

Pre-training on High-Resource Speech Recognition Improves Low-Resource Speech-to-Text Translation

Sameer Bansal¹ Herman Kamper² Karen Livescu³ Adam Lopez¹ Sharon Goldwater¹

¹School of Informatics, University of Edinburgh, UK

²E&E Engineering, Stellenbosch University, South Africa

³Toyota Technological Institute at Chicago, USA

{sameer.bansal, sgwater, alopez}@inf.ed.ac.uk
kamperh@sun.ac.za
klivescu@ttic.edu

Abstract

We present a simple approach to improve direct speech-to-text translation (ST) when the source language is low-resource: we pre-train the model on a high-resource automatic speech recognition (ASR) task, and then fine-tune its parameters for ST. We demonstrate that our approach is effective by pre-training on 300 hours of English ASR data to improve Spanish-English ST from 10.8 to 20.2 BLEU when only 20 hours of Spanish-English ST training data are available. Through an ablation study, we find that the pre-trained encoder (acoustic model) accounts for most of the improvement, despite the fact that the shared language in these tasks is the target language text, not the source language audio. Applying this insight, we show that pre-training on ASR helps ST even when the ASR language differs from both source and target ST languages: pre-training on French ASR also improves Spanish-English ST. Finally, we show that the approach improves performance on a true low-resource task: pre-training on a combination of English ASR and French ASR improves Mboshi-French ST, where only 4 hours of data are available, from 3.5 to 7.1 BLEU.

1 Introduction

Speech-to-text Translation (ST) has many potential applications for low-resource languages: for example in language documentation, where the source language is often unwritten or endangered (Besacier et al., 2006; Martin et al., 2015; Adams et al., 2016a,b; Anastasopoulos and Chiang, 2017); or in crisis relief, where emergency workers might need to respond to calls or requests in a foreign language (Munro, 2010). Traditional ST is a pipeline of automatic speech recognition (ASR) and machine translation (MT), and thus requires transcribed source audio to train ASR and parallel text to train MT. These resources are often unavailable

for low-resource languages, but for our potential applications, there may be some source language audio paired with target language text translations. In these scenarios, end-to-end ST is appealing.

Recently, Weiss et al. (2017) showed that end-to-end ST can be very effective, achieving an impressive BLEU score of 47.3 on Spanish-English ST. But this result required over 150 hours of translated audio for training, still a substantial resource requirement. By comparison, a similar system trained on only 20 hours of data for the same task achieved a BLEU score of 5.3 (Bansal et al., 2018). Other low-resource systems have similarly low accuracies (Anastasopoulos and Chiang, 2018; Bérard et al., 2018).

To improve end-to-end ST in low-resource settings, we can try to leverage other data resources. For example, if we have transcribed audio in the source language, we can use multi-task learning to improve ST (Anastasopoulos and Chiang, 2018; Weiss et al., 2017; Bérard et al., 2018). But source language transcriptions are unlikely to be available in our scenarios of interest.

Could we improve low-resource ST by leveraging data from a high-resource language? For ASR, training a single model on multiple languages can be effective for all of them (Toshniwal et al., 2018b; Deng et al., 2013). For MT, *transfer learning* (Thrun, 1995) has been very effective: pre-training a model for a high-resource language pair and transferring its parameters to a low-resource language pair when the target language is shared (Zoph et al., 2016; Johnson et al., 2017). Inspired by these successes, we show that low-resource ST can leverage transcribed audio in a high-resource target language, or even a different language altogether, simply by pre-training a model for the high-resource ASR task, and then transferring and fine-tuning some or all of the model’s parameters for low-resource ST.

We first test our approach using Spanish as the source language and English as the target. After training an ASR system on 300 hours of English, fine-tuning on 20 hours of Spanish-English yields a BLEU score of 20.2, compared to only 10.8 for an ST model without ASR pre-training. Analyzing this result, we discover that the main benefit of pre-training arises from the transfer of the *encoder* parameters, which model the input acoustic signal. In fact, this effect is so strong that we also obtain improvements by pre-training on a language that differs from both the source and the target: pre-training on French and fine-tuning on Spanish-English. We hypothesize that pre-training the encoder parameters, even on a different language, allows the model to better learn about linguistically meaningful phonetic variation while normalizing over acoustic variability such as speaker and channel differences. We conclude that the acoustic-phonetic learning problem, rather than translation itself, is one of the main difficulties in low-resource ST. A final set of experiments confirm that ASR pre-training also helps on another language pair where the input is truly low-resource: Mboshi-French.

2 Method

For both ASR and ST, we use an encoder-decoder model with attention adapted from Weiss et al. (2017), Bérard et al. (2018) and Bansal et al. (2018), as shown in Figure 1. We use the same model architecture for all our models, allowing us to conveniently transfer parameters between them. We also constrain the hyper-parameter search to fit a model into a single Titan X GPU, allowing us to maximize available compute resources.

We use a pre-trained English ASR model to initialize training of Spanish-English ST models, and a pre-trained French ASR model to initialize training of Mboshi-French ST models. During ST training, all model parameters are updated. In these configurations, the decoder shares the same vocabulary across the ASR and ST tasks. This is practical for settings where the target text language is high-resource with ASR data available.

In settings where both ST languages are low-resource, ASR data may only be available in a third language. To test whether transfer learning will help in this setting, we use a pre-trained French ASR model to train Spanish-English ST models; and English ASR for Mboshi-French models. In these cases, the ST languages are different from the

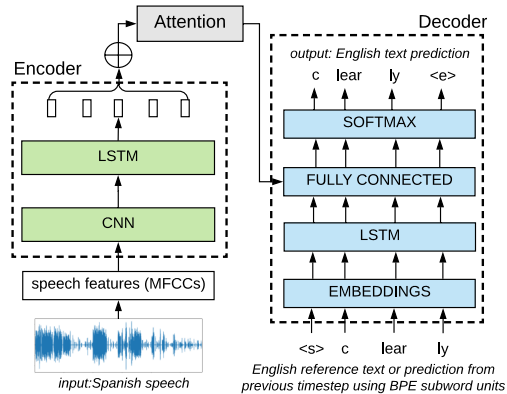


Figure 1: Encoder-decoder with attention model architecture for both ASR and ST. The encoder input is the Spanish speech utterance *claro*, translated as *clearly*, represented as BPE (subword) units.

ASR language, so we can only transfer the encoder parameters of the ASR model, since the dimensions of the decoder’s output softmax layer are indexed by the vocabulary, which is not shared.¹ Sharing only the speech encoder parameters is much easier, since the speech input can be preprocessed in the same manner for all languages. This form of transfer learning is more flexible, as there are no constraints on the ASR language used.

3 Experimental Setup

3.1 Data sets

English ASR. We use the Switchboard Telephone speech corpus (Godfrey and Holliman, 1993), which consists of around 300 hours of English speech and transcripts, split into 260k utterances. The development set consists of 5 hours that we removed from the training set, split into 4k utterances.

French ASR. We use the French speech corpus from the GlobalPhone collection (Schultz, 2002), which consists of around 20 hours of high quality read speech and transcripts, split into 9k utterances. The development set consists of 2 hours, split into 800 utterances.

Spanish-English ST. We use the Fisher Spanish speech corpus (Graff et al., 2010), which consists of 160 hours of telephone speech in a variety of Spanish dialects, split into 140K utterances. To simulate low-resource conditions, we construct smaller train-

¹Using a shared vocabulary of characters or subwords is an interesting direction for future work, but not explored here.

ing corpora consisting of 50, 20, 10, 5, or 2.5 hours of data, selected at random from the full training data. The development and test sets each consist of around 4.5 hours of speech, split into 4K utterances. We do not use the corresponding Spanish transcripts; our target text consists of English translations that were collected through crowdsourcing (Post et al., 2013, 2014).

Mboshi-French ST. Mboshi is a Bantu language spoken in the Republic of Congo, with around 160,000 speakers.² We use the Mboshi-French parallel corpus (Godard et al., 2018), which consists of around 4 hours of Mboshi speech, split into a training set of 5K utterances and a development set of 500 utterances. Since this corpus does not include a designated test set, we randomly sampled and removed 200 utterances from training to use as a development set, and use the designated development data as a test set.

3.2 Preprocessing

Speech. We convert raw speech input to 13-dimensional MFCCs using Kaldi (Povey et al., 2011).³ We also perform speaker-level mean and variance normalization.

Text. The target text of the Spanish-English data set contains 1.5M word tokens and 17K word types. If we model text as sequences of words, our model cannot produce any of the unseen word types in the test data and is penalized for this, but it can be trained very quickly (Bansal et al., 2018). If we instead model text as sequences of characters as done by Weiss et al. (2017), we would have 7M tokens and 100 types, resulting in a model that is open-vocabulary, but very slow to train (Bansal et al., 2018). As an effective middle ground, we use byte pair encoding (BPE; Sennrich et al., 2016) to segment each word into subwords, each of which is a character or a high-frequency sequence of characters—we use 1000 of these high-frequency sequences. Since the set of subwords includes the full set of characters, the model is still open vocabulary; but it results in a text with only 1.9M tokens and just over 1K types, which can be trained almost as fast as the word-level model.

The vocabulary for BPE depends on the fre-

²ethnologue.com/language/mdw

³In preliminary experiments, we did not find much difference between MFCCs and more raw spectral representations like Mel filterbank features.

quency of character sequences, so it must be computed with respect to a specific corpus. For English, we use the full 160-hour Spanish-English ST target training text. For French, we use the Mboshi-French ST target training text.

3.3 Model architecture for ASR and ST

Speech encoder. As shown schematically in Figure 1, MFCC feature vectors, extracted using a window size of 25 ms and a step size of 10ms, are fed into a stack of two CNN layers, with 128 and 512 filters with a filter width of 9 frames each. In each CNN layer we stride with a factor of 2 along time, apply a ReLU activation (Nair and Hinton, 2010), and apply batch normalization (Ioffe and Szegedy, 2015). The output of the CNN layers is fed into a three-layer bi-directional long short term memory network (LSTM; Hochreiter and Schmidhuber, 1997); each hidden layer has 512 dimensions.

Text decoder. At each time step, the decoder chooses the most probable token from the output of a softmax layer produced by a fully-connected layer, which in turn receives the current state of a recurrent layer computed from previous time steps and an attention vector computed over the input. Attention is computed using the *global attentional model* with *general* score function and *input-feeding*, as described in Luong et al. (2015). The predicted token is then fed into a 128-dimensional embedding layer followed by a three-layer LSTM to update the recurrent state; each hidden state has 256 dimensions. While training, we use the predicted token 20% of the time as input to the next decoder step and the training token for the remaining 80% of the time (Williams and Zipser, 1989). At test time we use beam decoding with a beam size of 5 and length normalization (Wu et al., 2016) with a weight of 0.6.

Training and implementation. Parameters for the CNN and RNN layers are initialized using the scheme from (He et al., 2015). For the embedding and fully-connected layers, we use Chainer’s (Tokui et al., 2015) default initialization.

We regularize using dropout (Srivastava et al., 2014), with a ratio of 0.3 over the embedding and LSTM layers (Gal, 2016), and a weight decay rate of 0.0001. The parameters are optimized using Adam (Kingma and Ba, 2015), with a starting alpha of 0.001.

Following some preliminary experimentation on our development set, we add Gaussian noise with standard deviation of 0.25 to the MFCC features during training, and drop frames with a probability of 0.10. After 20 epochs, we corrupt the true decoder labels by sampling a random output label with a probability of 0.3.

Our code is implemented in Chainer (Tokui et al., 2015) and is freely available.⁴

3.4 Evaluation

Metrics. We report BLEU (Papineni et al., 2002) for all our models.⁵ In low-resource settings, BLEU scores tend to be low, difficult to interpret, and poorly correlated with model performance. This is because BLEU requires exact four-gram matches only, but low four-gram accuracy may obscure a high unigram accuracy and inexact translations that partially capture the semantics of an utterance, and these can still be very useful in situations like language documentation and crisis response. Therefore, we also report word-level unigram precision and recall, taking into account *stem*, *synonym*, and *paraphrase* matches. To compute these scores, we use METEOR (Lavie and Agarwal, 2007) with default settings for English and French.⁶ For example, METEOR assigns “eat” a recall of 1 against reference “eat” and a recall of 0.8 against reference “feed”, which it considers a synonym match.

Naive baselines. We also include evaluation scores for a naive baseline model that predicts the K most frequent words of the training set as a bag of words for each test utterance. We set K to be the value at which precision/recall are most similar, which is always between 5 and 20 words. This provides an empirical lower bound on precision and recall, since we would expect any usable model to outperform a system that does not even depend on the input utterance. We do not compute BLEU for these baselines, since they do not predict sequences, only bags of words.

4 ASR results

Using the experimental setup of Section 3, we pre-trained ASR models in English and French, and report their word error rates (WER) on develop-

	en-100h	en-300h	fr-20h
WER	35.4	27.3	29.6

Table 1: Word Error Rate (WER, in %) for the ASR models used as pretraining, computed on Switchboard *train-dev* for English and Globalphone dev for French.

ment data in Table 1.⁷ We denote each ASR model by $L-Nh$, where L is a language code and N is the size of the training set in hours. For example, *en-300h* denotes an English ASR model trained on 300 hours of data.

Training ASR models for state-of-the-art performance requires substantial hyper-parameter tuning and long training times. Since our goal is simply to see whether pre-training is useful, we stopped pre-training our models after around 30 epochs (3 days) to focus on transfer experiments. As a consequence, our ASR results are far from state-of-the-art: current end-to-end Kaldi systems obtain 16% WER on Switchboard *train-dev*, and 22.7% WER on the French Globalphone dev set.⁸ We believe that better ASR pre-training may produce better ST results, but we leave this for future work.

5 Spanish-English ST

In the following, we denote an ST model by $S-T-Nh$, where S and T are source and target language codes, and N is the size of the training set in hours. For example, *sp-en-20h* denotes a Spanish-English ST model trained using 20 hours of data. We use the code *mb* for Mboshi and *fr* for French.

5.1 Using English ASR to improve ST

Figure 2 shows the BLEU and unigram precision/recall scores on the development set for baseline Spanish-English ST models and those trained after initializing with the *en-300h* model. Corresponding results on the test set (Table 2) reveal very similar patterns. The remainder of our analysis is confined to the development set. The naive baseline, which predicts the 15 most frequent English words in the training set, achieves a precision/recall of around 20%, setting a performance lower bound.

Low-resource: 20-50 hours of ST training data. Our baseline ST models substantially improve over

⁴github.com/0xSameer/ast

⁵We compute BLEU with `multi-bleu.pl` from the Moses toolkit (Koehn et al., 2007).

⁶cs.cmu.edu/~alavie/METEOR

⁷We computed WER with the NIST `sclite` script.

⁸These WER results taken from respective Kaldi recipes on GitHub, and may not represent the very best results on these data sets.

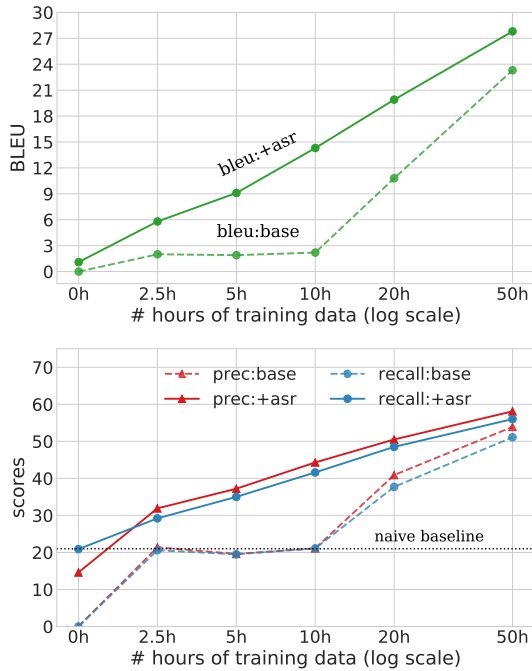


Figure 2: (top) BLEU and (bottom) Unigram precision/recall for Spanish-English ST models computed on Fisher dev set. base indicates no transfer learning; +asr are models trained by fine-tuning *en-300h* model parameters. *naive baseline* indicates the score when we predict the 15 most frequent English words in the training set.

previous results (Bansal et al., 2018) using the same train/test splits, primarily due to better regularization and modeling of subwords rather than words. Yet transfer learning still substantially improves over these strong baselines. For *sp-en-20h*, transfer learning improves dev set BLEU from 10.8 to 19.9, precision from 41% to 51%, and recall from 38% to 49%. For *sp-en-50h*, transfer learning improves BLEU from 23.3 to 27.8, precision from 54% to 58%, and recall from 51% to 56%.

Very low-resource: 10 hours or less of ST training data. Figure 2 shows that without transfer learning, ST models trained on less than 10 hours of data struggle to learn, with precision/recall scores close to or below that of the naive baseline. But with transfer learning, we see gains in precision and recall of between 10 and 20 points.

We also see that with transfer learning, a model trained on only 5 hours of ST data achieves a BLEU of 9.1, nearly as good as the 10.8 of a model trained on 20 hours of ST data without transfer learning. In other words, fine-tuning an English ASR model—which is relatively easy to obtain—produces similar results to training an ST model on four times as

$N =$	0	2.5	5	10	20	50
base	0	2.1	1.8	2.1	10.8	22.7
+asr	0.5	5.7	9.1	14.5	20.2	28.2

Table 2: BLEU scores for Spanish-English ST on the Fisher test set, using N hours of training data. base: no transfer learning. +asr: using model parameters from English ASR (en-300h).

<i>Spanish</i>	super caliente pero muy bonito
<i>English</i>	super hot but very nice
<i>20h</i>	you support it <u>but</u> it was very nice
<i>20h+asr</i>	you can get alright <u>but</u> it's very nice
<i>50h</i>	<u>super</u> expensive <u>but</u> very nice
<i>50h+asr</i>	super hot but it's very nice
<i>Spanish</i>	sí y usted hace mucho tiempo que que vive aquí
<i>English</i>	yes and have you been living here a long time
<i>20h</i>	<u>yes</u> i've <u>been</u> a long time what did you come <u>here</u>
<i>20h+asr</i>	<u>yes</u> and you <u>have</u> a long time that you <u>live</u> <u>here</u>
<i>50h</i>	<u>yes</u> you are a long time that you <u>live</u> <u>here</u>
<i>50h+asr</i>	yes and have you been <u>here</u> <u>long</u>


Table 3: Example translations on selected sentences from the Fisher development set, with stem-level n -gram matches to the reference sentence underlined. 20h and 50h are Spanish-English models without pre-training; 20h+asr and 50h+asr are pre-trained on 300 hours of English ASR.

much data, which may be difficult to obtain.


We even find that in the very low-resource setting of just 2.5 hours of ST data, with transfer learning the model achieves a precision/recall of around 30% and improves by more than 10 points over the naive baseline. In very low-resource scenarios with time constraints—such as in disaster relief—it is possible that even this level of performance may be useful, since it can be used to spot keywords in speech and can be trained in just three hours.

Sample translations. Table 3 shows example translations for models *sp-en-20h* and *sp-en-50h* with and without transfer learning using *en-300h*.

Figure 3 shows the attention weights for the last sample utterance in Table 3. For this utterance, the Spanish and English text have a different word order: *mucho tiempo* occurs in the middle of the speech utterance, and its translation, *long time*, is at the end of the English reference. Similarly, *vive aquí* occurs at the end of the speech utterance, while the translation, *living here*, is in the middle of the English reference. The baseline *sp-en-50h* model translates the words correctly but doesn't get



(a) 50h:baseline



(b) 50h:asr

Figure 3: Attention plots for the final example in Table 3, using 50h models with and without pre-training. The x -axis shows the reference Spanish word positions in the input; the y -axis shows the predicted English sub-words. In the reference, *mucho tiempo* is translated to *long time*, and *vive aquí* to *living here*, but their order is reversed, and this is reflected in (b).

the English word order right. With transfer learning, the model produces a shorter but still accurate translation in the correct word order.

5.2 Analysis

To understand the source of these improvements, we carried out a set of ablation experiments. For most of these experiments, we focus on Spanish-English ST with 20 hours of training data, with and without transfer learning.

Transfer learning with selected parameters. In our first set of experiments, we transferred all parameters of the *en-300h* model, including the speech encoder CNN and LSTM; the text decoder embedding, LSTM and output layer parameters; and attention parameters. To see which set of parameters has the most impact, we train the *sp-en-20h* model by transferring only selected parameters from *en-300h*, and randomly initializing the rest.

The results (Figure 4) show that transferring all




Figure 4: Fisher development set training curves (reported using BLEU) for *sp-en-20h* using selected parameters from *en-300h*: none (base); encoder CNN only (+asr:cnn); encoder CNN and LSTM only (+asr:enc); decoder only (+asr:dec); and all: encoder, attention, and decoder (+asr:all). These scores do not use beam search and are therefore lower than the best scores reported in Figure 2.

parameters is most effective, and that the speech encoder parameters account for most of the gains. We hypothesize that the encoder learns transferable low-level acoustic features that normalize across variability like speaker and channel differences to better capture meaningful phonetic differences, and that much of this learning is language-independent. This hypothesis is supported by other work showing the benefits of cross-lingual and multilingual training for speech technology in low-resource target languages (Carlin et al., 2011; Jansen et al., 2010; Deng et al., 2013; Vu et al., 2012; Thomas et al., 2012; Cui et al., 2015; Alumäe et al., 2016; Yuan et al., 2016; Renshaw et al., 2015; Hermann and Goldwater, 2018).

By contrast, transferring only decoder parameters does not improve accuracy. Since decoder parameters help when used in tandem with encoder parameters, we suspect that the dependency in parameter training order might explain this: the transferred decoder parameters have been trained to expect particular input representations from the encoder, so transferring only the decoder parameters without the encoder might not be useful.

Figure 4 also suggests that models make strong gains early on in the training when using transfer learning. The *sp-en-20h* model initialized with all model parameters (+asr:all) from *en-300h* reaches a higher BLEU score after just 5 epochs (2 hours) of training than the model without transfer learning trained for 60 epochs/20 hours. This again can be useful in disaster-recovery scenarios, where the




Figure 5: Spanish-to-English BLEU scores on Fisher dev set, with 0h (no transfer learning), 100h and 300h of English ASR data used.

time to deploy a working system must be minimized.

Amount of ASR data required. Figure 5 shows the impact of increasing the amount of English ASR data used on Spanish-English ST performance for two models: *sp-en-20h* and *sp-en-50h*.

For *sp-en-20h*, we see that using *en-100h* improves performance by almost 6 BLEU points. By using more English ASR training data (*en-300h*) model, the BLEU score increases by almost 9 points. However, for *sp-en-50h*, we only see improvements when using *en-300h*. This implies that transfer learning is most useful when only a few tens of hours of training data are available for ST. As the amount of ST training data increases, the benefits of transfer learning tail off, although it’s possible that using even more monolingual data, or improving the training at the ASR step, could extend the benefits to larger ST data sets.

Impact of code-switching. We also tried using the *en-300h* ASR model without any fine-tuning to translate Spanish audio to English text. This model achieved a BLEU score of 1.1, with a precision of 15 and recall of 21. The non-zero BLEU score indicates that the model is matching *some* 4-grams in the reference. This seems to be due to code-switching in the Fisher-Spanish speech data set. Looking at the dev set utterances, we find several examples where the Spanish transcriptions match the English translations, indicating that the speaker switched into English. For example, there is an utterance whose Spanish transcription and English translation are both “right yeah”, and this English expression is indeed present in the source audio. The English ASR model correctly translates this utterance, which is unsurprising since the phrase “right yeah” occurs nearly 500 times in Switchboard.

Overall, we find that in nearly 500 of the 4,000 development set utterances (14%), the Spanish transcription and English translations share more than half of their tokens, indicating likely code-switching. This suggests that transfer learning from English ASR models might help more than from other languages. To isolate this effect from transfer learning of language-independent speech features, we carried out a further experiment.

5.3 Using French ASR to improve Spanish-English ST

In this experiment, we pre-train using French ASR data for a Spanish-English translation task. Here, we can only transfer the speech encoder parameters, and there should be little if any benefit due to code-switching.

Because our French data set (20 hours) is much smaller than our English one (300 hours), for a fair comparison we used a 20 hour subset of the English data for pre-training in this experiment. For both the English and French models, we transferred only the encoder parameters.

Table 4 shows that both the English and French 20-hour pre-trained models improve performance on Spanish-English ST. The English model works slightly better, as would be predicted given our discussion of code-switching, but the French model is also useful, improving BLEU from 10.8 to 12.5. This result strengthens the claim that ASR pre-training on a completely distinct third language can help low-resource ST. Presumably benefits would be much greater if we used a larger ASR data set, as we did with English above.

In this experiment, the French pre-trained model used a French BPE output vocabulary, distinct from the English BPE vocabulary used in the ST system. In the future it would be interesting to try combining the French and English text to create a combined output vocabulary, which would allow transferring both the encoder and decoder parameters, and may be useful for translating names or cognates. More generally, it would also be possible to pre-train on multiple languages simultaneously using a shared BPE vocabulary. There is evidence that speech features trained on multiple languages transfer better than those trained on the same amount of data from a single language (Hermann and Goldwater, 2018), so multilingual pre-training for ST could improve results.

	baseline	+fr-20h	+en-20h
sp-en-20h	10.8	12.5	13.2

Table 4: Fisher dev set BLEU scores for *sp-en-20h*. baseline: model without transfer learning. Last two columns: Using encoder parameters from French ASR (+fr-20h), and English ASR (+en-20h).

model	pretrain	BLEU	Pr.	Rec.
fr-top-8w	–	0	23.5	22.2
fr-top-10w	–	0	20.6	24.5
en-300h	–	0	0.2	5.7
fr-20h	–	0	4.1	3.2
	–	3.5	18.6	19.4
mb-fr-4h	fr-20h	5.9	23.6	20.9
	en-300h	5.3	23.5	22.6
	en + fr	7.1	26.7	23.1

Table 5: Mboshi-to-French translation scores, with and without ASR pre-training. Pr. is the precision, and Rec. the recall score. fr-top-8w and fr-top-10w are *naive baselines* that, respectively, predict the 8 or 10 most frequent training words. For en + fr, we use encoder parameters from *en-300h* and attention+decoder parameters from *fr-20h*

6 Mboshi-French ST

Our final set of experiments test our transfer method on ST for the low-resource language Mboshi, where we have only 4 hours of ST training data: Mboshi speech input paired with French text output.

Table 5 shows the ST model scores for Mboshi-French with and without using transfer learning. The first two rows *fr-top-8w*, *fr-top-10w*, show precision and recall scores for the *naive baselines* where we predict the top 8 or 10 most frequent French words in the Mboshi-French training set. These show that a precision/recall in the low 20s is easy to achieve, although with no n-gram matches (0 BLEU). The pre-trained ASR models by themselves (next two lines) are much worse.

The baseline model trained only on ST data actually has lower precision/recall than the naive baseline, although its non-zero BLEU score indicates that it is able to correctly predict some n-grams. We see comparable precision/recall to the naive baseline with improvements in BLEU by transferring either French ASR parameters (both encoder

and decoder, *fr-20h*) or English ASR parameters (encoder only, *en-300h*).

Finally, to achieve the benefits of both the larger training set size for the encoder and the matching language of the decoder, we tried transferring the encoding parameters from the *en-300h* model and the decoding parameters from the *fr-20h* model. This configuration (*en+fr*) gives us the best evaluation scores on all metrics, and highlights the flexibility of our framework. Nevertheless, the 4-hour scenario is clearly a very challenging one.

7 Conclusion

This paper introduced the idea of pre-training an end-to-end speech translation system involving a low-resource language using ASR training data from a higher-resource language. We showed that large gains are possible: for example, we achieved an improvement of 9 BLEU points for a Spanish-English ST model with 20 hours of parallel data and 300 hours of English ASR data. Moreover, the pre-trained model trains faster than the baseline, achieving higher BLEU in only a couple of hours, while the baseline trains for more than a day.

We also showed that these methods can be used effectively on a real low-resource language, Mboshi, with only 4 hours of parallel data. The very small size of the data set makes the task challenging, but by combining parameters from an English encoder and French decoder, we outperformed baseline models to obtain a BLEU score of 7.1 and precision/recall of about 25%. We believe ours is the first paper to report word-level BLEU scores on this data set.

Our analysis indicates that, other things being equal, transferring both encoder and decoder parameters works better than just transferring one or the other. However, transferring the encoder parameters is where most of the benefit comes from. Pre-training using a large ASR corpus from a mismatched language will therefore probably work better than using a smaller ASR corpus that matches the output language.

Our analysis suggests several avenues for further exploration. On the speech side, it might be even more effective to use multilingual training; or to replace the MFCC input features with pre-trained multilingual features, or features that are targeted to low-resource multispeaker settings (Kamper et al., 2015, 2017; Thomas et al., 2012; Cui et al., 2015; Yuan et al., 2016; Renshaw et al., 2015). On

the language modeling side, simply transferring decoder parameters from an ASR model did not work; it might work better to use pre-trained decoder parameters from a language model, as proposed by Ramachandran et al. (2017), or *shallow fusion* (Gülçehre et al., 2015; Toshniwal et al., 2018a), which interpolates a pre-trained language model during beam search. In these methods, the decoder parameters are independent, and can therefore be used on their own. We plan to explore these strategies in future work.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable feedback. This work was supported in part by a James S McDonnell Foundation Scholar Award, a Google faculty research award, and NSF grant 1816627. We thank Ida Szubert and Clara Vania for helpful comments on previous drafts of this paper and Antonios Anastasopoulos for tips on experimental setup.

References

- Oliver Adams, Graham Neubig, Trevor Cohn, and Steven Bird. 2016a. Learning a translation model from word lattices. In *Proc. Interspeech*.
- Oliver Adams, Graham Neubig, Trevor Cohn, Steven Bird, Quoc Truong Do, and Satoshi Nakamura. 2016b. Learning a lexicon and translation model from phoneme lattices. In *Proc. EMNLP*.
- Tanel Alumäe, Stavros Tsakalidis, and Richard M Schwartz. 2016. Improved multilingual training of stacked neural network acoustic models for low resource languages. In *Proc. Interspeech*.
- Antonios Anastasopoulos and David Chiang. 2017. A case study on using speech-to-translation alignments for language documentation. In *Proc. ACL*.
- Antonios Anastasopoulos and David Chiang. 2018. Tied multitask learning for neural speech translation. In *Proc. NAACL HLT*.
- Sameer Bansal, Herman Kamper, Karen Livescu, Adam Lopez, and Sharon Goldwater. 2018. Low-resource speech-to-text translation. In *Proc. Interspeech*.
- Alexandre Bérard, Laurent Besacier, Ali Can Kobayiyikoglu, and Olivier Pietquin. 2018. End-to-end automatic speech translation of audiobooks. In *Proc. ICASSP*.
- Laurent Besacier, Bowen Zhou, and Yuqing Gao. 2006. Towards speech translation of non written languages. In *Proc. SLT*.
- Michael A Carlin, Samuel Thomas, Aren Jansen, and Hynek Hermansky. 2011. Rapid evaluation of speech representations for spoken term discovery. In *Proc. Interspeech*.
- Jia Cui, Brian Kingsbury, Bhuvana Ramabhadran, Abhinav Sethy, Kartik Audhkhasi, Xiaodong Cui, Ellen Kislal, Lidia Mangu, Markus Nussbaum-Thom, Michael Picheny, et al. 2015. Multilingual representations for low resource speech recognition and keyword search. In *Proc. ASRU*.
- Li Deng, Jinyu Li, Jui-Ting Huang, Kaisheng Yao, Dong Yu, Frank Seide, Mike Seltzer, Geoff Zweig, Xiaodong He, Jason Williams, Yifan Gong, and Alex Acero. 2013. Recent advances in deep learning for speech research at Microsoft. In *Proc. ICASSP*.
- Yarin Gal. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Proc. NIPS*.
- Pierre Godard, Gilles Adda, Martine Adda-Decker, Juan Benjumea, Laurent Besacier, Jamison Cooper-Leavitt, Guy-Noël Kouarata, Lori Lamel, Hélène Maynard, Markus Müller, Annie Rialland, Sebastian Stüker, François Yvon, and Marcely Zanon Boito. 2018. A very low resource language speech corpus for computational language documentation experiments. In *Proc. LREC*.
- John Godfrey and Edward Holliman. 1993. Switchboard-1 Release 2 (LDC97S62). <https://catalog.ldc.upenn.edu/ldc97s62>.
- David Graff, Shudong Huang, Ingrid Cartagena, Kevin Walker, and Christopher Cieri. 2010. Fisher Spanish Speech (LDC2010S01). <https://catalog.ldc.upenn.edu/ldc2010s01>.
- Caglar Gülçehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loïc Barrault, Huei-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2015. On using monolingual corpora in neural machine translation. *CoRR*, abs/1503.03535.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *Proc. ICCV*.
- Enno Hermann and Sharon Goldwater. 2018. Multilingual bottleneck features for subword modeling in zero-resource languages. In *Proc. Interspeech*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. ICML*.
- Aren Jansen, Kenneth Church, and Hynek Hermansky. 2010. Towards spoken term discovery at scale with zero resources. In *Proc. Interspeech*.

- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Gregory S. Corrado, Macduff Hughes, and Jeffrey Dean. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Trans. ACL*, 5:339–351.
- Herman Kamper, Micha Elsner, Aren Jansen, and Sharon Goldwater. 2015. Unsupervised neural network based feature extraction using weak top-down constraints. In *Proc. ICASSP*.
- Herman Kamper, Aren Jansen, and Sharon Goldwater. 2017. A segmental framework for fully-unsupervised large-vocabulary speech recognition. *Comput. Speech Lang.*, 46:154–174.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc. ICLR*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. ACL*.
- Alon Lavie and Abhaya Agarwal. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proc. WMT*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proc. EMNLP*.
- Lara J Martin, Andrew Wilkinson, Sai Sumanth Miryala, Vivian Robison, and Alan W Black. 2015. Utterance classification in speech-to-speech translation for zero-resource languages in the hospital administration domain. In *Proc. ASRU*.
- Robert Munro. 2010. Crowdsourced translation for emergency response in Haiti: The global collaboration of local knowledge. In *AMTA Workshop Collaborative Crowdsourcing Transl.*
- Vinod Nair and Geoffrey E. Hinton. 2010. Rectified linear units improve restricted Boltzmann machines. In *Proc. ICML*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. ACL*.
- Matt Post, Gaurav Kumar, Adam Lopez, Damianos Karakos, Chris Callison-Burch, and Sanjeev Khudanpur. 2013. Improved speech-to-text translation with the Fisher and Callhome Spanish-English speech translation corpus. In *Proc. IWSLT*.
- Matt Post, Gaurav Kumar, Adam Lopez, Damianos Karakos, Chris Callison-Burch, and Sanjeev Khudanpur. 2014. Fisher and CALLHOME Spanish-English Speech Translation. <https://catalog.ldc.upenn.edu/ldc2014t23>.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely. 2011. The Kaldi Speech Recognition Toolkit. In *Proc. ASRU*.
- Prajit Ramachandran, Peter J Liu, and Quoc V Le. 2017. Unsupervised pretraining for sequence to sequence learning. In *Proc. EMNLP*.
- Daniel Renshaw, Herman Kamper, Aren Jansen, and Sharon Goldwater. 2015. A comparison of neural network methods for unsupervised representation learning on the zero resource speech challenge. In *Proc. Interspeech*.
- Tanja Schultz. 2002. Globalphone: a multilingual speech and text database developed at karlsruhe university. In *Seventh International Conference on Spoken Language Processing*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proc. ACL*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*
- Samuel Thomas, Sriram Ganapathy, and Hynek Hermansky. 2012. Multilingual mlp features for low-resource LVCSR systems. In *Proc. ICASSP*.
- Sebastian Thrun. 1995. Is learning the n-th thing any easier than learning the first? In *Proc. NIPS*.
- Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. 2015. Chainer: A next-generation open source framework for deep learning. In *Proc. LearningSys*.
- Shubham Toshniwal, Anjuli Kannan, Chung-Cheng Chiu, Yonghui Wu, Tara N Sainath, and Karen Livescu. 2018a. A Comparison of Techniques for Language Model Integration in Encoder-Decoder Speech Recognition. In *Proc. SLT*.
- Shubham Toshniwal, Tara N. Sainath, Ron J. Weiss, Bo Li, Pedro Moreno, Eugene Weinstein, and Kanishka Rao. 2018b. Multilingual Speech Recognition with A Single End-To-End Model. In *Proc. ICASSP*.
- Ngoc Thang Vu, Wojtek Breiter, Florian Metze, and Tanja Schultz. 2012. An investigation on initialization schemes for multilayer perceptron training using multilingual data and their effect on ASR performance. In *Proc. Interspeech*.
- Ron J Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. Sequence-to-sequence models can directly transcribe foreign speech. In *Proc. Interspeech*.

Ronald J. Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural Comput.*

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Gregory S. Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.

Yougen Yuan, Cheung-Chi Leung, Lei Xie, Bin Ma, and Haizhou Li. 2016. Learning neural network representations using cross-lingual bottleneck features with word-pair information. In *Proc. Interspeech*.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. In *Proc. EMNLP*.