



Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation

Alexandre Bérard, Olivier Pietquin, Laurent Besacier, Christophe Servan

► To cite this version:

Alexandre Bérard, Olivier Pietquin, Laurent Besacier, Christophe Servan. Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation. NIPS Workshop on end-to-end learning for speech and audio processing, Dec 2016, Barcelona, Spain. hal-01408086

HAL Id: hal-01408086

<https://hal.science/hal-01408086>

Submitted on 6 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation

Alexandre Bérard * Olivier Pietquin †

Univ. Lille, CNRS, Centrale Lille, Inria, UMR 9189 CRISTAL - Lille, France

alexandre.berard@ed.univ-lille1.fr

olivier.pietquin@univ-lille1.fr

Laurent Besacier ‡ Christophe Servan ‡

LIG, Univ. Grenoble Alpes, CNRS, Grenoble, France

firstname.lastname@imag.fr

1 Introduction and background

Current speech translation systems integrate (loosely or closely) two main modules: source language speech recognition (ASR) and source-to-target text translation (MT). In these approaches, source language text transcript (as a sequence or as a graph) appears as mandatory to produce a text hypothesis in the target language. In the meantime, deep neural networks have yielded breakthroughs in different areas including machine translation and speech recognition. Current systems have an “encoder-decoder” architecture where a sequence of input symbols (or an input signal) is projected into a continuous low dimensional space and the output sequence is generated from this representation. This type of architecture has been proposed for machine translation (Sutskever et al., 2014; Bahdanau et al., 2015) and for automatic speech recognition (Chorowski et al., 2015; Chan et al., 2016; Zweig et al., 2016). If we exclude some pioneering attempts to directly translate from a source phone sequence to a target word sequence (Besacier et al., 2006) or to discover a bilingual lexicon from parallel phone-word sequences (Godard et al., 2016), the only attempt to translate directly a source speech signal into target language text is that of Duong et al. (2016). However, the authors focus on the alignment between source speech utterances and their text translation without proposing a complete end-to-end translation system. The same authors (Anastasopoulos et al., 2016) also propose to jointly use IBM translation models (IBM Model 2), and dynamic time warping (DTW) to align source speech and target text, but again, only the alignment performance is measured in their work.

Consequently, this paper proposes a first attempt to build an end-to-end speech-to-text translation system, which does not use source language text during learning or decoding. Relaxing the need for source language transcription would drastically change the data collection methodology in speech translation, especially in under-resourced scenarios. For instance, in the former project DARPA TRANSTAC (speech translation from spoken Arabic dialects), a large effort was devoted to the collection of speech transcripts (and a prerequisite to obtain transcripts was often a detailed transcription guide for languages with little standardized spelling). Now, if end-to-end approaches for speech-to-text translation are successful, one might consider collecting data by asking bilingual speakers to directly utter speech in the source language from target language text utterances. Such an approach has the advantage to be applicable to any unwritten (source) language.

*also at LIG, Univ. Grenoble Alpes

†now with Google DeepMind - London, UK

‡now with Systran, Paris, France

2 Model

Description We compare two end-to-end models: one for standard machine translation, and one for speech translation. Both models are attention-based encoder-decoder neural networks (Sutskever et al., 2014; Bahdanau et al., 2015).

Given an input sequence (x_1, \dots, x_A) , the encoder (a multi-layered bidirectional LSTM) produces a sequence of outputs $h = (h_1, \dots, h_A)$. We initialize the decoder state with $s_0 = \tanh(W_{init}s'_A)$ where s'_A is the last state of the encoder. The next states of the decoder are computed as follows:

$$s_t, o_t = LSTM(s_{t-1}, Ez_{t-1}) \quad (1)$$

s_t is the new state of the LSTM, while o_t is its output. During training, z_{t-1} is the previous symbol in the target (reference) sequence. During evaluation, z_{t-1} is the predicted symbol at the previous time-step. E is the target embedding matrix.

We compute an attention vector d_t (over the encoder output h), which is concatenated to the LSTM’s output o_t :

$$y_t = W_{proj}(o_t \oplus d_t) + b_{proj} \quad d_t = attention(h, s_t) \quad (2)$$

This vector is then mapped to a vector of the same size as the target vocabulary, to compute a probability distribution over all target words. At each time-step t , the greedy decoder picks the word w_t with the highest probability according to this distribution.

$$p(w_t | y_{t-1}, s_{t-1}) = softmax(W_{out}y_t + b_{out}) \quad (3)$$

Attention mechanism For text input, we use the same attention model as Vinyals et al. (2015):

$$attention(h, s_t) = \sum_{i=1}^A a_i^t h_i \quad (4)$$

$$a_i^t = softmax(u_i^t) \quad u_i^t = v^T \tanh(W_1 h_i + W_2 s_t + b_2) \quad (5)$$

Convolutional attention For speech input, we use the same attention model as Chorowski et al. (2015), which uses a convolution filter F to take into account the attention weights at the previous time-step. This allows the attention model to “remember” and not translate the same part of the signal twice. This is especially useful for speech as the input sequence can be quite long.

$$u_i^t = v^T \tanh(W_1 h_i + W_2 s_t + b_2 + f_{t,i} \mu) \quad f_t = F * a_{t-1} \quad (6)$$

$W_{init} \in \mathbb{R}^{2m \times 2m}$, $E \in \mathbb{R}^{n \times \bar{V}}$, $W_{proj} \in \mathbb{R}^{m \times 2m}$, $W_{out} \in \mathbb{R}^{\bar{V} \times m}$, $W_1 \in \mathbb{R}^{m \times m}$, $W_2 \in \mathbb{R}^{m \times 2m}$, $\mu \in \mathbb{R}^m$, $F \in \mathbb{R}^k$, where n is the embedding size, m the size of the LSTM layers (the state size is $2m$), k the size of the convolution filter, and \bar{V} the target vocabulary size.

Implementation Our implementation is based on the *seq2seq* model implemented by *TensorFlow* (Abadi et al., 2015). It reuses some of its components, while adding a number of features, like a bidirectional encoder (Bahdanau et al., 2015), a beam-search decoder, a convolutional attention model and a hierarchical encoder (Chorowski et al., 2015).⁴ Our implementation also allows pre-training, multiple encoders (Zoph and Knight, 2016) and multi-task training (Luong et al., 2016). Contrary to the reference implementation by *TensorFlow*, which uses a bucket-based method with static (unrolled) RNNs, we use symbolic loops. This makes model creation much faster, and removes the need to manually configure bucket sizes.

Model settings We describe two models: one for text translation, and one for speech translation. Both models use a decoder with 2 layers of 256 LSTM units, with word embeddings of size 256. The *speech* model does not have source embeddings, as the input is already a sequence of vectors (of 40 MFCCs + frame energy). Both encoders are multi-layer bidirectional LSTMs, with 256 units. The *text* model’s encoder has 2 layers. The *speech* model has a hierarchical encoder with 3 layers, i.e., like Chorowski et al. (2015), the 2nd and 3rd layers read every other output from the previous layer, resulting in a sequence of 1/4th the size of the input. For the *text* model, we use a standard attention model, while the *speech* model uses convolutional attention (with a single filter of size 25). The latter also has two fully connected layers of size 256 between the input and the first layer of the encoder.

For training, we use the *Adam* algorithm with an initial learning rate of 0.001 (Kingma and Ba, 2014), and a mini-batch size of 64. We apply dropout (with a rate of 0.5) during training on the connections between LSTM layers in the encoder and decoder (Zaremba et al., 2014).

⁴Our source code is available online at github.com/eske/seq2seq

Corpus	Lines	French				English	
		Words	Chars	Frames	Duration (h)	Words	Chars
BTEC train	19972	201k	1020k	5715k	15:51	189k	842k
BTEC dev	1512	12.2k	61.9k	357k	0:59	11.5k	50.3k
BTEC test1	469	3.9k	19.5k	111k	0:18	3.6k	15.6k
BTEC test2	464	3.8k	19.2k	110k	0:18	3.6k	15.8k

Table 1: On the full corpus, there is on average 9.8 words/sent on the French side, and 9.3 on the English side. On the French audio side (with *Agnes* as a speaker), there is on average 281 frames/sent (2.8 sec/sent).

Corpus	BLEU score				
	Greedy	+BeamSearch	+LM	+Ensemble	SMT
dev	42.5	43.6	45.1	51.6	54.3
test1	39.8	41.7	43.0	47.1	49.4
test2	41.0	41.0	42.8	47.9	45.8
test1 (7 refs)	50.7	53.4	54.7	60.3	60.0
test2 (7 refs)	47.0	47.3	49.6	55.3	53.2
train.1000	78.4	80.5	82.9	89.6	76.6

Table 2: Results of the Text Translation task on BTEC, under different settings of the decoder. “Greedy” uses a greedy decoder. “BeamSearch” adds a beam-search decoder with a beam size of 8. “LM” adds an external trigram language model estimated on the target side of *BTEC train*. “Ensemble” uses a log-linear model with 5 NMT models trained independently (+ language model). The *train* corpus used in evaluation is a subset of 1000 sentences from *BTEC train*. For all the non-ensemble configurations, the best model out of 5 is used (according to its score on *dev*). The SMT baseline is a phrase-based model (Moses) trained on *BTEC train*, and tuned on *BTEC dev*.

3 Experimental Settings

As we do not have any good-sized speech translation corpus, we generated one with speech synthesis. We decided to use the French→English BTEC corpus, which is a small parallel corpus (see Table 1). BTEC (Basic Travel Expression Corpus) is very clean and contains short sentences (10 words/sentence on average), with a small vocabulary (9218 unique tokens on the French side, 7186 on the English side). This corpus also has multiple references (7 references per source sentence for *test1* and *test2*), which we use for BLEU evaluation. We lowercase and tokenize all the data with the *Moses* tokenizer.

Synthetic speech was generated using *Voxxygen*⁵, a commercial speech synthesis system, for 4 different female voices (Agnes, Marion, Helene, Fabienne) and 3 different male voices (Michel, Loic, Philippe). It is important to note that this is corpus-based concatenative speech synthesis (Schwarz, 2007) and not parametric synthesis. So, for each speaker’s voice, speech utterances are generated by concatenation of units mined from a large speech corpus (generally around 3000 sentences/speaker). This means that despite having very little intra-speaker variability in our speech data, there is a realistic level of inter-speaker variability. To challenge the robustness of our system to inter-speaker variability, we concatenate the training data for 6 speakers, and leave the last speaker (*Agnes*) for evaluation only.

Like Chan et al. (2016), as input for our *speech* model, we segment the speech raw data into frames of 40 ms, with a step-size of 10 ms, and extract 40 MFCCs for each frame, along with the frame energy.

4 Experiments

Table 2 shows the results of our Machine Translation experiments. The main reason why we perform these experiments, is to assess whether it makes sense to train NMT models on such a small and specialized corpus. Our best system (an ensemble of 5 models with a language model) gives similar results in terms of BLEU score to a baseline SMT system. This is rather surprising, considering the small amount of data used. Also, even though the models seem to overfit quite a lot (see *train.1000* scores), this does not seem to be a problem during evaluation.

⁵www.voxxygen.fr

Corpus	Speaker	BLEU score				
		Greedy	+BeamSearch	+LM	+Ensemble	Baseline
dev	Agnes	30.1	32.3	33.5	40.2	43.7
test1		29.2	31.9	32.7	38.6	42.1
test2		29.0	30.6	31.0	37.1	39.6
test1 (7 refs)	Agnes	37.6	40.1	41.8	48.8	52.7
	Michel	40.6	43.1	44.4	50.4	51.9
test2 (7 refs)	Agnes	33.6	35.6	36.7	43.1	46.7
	Michel	37.3	39.6	41.0	46.7	46.4
train.1000	Michel	53.8	60.8	61.8	82.6	60.5

Table 3: Results of the speech translation experiments on BTEC. The models are trained with 6 different speakers (including *Michel*). *Agnes* is not used for training. The *Ensemble* configuration uses 5 models trained independently. The baseline system uses a pipeline Google Speech ASR + SMT system trained on BTEC (we trained a new system on French without punctuation, as Google Speech API does not output any punctuation symbol, except when actually pronounced). The WER scores obtained by the baseline ASR system range between 23% and 26%. The last column from Table 2 provides a baseline assuming a perfect ASR.

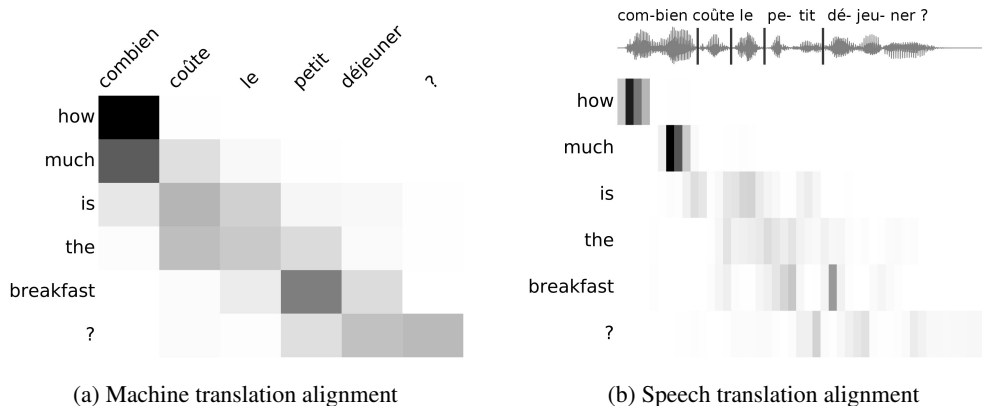


Figure 1: Alignments performed by the attention model during training

Table 3 shows the results of our Speech Translation experiments. Our best results are behind, but not very far from those of the baseline (which uses Google Speech API and a specialized SMT system). The results on a new speaker (*Agnes*) are also encouraging, considering that we did not use any speaker adaptation technique.

Figure 1 shows an example of alignment performed by the attention mechanism during training, for text and for speech. We trained our models for 20k steps, which takes less than 2 hours for the *text* models, and roughly 8 hours for the *speech* models (on a single GTX 1070).

5 Conclusion

We proposed a model for end-to-end speech translation, which gives promising results on a small French-English synthetic corpus. These results, however, are to be mitigated due to the synthetic nature of the data: there is not much variability (except across speakers), which makes it easier for the network to learn. In any case, the experiments show that our architecture is able to generalize quite well to a new speaker.

Future work includes the evaluation of this model on non-synthetic data. Some datasets that could be used are: TED talks (there exists a parallel corpus and a speech corpus, but their alignment is not straightforward); or books from the Gutenberg project⁶ (public-domain books in many languages, along with some human-read audiobooks). We also wish to evaluate this model in a multi-source setting: does it help when translating text to have access to a human-read version of it, or other kinds of information such as part-of-speech or images?

⁶www.gutenberg.org

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, Ł., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org.
- Anastasopoulos, A., Chiang, D., and Duong, L. (2016). An Unsupervised Probability Model for Speech-to-Translation Alignment of Low-Resource Languages. *arXiv preprint arXiv:1609.08139*.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR 2015*, pages 3104–3112, San Diego, California, USA.
- Besacier, L., Zhou, B., and Gao, Y. (2006). Towards speech Translation of Non Written Languages. In *Spoken Language Technology Workshop (SLT 2006)*, pages 222–225, Palm Beach, Aruba. IEEE.
- Chan, W., Jaitly, N., Le, Q. V., and Vinyals, O. (2016). Listen, Attend and Spell: A Neural Network for Large Vocabulary Conversational Speech Recognition. In *ICASSP 2016*, pages 4960–4964, Shanghai, China. IEEE.
- Chorowski, J. K., Bahdanau, D., Serdyuk, D., Cho, K., and Bengio, Y. (2015). Attention-Based Models for Speech Recognition. In *Advances in Neural Information Processing Systems (NIPS 2015)*, pages 577–585, Montréal, Canada.
- Duong, L., Anastasopoulos, A., Chiang, D., Bird, S., and Cohn, T. (2016). An Attentional Model for Speech Translation Without Transcription. In *NAACL-HLT 2016*, Denver, Colorado, USA.
- Godard, P., Adda, G., Adda-Decker, M., Allauzen, A., Besacier, L., Bonneau-Maynard, H., Kouarata, G.-N., Löser, K., Rialland, A., and Yvon, F. (2016). Preliminary Experiments on Unsupervised Word Discovery in Mboshi. In *Interspeech 2016*, San Francisco, California, USA.
- Kingma, D. and Ba, J. (2014). Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- Luong, M.-T., Le, Q. V., Sutskever, I., Vinyals, O., and Kaiser, Ł. (2016). Multi-task Sequence to Sequence Learning. In *ICLR 2016*, San Juan, Puerto Rico.
- Schwarz, D. (2007). Corpus-Based Concatenative Synthesis: Assembling sounds by content-based selection of units from large sound databases. *IEEE Signal Processing Magazine*, pages 24–2.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to Sequence Learning with Neural Networks. In *Advances in Neural Information Processing Systems (NIPS 2014)*, pages 3104–3112, Montréal, Canada.
- Vinyals, O., Kaiser, Ł., Koo, T., Petrov, S., Sutskever, I., and Hinton, G. (2015). Grammar as a Foreign Language. In *Advances in Neural Information Processing Systems (NIPS 2015)*, pages 2773–2781, Montréal, Canada.
- Zaremba, W., Sutskever, I., and Vinyals, O. (2014). Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*.
- Zoph, B. and Knight, K. (2016). Multi-Source Neural Translation. In *NAACL-HLT 2016*, Denver, Colorado, USA.
- Zweig, G., Yu, C., Droppo, J., and Stolcke, A. (2016). Advances in All-Neural Speech Recognition. *arXiv preprint arXiv:1609.05935*.