



HAL
open science

ON-TRAC Consortium for End-to-End and Simultaneous Speech Translation Challenge Tasks at IWSLT 2020

Maha Elbayad, Ha Nguyen, Fethi Bougares, Natalia Tomashenko, Antoine Caubrière, Benjamin Lecouteux, Yannick Estève, Laurent Besacier

► **To cite this version:**

Maha Elbayad, Ha Nguyen, Fethi Bougares, Natalia Tomashenko, Antoine Caubrière, et al.. ON-TRAC Consortium for End-to-End and Simultaneous Speech Translation Challenge Tasks at IWSLT 2020. Proceedings of the 17th International Conference on Spoken Language Translation, Jul 2020, Seattle, WA, United States. pp.35-43, 10.18653/v1/2020.iwslt-1.2 . hal-02895893

HAL Id: hal-02895893

<https://hal.science/hal-02895893v1>

Submitted on 10 Jul 2020

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.

ON-TRAC Consortium for End-to-End and Simultaneous Speech Translation Challenge Tasks at IWSLT 2020

Maha Elbayad^{1,4*}, Ha Nguyen^{1,2*}, Fethi Bougares³, Natalia Tomashenko²,
Antoine Caubrière³, Benjamin Lecouteux¹, Yannick Estève², Laurent Besacier¹

¹LIG - Université Grenoble Alpes, France

²LIA - Avignon Université, France

³LIUM - Le Mans Université, France

⁴Inria - Grenoble, France

Abstract

This paper describes the ON-TRAC Consortium translation systems developed for two challenge tracks featured in the Evaluation Campaign of IWSLT 2020, offline speech translation and simultaneous speech translation. ON-TRAC Consortium is composed of researchers from three French academic laboratories: LIA (Avignon Université), LIG (Université Grenoble Alpes), and LIUM (Le Mans Université). Attention-based encoder-decoder models, trained end-to-end, were used for our submissions to the offline speech translation track. Our contributions focused on data augmentation and ensembling of multiple models. In the simultaneous speech translation track, we build on Transformer-based *wait-k* models for the text-to-text subtask. For speech-to-text simultaneous translation, we attach a *wait-k* MT system to a hybrid ASR system. We propose an algorithm to control the latency of the ASR+MT cascade and achieve a good latency-quality trade-off on both subtasks.

1 Introduction

While cascaded speech-to-text translation (AST) systems (combining source language speech recognition (ASR) and source-to-target text translation (MT)) remain state-of-the-art, recent works have attempted to build end-to-end AST with very encouraging results (Bérard et al., 2016; Weiss et al., 2017; Bérard et al., 2018; Jia et al., 2019; Sperber et al., 2019). This year, IWSLT 2020 *offline translation track* attempts to evaluate if end-to-end AST will close the gap with cascaded AST for the English-to-German language pair.

Another increasingly popular topic is simultaneous (online) machine translation which consists in generating an output hypothesis before the entire

input sequence is available. To deal with this low latency constraint, several strategies were proposed for neural machine translation with input text (Ma et al., 2019; Arivazhagan et al., 2019; Ma et al., 2020). Only a few works investigated low latency neural speech translation (Niehues et al., 2018). This year, IWSLT 2020 *simultaneous translation track* attempts to stimulate research on this challenging task. This paper describes the ON-TRAC consortium automatic speech translation (AST) systems for the IWSLT 2020 Shared Task (Ansari et al., 2020). ON-TRAC Consortium is composed of researchers from three French academic laboratories: LIA (Avignon Université), LIG (Université Grenoble Alpes), and LIUM (Le Mans Université).

We participated in:

- IWSLT 2020 *offline translation track* with end-to-end models for the English-German language pair,
- IWSLT 2020 *simultaneous translation track* with a cascade of an ASR system trained using Kaldi (Povey et al., 2011) and an online MT system with *wait-k* policies (Dalvi et al., 2018; Ma et al., 2019).

This paper goes as follows: we review the systems built for the offline speech translation track in §2. Then, we present our approaches to the simultaneous track for both text-to-text and speech-to-text subtasks in §3. We ultimately conclude this work in §4.

2 Offline Speech translation Track

In this work, we developed several end-to-end speech translation systems, using a similar architecture as last year (Nguyen et al., 2019) and adapting it for translating English speech into German text (En-De). All the systems were developed using the

*Equal contribution.

Name	#segments	Total length (in hours)
MuST-C train	229.703	400
MuST-C dev	1.423	2.5
MuST-C tst-COMMON	2.641	4.1
MuST-C tst-HE	600	1.2
Europarl train	32.628	77
Europarl dev	1.320	3.1
How2 synthetic	176.564	285.5
tst2019	2.813	5.1
tst2020	2.263	4.1

Table 1: Statistics of training and evaluation data. The statistics of tst2019 and tst2020 are measured on the segmented version provided by IWSLT2020 organizers.

ESPnet (Watanabe et al., 2018) end-to-end speech processing toolkit.

2.1 Data and pre-processing

Data. We relied on MuST-C (Di Gangi et al., 2019) English-to-German (hereafter called MuST-C original), and Europarl (Iranzo-Sánchez et al., 2020) English-to-German as our main corpora. Besides, we automatically translated (into German) the English transcription of MuST-C and How2 (Sanabria et al., 2018) in order to augment training data. This resulted in two synthetic corpora, which are called *MuST-C synthetic* and *How2 synthetic* respectively. The statistics of these corpora, along with the provided evaluation data, can be found in Table 1. We experimented with different ways of combining those corpora. The details of these experiments are presented later in this section.

Speech features and data augmentation. 80-dimensional Mel filter-bank features, concatenated with 3-dimensional pitch features¹ are extracted from windows of 25ms with a frame shift of 10ms. We computed mean and variance normalization on these raw features of the training set, then applied it on all the data. Beside speed perturbation with factors of 0.9, 1.0, and 1.1, SpecAugment (Park et al., 2019) is applied to the training data (Ko et al.,

¹Pitch-features are computed using the Kaldi toolkit (Povey et al., 2011) and consist of the following values (Ghahremani et al., 2014): (1) probability of voicing (POV-feature), (2) pitch-feature and (3) delta-pitch feature. For details, see http://kaldi-asr.org/doc/process-kaldi-pitch-feats_8cc.html




Figure 1: Architecture of the speech encoder: a stack of two VGG blocks followed by 5 BLSTM layers.

2015). All three SpecAugment methods were used, including time warping ($W = 5$), frequency masking ($F = 30$), and time masking ($T = 40$).

Text preprocessing. The same as last year, we normalize punctuation, and tokenize all the German text using Moses.² Texts are case-sensitive and contain punctuation. Moreover, the texts of the MuST-C corpus contain multiple non speech events (*i.e.* 'Laughter', 'Applause' etc.). All these marks are removed from the texts before training our models. This results in a vocabulary of 201 characters. We find that some of these characters should not appear in the German text, for example, ♪, 你, 葱, 送, etc. Therefore, we manually exclude them from the vocabulary. In the end, we settle with an output vocabulary of 182 characters.

2.2 Architecture

We reuse our last year attention-based encoder-decoder architecture. As illustrated in Figure 1, the encoder has two VGG-like (Simonyan and Zisserman, 2015) CNN blocks followed by five stacked 1024-dimensional BLSTM layers. Each VGG block is a stack of two 2D-convolution layers followed by a 2D-maxpooling layer aiming to reduce both time (T) and frequency (D) dimensions of the input speech features by a factor of 2. After these two VGG blocks, input speech features' shape is transformed from $(T \times D)$ to $(T/4 \times D/4)$. We used Bahdanau's attention mechanism (Bahdanau et al., 2015) in all our experiments. The decoder is a stack of two LSTM layers 1024 dimensional memory cells. We would like to men-

²<http://www.statmt.org/moses/>

No.	Experiment	MuST-C tst-COMMON	MuST-C tst-HE	tst2015 (iwslt seg)	tst2015 (ASR seg)
1	MuST-C original + EuroParl	20.18	19.82	12.59	14.85
2	MuST-C original + Europarl + How2 synthetic	20.51	20.10	12.10	13.66
3*	MuST-C original + Europarl + How2 synthetic	23.55	22.35	13.00	15.30
4*	MuST-C original + Europarl + How2 synthetic + MuST-C synthetic	22.75	21.31	14.00	16.45
5*	Finetune 3* on MuST-C original	23.60	22.26	13.71	15.30
6*	Finetune 3* on MuST-C original+ MuST-C synthetic	23.64	22.23	13.67	15.29
7	Ensemble (1 to 6)	25.22	23.80	15.20	16.53

Table 2: Detokenized case-sensitive BLEU scores for different experiments - * represents experiments that apply SpecAugment.

Model	iwslt seg	ASR seg
3*	constrastive5	constrastive3
4*	constrastive4	constrastive2
Ensemble	constrastive1	primary

Table 3: The ranking of out submitted systems. Model 3* and 4* are respectively corresponding to No.3* and No.4* of Table 2.

tion that Transformer based models have also been tested using the default ESPnet architecture and showed weaker results compared to the LSTM-based encoder-decoder architecture.

Hyperparameters’ details. All of our models are trained in maximum 20 epochs, with early stopping after 3 epochs if the accuracy on the development set does not improve. Dropout is set to 0.3 on the encoder part, and Adadelata is chosen as our optimizer. During decoding time, the beam size is set to 10. We prevent the models from generating too long sentences by setting a $maxlenratio^3 = 1.0$. All our end-to-end models are similar in terms of architecture. They are different mainly in the following aspects: (1) training corpus; (2) type of

tokenization units;⁴ (3) fine-tuning and pretraining strategies. Description of different models and evaluation results are given in Section 2.4.

2.3 Speech segmentation

Two types of segmentation of evaluation and development data were used for experiments and submitted systems: segmentation provided by the IWSLT organizers and automatic segmentation based on the output of an ASR system.

The ASR system, used to obtain automatic segmentation, was trained with the Kaldi speech recognition toolkit (Povey et al., 2011). An acoustic model was trained using the TED-LIUM 3 corpus (Hernandez et al., 2018).⁵ This ASR system produces recognized words with timecodes (start time and duration for each word). Then we form the speech segments based on this output following the rules: (1) if silence duration between two words is longer than a given threshold $\Theta = 0.65$ seconds, we split the audio file; (2) if the number of words in the current speech segment exceeds 40, then Θ is reduced to 0.15 seconds in order to avoid too long segments. These thresholds have been optimised to get segment duration distribu-

⁴All systems use 182 output character tokens except system 1 which has 201

⁵The off-limit TED talks from IWSLT-2019 were excluded from the training subset

³ $maxlenratio = \frac{maximum_output_length}{encoder_hidden_state_length}$

No.	Set	BLEU	TER	BEER	CharacTER	BLEU(ci)	TER(ci)
1	2019.contrastive1	17.57	71.68	47.24	58.03	18.64	69.66
2	2019.contrastive2	17.83	71.60	48.66	53.49	18.9	69.26
3	2019.contrastive3	19.03	66.96	49.12	54.10	19.97	65.01
4	2019.contrastive4	15.08	78.79	45.87	59.06	16.06	76.62
5	2019.contrastive5	15.87	74.17	46.18	59.96	16.86	72.15
6	2019.primary	20.19	66.38	49.89	52.51	21.23	64.26
7	2020.contrastive1	18.47	71.85	48.92	55.83	19.46	69.88
8	2020.contrastive2	19.31	69.30	49.55	52.68	20.36	67.14
9	2020.contrastive3	20.51	64.88	50.19	53.06	21.5	62.99
10	2020.contrastive4	15.48	83.45	46.68	57.56	16.42	81.33
11	2020.contrastive5	16.5	75.15	47.23	57.90	17.42	73.22
12	2020.primary	22.12	63.87	51.20	51.46	23.25	61.85

Table 4: IWSLT 2020 official results (offline track) on tst2019 and tst2020.

tion in the development and evaluation data that is similar to the one observed in the training data. It will be shown in next subsection that this ASR segmentation improves results over the provided segmentation when the latter is noisy (see experimental results on iwslt/tst2015).

2.4 Experiments and results

After witnessing the benefit of merging different corpora from our submission last year (Nguyen et al., 2019), we continue exploring different combinations of corpora in this submission. As shown in the first two rows of Table 2, merging How2 synthetic with the baseline (MuST-C original + Europarl) does not bring significant improvement. It is noticeable that this pool is worse than the baseline on both tst2015 (iwslt seg) and tst2015 (ASR seg). However, we find that applying data augmentation (SpecAugment) on this same combination helps outperform the baseline on every investigated testset, most significantly on MuST-C tst-COMMON, and MuST-C tst-HE. Therefore, SpecAugment is consistently applied to all the experiments that follow. Adding MuST-C synthetic to this pool surprisingly decreases BLEU scores on both MuST-C testsets, while significantly increases the scores on both tst2015 (iwslt seg) and tst2015 (ASR seg). Not being able to investigate further on this matter due to time constraint, instead of fine-tuning 4*, we decided to fine-tune 3*, which performs reasonably well among all the testsets, on MuST-C original and MuST-C original+synthetic. We witness that the impact of fine tuning is very limited. One can also see once again that adding MuST-C synthetic does not make much difference. Finally, the last

row of the table shows the results of ensembling all six models at decoding time. It is clear from the table that ensembling yields the best BLEU scores across all the testsets.

2.5 Overview of systems submitted

Two conclusions that can be drawn from Table 2 are (1) ensembling all six models is the most promising among all presented models, (2) our own segmentation (tst2015 ASR segmentation) is better than the default one. Therefore, we choose as our primary submission the translations of the ASR segmentations generated by the ensemble of all six models. Model 3* and 4* (Table 2) are also used to translate our contrastive submission runs, whose ranks are shown in Table 3. The official results for all our submitted systems can be found in Table 4. They confirm that our segmentation approach proposed is beneficial.

3 Simultaneous Speech Translation Track

In this section, we describe our submission to the Simultaneous Speech Translation (SST) track. Our pipeline consists of an automatic speech recognition (ASR) system followed by an online machine translation (MT) system. We first define our online ASR and MT models in §3.1 and §3.2 respectively. Then, we outline in §3.3 how we arrange the two systems for the speech-to-text subtask. We detail our experimental setup and report our results on the text-to-text subtask in §3.4 and on the speech-to-text in §3.5.

3.1 Online ASR

Our ASR system is a hybrid HMM/DNN system trained with lattice-free MMI (Povey et al., 2016), using the Kaldi speech recognition toolkit (Povey et al., 2011). The acoustic model (AM) topology consists of a Time Delay Neural Network (TDNN) followed by a stack of 16 factorized TDNNs (Povey et al., 2018). The acoustic feature vector is a concatenation of 40-dimensional MFCCs without cepstral truncation (MFCC-40) and 100-dimensional i-vectors for speaker adaptation (Dehak et al., 2010). Audio samples were randomly perturbed in speed and amplitude during the training process. This approach is commonly called audio augmentation and is known to be beneficial for speech recognition (Ko et al., 2015).

Online decoding with Kaldi. The online ASR system decodes under a set of rules to decide when to stop decoding and output a transcription. An endpoint is detected if either of the following conditions is satisfied:

- (a) After t seconds of silence even if nothing was decoded.
- (b) After t seconds of silence after decoding something, if the final-state was reached with $\text{cost}_{\text{relative}} < c$.
- (c) After t seconds of silence after decoding something, even if no final-state was reached.
- (d) After the utterance is t seconds long regardless of anything else.

Each rule has an independent characteristic time t and condition (b) can be duplicated with different times and thresholds (t, c). The value of $\text{cost}_{\text{relative}}$ reflects the quality of the output, it is *null* if a final-state of the decoding graph had the best cost at the final frame, and *infinite* if no final-state was active.

3.2 Online MT

Our MT systems are Transformer-based (Vaswani et al., 2017) wait- k decoders with unidirectional encoders. Wait- k decoding starts by reading k source tokens, then alternates between reading and writing a single token at a time, until the source is depleted, or the target generation is terminated. With a source-target pair (\mathbf{x}, \mathbf{y}) , the number of source tokens read when decoding y_t following a wait- k policy is $z_t^k = \min(k + t - 1, |\mathbf{x}|)$. To stop leaking signal from future source tokens, the energies of the encoder-decoder multihead-attention

are masked to only include the z_t tokens read so far.

Unlike Transformer wait- k models introduced in Ma et al. (2019) where the source is processed with a bidirectional encoder, we opt for a unidirectional encoding of the source. In fact, this change alleviates the cost of re-encoding the source sequence after each read operation. Contrary to offline task, where bidirectional encoders are superior, unidirectional encoder achieve better quality-lagging trade-offs in online MT.

Ma et al. (2019) optimize their models with maximum likelihood estimation w.r.t. a single wait- k decoding path z^k :

$$\log p(\mathbf{y} | \mathbf{x}, z^k) = \sum_{t=1}^{|\mathbf{y}|} \log p_{\theta}(y_t | \mathbf{y}_{<t}, \mathbf{x}_{\leq z_t^k}). \quad (1)$$

Instead of optimizing a single decoding path, we jointly optimize across multiple wait- k paths. The additional loss terms provide a richer training signal, and potentially yield models that could perform well under different lagging constraints. Formally, we consider an exhaustive set of wait- k paths and in each training epoch we encode the source sequence then uniformly sample a path to decode with. As such, we optimize:

$$\mathcal{Z} = \{z^k | k \in \{1, \dots, |\mathbf{x}|\}\}, \quad (2)$$

$$\mathbb{E}_z[\log p(\mathbf{y} | \mathbf{x}, z)] \approx \sum_{z \in \mathcal{Z}} \log p_{\theta}(\mathbf{y} | \mathbf{x}, z). \quad (3)$$

We will refer to this training with *multi-path*.

3.3 Cascaded ASR+MT

For speech-to-text online translation we pair an ASR system with our online MT system and decode following the algorithm described in Algorithm 1.

In this setup, the lagging is controlled by the endpointing of the ASR system. The online MT system follows the lead of the ASR and translates prefix-to-prefix. Since the MT system is not trained to detect end of segments and can only halt the translation by emitting $\langle /s \rangle$, we constrain it to decode $\alpha |\mathbf{x}^{\text{asr}}| + \beta$ tokens, where \mathbf{x}^{asr} is the partial transcription and (α, β) two hyper-parameters.

Along with the hyper-parameters of the ASR’s endpointing rules, we tune (α, β) on a development set to achieve good latency-quality trade-offs.

3.4 Text-to-text translation subtask

Training MT. We train our online MT systems on English-to-German MuST-C (Di Gangi et al., 2019)

Algorithm 1 ASR+MT decoding algorithm

Input: source audio blocks \mathbf{x} .
Output: translation hypothesis \mathbf{y} .
Initialization: action=READ, $z=0$, $t=1$,
 $\mathbf{x}^{\text{asr}}=(\text{ })$, $\mathbf{y}=(\langle s \rangle)$
Hyper-parameters sz, α, β .
while $y_t \neq \langle /s \rangle$ **do**
 while action = READ $\wedge z < |\mathbf{x}|$ **do**
 Read sz elements from \mathbf{x} . $z += sz$
 Feed the new audio blocks to the ASR system.
 if Endpoint detected $\vee z = |\mathbf{x}|$ **then**
 Output transcription and append it to \mathbf{x}^{asr} .
 action = WRITE
 end if
end while
if $|\mathbf{y}| < \alpha|\mathbf{x}^{\text{asr}}| + \beta$ **then**
 Given \mathbf{y} and \mathbf{x}^{asr} , predict the next token y_{t+1}
 $t += 1$
else
 action = READ
end if
end while

	Pairs	English words	German words
Europarl	1,730K	43,7M	41,1M
Common Crawl	1,543K	31,0M	30,0M
News Commentary	320K	7,0M	7,2M
MuST-C	214K	3,9M	3,7M

Table 5: Parallel training data for the MT systems.

and WMT’19 data,⁶ namely, Europarl (Koehn, 2005), News Commentary (Tiedemann, 2012) and Common Crawl (Smith et al., 2013). We remove pairs with a length-ratio exceeding 1.3 from Common Crawl and pairs exceeding a length-ratio of 1.5 from the rest. We develop on MuST-C dev and report results on MuST-C tst-COMMON. For open-vocabulary translation, we use *SentencePiece* (Kudo and Richardson, 2018) to segment the bi-texts with byte pair encoding (Sennrich et al., 2016). This results in a joint vocabulary of 32K types. Details of the training data are provided in Table 5.

We train Transformer *big* architectures and tie the embeddings of the encoder with the decoder’s input and output embeddings. We optimize our models with label-smoothed maximum likelihood (Szegedy et al., 2016) with a smoothing rate $\epsilon = 0.1$. The parameters are updated using


⁶<http://www.statmt.org/wmt19/>

Figure 2: [Text-to-Text] Latency-quality trade-offs evaluated on MuST-C tst-COMMON with greedy decoding. Offline systems have an AL of 18.55 words. The red vertical bars correspond to the AL evaluation thresholds.

Adam (Kingma and Ba, 2015) ($\beta_1, \beta_2 = 0.9, 0.98$) with a learning rate that follows an inverse square-root schedule. We train for a total of 50K updates and evaluate with the check-pointed weights corresponding to the lowest (best) loss on the development set. Our models are implemented with Fairseq (Ott et al., 2019). We generate translation hypotheses with greedy decoding and evaluate the latency-quality trade-off by measuring case-sensitive detokenized BLEU (Papineni et al., 2002) and word-level Average Lagging (AL) (Ma et al., 2019).

Results. We show in Figure 2 the performance of our systems on the test set (MuST tst-COMMON) measured with the provided evaluation server.⁷ We denote with $k_{\text{train}}=\infty$ a unidirectional model trained for wait-until-end decoding *i.e.* reading the full source before writing the target. We evaluate four *wait-k* systems each trained with a value of k_{train} in $\{5, 7, 9, \infty\}$ and decoded with k_{eval} ranging from 2 to 11. We then ensemble the aforementioned *wait-k* models and evaluate a *multi-path* model that jointly optimizes a large set of *wait-k* paths. The results demonstrate that *multi-path* is competitive with *wait-k* without the need to select which path to optimize (some values of k , *e.g.* 5, underperform in comparison). Ensembling the *wait-k* models gives a boost of 1.43 BLEU points on average.

⁷https://github.com/pytorch/fairseq/blob/simulastsharedtask/examples/simultaneous_translation

Corpus	#hours	#words	#speakers
TED-LIUM 3	452	5.05M	2,028
How2	365	3.31M	13,147
Europarl	94	0.75M	171

Table 6: Corpora used for the acoustic model.

Decoding	Corpus	WER
Offline	TED-LIUM 3 dev	7.65
Offline	TED-LIUM 3 test	7.84
Offline	MuST-C tst-COMMON	14.2
Online	MuST-C tst-COMMON	16.3

Table 7: WERs for the ASR system with offline and online decoding (AL=5s for online)

3.5 Speech-to-text translation subtask

Training ASR. We train our system following the *tedlium* recipe⁸ while adapting it for the IWSLT task. The TDNN layers have a hidden dimension of 1536 with a linear bottleneck dimension of 160 in the factorized layers. The i-vector extractor is trained on all acoustic data (speech perturbed + speech) using a 10s window. The acoustic training data includes TED-LIUM 3, How2 and Europarl. These corpora are detailed in Table 6 and represent about 900 hours of audio.

As a language model, we use the 4-grams *small* model provided with TED-LIUM 3. The vocabulary size is 152K, with 1.2 million of 2-grams, 622K 3-grams and 70K 4-grams.

The final system is tuned on TED-LIUM 3 dev and tested with TED-LIUM 3 test and MuST-C tst-COMMON. Results are shown in Table 7.

Training MT. To train the MT system for the ASR+MT cascade we process source-side data (English) to match transcriptions of the ASR. This consists of lower-casing, removing punctuation and converting numbers into letters. For this task we use two distinct English and German vocabularies of 32K BPE tokens each. We train Transformer *big* architectures with tied input-output decoder embeddings following the setup described in §3.4.

Results. Similar to the text-to-text subtask, we show our results in a plot of BLEU-to-AL in Figure 3. The systems are evaluated on the test via the provided evaluation server where MuST-

⁸<https://github.com/kaldi-asr/kaldi/tree/master/egs/tedlium>




Figure 3: [Speech-to-Text] Latency-quality trade-offs evaluated on MuST-C tst-COMMON with greedy decoding. Offline systems have an AL of 5806 ms. The red vertical bars correspond to the AL evaluation thresholds.

C’s sentence-level aligned segments are streamed and decoded online and the lagging is measured in milliseconds. Note that in this task we use a single ASR model and only ensemble the MT *wait-k* models. The cascade of an online ASR with *wait-k* MT follows the same trends as the text-to-text models. In particular, *multi-path* is competitive with specialized *wait-k* models and ensembling boosts the BLEU scores by 0.67 points on average.

4 Conclusion

This paper described the ON-TRAC consortium submission to the IWSLT 2020 shared task. In the continuity of our 2019 participation, we have submitted several end-to-end systems to the offline speech translation track. A significant part of our efforts was also dedicated to the new simultaneous translation track: we improved *wait-k* models with unidirectional encoders and multi-path training and cascaded them with a strong ASR system. Future work will be dedicated to simultaneous speech translation using end-to-end models.

5 Acknowledgements

This work was funded by the French Research Agency (ANR) through the ON-TRAC project under contract number ANR-18-CE23-0021.

References

- Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondrej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, and Changan Wang. 2020. Findings of the IWSLT 2020 Evaluation Campaign. In *Proceedings of the 17th International Conference on Spoken Language Translation (IWSLT 2020)*, Seattle, USA.
- Naveen Arivazhagan, Colin Cherry, Wolfgang Macherey, Chung-Cheng Chiu, Semih Yavuz, Ruoming Pang, Wei Li, and Colin Raffel. 2019. [Monotonic infinite lookback attention for simultaneous machine translation](#). In *Proc. of ACL*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proc. of ICLR*.
- Alexandre Bérard, Laurent Besacier, Ali Can Kocabiyikoglu, and Olivier Pietquin. 2018. [End-to-End Automatic Speech Translation of Audiobooks](#). In *Proc. of ICASSP*.
- Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. 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*.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, and Stephan Vogel. 2018. [Incremental decoding and training methods for simultaneous translation in neural machine translation](#). In *Proc. of NAACL-HLT*.
- Najim Dehak, Patrick J Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet. 2010. Front-end factor analysis for speaker verification. *IEEE Transactions on Audio, Speech, and Language Processing*.
- Mattia Antonino Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. Must-c: a multilingual speech translation corpus. In *Proc. of NAACL-HLT*.
- Pegah Ghahremani, Bagher BabaAli, Daniel Povey, Korbinian Riedhammer, Jan Trmal, and Sanjeev Khudanpur. 2014. A pitch extraction algorithm tuned for automatic speech recognition. In *Proc. of ICASSP*.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. 2018. TED-LIUM 3: twice as much data and corpus repartition for experiments on speaker adaptation. In *International Conference on Speech and Computer*.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Adrià Giménez, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. Europarl-ST: A multilingual corpus for speech translation of parliamentary debates. In *Proc. of ICASSP*.
- Ye Jia, Melvin Johnson, Wolfgang Macherey, Ron J. Weiss, Yuan Cao, Chung-Cheng Chiu, Naveen Ari, Stella Laurenzo, and Yonghui Wu. 2019. [Leveraging weakly supervised data to improve end-to-end speech-to-text translation](#). In *Proc. of ICASSP*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proc. of ICLR*.
- Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. 2015. [Audio augmentation for speech recognition](#). In *Proc. of INTERSPEECH*.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *MT summit*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proc. of EMNLP: System Demonstrations*.
- Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. 2019. [STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework](#). In *Proc. of ACL*.
- Xutai Ma, Juan Pino, James Cross, Liezl Puzon, and Jiatao Gu. 2020. [Monotonic multihead attention](#). In *Proc. of ICLR*.
- Ha Nguyen, Natalia Tomashenko, Marcelly Zanon Boito, Antoine Caubriere, Fethi Bougares, Mickael Rouvier, Laurent Besacier, and Yannick Esteve. 2019. ON-TRAC consortium end-to-end speech translation systems for the IWSLT 2019 shared task. In *Proc. of IWSLT*.
- Jan Niehues, Ngoc-Quan Pham, Thanh-Le Ha, Matthias Sperber, and Alex Waibel. 2018. [Low-latency neural speech translation](#). In *Proc. of INTERSPEECH*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proc. of NAACL-HLT: Demonstrations*.
- Kishore Papineni, Salim Roukos, Todd Ard, and Weijing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proc. of ACL*.
- Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. 2019. [SpecAugment: A simple data augmentation method for automatic speech recognition](#). In *Proc. of INTERSPEECH*.
- Daniel Povey, Gaofeng Cheng, Yiming Wang, Ke Li, Hainan Xu, Mahsa Yarmohammadi, and Sanjeev Khudanpur. 2018. [Semi-orthogonal low-rank matrix factorization for deep neural networks](#). In *Proc. of INTERSPEECH*.

- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Nandragoel, Mirko Hannemann, Yanmin Qian, Petr Schwarz, Georg Stemmer, et al. 2011. The Kaldi speech recognition toolkit. In *In IEEE 2011 workshop*.
- Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahremani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur. 2016. Purely sequence-trained neural networks for ASR based on lattice-free MMI. In *Proc. of INTERSPEECH*.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. How2: a large-scale dataset for multimodal language understanding. In *ViGIL Workshop, NeurIPS*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proc. of ACL*.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *Proc. of ICLR*.
- Jason Smith, Herve Saint-Amand, Magdalena Plamadă, Philipp Koehn, Chris Callison-Burch, and Adam Lopez. 2013. [Dirt cheap web-scale parallel text from the common crawl](#). In *Proc. of ACL*.
- Matthias Sperber, Graham Neubig, Jan Niehues, and Alex Waibel. 2019. [Attention-passing models for robust and data-efficient end-to-end speech translation](#). *TACL*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. [Rethinking the inception architecture for computer vision](#). In *Proc. of CVPR*.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proc. of LREC*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proc. of NeurIPS*.
- Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, et al. 2018. [Espnet: End-to-end speech processing toolkit](#). In *Proc. of 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. of INTERSPEECH*.