

Harnessing Indirect Training Data for End-to-End Automatic Speech Translation: Tricks of the Trade

Juan Pino¹, Liezl Puzon¹, Jiatao Gu¹, Xutai Ma^{1,2}, Arya D. McCarthy^{1,2}, Deepak Gopinath¹

¹Facebook

²Johns Hopkins University

{juancarabina, lie, jgu, dgopinath}@fb.com {xutai_ma, arya}@jhu.edu

Abstract

For automatic speech translation (AST), end-to-end approaches are outperformed by cascaded models that transcribe with automatic speech recognition (ASR), then translate with machine translation (MT). A major cause of the performance gap is that, while existing AST corpora are small, massive datasets exist for both the ASR and MT subsystems. In this work, we evaluate several data augmentation and pretraining approaches for AST, by comparing all on the same datasets. Simple data augmentation by translating ASR transcripts proves most effective on the English–French augmented LibriSpeech dataset, closing the performance gap from 8.2 to 1.4 BLEU, compared to a very strong cascade that could directly utilize copious ASR and MT data. The same end-to-end approach plus fine-tuning closes the gap on the English–Romanian MuST-C dataset from 6.7 to 3.7 BLEU. In addition to these results, we present practical recommendations for augmentation and pretraining approaches. Finally, we decrease the performance gap to 0.01 BLEU using a Transformer-based architecture.

1. Introduction

Automatic speech-to-text translation (AST) is the task of transforming speech input to its corresponding textual translation. Traditionally, AST has been conducted via a cascade approach [1, 2]: an automatic speech recognition (ASR) system creates a transcript of the speech signal, which is then translated by a machine translation (MT) system. Recent approaches use end-to-end neural network models [3, 4], which are directly inspired by end-to-end models for ASR [5, 6]. Cascade approaches can benefit from large amounts of training data available to its components for certain language pairs. For example, an English ASR model can be trained on 960 hours of speech [7], and an English–French translation model can be trained on about 40 million sentence pairs [8].

End-to-end approaches have very limited data available [9, 1, 10]; nevertheless, they present several benefits. First, end-to-end models can enable lower inference latency since they involve only one prediction. Second, it may be easier to reduce the model size for a single integrated model. Finally, end-to-end approaches avoid compounding errors

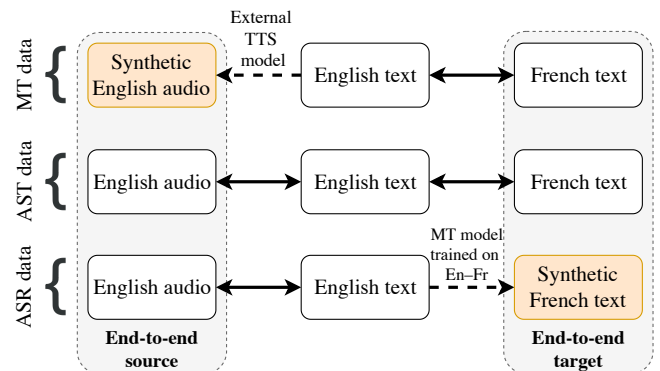


Figure 1: AST datasets typically include three parts: recorded speech, transcripts, and translations. While cascaded models can leverage these, an advantage over end-to-end systems is that they can also leverage data that provide adjacent pairs—which are far more prevalent. The approaches we investigate involve completing the triplets, so end-to-end systems can also benefit from incomplete triplets.

from the ASR and MT models.

End-to-end models for AST have been shown to perform better than or on par with cascade models [4, 11] when both are trained only on speech translation parallel corpora. However, when additional data are used to train its ASR and MT subsystems, the cascade outperforms the vanilla end-to-end approach [3, 12, 13]. In this work, we explore several techniques that leverage the wealth of ASR and MT data to aid end-to-end systems, by means of data augmentation.

The major contributions of this paper are:

1. **We confirm that end-to-end models underperform cascade models by a large margin**, when the components can be trained on additional ASR and MT training data while the end-to-end model is constrained to be trained on AST training data. In particular, we build a very strong cascade model that outperforms a previously reported system [11] by 5.5 BLEU.
2. **We investigate the strategies that can improve end-to-end AST models.** We augment the data by leveraging ASR training data with MT and MT training data with

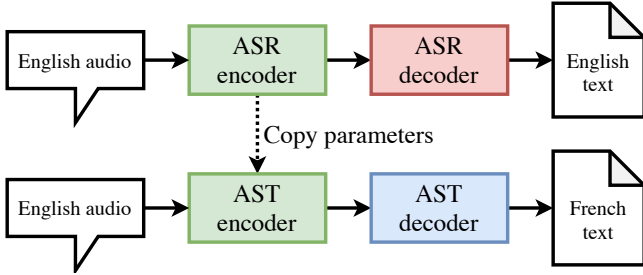


Figure 2: Pre-training end-to-end AST encoder on ASR task.

text-to-speech synthesis (TTS). We also study the effect of pretraining the ASR encoder as well as how to better utilize out-of-domain augmented data with fine-tuning. In the case of TTS-augmented data, we analyze the effect of the amount of data added, which TTS engine is used, and whether one speaker or multiple speakers are used to generate the data.

3. **We benchmark the performance of several architectures on AST task on public available datasets.** We first propose an extension to the Bérard model [11] that increases its capacity for training on larger data settings. We also benchmark models on the AST task that have been previously applied to the ASR task only: VGG LSTM [14] and VGG Transformer [15]. To our knowledge, this is the first time the VGG Transformer architecture has been applied to the AST task. For better reproducibility, experiments are conducted on two publicly available datasets, AST Librispeech [9] and MuST-C [10]. With data augmentation, pretraining, fine-tuning and careful architecture selection, we obtain competitive end-to-end models on the corresponding English–French (En–Fr) and English–Romanian (En–Ro) tasks.

2. Approach

For certain language pairs, cascade models can access large amounts of training data. In this section, we present our strategies to leverage this additional data for end-to-end models.

The first prong of our approach involves generating synthetic data to augment the existing AST data (§5.2), and it is summarised in Figure 1. MT models can be used to generate synthetic AST training data by automatically translating the transcript portion of ASR training data. (For our languages of interest, the MT model can be trained on a large amount of data and will be able to generate high quality translations.) In addition, the MT model itself would not be directly used in training the end-to-end AST model, which will avoid reinforcing errors produced by that model.

Similarly, we can generate additional synthetic AST training data by generating speech from the source side of an MT parallel corpus. This technique is similar to backtranslation [16] and has been previously applied to end-to-end ASR

training [17].

Another important aspect to consider is that the additional synthetic data we generate may not be in-domain. This can be problematic when there is a large gap between the amount of available gold data compared to the amount of synthetic data. We investigate fine-tuning techniques to address this issue (§5.4).

As a second prong to the approach, we examine pre-training (§5.3): we can use large ASR corpus to pretrain the speech encoder of an AST model; see Figure 2.

3. Models

In this section, we describe the various model architectures used for ASR and AST experimentation.

3.1. Bérard Model and Extension

We use a similar architecture to [11] with a speech encoder consisting of two non-linear layers followed by two convolutional layers and three bidirectional LSTM layers, and a custom LSTM decoder. In addition, the input state to an LSTM layer is the state emitted at the current timestep by the layer beneath. The input state to the bottom layer is the state emitted by the top LSTM layer at the previous timestep. Preliminary experiments showed that passing the state from the previous timestep to the LSTM layer one level above at the current timestep was not as effective. Finally, we extend the architecture to an arbitrary number N of decoder layers as shown in Equation 1:

$$\begin{aligned}
 s_t^1, o_t^1 &= \text{LSTM}^1(s_{t-1}^N, e(y_t)) \\
 c_t &= \text{attention}(o_t^1, h) \\
 s_t^n, o_t^n &= \text{LSTM}^n(s_t^{n-1}, c_t)
 \end{aligned} \tag{1}$$

where subscript t indicates the timestep, superscript n indicates the position in the stack of LSTMs, s is the state, o is the output, e is the embedding function, y is a target token and c is a context vector.

3.2. VGG LSTM

We investigate the performance of ASR and AST with a model similar to the ESPnet¹ implementation [14]. The encoder is composed of two blocks of VGG layers [18] followed by bidirectional LSTM layers. We use a hybrid attention mechanism from [5] that takes into account both location and content information. The decoder is an LSTM, following [19]. This model will subsequently be called VGGLSTM.

3.3. VGG Transformer

We also investigate the performance of a Transformer model. To our knowledge, this is the first time this type of model is applied to the AST task. We use a variant of the Transformer network [20] that has been shown to perform well on the

¹<https://github.com/espnet/espnet>

ASR task [15] by replacing the sinusoidal positional embedding with input representations learned by convolutions. This model will subsequently be called VGGTRANSFORMER.

4. Experimental Setup

4.1. Datasets

For both the En–Fr and En–Ro language pairs, we use three datasets corresponding to the AST, ASR and MT tasks. We choose some of the largest publicly available datasets for ASR and MT in order to have the ASR and MT models in an unconstrained-like setting and make the comparison between end-to-end and cascade models more realistic. Dataset statistics are summarized in Table 1.

For the En–Fr AST task, we use the publicly available augmented Librispeech corpus [9] (*AST Librispeech*), which is the second largest dataset publicly available for this task. We compare our results to the most recent work on this dataset [11]. For En–Ro AST, we use the recently released MuST-C corpus [10], which is the largest publicly available dataset, and compare our results to the original dataset benchmarks.

For the En ASR task, we use Librispeech [7] (*ASR Librispeech*) which is in the same domain as AST Librispeech and also the largest publicly available dataset for ASR. Since the validation and test set from AST Librispeech come from the training portion of ASR Librispeech, we filter all training utterances from ASR Librispeech that contain any of the validation or test utterances from AST Librispeech with at least two words. We allow ASR Librispeech training utterances to contain validation and test utterances from AST Librispeech with only one word, otherwise we would for example remove all training utterances containing the word “no”. The ASR Librispeech corpus is used for the ASR task for both En–Fr and En–Ro experiments.

For the En–Fr MT task, we use the En–Fr parallel data available as part of the WMT14² competition [8]. For En–Ro, we use the En–Ro parallel data available for the WMT16³ competition [21].

4.2. Preprocessing Settings

For the En–Fr task, we follow the same preprocessing as [11]. The English text is simply lowercased as it does not contain punctuation. The French text is punctuation-normalized, tokenized and lowercased. For the En–Ro task, the English text is tokenized, punctuation-stripped and lowercased. The Romanian text is punctuation-normalized and tokenized but the casing is preserved, following [10]. We do not limit the number of frames in the training data except to avoid GPU out-of-memory errors.

When training on AST Librispeech only, we use a character-level decoder. Otherwise, we use a unigram

Dataset	# utterances	# hours
En–Fr		
AST Librispeech	94,542	100
ASR Librispeech	265,754	902
WMT14	29,362,441	-
dev	1071	2
test	2048	3:44
En–Ro		
MuST-C	236,168	432
WMT16	612,422	-
dev	1370	2:33
tst-COMMON	2556	4:10

Table 1: Dataset statistics for the AST, ASR and MT tasks for En–Fr and En–Ro. The development and test sets are common to the ASR and AST tasks. tst-COMMON is the common test set for MuST-C.

model with size 10,000 using the SentencePiece implementation [22] as training on larger datasets with a character-level decoder would be prohibitively slow.

4.3. Model Settings

We use two sets of hyperparameters for the Bérard architecture. When training on AST Librispeech (En–Fr), we reuse the same parameters as [11]. In all other settings—for En–Fr with additional data and En–Ro—we use 3 decoder layers based on the extended model presented in §3 in order to give more capacity to the model.

The VGG LSTM encoder uses 80 log-scaled mel spectrogram features, 2 VGG blocks with 64 and 128 channels, filter size 3, pooling size 2, 2 convolutional layers and layer normalization and 5 bidirectional LSTM layers of size 1024. The decoder uses embeddings of size 1024 and 2 LSTM layers of size 1024. The attention has dimension 1024 and 10 channels with filter size 201. VGG LSTM uses no dropout.

Our VGG TRANSFORMER also uses 80 features, the same VGG block configuration as VGG LSTM, 14 transformer encoder layers and 4 transformer decoder layers with size 1024, 16 heads, a feed forward network of size 4096 and dropout with probability 0.15. The VGG TRANSFORMER decoder uses target embeddings of size 128 and 4 convolutional layers with 256 channels, filter size 3 and layer normalization.

Table 2 gives the number of parameters for these four models as well as the Transformer model used for MT.

4.4. Training Settings

For the Bérard architecture, we use the Adam optimizer [23] with a learning rate of 0.001. For the smaller AST Librispeech task, we use a minibatch size of 16000 frames to help convergence. For other tasks, we use a minibatch size of 96,000 frames except for the VGG TRANSFORMER where

²<http://www.statmt.org/wmt14/translation-task.html>

³<http://statmt.org/wmt16/translation-task.html>

Architecture	# Parameters
Bérard	8.9M
Bérard with 3 decoder layers	13.5M
VGG LSTM	176M
VGG Transformer	258.9M
Transformer (MT)	214.2M

Table 2: Number of parameters for each model architecture.

we use 72,000 frames (to avoid memory issues). We also use delayed updates [24] in order to keep the same effective batch size and avoid GPU out-of-memory errors. All experiments are conducted on 8 GPUs. For other architectures than Bérard, we use ADADELTA [25] with a learning rate of 1 and we normalize the loss per utterance instead of per token. These hyperparameters were chosen based on preliminary experimentation on the ASR Librispeech task.

5. Experiments

5.1. Cascade Baselines

The baseline approach, *CASCADE*, involves two steps: first, transcribe input speech with an ASR model, then translate the transcript with an MT model. Both models are trained separately on large training datasets.

The ASR models for En–Fr use the same architectures from §3 and are trained on the full Librispeech corpus, which is much larger than the available AST data. For the En–Ro task, ASR models are trained on the MuST-C and the Librispeech datasets. We use a Transformer [20] as the basic MT architecture. More precisely, for En–Fr, we first pretrain a large Transformer model (*transformer big*) over the entire WMT14 corpus, then fine-tune this model on the AST Librispeech data. For En–Ro, we merge the MuST-C and the WMT16 corpora since they have comparable sizes and train a smaller Transformer (*transformer base*) on the joint corpus.

5.2. Data Augmentation

5.2.1. MT: Producing AST data from ASR data.

The MT models described in §5.1 are used to automatically translate the English transcript of ASR Librispeech into French and Romanian. The resulting synthetic data can directly be used as additional training data for end-to-end AST models.

5.2.2. TTS: Producing AST data from MT data.

We also explore augmenting the MT training data with TTS. This technique is similar to backtranslation [16] and has been previously applied to end-to-end ASR training [17]. We use two pretrained TTS engines. **TTS1** uses the OpenSeq2Seq framework [26] to generate speech samples in five different voices. The TTS model is based on an extension of the

Task	Model	Data	En–Fr	En–Ro
ASR	Bérard [11]	AST	15.1%	27.6%
	Bérard	ASR	10.0%	19.7%
MT	Bérard [11]	AST	19.3	28.2
	Transformer	AST + MT	24.7	28.7
	oracle BLEU [11]	—	19.3	
	Cascade [11]	ASR + AST + MT	15.8	
	Our cascade	ASR + AST + MT	21.3	21.0
	Bérard	AST	13.1	14.3
	Bérard		19.9	16.4
	+ pretraining	AST + MT	19.2	17.3
	Bérard		12.6	10.3
	+ fine-tuning		14.3	12.5
+ pretraining	AST + TTS	15.2	13.0	
+ pretraining + fine-tuning		16.4	15.2	
Bérard		19.0	10.0	
+ fine-tuning		19.3	12.3	
+ pretraining	AST + MT + TTS	19.2	14.3	
+ pretraining + fine-tuning		19.2	15.7	

Table 3: En–Fr results on AST Librispeech test set and En–Ro results on MuST-C test set with Bérard model. For ASR, we use the same AST subset of Librispeech as the ASR test set. ASR performance is measured by word error rate (WER; lower is better) and MT and AST performance is measured by BLEU (higher is better). MT, TTS, fine-tuning and pre-training help bridge the gap to a very strong cascade model.

Tacotron 2 model [27] with Global Style Tokens [28]. **TTS2** is trained on about 15 hours of single speaker data. The text comes from several domains such as Wikipedia, news articles, parliament speech, and novels. We use TTS1 to generate speech from a random sample of WMT14 with the same size as ASR Librispeech (265,754 utterances) and TTS2 to generate speech from WMT16 (612,422 utterances).

5.3. Speech Encoder Pretraining

Speech encoder pretraining is another way to use the full ASR Librispeech dataset [29]. We first pretrain an English ASR model on ASR Librispeech plus the TTS1 corpus generated in §5.2—the parallel corpus built from the generated TTS and WMT14 English text. We then take the encoder of the ASR model to initialize the encoder of an AST model with the same architecture.

5.4. Fine-tuning

On the En–Fr task, the TTS data is generated from WMT14, which is out-of-domain with respect to Librispeech. On the En–Ro task, both the TTS and the MT data are out-of-domain. We investigate fine-tuning as a technique to mitigate the domain shift. We fine-tune by continuing training on only AST Librispeech or MuST-C, starting from the best checkpoint on development data after convergence.

Architecture	BLEU
Cascade (ASR architecture; Transformer used for MT)	
BÉRARD	21.3
VGGLSTM	21.8
VGGTRANSFORMER	21.7
End-to-end AST	
BÉRARD	19.9
VGGLSTM	19.8
VGGTRANSFORMER	21.7

Table 4: Performance of the architectures we consider on the (En–Fr) AST Librispeech test set in the AST + MT data setting, which performed best in Table 3. The large VGGTRANSFORMER outperforms other end-to-end approaches, now within 0.15 BLEU of the best cascade model.

6. Results and Analysis

We first investigate techniques to improve end-to-end AST on Bérard model. Results for the En–Fr and En–Ro tasks are summarized in Table 3. The ASR and MT components we trained with additional data are very strong: the cascade model outperforms the vanilla end-to-end model by 8.2 BLEU (21.3 vs. 13.1) on the AST task. It is also important to note that our cascade baseline is greater than the best reported result on this task by 5.5 BLEU and even better than the previously reported oracle BLEU of 19.3 by 2 BLEU, where the gold transcript is passed to the translation system.

6.1. Effect of Data Augmentation

By augmenting ASR Librispeech with automatic translations (AST + MT), we show an improvement of 6.8 BLEU for En–Fr and 3.0 BLEU for En–Ro. Under this setting, we observe that pretraining is not beneficial for En–Fr but is for En–Ro.⁴ In this setting on En–Ro, we close the gap between cascade and end-to-end BLEU from 6.7 to 3.7.

Additional TTS-augmented data (AST + TTS) initially hurts performance for both language pairs. With pretraining and fine-tuning, TTS data provides a gain of 3.3 BLEU over the vanilla end-to-end baseline. However, it still underperforms the model using MT-augmented data only.

We also augmented the data with MT and TTS (AST + MT + TTS) at the same time. We find that it does not provide additional gain over using MT data only. In general, MT data can efficiently help the model, while TTS data is less efficient and can be hurtful. We analyze TTS in more detail in §6.4.

We present the performance of different architectures on the higher-resource En–Fr task in Table 4. In the AST + MT setting, we obtain BLEU of 19.9 with Bérard, 19.8 with VGGLSTM and 21.7 with VGGTRANSFORMER, showing that

⁴Data setting experiments were conducted on the Bérard architecture. We do not expect the conclusions on data augmentation, pretraining and fine-tuning to change on the VGG architectures.

Dataset	Bérard	Δ PT	+ FT	Δ PT
AST + TTS	12.6	2.6	14.3	2.0
AST + MT + TTS	19.0	0.1	19.2	0.0
MuST-C + TTS	10.3	2.7	13.0	2.2
MuST-C + MT	16.4	-0.2	15.7	0.8
MuST-C + MT + TTS	10.0	4.2	14.3	1.4

Table 5: Effect of pretraining (PT), when added to the baseline or the fine-tuned (FT) baseline.

these architectures are effective with additional data. VGGTRANSFORMER in particular (obtaining the best end-to-end AST score) is on par with the VGGLSTM cascade and -0.1 BLEU behind the VGGLSTM cascade.

6.2. Effect of Encoder Pretraining

Results are summarized in Table 5. Pretraining improves the AST + TTS system by +2.6 BLEU and +2.0 BLEU with fine-tuning, improves the MuST-C + TTS system by +2.7 BLEU and +2.2 BLEU with fine-tuning, and improves the MuST-C + MT + TTS system by +4.2 BLEU and +1.4 BLEU.

However, gains from pretraining ASR on the full Librispeech dataset do not compound with gains from MT augmentation. Pretraining does not help as much in the AST + MT + TTS setup, showing a negligible change in BLEU score. Pretraining has mixed results for the MuST-C + MT case, with -0.2 BLEU and +0.8 BLEU with pretraining.

Pretraining on in-domain ASR data is not a good substitute for MT-augmenting the ASR data. However we note that using a pretrained speech encoder will speed up convergence of the AST model. Thus, pretraining could be used in experiments with the same architecture and provides a good starting point for more rapid iteration.

6.3. Effect of Fine-tuning

Table 6 summarizes the fine-tuning results. We apply fine-tuning whenever TTS-augmentation is used. Fine-tuning seems to mitigate the effect of domain shift introduced by the additional out-of-domain TTS data: in the Librispeech AST + TTS setup, fine-tuning improves by +1.7 (+1.2 BLEU on top of pretraining). For the MuST-C + MT setup, we see +2.3 BLEU (+1.4 on top of pretraining).

Fine-tuning does not improve the AST Librispeech model on top of MT-augmentation though, likely because the MT-augmented data is already in-domain. For the AST + TTS + MT setup, we see neutral results: +0.3 BLEU and no effect on top of pretraining. However, for the MuST-C + MT setup we see a gain of +0.3–0.9 BLEU because the MT-augmented data is out-of-domain for the MuST-C dataset.

6.4. TTS Data: Quantity, Quality and Diversity

How does augmenting the AST training data with TTS affect performance in the En–Fr task? First, in Figure 3, we see that

Dataset	Bérard	Δ FT	+ preT	Δ FT
AST + TTS	12.6	1.8	15.2	1.2
AST + MT + TTS	19.0	0.2	19.2	0.0
MuST-C + TTS	10.3	2.2	13.0	2.2
MuST-C + MT	16.4	0.9	16.2	0.3
MuST-C + MT + TTS	10.0	2.3	14.3	1.4

Table 6: Effect of fine-tuning (FT), both when added to the baseline or the pretrained (preT) baseline.

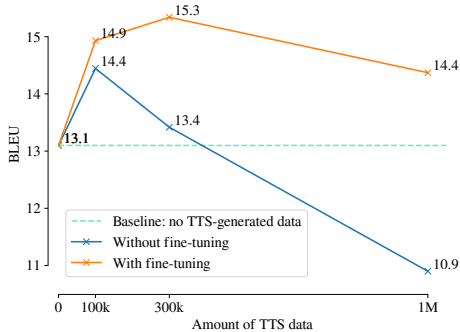


Figure 3: Varying amounts of TTS data for the En-Fr task, with and without fine-tuning.

while adding TTS data up to 100,000 utterances improves the performance, the performance degrades beyond that. We hypothesize that this is because the additional TTS data is out of domain. With fine-tuning, adding up to 300,000 utterances improves performance, which confirms our hypothesis; however, adding 1M utterances starts degrading performance. In the future, we will investigate how to make more effective use of larger quantities of TTS-generated data.

We also study the effect of using single-speaker or multi-speaker TTS. We use a sample of size 300k utterances from WMT14 and generate speech with the TTS1 engine using the first speaker Speaker 0, the second speaker Speaker 1, and all five speakers in a round-robin fashion. Finally, we investigate whether the quality of the TTS engine matters. For the same sample of 300k utterances, we generate speech using the TTS2 engine both from the English text and the corresponding French translations. The latter is analogous to copying the target to the source in machine translation [30].

Results are reported in Table 7. Comparing the first two rows, we conclude that performance may vary depending on the speaker. The third row shows that using multiple speakers performs on par with choosing Speaker 0 (+0.2 BLEU) but outperforms choosing Speaker 1 (+0.9 BLEU). In the future, we will investigate whether the multi speaker approach can mitigate the effect observed in Figure 3. Finally, comparing rows 1 and 4, we conclude that the quality of the TTS may matter marginally, with TTS2 slightly outperforming TTS1. The last row shows that the analogue of copying target to

Row #	Configuration	TTS Engine	BLEU
1	Speaker 0	TTS1	13.4
2	Speaker 1	TTS1	12.7
3	Multi speaker	TTS1	13.6
4	English	TTS2	13.6
5	French	TTS2	13.6

Table 7: Effect of the number of speakers and the TTS engine when adding TTS data.

the source in machine translation is an interesting avenue for further investigation.

7. Related Work

Initial attempts at speech translation [2] incorporate lattices from ASR systems as inputs to statistical MT models [31, 32]. More recent approaches have focused on end-to-end models. [3] demonstrate the viability of this approach on a small synthetic corpus. [4] outperform a cascade model using a similar architecture to an attention-based ASR model. [4] and [11] show that multi-task learning can further improve an end-to-end model. Pretraining has also been shown to improve end-to-end models [11, 29].

[12] note that cascaded models are at a disadvantage when constrained to be trained on speech translation data only. They show how to leverage additional ASR and MT training data with an attention-passing mechanism. [13] improve an end-to-end model using MT-augmented and TTS-augmented data but do so *on a proprietary dataset*. In contrast, we experiment on two public datasets where we obtain new state-of-the-art performance; further, we provide additional analyses on network architectures and recommendations on how to better leverage TTS-augmented data.

8. Conclusion

We have demonstrated that cascaded models are very competitive when not constrained to only train on AST data. We have studied several techniques aimed at bridging the gap between end-to-end and cascade models. With data augmentation, pretraining, fine-tuning and architecture selection, we trained end-to-end models that show competitive performance when compared to cascade approach. Our approaches reduced the performance gap between end-to-end and strong cascade models, from 8.2 to 1.4 BLEU on En-Fr Librispeech AST data and from 6.7 to 3.7 on the En-Ro MuST-C corpus. We also analyzed the effect of TTS data in terms of quality, quantity, and the use of single speaker vs. multiple speakers, and we provide recommendations on how to harness this type of data. In the future, we would like to investigate how to better use larger-scale TTS-generated data.

9. References

- [1] M. Post, G. Kumar, A. Lopez, D. Karakos, C. Callison-Burch, and S. Khudanpur, “Improved speech-to-text translation with the Fisher and Callhome Spanish–English speech translation corpus,” in *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Heidelberg, Germany, December 2013.
- [2] H. Ney, “Speech translation: Coupling of recognition and translation,” in *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No. 99CH36258)*, vol. 1. IEEE, 1999, pp. 517–520.
- [3] A. Bérard, O. Pietquin, L. Besacier, and C. Servan, “Listen and translate: A proof of concept for end-to-end speech-to-text translation,” in *NIPS Workshop on end-to-end learning for speech and audio processing*, 2016.
- [4] R. J. Weiss, J. Chorowski, N. Jaitly, Y. Wu, and Z. Chen, “Sequence-to-sequence models can directly translate foreign speech,” in *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, 2017, pp. 2625–2629. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-503>
- [5] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Advances in neural information processing systems*, 2015, pp. 577–585.
- [6] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [7] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [8] O. Bojar, C. Buck, C. Federmann, B. Haddow, P. Koehn, J. Leveling, C. Monz, P. Pecina, M. Post, H. Saint-Amant, R. Soricut, L. Specia, and A. Tamchyna, “Findings of the 2014 workshop on statistical machine translation,” in *Proceedings of the Ninth Workshop on Statistical Machine Translation*. Baltimore, Maryland, USA: Association for Computational Linguistics, June 2014, pp. 12–58. [Online]. Available: <http://www.aclweb.org/anthology/W/W14/W14-3302>
- [9] A. Kocabiyikoglu, L. Besacier, and O. Kraif, “Augmenting Librispeech with French translations: A multimodal corpus for direct speech translation evaluation,” in *LREC (Language Resources and Evaluation Conference)*, 2018.
- [10] M. A. Di Gangi, R. Cattoni, L. Bentivogli, M. Negri, and M. Turchi, “MuST-C: a multilingual speech translation corpus,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, Minneapolis, MN, USA, June 2019.
- [11] A. Bérard, L. Besacier, A. C. Kocabiyikoglu, and O. Pietquin, “End-to-end automatic speech translation of audiobooks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6224–6228.
- [12] M. Sperber, G. Neubig, J. Niehues, and A. Waibel, “Attention-passing models for robust and data-efficient end-to-end speech translation,” *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 313–325, Mar. 2019.
- [13] Y. Jia, M. Johnson, W. Macherey, R. J. Weiss, Y. Cao, C.-C. Chiu, N. Ari, S. Laurenzo, and Y. Wu, “Leveraging weakly supervised data to improve end-to-end speech-to-text translation,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7180–7184.
- [14] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, “ESPnet: End-to-end speech processing toolkit,” in *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018.*, 2018, pp. 2207–2211. [Online]. Available: <https://doi.org/10.21437/Interspeech.2018-1456>
- [15] A. Mohamed, D. Okhonko, and L. Zettlemoyer, “Transformers with convolutional context for ASR,” *CoRR*, vol. abs/1904.11660, 2019. [Online]. Available: <http://arxiv.org/abs/1904.11660>
- [16] R. Sennrich, B. Haddow, and A. Birch, “Improving neural machine translation models with monolingual data,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 86–96.
- [17] T. Hayashi, S. Watanabe, Y. Zhang, T. Toda, T. Hori,

- R. Astudillo, and K. Takeda, “Back-translation-style data augmentation for end-to-end ASR,” in *2018 IEEE Spoken Language Technology Workshop, SLT 2018 - Proceedings*, ser. 2018 IEEE Spoken Language Technology Workshop, SLT 2018 - Proceedings. Institute of Electrical and Electronics Engineers Inc., 2019, pp. 426–433.
- [18] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [19] T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sept. 2015, pp. 1412–1421. [Online]. Available: <https://www.aclweb.org/anthology/D15-1166>
- [20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [21] O. Bojar, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, M. Huck, A. Jimeno Yepes, P. Koehn, V. Logacheva, C. Monz, M. Negri, A. Neveol, M. Neves, M. Popel, M. Post, R. Rubino, C. Scarton, L. Specia, M. Turchi, K. Verspoor, and M. Zampieri, “Findings of the 2016 conference on machine translation,” in *Proceedings of the First Conference on Machine Translation*. Berlin, Germany: Association for Computational Linguistics, August 2016, pp. 131–198. [Online]. Available: <http://www.aclweb.org/anthology/W/W16/W16-2301>
- [22] T. Kudo and J. Richardson, “SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Brussels, Belgium: Association for Computational Linguistics, Nov. 2018, pp. 66–71. [Online]. Available: <https://www.aclweb.org/anthology/D18-2012>
- [23] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [24] D. Saunders, F. Stahlberg, A. de Gispert, and B. Byrne, “Multi-representation ensembles and delayed SGD updates improve syntax-based NMT,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Melbourne, Australia: Association for Computational Linguistics, July 2018, pp. 319–325. [Online]. Available: <https://www.aclweb.org/anthology/P18-2051>
- [25] M. D. Zeiler, “ADADELTA: an adaptive learning rate method,” *CoRR*, vol. abs/1212.5701, 2012. [Online]. Available: <http://arxiv.org/abs/1212.5701>
- [26] O. Kuchaiev, B. Ginsburg, I. Gitman, V. Lavrukhin, C. Case, and P. Micikevicius, “Openseq2seq: extensible toolkit for distributed and mixed precision training of sequence-to-sequence models,” *CoRR*, vol. abs/1805.10387, 2018. [Online]. Available: <http://arxiv.org/abs/1805.10387>
- [27] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, et al., “Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [28] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. Stockholmsmässan, Stockholm Sweden: PMLR, 10–15 Jul 2018, pp. 5180–5189. [Online]. Available: <http://proceedings.mlr.press/v80/wang18h.html>
- [29] S. Bansal, H. Kamper, K. Livescu, A. Lopez, and S. Goldwater, “Pre-training on high-resource speech recognition improves low-resource speech-to-text translation,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 58–68.
- [30] A. Currey, A. V. Miceli Barone, and K. Heafield, “Copied monolingual data improves low-resource neural machine translation,” in *Proceedings of the Second Conference on Machine Translation*. Copenhagen, Denmark: Association for Computational Linguistics, Sept. 2017, pp. 148–156. [Online]. Available: <https://www.aclweb.org/anthology/W17-4715>
- [31] E. Matusov, S. Kanthak, and H. Ney, “On the integra-

tion of speech recognition and statistical machine translation,” in *Ninth European Conference on Speech Communication and Technology*, 2005.

- [32] F. Casacuberta, M. Federico, H. Ney, and E. Vidal, “Recent efforts in spoken language translation,” *IEEE Signal Processing Magazine*, vol. 25, no. 3, pp. 80–88, 2008.