

Direct Models for Simultaneous Translation and Automatic Subtitling: FBK@IWSLT2023

Sara Papi^{◇,□}, Marco Gaido[◇], Matteo Negri[◇]

[◇]Fondazione Bruno Kessler

[□]University of Trento

{spapi,mgaido,negri}@fbk.eu

Abstract

This paper describes the FBK’s participation in the Simultaneous Translation and Automatic Subtitling tracks of the IWSLT 2023 Evaluation Campaign. Our submission focused on the use of direct architectures to perform both tasks: for the simultaneous one, we leveraged the knowledge already acquired by offline-trained models and directly applied a policy to obtain the real-time inference; for the subtitling one, we adapted the direct ST model to produce well-formed subtitles and exploited the same architecture to produce timestamps needed for the subtitle synchronization with audiovisual content. Our English-German SimulST system shows a reduced computational-aware latency compared to the one achieved by the top-ranked systems in the 2021 and 2022 rounds of the task, with gains of up to 3.5 BLEU. Our automatic subtitling system outperforms the only-existing solution based on a direct system by 3.7 and 1.7 SubER in English-German and English-Spanish respectively.

1 Introduction

In recent years, the advances in natural language processing and machine learning led to a surge of interest in developing speech translation (ST) systems that can translate speech from one language into text in another language without human intervention. Significant progress has been specially made toward end-to-end ST models (Bérard et al., 2016; Weiss et al., 2017) trained to directly translate speech without the intermediate steps of transcription (through automatic speech recognition - ASR) and translation (through machine translation - MT). Along with this growing interest in direct ST, also accompanied by a reduction of the performance gap with respect to cascaded architectures (Bentivogli et al., 2021), other trends have emerged thanks to deep learning advancements, which made it possible to deploy direct solutions to perform the task in real-time (i.e. to produce partial translations

while continuing to process the input audio) or to automatically generate subtitles for audiovisual content (i.e. pieces of translated text which have to conform to specific spatiotemporal constraints and be synchronized with the video).

The International Workshop on Spoken Language Translation (IWSLT) is playing an important role in advancing the state-of-the-art in these fields by organizing a series of evaluation campaigns (Ansari et al., 2020; Anastasopoulos et al., 2021, 2022) focused on simultaneous speech translation (SimulST) and, this year for the first time, automatic subtitling. These campaigns provide a unique opportunity for researchers to compare their systems against others, share their findings, and identify areas for further improvement.

In this paper, we describe FBK’s participation in the IWSLT 2023 Evaluation Campaigns (Agarwal et al., 2023) for simultaneous translation and automatic subtitling. Motivated by the promising results reported in previous works (Ren et al., 2020; Papi et al., 2022a), our approach is characterized by the use of direct ST models to address both tasks.

For the simultaneous speech-to-text translation (SimulST) task, we participated in the English→German track and leveraged an offline-trained direct model without performing any adaptation to the real-time scenario, as this has recently been shown not to be necessary to achieve competitive results (Papi et al., 2022b). For the automatic subtitling task, we participated in both the English→German and English→Spanish tracks by adapting a direct ST model to produce well-formed subtitles and exploiting the same architecture to produce the timestamps needed for their synchronization with audiovisual contents, as in (Papi et al., 2022a).

Our results demonstrate the effectiveness of our approach. In SimulST, the computational-aware latency of our models is lower compared to the winning systems of the last two rounds (2021, and

2022) of the IWSLT SimulST Evaluation Campaign, with gains up to 3.5 BLEU. In automatic subtitling, our systems improve the results reported in (Papi et al., 2022a) which, to the best of our knowledge, represents the only-existing solution based on a direct model. Specifically, on average among the various dev sets available for the task, we achieve 3.7 SubER on en-de and 1.7 SubER on en-es.

2 Applied Direct Models

For this year’s submission, we applied the direct ST models to the two different scenarios of simultaneous translation and automatic subtitling.

2.1 Simultaneous Translation

Recent trends in SimulST consist of using offline-trained models for simultaneous inference (Papi et al., 2022b). There are several motivations for this choice: *i*) it avoids re-training or building specific architectures for SimulST, saving time and computational resources; *ii*) only one model has to be trained and maintained to perform both offline and simultaneous ST; and *iii*) there is no need to train several models, each specialized to support different latency regimes.

A key aspect of SimulST, also critical when approaching the task with offline models at inference time, is the so-called *decision policy*: the mechanism that is in charge of deciding whether to read more information or to emit a partial hypothesis. One of the first and most popular policies is the wait- k (Ma et al., 2019), initially introduced for simultaneous MT, and then applied to the speech scenario (Ma et al., 2020b; Chen et al., 2021; Zeng et al., 2021; Karakanta et al., 2021b). The wait- k , which prescribes waiting for an initial number of k words before starting to translate, is defined as a “fixed” policy (Zheng et al., 2020) because the decision is taken independently from the source input content. However, as the actual information contained in the input (e.g. in terms of ambiguity, completeness, and syntactic/semantic cohesion) is also important for the sake of good-quality incremental translations, several “adaptive” policies have been introduced, which instead adapt their decisions to the input content. Some adaptive policies require system re-training or the development of ad-hoc modules (Liu et al., 2021b; Chang and Lee, 2022; Zhang and Feng, 2022), while some others do not (Liu et al., 2020; Nguyen et al., 2021; Papi et al.,

2022e). Since our objective is to avoid any modifications to the offline-trained model, we pointed our attention to the latter, more conservative category. Among these policies, we analyzed the three following alternatives:

- **Local Agreement (LA)** (Liu et al., 2020): this policy prescribes generating a partial hypothesis from scratch at every newly received audio segment, and emitting it (or only a part of it) if it coincides with one of those generated in the previous time step;
- **Encoder-Decoder Attention (EDATT)** (Papi et al., 2022e): it exploits the cross-attention scores modeling the audio-translation relation to decide whether to emit the words of a partial hypothesis or not. If, for the current word, the sum of the attention scores of the last λ received speech frames exceeds a certain threshold α (both λ and α are hyperparameters), the emission is delayed because the system needs more context to translate that word. Otherwise, the word is emitted and we proceed to the next word of the hypothesis;
- **ALIGNATT** (Papi et al., 2023b): as for EDATT, the cross-attention scores are leveraged to decide what to emit but, in this case, instead of summing the attention scores of the last speech frames, each word is uniquely assigned (or aligned) to the frame having the maximum attention score. If the aligned frame corresponds to one of the last f frames (f being a hyperparameter that controls the latency) the emission is stopped. Otherwise, we proceed to the next word.

2.2 Automatic Subtitling

So far, the adoption of direct ST architectures to address the automatic subtitling task has only been explored in (Papi et al., 2022a). As a matter of fact, all previous works on the topic (Piperidis et al., 2004; Melero et al., 2006; Matusov et al., 2019; Koponen et al., 2020; Bojar et al., 2021) rely on cascade architectures that usually involve an ASR component to transcribe the input speech, a subtitle segmenter that segments the transcripts into subtitles, a timestamp estimator that predicts the start and times of each subtitle, and an MT model that translates the subtitle transcripts.

Cascaded architectures, however, cannot access information contained in the speech, such as

prosody, which related works proved to be an important source of information for the segmentation into subtitles (Öktem et al., 2019; Federico et al., 2020; Virkar et al., 2021; Tam et al., 2022). The importance of such information has been further verified in (Karakanta et al., 2020a), which proved that the direct ST models are better in subtitle segmentation compared to the cascade ones. Another study by Karakanta et al. 2021a, also pointed out the importance of consistency between captions (segmented transcripts) and subtitles (segmented translations), showing that the predicted caption content can also be useful for the translation. Specifically, the authors obtained significant improvements by using a Triangle Transformer-based architecture (Anastasopoulos and Chiang, 2018) composed of one encoder and two decoders: the first decoder is in charge of emitting the transcripts and the second one is in charge of emitting the translation by also attending to the output embeddings of the predicted transcript. Therefore, in our submission, based on the findings of the aforementioned work, we inspected the use of both a classic single encoder-single decoder architectures, as in (Papi et al., 2022a), and of the Triangle architecture for automatic subtitling.

3 Experimental Setup

3.1 Data

Simultaneous We developed a pure offline model trained on the same data used for our last year’s (constrained) submission (Gaido et al., 2022b).

Subtitling We used the same data settings of (Papi et al., 2022a), for which we leverage the multimodal segmenter by Papi et al. (2022d) to segment into subtitles ST and machine-translated ASR corpora as per (Gaido et al., 2021b, 2022a).¹ No OpenSubtitles or text-only data were used to train our models.

3.2 Training Settings

All the models used for our participation were implemented using the newly released implementation of the Conformer architecture by Papi et al.

¹All the corpora used in (Papi et al., 2022a) are allowed ASR and ST training data for the Subtitling task (<https://iwslt.org/2023/subtitling#training-and-data-conditions>). Therefore, our submission has to be considered “Constrained”.

(2023a)² based on Fairseq-ST (Wang et al., 2020). In their paper, the authors analyzed the most popular open-source libraries for speech recognition/translation and found at least one bug affecting all the existing Conformer implementations, therefore claiming the importance of testing code to avoid the propagation of unreliable findings masked by good results.

Simultaneous We tested a Conformer-based architecture (Gulati et al., 2020) with two configurations: 12 encoder layers and 16 encoder layers. The number of Transformer decoder layers is 6, we set 512 features for the attention layers and 2,048 hidden units for the feed-forward layers. We used 0.1 dropout for the feed-forward layers, attention layers, and convolutional modules. The kernel size was set to 31 for the point- and depth-wise convolutions. We trained with the Adam optimizer (Kingma and Ba, 2015) by setting $\beta_1 = 0.9$ and $\beta_2 = 0.98$, a weight decay of 0.001, the learning rate to 0.002 using the inverse square-root scheduler with 25,000 warm-up steps. Label smoothed cross-entropy loss (0.1 smoothing factor) was used together with the CTC loss (Graves et al., 2006) with weight 0.5. We experimented also by applying the CTC compression mechanism (Gaido et al., 2021a) to the source input to shrink its dimension and reduce RAM consumption. Utterance Cepstral Mean and Variance Normalization (CMVN) was applied during training. Also, we leveraged SpecAugment (Park et al., 2019) with frequency mask ($F = 27$, and $N = 2$), time mask ($N = 10$, $T = 300$, and $p = 0.05$), and no time warp. Both ST training and ASR pre-training were performed with the same settings. The target vocabulary is of size 16,000, and the source vocabulary is of size 10,000, and are both based on SentencePiece (Kudo and Richardson, 2018). We differentiate between original and machine-translated training data by pre-pending a tag (nomt and mt, respectively) to the target text as in all our last years’ submissions (Gaido et al., 2020; Papi et al., 2021; Gaido et al., 2022b). The total batch size was set to 1,280,000 and was performed on 4 NVIDIA A40 GPUs with 40GB of RAM by setting the mini-batch update frequency to 8 and 40,000 maximum tokens. Maximum updates were set to 100,000.

²Code available at <https://github.com/hlt-mt/FBK-fairseq>

Automatic Subtitling Both the classic encoder-decoder architecture and the triangle architecture are composed of 12 layers of Conformer encoder and 6 layers of Transformer decoder (which is replicated twice in the triangle model). The dimension of the feed-forward layers is 2,048 and $d = 512$ in the attention. The kernel size of the point- and depth-wise convolutions in the convolutional modules is 31. The dropout was set to 0.1. CTC loss with compression is added with weight 0.5 to the cross entropy loss with label smoothing (0.1 of smoothing factor) and optimized with Adam ($\beta_1 = 0.9, \beta_2 = 0.98$). The source vocabulary is of size 8,000 and the target vocabulary of size 16,000 (`<eob>` and `<eol>` included); both are obtained by SentencePiece models. The ST pre-training was done by setting the learning rate to 0.002 with inverse square-root scheduler and 25,000 warm-up updates. The SubST fine-tuning was done by setting a constant learning rate of 0.001. A second fine-tuning was done with the same setting of (Papi et al., 2022a), but we restored the punctuation of the ASR datasets which do not contain any (i.e., the TEDLIUM corpus (Hernandez et al., 2018)) by using `bert-restore-punctuation`,³ before machine-translating and segmenting the target texts into subtitles. We trained the standard architecture with 40,000 maximum tokens on 4 NVIDIA A100 GPUs with 40GB of RAM and we set the update frequency to 2. For the triangle architecture, we set maximum tokens to 20,000 to fit the architecture in memory and the update frequency to 4 to hold the same total batch size of 320,000 tokens. Maximum updates were set to 100,000 for both the pre-training and training phases.

3.3 Evaluation Settings

Simultaneous We exploit the SimulEval tool (Ma et al., 2020a). To be comparable with the previous years, all the results except this year’s submission are shown for the SimulEval v1.0.2, which adopts BLEU (Post, 2018)⁴ to measure translation quality and Average Lagging or AL (Ma et al., 2019) to measure latency. Instead, for this year’s submission, we adopt the latest version of SimulEval (1.1.0) with BLEU measured with `sacrebleu` 2.3.0 and we also report Length-Adaptive Average Lagging or LAAL (Papi et al., 2022c) and Average Token Delay or ATD (Kano

³<https://huggingface.co/felflare/bert-restore-punctuation>

⁴case:mixedleff:noltok:13alsmooth:explversion:1.5.1

et al., 2022) as additional latency metrics. All the evaluations were run on a single NVIDIA K80 with 12GB of RAM, by applying global CMVN to audio input, whose features were estimated on the MuST-C v2 training set. Computational aware metrics (“_CA”) refer to the single NVIDIA K80 setting and consider also the model computational time in the delay calculation.

Automatic Subtitling We adopt the following metrics: SubER-cased (henceforth, SubER) (Wilken et al., 2022) for overall subtitle quality, Sigma (Karakanta et al., 2022) for the subtitle segmentation quality, and BLEU⁵ for translation quality. We also compute the conformity percentage of 42 characters per line (CPL) and 21 characters per second (CPS) or reading speed, as suggested on the track website.⁶ We neglected the conformity computation of the subtitles with more than two lines since our model only produces subtitles with two lines or less, thus being always 100% conform. Conformity scores are computed by using the script released for the paper (Papi et al., 2022a).⁷ Dev/test audios are segmented with SHAS (Tsiamas et al., 2022). No audio cleaning is applied.

4 Results

4.1 Simultaneous Translation

Since we directly employ an offline model for the simultaneous inference, we show in Table 1 the results of the offline ASR pre-training and ST training. Although the model with 12 encoder layers (row 0) obtains lower – hence better – WER compared to the 16 encoder-layers model (row 1), the highest – hence better – BLEU in ST is achieved by the bigger architecture. The performance is also slightly enhanced by adding the CTC compression (row 3) during training, which is particularly useful also for the SimulST scenario since it speeds up inference (of about 12/15%). Therefore, we select this model for the final submission. Compared to our last year’s submission (row 5), our 16 encoder-layers model scores +0.4 BLEU even if, at this time, we have not fine-tuned it on the in-domain (TED talks) datasets. Our model also performs

⁵case:mixedleff:noltok:13alsmooth:explversion:2.0.0

⁶<https://iwslt.org/2023/subtitling#automatic-evaluation>

⁷Script available at: https://github.com/hlt-mt/FBK-fairseq/blob/master/examples/speech_to_text/scripts/subtitle_compliance.py

better than the NAIST last year’s system (+11.1 BLEU) while is worse (-1.0 BLEU) compared to the last year’s SimulST task winner CUNI-KIT whose model, however, leveraged large pre-trained models such as wav2vec 2.0 and mBART50. Compared to last year’s cascade model by UPV, we score -1.7 BLEU. This system, however, also outperformed the CUNI-KIT system by 0.7 BLEU points, indicating that a gap between direct and cascade architectures still exists.

id	Model	WER% (\downarrow)	BLEU (\uparrow)
1	12 encoder layers	9.7	31.6
2	16 encoder layers	9.9	31.9
3	+ CTC compress.	-	32.1
4	CUNI-KIT 2022 [†]	-	33.1
5	FBK 2022	-	31.7
6	NAIST 2022 [‡]	-	21.0
7	UPV 2022 (Cascade)*	9.5	33.8

Table 1: Offline results of our Conformer-based architectures on MuST-C v2 tst-COMMON together with the available results of the last year’s SimulST competitors. [†](Polák et al., 2022), [‡](Fukuda et al., 2022), *(Iranzo-Sánchez et al., 2022).

In Figure 1, we show the simultaneous results of the different policies mentioned in Section 2.1 applied to our offline model. The differences in terms of quality-latency trade-off between the LA and both EDATT and ALIGNATT are evident: the last ones outperform the former with an improvement peak of 1.5 BLEU at lower latency (approximately 1s). Moreover, when the computationally aware AL is considered, EDATT and ALIGNATT are the only policies able to reach a latency $\leq 2s$. Regarding the comparison between EDATT and ALIGNATT, ALIGNATT can span a latency between 1 and 2.6s ideally (when unlimited computing resources are available), and between 1.8 and 3.7s computationally aware, while EDATT is limited to a latency of 1.4 to 2.5s ideally, and 2.3 to 3.6s computationally aware. We hence select ALIGNATT as it is able to reach a wider range of latency.

Lastly, we compare our policy with the two winning systems of the last two years (2021, and 2022). The 2021 winner (Liu et al., 2021a) was based on an architecture named Cross Attention Augmented Transducer (CAAT), which was specifically tailored for the SimulST task (Liu et al., 2021b) and still represents the state of the art in terms of low latency (considering ideal AL only). The 2022 winner (CUNI-KIT (Polák et al., 2022)) was based on the wav2vec 2.0 + mBART50 offline architecture




Figure 1: Comparison between the LA, EDATT, and ALIGNATT policies described in Section 2.1 on MuST-C v2 en \rightarrow de tst-COMMON. Solid curves represent AL, dashed curves represent AL_CA.

reported in Table 1, row 4. They applied the LA policy, the same we analyze in Figure 1, to the aforementioned architecture for simultaneous inference. The comparison is reported in Figure 2. As we can see, there is a 1.0-2.0 BLEU difference between our approach and the IWSLT 2022 winner, which is expected since their offline system is superior compared to ours, as already observed in Table 1. Compared to the IWSLT 2021 winner, we observe a performance drop in our system with $AL \leq 1.5s$, while the situation is opposite with $AL > 1.5s$. However, when we look at the computationally aware metrics, the results completely change. Our system clearly outperforms the 2021 winner, with a maximum improvement of about 2 BLEU points. Moreover, our system is the only one able to reach a computational aware latency of 2s while, instead, the IWSLT 2022 winner curve starts only at around 3s. Therefore, our system is significantly faster and, at around 3s, we achieve a relative improvement of more than 3.5 BLEU compared to the IWSLT 2022 winner.

To sum up, when the computationally aware metric is considered, our approach outperforms the winners of both the 2021 and 2022 rounds of the SimulST task. In addition, in this year’s round, the systems are evaluated with the threshold $AL = 2s$ and with the new version of SimulEval.⁸ With respect to these settings, our submitted system scores 30.7 BLEU with $AL = 1.89s$ ($LAAL = 2.07s$, $ATD = 1.80s$).

⁸<https://iwslt.org/2023/simultaneous#ranking>




Figure 2: Comparison with the 2021 and 2022 winners of the SimulST Evaluation Campaigns MuST-C v2 en→{de, es} tst-COMMON. Solid curves represent AL, dashed curves represent AL_CA.

4.2 Automatic Subtitling

In Table 2, we show a comparison between the standard encoder-decoder and the Triangle architectures for automatic subtitling. The results are computed on MuST-Cinema (Karakanta et al., 2020b), the only existing corpus for SubST. Unfortunately, in contrast with the results achieved by (Karakanta et al., 2021a), we found that the standard architectures perform better on all the considered metrics. While the differences in terms of translation quality are not so big (0.8-9 BLEU drop in both languages), there is a huge gap in the quality of the segmentation into subtitles, with the standard model improving by 3.3 and 4.7 Sigma the scores obtained by the Triangle respectively on en-de and en-es. This is also reflected by a worse SubER score (the lower, the better) of the Triangle, exhibiting a performance drop of, respectively, 0.9 and 1.6 SubER for en-de and en-es compared to the standard architecture. Therefore, we can conclude that the generated captions seem not to help with subtitle generation. Rather, they negatively affect subtitle generation to the detriment of segmentation quality. For this reason, we decided to employ the standard encoder-decoder architecture for our participation in the automatic subtitling task.

In the following, we present the results of our model on the four dev sets released for the task,⁹ namely: **MuST-Cinema or TED** containing TED talks videos, **EuroParITV or EPTV** containing recordings related to the European Parliament ac-

⁹<https://iwslt.org/2023/subtitling#development-and-evaluation-data>

en-de					
Model	SubER	BLEU	Sigma	CPL	CPS
(Papi et al., 2022a)	59.9	23.4	77.9	86.9	68.9
Triangle	60.8	22.6	74.6	84.5	67.7

en-es					
Model	SubER	BLEU	Sigma	CPL	CPS
(Papi et al., 2022a)	46.8	37.4	81.6	93.2	74.6
Triangle	48.4	36.5	76.9	90.3	71.7

Table 2: Results of the direct ST models standard and Triangle architectures described in Section 2.2 on MuST-Cinema test set for en→{de, es}.

tivities, **Peloton** containing online fitness classes, and **ITV Studios or ITV** containing videos from a broad range of programming (drama, entertainment, factual). For both language pairs (en-de and en-es), Table 3 shows the results computed with SubER, which is the primary metric used for the task.¹⁰ As we can see, the models fine-tuned on data with restored punctuation score the best results in both languages. Across the four dev sets, there is a 3.7 SubER improvement for en-de, and 1.7 for en-es. Moreover, coherently among languages, the TED talks scenario results in the easiest one for our model, as it is in-domain (e.g., MuST-Cinema, based on TED talks, was used to train the model). Conversely, the ITV scenario is the most difficult one since it contains TV series, which is a completely unseen domain for our model. Indeed, its data contain a larger amount of background music/noise, as well as dialogues with multiple speakers which are not present in our training data. In light of the results obtained by the fine-tuned models, we select them for our submission to the automatic subtitling task.

en-de					
Model	TED	EPTV	Peloton	ITV	Avg
(Papi et al., 2022a)	72.7	82.3	84.7	88.0	81.9
+ fine-tuning	69.4	80.6	79.1	83.7	78.2

en-es					
Model	TED	EPTV	Peloton	ITV	Avg
(Papi et al., 2022a)	54.8	75.3	82.3	84.1	74.1
+ fine-tuning	52.5	73.7	80.3	82.2	72.4

Table 3: SubER (↓) scores for en→{de, es} of the direct ST models on the four dev sets of the competition. “fine-tuning” represents the second fine-tuning on data with restored punctuation mentioned in Section 3.2.

¹⁰<https://iwslt.org/2023/subtitling#automatic-evaluation>

5 Conclusions

We presented the FBK’s systems built to participate in the IWSLT 2023 Evaluation Campaigns for simultaneous speech translation (en-de) and automatic subtitling (en- $\{de, es\}$). Our submissions are characterized by the use of direct speech translation models to address both tasks, without any further modification nor adaptation for the simultaneous task, and with a fine-tuning on subtitle-like translations for the automatic subtitling task. Our SimulST system achieves a lower computational-aware latency with up to 3.5 BLEU gain compared to the last two years’ winners. Our automatic subtitling system achieves 3.7 and 1.7 SubER improvement on en-de and en-es respectively, compared to the only solution published in the literature based on a direct system.

Acknowledgements

This work has been supported by the project “AI@TN” funded by the Autonomous Province of Trento, Italy.

References

- Milind Agarwal, Sweta Agrawal, Antonios Anastasopoulos, Ondřej Bojar, Claudia Borg, Marine Carpuat, Roldano Cattoni, Mauro Cettolo, Mingda Chen, William Chen, Khalid Choukri, Alexandra Chronopoulou, Anna Currey, Thierry Declerck, Qianqian Dong, Yannick Estève, Kevin Duh, Marcello Federico, Souhir Gahbiche, Barry Haddow, Benjamin Hsu, Phu Mon Htut, Hirofumi Inaguma, Dávid Javorský, John Judge, Yasumasa Kano, Tom Ko, Rishu Kumar, Pengwei Li, Xutai Ma, Prashant Mathur, Evgeny Matusov, Paul McNamee, John P. McCrae, Kenton Murray, Maria Nadejde, Satoshi Nakamura, Matteo Negri, Ha Nguyen, Jan Niehues, Xing Niu, Atul Ojha Kr., John E. Ortega, Proyag Pal, Juan Pino, Lonneke van der Plas, Peter Polák, Elijah Rippeth, Elizabeth Salesky, Jiatong Shi, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Yun Tang, Brian Thompson, Kevin Tran, Marco Turchi, Alex Waibel, Mingxuan Wang, Shinji Watanabe, and Rodolfo Zevallos. 2023. Findings of the IWSLT 2023 Evaluation Campaign. In *Proceedings of the 20th International Conference on Spoken Language Translation (IWSLT 2023)*.
- Antonios Anastasopoulos, Loïc Barrault, Luisa Bentivogli, Marceley Zanon Boito, Ondřej Bojar, Roldano Cattoni, Anna Currey, Georgiana Dinu, Kevin Duh, Maha Elbayad, Clara Emmanuel, Yannick Estève, Marcello Federico, Christian Federmann, Souhir Gahbiche, Hongyu Gong, Roman Grundkiewicz, Barry Haddow, Benjamin Hsu, Dávid Javorský, Věra Kloudová, Surafel Lakew, Xutai Ma, Prashant Mathur, Paul McNamee, Kenton Murray, Maria Nadejde, Satoshi Nakamura, Matteo Negri, Jan Niehues, Xing Niu, John Ortega, Juan Pino, Elizabeth Salesky, Jiatong Shi, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Yogesh Virkar, Alexander Waibel, Changhan Wang, and Shinji Watanabe. 2022. [Findings of the IWSLT 2022 evaluation campaign](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 98–157, Dublin, Ireland (in-person and online).
- Antonios Anastasopoulos, Ondřej Bojar, Jacob Breermann, Roldano Cattoni, Maha Elbayad, Marcello Federico, Xutai Ma, Satoshi Nakamura, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Alexander Waibel, Changhan Wang, and Matthew Wiesner. 2021. [FINDINGS OF THE IWSLT 2021 EVALUATION CAMPAIGN](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 1–29, Bangkok, Thailand (online).
- Antonios Anastasopoulos and David Chiang. 2018. [Tied multitask learning for neural speech translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 82–91, New Orleans, Louisiana.
- Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, Alexander Waibel, and Changhan Wang. 2020. [FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34, Online.
- Luisa Bentivogli, Mauro Cettolo, Marco Gaido, Alina Karakanta, Alberto Martinelli, Matteo Negri, and Marco Turchi. 2021. [Cascade versus direct speech translation: Do the differences still make a difference?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2873–2887, Online.
- Alexandre Bérard, Olivier Pietquin, Christophe Serivan, 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*, Barcelona, Spain.
- Ondřej Bojar, Dominik Macháček, Sangeet Sagar, Otakar Smrž, Jonáš Kratochvíl, Peter Polák, Ebrahim Ansari, Mohammad Mahmoudi, Rishu Kumar, Dario

- Franceschini, Chiara Canton, Ivan Simonini, Thai-Son Nguyen, Felix Schneider, Sebastian Stüker, Alex Waibel, Barry Haddow, Rico Sennrich, and Philip Williams. 2021. [ELITR multilingual live subtitling: Demo and strategy](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 271–277, Online.
- Chih-Chiang Chang and Hung-Yi Lee. 2022. [Exploring Continuous Integrate-and-Fire for Adaptive Simultaneous Speech Translation](#). In *Proc. Interspeech 2022*, pages 5175–5179.
- Junkun Chen, Mingbo Ma, Renjie Zheng, and Liang Huang. 2021. [Direct simultaneous speech-to-text translation assisted by synchronized streaming ASR](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4618–4624, Online.
- Marcello Federico, Yogesh Virkar, Robert Enyedi, and Roberto Barra-Chicote. 2020. [Evaluating and Optimizing Prosodic Alignment for Automatic Dubbing](#). In *Proc. Interspeech 2020*, pages 1481–1485.
- Ryo Fukuda, Yuka Ko, Yasumasa Kano, Kosuke Doi, Hirotaka Tokuyama, Sakriani Sakti, Katsuhito Sudoh, and Satoshi Nakamura. 2022. [NAIST simultaneous speech-to-text translation system for IWSLT 2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 286–292, Dublin, Ireland (in-person and online).
- Marco Gaido, Mauro Cettolo, Matteo Negri, and Marco Turchi. 2021a. [CTC-based compression for direct speech translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 690–696, Online.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2020. [End-to-end speech-translation with knowledge distillation: FBK@IWSLT2020](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 80–88, Online.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2021b. [On Knowledge Distillation for Direct Speech Translation](#). In *Proceedings of CLiC-IT 2020*, Online.
- Marco Gaido, Matteo Negri, and Marco Turchi. 2022a. [Direct speech-to-text translation models as students of text-to-text models](#). *Italian Journal of Computational Linguistics*.
- Marco Gaido, Sara Papi, Dennis Fucci, Giuseppe Fiameni, Matteo Negri, and Marco Turchi. 2022b. [Efficient yet competitive speech translation: FBK@IWSLT2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 177–189, Dublin, Ireland (in-person and online).
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006. [Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks](#). In *Proceedings of the 23rd international conference on Machine learning (ICML)*, pages 369–376, Pittsburgh, Pennsylvania.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented Transformer for Speech Recognition](#). In *Proc. Interspeech 2020*, pages 5036–5040.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. 2018. [Tedlium 3: Twice as much data and corpus repartition for experiments on speaker adaptation](#). In *Speech and Computer*, pages 198–208, Cham. Springer International Publishing.
- Javier Iranzo-Sánchez, Javier Jorge Cano, Alejandro Pérez-González-de Martos, Adrián Giménez Pastor, Gonçal Garcés Díaz-Munío, Pau Baquero-Arnal, Joan Albert Silvestre-Cerdà, Jorge Civera Saiz, Albert Sanchis, and Alfons Juan. 2022. [MLLP-VRAIN UPV systems for the IWSLT 2022 simultaneous speech translation and speech-to-speech translation tasks](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 255–264, Dublin, Ireland (in-person and online).
- Yasumasa Kano, Katsuhito Sudoh, and Satoshi Nakamura. 2022. [Average token delay: A latency metric for simultaneous translation](#). *arXiv preprint arXiv:2211.13173*.
- Alina Karakanta, François Buet, Mauro Cettolo, and François Yvon. 2022. [Evaluating Subtitle Segmentation for End-to-end Generation Systems](#). In *Proceedings of the 13th Language Resources and Evaluation Conference (LREC)*, pages 3069–3078, Marseilles, France.
- Alina Karakanta, Marco Gaido, Matteo Negri, and Marco Turchi. 2021a. [Between flexibility and consistency: Joint generation of captions and subtitles](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 215–225, Bangkok, Thailand (online).
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2020a. [Is 42 the answer to everything in subtitling-oriented speech translation?](#) In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 209–219, Online.
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2020b. [MuST-cinema: a speech-to-subtitles corpus](#). In *Proc. of the 12th Language Resources and Evaluation Conference*, pages 3727–3734, Marseille, France.

- Alina Karakanta, Sara Papi, Matteo Negri, and Marco Turchi. 2021b. [Simultaneous speech translation for live subtitling: from delay to display](#). In *Proceedings of the 1st Workshop on Automatic Spoken Language Translation in Real-World Settings (ASLTRW)*, pages 35–48, Virtual.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Maarit Koponen, Umut Sulubacak, Kaisa Vitikainen, and Jörg Tiedemann. 2020. [MT for subtitling: User evaluation of post-editing productivity](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 115–124, Lisboa, Portugal.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium.
- Dan Liu, Mengge Du, Xiaoxi Li, Yuchen Hu, and Lirong Dai. 2021a. [The USTC-NELSLIP systems for simultaneous speech translation task at IWSLT 2021](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 30–38, Bangkok, Thailand (online).
- Dan Liu, Mengge Du, Xiaoxi Li, Ya Li, and Enhong Chen. 2021b. [Cross attention augmented transducer networks for simultaneous translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 39–55, Online and Punta Cana, Dominican Republic.
- Danni Liu, Gerasimos Spanakis, and Jan Niehues. 2020. [Low-Latency Sequence-to-Sequence Speech Recognition and Translation by Partial Hypothesis Selection](#). In *Proc. Interspeech 2020*, pages 3620–3624.
- 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 *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3036, Florence, Italy.
- Xutai Ma, Mohammad Javad Dousti, Changhan Wang, Jiatao Gu, and Juan Pino. 2020a. [SIMULEVAL: An evaluation toolkit for simultaneous translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 144–150, Online.
- Xutai Ma, Juan Pino, and Philipp Koehn. 2020b. [SimulMT to SimulST: Adapting simultaneous text translation to end-to-end simultaneous speech translation](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 582–587, Suzhou, China.
- Evgeny Matusov, Patrick Wilken, and Yota Georgakopoulou. 2019. [Customizing neural machine translation for subtitling](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 82–93, Florence, Italy.
- Maite Melero, Antoni Oliver, and Toni Badia. 2006. [Automatic Multilingual Subtitling in the eTITLE Project](#). In *Proceedings of ASLIB Translating and the Computer 28*.
- Ha Nguyen, Yannick Estève, and Laurent Besacier. 2021. An empirical study of end-to-end simultaneous speech translation decoding strategies. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7528–7532. IEEE.
- Alp Öktem, Mireia Farrús, and Antonio Bonafonte. 2019. Prosodic phrase alignment for machine dubbing. *ArXiv*, abs/1908.07226.
- Sara Papi, Marco Gaido, Alina Karakanta, Mauro Cettolo, Matteo Negri, and Marco Turchi. 2022a. Direct speech translation for automatic subtitling. *arXiv preprint arXiv:2209.13192*.
- Sara Papi, Marco Gaido, Matteo Negri, and Andrea Pilzer. 2023a. [Reproducibility is Nothing without Correctness: The Importance of Testing Code in NLP](#). *arXiv preprint arXiv:2303.16166*.
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2021. [Dealing with training and test segmentation mismatch: FBK@IWSLT2021](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 84–91, Bangkok, Thailand (online).
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2022b. [Does simultaneous speech translation need simultaneous models?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 141–153, Abu Dhabi, United Arab Emirates.
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2022c. [Over-generation cannot be rewarded: Length-adaptive average lagging for simultaneous speech translation](#). In *Proceedings of the Third Workshop on Automatic Simultaneous Translation*, pages 12–17, Online.
- Sara Papi, Alina Karakanta, Matteo Negri, and Marco Turchi. 2022d. [Dodging the data bottleneck: Automatic subtitling with automatically segmented ST corpora](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint*

- Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 480–487, Online only.
- Sara Papi, Matteo Negri, and Marco Turchi. 2022e. Attention as a guide for simultaneous speech translation. *arXiv preprint arXiv:2212.07850*.
- Sara Papi, Matteo Negri, and Marco Turchi. 2023b. Alignatt: Using attention-based audio-translation alignments as a guide for simultaneous speech translation. In *Proc. of Interspeech 2023*, Dublin, Ireland.
- 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. Interspeech 2019*, pages 2613–2617.
- Stelios Piperidis, Iason Demiros, Prokopis Prokopidis, Peter Vanroose, Anja Hoethker, Walter Daelemans, Elsa Sklavounou, Manos Konstantinou, and Yannis Karavidas. 2004. [Multimodal, multilingual resources in the subtitling process](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal.
- Peter Polák, Ngoc-Quan Pham, Tuan Nam Nguyen, Danni Liu, Carlos Mullov, Jan Niehues, Ondřej Bojar, and Alexander Waibel. 2022. [CUNI-KIT system for simultaneous speech translation task at IWSLT 2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 277–285, Dublin, Ireland (in-person and online).
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium.
- Yi Ren, Jinglin Liu, Xu Tan, Chen Zhang, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2020. [SimulSpeech: End-to-end simultaneous speech to text translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3787–3796, Online.
- Derek Tam, Surafel M. Lakew, Yogesh Virkar, Prashant Mathur, and Marcello Federico. 2022. [Isochrony-Aware Neural Machine Translation for Automatic Dubbing](#). In *Proc. Interspeech 2022*, pages 1776–1780.
- Ioannis Tsiamas, Gerard I. Gállego, José A. R. Fonollosa, and Marta R. Costa-jussà. 2022. [SHAS: Approaching optimal Segmentation for End-to-End Speech Translation](#). In *Proc. Interspeech 2022*, pages 106–110.
- Yogesh Virkar, Marcello Federico, Robert Enyedi, and Roberto Barra-Chicote. 2021. [Improvements to prosodic alignment for automatic dubbing](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7543–7574.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020. [fairseq s2t: Fast speech-to-text modeling with fairseq](#). In *Proceedings of the 2020 Conference of the Asian Chapter of the Association for Computational Linguistics (ACL): System Demonstrations*.
- Ron J. Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. [Sequence-to-Sequence Models Can Directly Translate Foreign Speech](#). In *Proceedings of Interspeech 2017*, pages 2625–2629, Stockholm, Sweden.
- Patrick Wilken, Panayota Georgakopoulou, and Evgeny Matusov. 2022. [SubER - a metric for automatic evaluation of subtitle quality](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 1–10, Dublin, Ireland (in-person and online).
- Xingshan Zeng, Liangyou Li, and Qun Liu. 2021. [Real-TranS: End-to-end simultaneous speech translation with convolutional weighted-shrinking transformer](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2461–2474, Online.
- Shaolei Zhang and Yang Feng. 2022. [Information-transport-based policy for simultaneous translation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 992–1013, Abu Dhabi, United Arab Emirates.
- Baigong Zheng, Kaibo Liu, Renjie Zheng, Mingbo Ma, Hairong Liu, and Liang Huang. 2020. [Simultaneous translation policies: From fixed to adaptive](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2847–2853, Online.