A System for Simultaneous Translation of Lectures and Speeches*. Ph.D. thesis, University of Karlsruhe.
- Yarin Gal and Zoubin Ghahramani. 2016. **Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning**. In *International Conference on Machine Learning (ICML)*.
- Antonino Mattia di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. **MuST-C : a Multilingual Speech Translation Corpus**. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, Minneapolis, USA.
- Xavier Glorot and Yoshua Bengio. 2010. **Understanding the difficulty of training deep feedforward neural networks**. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, Sardinia, Italy.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.
- Roger Hsiao, Ashish Venugopal, Thilo Köhler, Ting Zhang, Paisarn Charoenpornasawat, Andreas Zollmann, Stephan Vogel, Alan W. Black, Tanja Schultz, and Alex Waibel. 2006. **Optimizing components for handheld two-way speech translation for an English-Iraqi Arabic system**. In *Annual Conference of the International Speech Communication Association (InterSpeech)*, pages 765–768, Pittsburgh, USA.
- Niehues Jan, Roldano Cattoni, Stuker Sebastian, Matteo Negri, Marco Turchi, Salesky Elizabeth, Sanabria Ramon, Barrault Loic, Specia Lucia, and Marcello Federico. 2019. The iwslt 2019 evaluation campaign. In *16th International Workshop on Spoken Language Translation 2019*.
- Diederik P. Kingma and Jimmy L. Ba. 2014. **Adam: A Method for Stochastic Optimization**. In *International Conference on Learning Representations (ICLR)*, Banff, Canada.
- Taku Kudo and John Richardson. 2018. **SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. Empirical Methods in Natural Language Processing (EMNLP)**, pages 66–71.
- Hang Le, Juan Pino, Changhan Wang, Jiatao Gu, Didier Schwab, and Laurent Besacier. 2020. **Dual-decoder transformer for joint automatic speech recognition and multilingual speech translation**. In

- Proceedings of the 28th International Conference on Computational Linguistics*, pages 3520–3533, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Danni Liu, Gerasimos Spanakis, and Jan Niehues. 2020. Low-latency sequence-to-sequence speech recognition and translation by partial hypothesis selection. *arXiv preprint arXiv:2005.11185*.
- Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, et al. 2019. Stacl: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3036.
- Graham Neubig, Matthias Sperber, Xinyi Wang, Matthieu Felix, Austin Matthews, Sarguna Padmanabhan, Ye Qi, Devendra Singh Sachan, Philip Arthur, Pierre Godard, John Hewitt, Rachid Riad, and Liming Wang. 2018. **XNMT: The eXtensible Neural Machine Translation Toolkit**. In *Conference of the Association for Machine Translation in the Americas (AMTA) Open Source Software Showcase*, Boston, USA.
- Jan Niehues, Ngoc-Quan Pham, Thanh-Le Ha, Matthias Sperber, and Alex Waibel. 2018. Low-latency neural speech translation. *arXiv preprint arXiv:1808.00491*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. **PyTorch: An Imperative Style, High-Performance Deep Learning Library**. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada.
- Ngoc-Quan Pham, Thai-Son Nguyen, Thanh-Le Ha, Juan Hussain, Felix Schneider, Jan Niehues, Sebastian Stüker, and Alexander Waibel. 2019. The iwslt 2019 kit speech translation system. In *Proceedings of the 16th International Workshop on Spoken Language Translation*.
- Matt Post. 2019. **A Call for Clarity in Reporting BLEU Scores**. In *Conference on Machine Translation (WMT)*, pages 186–191, Brussels, Belgium.
- Yi Ren, Jinglin Liu, Xu Tan, Chen Zhang, QIN Tao, Zhou Zhao, and Tie-Yan Liu. 2020. Simulspeech: End-to-end simultaneous speech to text translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3787–3796.
- Nicholas Ruiz and Marcello Federico. 2014. Assessing the impact of speech recognition errors on machine translation quality. In *11th Conference of the Association for Machine Translation in the Americas (AMTA)*, Vancouver, BC, Canada.
- Matthias Sperber and Matthias Paulik. 2020. Speech Translation and the End-to-End Promise: Taking Stock of Where We Are. In *Association for Computational Linguistic (ACL)*, Seattle, USA.
- Matthias Sperber, Hendra Setiawan, Christian Gollan, Udhaykumar Nallasamy, and Matthias Paulik. 2020. **Consistent Transcription and Translation of Speech**. *Transactions of the Association for Computational Linguistics (TACL)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. **Attention Is All You Need**. In *Neural Information Processing Systems Conference (NIPS)*, pages 5998–6008, Long Beach, USA.

## A Model Details

In this section we will describe implementation details of the model architectures (shown in Figure 3) and training processes.

**Model Architectures** Unless otherwise noted, the same hyperparameters are used for all models. Weights for the speech encoder are initialized based on a pre-trained attentional ASR task that is identical to the ASR part of the direct multitask model. Other weights are initialized according to [Glorot and Bengio \(2010\)](#). The speech encoder is a 5-layer bidirectional LSTM with 700 dimensions per direction. Attentional decoders consist of 2 Transformer blocks ([Vaswani et al., 2017](#)) but use 1024-dimensional unidirectional LSTMs instead of self-attention, except for the *CONCAT* model, which uses 3 layers.

For the cascade’s MT model, encoder/decoder both contain 6 layers with 1024-dimensional LSTMs. Subword embeddings are size 1024. We regularize using LSTM dropout with  $p = 0.3$ , decoder input word-type dropout ([Gal and Ghahramani, 2016](#)), and attention dropout, both  $p = 0.1$ . We apply label smoothing with strength  $\epsilon = 0.1$ .

**Training** We optimize using Adam ([Kingma and Ba, 2014](#)) with  $\alpha = 0.0005$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ , 4000 warm-up steps, and learning rate decay by using the inverse square root of the iteration of each instance. We set the batch size dynamically based on the sentence length, such that the average batch size is 128 utterances. The training is stopped when the validation score has not improved over 10 epochs, where the validation score is corpus-level translation BLEU score (for the E2E and MT

models) and corpus-level WER for the cascade’s ASR model.

For decoding and generating n-best lists, we use beam size 5 and polynomial length normalization with exponent 1.5. Our implementation is based on PyTorch (Paszke et al., 2019) and XNMT (Neubig et al., 2018), and all models are trained in single-GPU environments, employing Tesla V100 GPUs with 32 GB memory. Most E2E and ASR models converged after approximately 30 epochs or 5 days of training. MT models converged after approximately 50 epochs or 2 days of training.

## B Normalized Erasure (Output Flicker)

We see similar curves for both the cascaded model and the E2E model when comparing normalized erasure in Figure 4. We see that most settings have an NE score of less than 0.2, while virtually all settings are less than 1. We note that a proportion of 0.2 for NE means that, on average, 1/5 of the tokens change once before they settle to their final state.

## C Prefix Training

We used prefix training to increase stability and reduce flickering in the streaming setting. We conducted this by utilizing each training instance twice in each epoch: one as normal and the other with only the prefix. The length of the prefixes were randomly sampled from  $[0, 1]$ . We found that this additional data augmentation was particularly helpful; without it, the models would hallucinate the rest of a partial sentence.

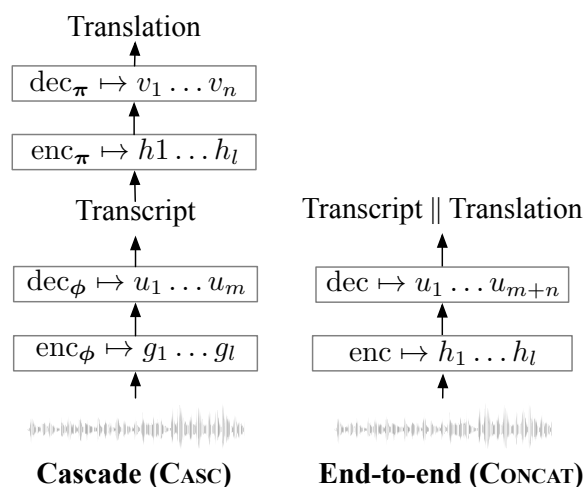


Figure 3: Architectures of the cascade and concatenated (E2E) model

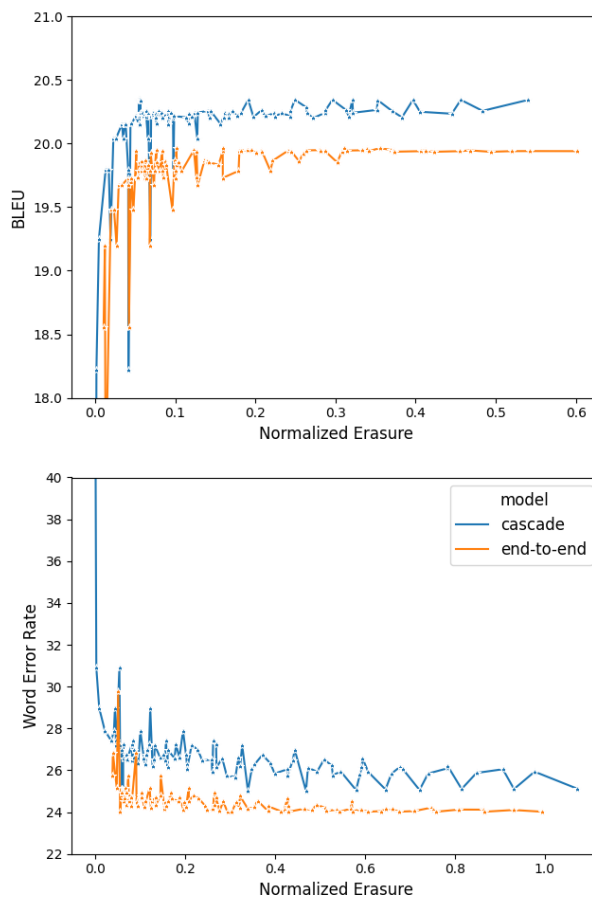


Figure 4: Left: average lag in seconds vs NE score. Right: NE vs WER score. All values are the mean of the results from the eight target languages.

We further found that starting the prefix sampling data augmentation too late in training was also negative. After testing initial models on the dev set, we found that starting this additional augmentation 15 epochs after training was best.

## D Utterance Segmentation

We follow the audio segments provided in the MuST-C corpus, created through a use of human alignment and XNMT (Neubig et al., 2018). We note that there exist a variety of methods for creating segments for such models, however, we leave additional exploration of E2E alignment methods as future work.