

Beyond Voice Activity Detection: Hybrid Audio Segmentation for Direct Speech Translation

Marco Gaido^{†,*}, Matteo Negri[†], Mauro Cettolo[†], Marco Turchi[†]

[†]Fondazione Bruno Kessler

^{*}University of Trento

{mgaido, cettolo, negri, turchi}@fbk.eu

Abstract

The audio segmentation mismatch between training data and those seen at run-time is a major problem in direct speech translation. Indeed, while systems are usually trained on manually segmented corpora, in real use cases they are often presented with continuous audio requiring automatic (and sub-optimal) segmentation. After comparing existing techniques (VAD-based, fixed-length and hybrid segmentation methods), in this paper we propose enhanced hybrid solutions to produce better results without sacrificing latency. Through experiments on different domains and language pairs, we show that our methods outperform all the other techniques, reducing by at least 30% the gap between the traditional VAD-based approach and optimal manual segmentation.

1 Introduction

Speech-to-text translation (ST) consists in translating utterances in one language into text in another language. From the architectural standpoint, ST systems are traditionally divided in cascade and direct. Cascade solutions first transcribe the audio via automatic speech recognition (ASR) and then translate the generated transcripts with a machine translation (MT) component. In direct ST, a single end-to-end model operates without intermediate representations. This allows reducing error propagation and latency, as well as exploiting more information (e.g. speaker’s vocal traits and prosody).

Different from MT, where sentence-level splits represent a natural (though not necessarily optimal) input segmentation criterion, handling audio data is more problematic. Existing training corpora (Cattoni et al., 2021; Iranzo-Sánchez et al., 2020) split continuous speech into utterances according to strong punctuation marks in the transcripts (which are known in advance), reflecting linguistic criteria

related to sentence well-formedness. This (*manual*) segmentation is optimal, as it allows ST systems to potentially generate correct outputs even for languages with different syntax and word order (e.g. subject-verb-object vs subject-object-verb). At run-time, though, audio transcripts are not known in advance and *automatic* segmentation techniques have to be applied. The traditional approach is to adopt a Voice Activity Detection (VAD) tool to break the audio on speaker silences (Sohn et al., 1999), considered as a proxy of clause boundaries. However, since the produced segmentation is not driven by syntactic information (unlike that of the training corpora), final performance on downstream tasks degrades considerably (Sinclair et al., 2014).

The impact of a syntax-unaware segmentation can be limited in cascade systems by means of dedicated components that re-segment the ASR transcripts, so to feed MT with well-formed sentences (Matusov et al., 2006). The absence of intermediate transcripts makes this solution unfeasible for direct systems, whose performance is therefore highly sensitive to sub-optimal audio segmentation. This has been shown in the 2020 IWSLT evaluation campaign (Ansari et al., 2020), where the best direct ST system had a key feature in the segmentation algorithm (Potapczyk and Przybysz, 2020). In the same evaluation setting, the second-best direct system (Bahar et al., 2020) exploited an external ASR model to segment the audio (with a +10% BLEU gain compared to its VAD-based counterpart). This solution, however, formally makes it closer to a cascade architecture, losing the advantage of the reduced latency of direct systems. For this reason, while in §4 we compare with the state-of-the-art method proposed in (Potapczyk and Przybysz, 2020), we will not consider approaches needing additional models (e.g. ASR) like the one in (Bahar et al., 2020).

So far, no work analyzed in depth the strengths

and weaknesses of different audio segmentation methods in the context of direct ST. To fill this gap, we study the behavior of the existing techniques and, based on the resulting observations, we propose improved hybrid methods that can also be applied to streaming audio. Through experiments in two domains (TED and European Parliament talks) and two target languages (German and Italian), we show that our solutions outperform the others in all conditions, reducing the gap with optimal manual segmentation by at least 30% compared to VAD systems.

2 Audio Segmentation Methods

2.1 Existing Methods

VAD systems. VAD tools are classifiers that determine whether a given audio frame contains speech or not. Based on this, a VAD-based segmentation considers a sequence of consecutive speech frames as a segment, filtering out non-speech frames. In this work, we evaluate two widely used open source VAD tools: LIUM (Meignier and Merlin, 2010) and WebRTC’s VAD.¹ For LIUM, we apply the configuration employed in the IWSLT campaign (Ansari et al., 2020). WebRTC takes as parameters the *frame size* (10, 20 or 30ms) and the *aggressiveness mode* (an integer in the range [0, 3], 3 being the most aggressive). We select three configurations based on the segmentation they produce on the MuST-C test set (Cattoni et al., 2021), one of the test sets used in our experiments (see §4). Specifically, we consider those not generating too many (more than two times the segments of the manual segmentation) or too long segments (more than 60s). They are: (3, 30ms), (2, 20ms) and (3, 20ms). The statistics computed for the two segmentation tools on the MuST-C test set are presented in Table 1, along with those corresponding to manual segmentation. To better understand the impact of different VADs on translation quality, the tools are compared on MuST-C and Europarl-ST data. Table 2 reports preliminary translation results for English-German (en-de) and English-Italian (en-it), obtained with the systems described in §3. LIUM and the most aggressive WebRTC configuration (3, 20ms) are significantly worse than the other two WebRTC configurations. As (2, 20ms) achieves comparable BLEU performance to (3, 30ms) on MuST-C and better on Europarl-ST, it is used in

¹<http://webrtc.org/>. We use the Python interface <http://github.com/wiseman/py-webrtcvad>.

the rest of the paper.

System	Man.	LIUM	WebRTC		
			3	2	3
Aggress.			30ms	20ms	20ms
Frame size					
% filtered	14.66	0.00	11.27	9.53	15.58
Num segm.	2,574	2,725	3,714	3,506	5,005
Max len (s)	51.97	18.63	48.84	58.62	46.76
Min len (s)	0.05	2.50	0.60	0.40	0.40
Avg len (s)	5.82	6.44	4.19	4.53	2.96

Table 1: Statistics for different segmentations of the MuST-C test set. “Man.” = sentence-based segmentation. “% filtered” = percentage of audio discarded.

VAD System	MuST-C		Europarl-ST	
	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)
English-German				
LIUM	19.55	76.21	15.39	94.06
WebRTC 3, 30ms	21.90	66.96	16.23	89.35
WebRTC 3, 20ms	19.48	72.25	14.07	99.32
WebRTC 2, 20ms	21.87	66.72	18.51	78.12
English-Italian				
LIUM	21.29	67.50	18.88	73.73
WebRTC 3, 30ms	22.46	64.99	19.85	72.28
WebRTC 3, 20ms	20.09	68.62	17.35	78.18
WebRTC 2, 20ms	22.34	66.12	20.90	69.54

Table 2: Results of the VAD systems on MuST-C and Europarl-ST for en-de and en-it.

Fixed-length. A simple approach is splitting the audio at a predefined segment length (Sinclair et al., 2014), without considering the content. In contrast with VAD, this naive method has the benefit of ensuring that the resulting segments are not too long or too short, which are typically hard conditions for ST systems. However, the split points are likely to break sentences in critical positions such as between a subject and a verb or even in the middle of a word. Unlike VAD, this method does not filter the non-speech frames from the input audio, which is entirely passed to the ST system. Fig. 1 shows that, with fixed segmentation, translation quality improves with the duration of the segments (slightly for values ≥ 16 s) up to 20s, after which it decreases. 20 seconds is the maximum segment length in our training data due to memory limits: we can conclude that longer segments produce better translations, but models can effectively translate only sequences whose length does not exceed the maximum observed in the training set.

SRPOL-like segmentation. The method described in (Potapczyk and Przybysz, 2020) takes into account both audio content (silences) and target segments’ length (i.e. the desired length of the generated segments) to split the audio. It recursively divides the audio segments on the longest




Figure 1: BLEU scores (Y axis) with different fixed-length segmentations (in seconds – X axis).

silence, until either there are no more silences in a segment, or the segment itself is shorter than a threshold. It is important to notice that, in (Potapczyk and Przybysz, 2020), the silences are detected with a manual operation, making the approach hard to reproduce and not scalable. In this paper, we replicate the logic, but we rely on Webrtc to automatically identify silences. For this reason, our results might be slightly different than the original ones, but the segmentation is automatic and easy to reproduce. Another major problem of this method is that it requires the full audio to be available for splitting it. So it is not applicable to audio streams and online use cases. Based on the previous considerations drawn from Fig. 1, in our experiments we set the maximum length threshold to 20s, so that the model is fed with sequences that are not longer than the maximum seen at training time. The resulting segments have an average length of 7-8s.

2.2 Proposed hybrid segmentation

Similar to (Potapczyk and Przybysz, 2020), our method is hybrid as it considers both the audio content and the target segments’ length. However, unlike (Potapczyk and Przybysz, 2020), we give more importance to the target segments’ length than to the detected pauses (we motivate this choice in §4). Specifically, we split on the longest pause in the interval (minimum and maximum length), if any, otherwise we split at maximum length. Maximum and minimum segment lengths are controlled by two hyper-parameters (MAX_LEN and MIN_LEN). Unlike the SRPOL-like approach, ours can operate on audio streams, as it does not require the full audio to start the segmentation procedure. Moreover, the latency is controlled by MAX_LEN and MIN_LEN , which can be tuned to trade translation quality for lower latency.

We tested different values for MIN_LEN and we chose 17s for our experiments, because it resulted in the best score on the MuST-C dev set. As in the other methods, and for the same reasons, MAX_LEN is set to 20s. The resulting segments have an average length slightly higher than 17s.

We also introduce a variant of this method that enforces splitting on pauses longer than 550ms. In (Karakanta et al., 2020), this threshold is shown to often represent a *terminal juncture*: a break between two utterances, usually corresponding to clauses. Splitting on such pauses should hence enforce separating different clauses. As a result, segments can be shorter than MIN_LEN , but we still ensure they are not longer than MAX_LEN . With this variant, the segments are much shorter, as their average length is 8s, similar to the SRPOL-like technique.

3 Experimental Settings

We use a Transformer (Vaswani et al., 2017) whose encoder is modified for ST. The encoder starts with two 2D convolutional layers that reduce the length of the Mel-filter-bank sequence by a factor of 4. The resulting tensors are passed to a linear layer that maps them into the dimension used by the following encoder Transformer layers. A logarithmic distance penalty (Di Gangi et al., 2019) is applied in all the encoder Transformer layers.

The ST models have 11 encoder Transformer layers and 4 decoder Transformer layers. We use 8 attention heads, 512 attention hidden units and 2,048 features in the FFNs’ hidden layer. We set dropout to 0.1. The optimizer is Adam (Kingma and Ba, 2015) with betas (0.9, 0.98). The learning rate is scheduled with inverse square root decay after 4,000 warm-up updates, during which it increases linearly from $3 \cdot 10^{-4}$ up to $5 \cdot 10^{-3}$. The update frequency is set to 8 steps; we train on 8

Segm. method	MuST-C en-de		Europarl en-de		MuST-C en-it		Europarl en-it	
	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)
Manual segm.	27.55	58.84	26.61	60.99	27.70	58.72	28.79	59.16
Best VAD	21.87	66.72	18.51	78.12	22.34	66.12	20.90	69.54
Best Fixed (20s)	23.86	61.29	23.27	64.01	23.20	64.24	22.28	64.57
SRPOL-like	22.26	71.10	20.49	77.61	23.12	66.27	23.26	66.19
Pause in 17-20s	24.39	61.35	23.78	63.15	23.50	63.76	22.86	63.44
+ force split	23.17	66.20	22.52	68.56	23.45	63.79	24.15	63.31

Table 3: Comparison between manual and automatic segmentations: VAD, fixed-length and hybrid approaches.

GPUs and each mini-batch is limited to 12,000 tokens or 8 sentences, so the resulting batch size is slightly lower than 512. We initialize the convolutional and the first encoder Transformer layers with the encoder of a model trained on ASR data. We pre-train our ST models on the ASR corpora with synthetic targets generated by an MT model fed with the known transcripts (Jia et al., 2019) and we fine-tune on the ST corpora. Both these trainings adopt knowledge distillation (Hinton et al., 2015) with the MT model as teacher (Liu et al., 2019; Gaido et al., 2020). Finally, we fine-tune on the ST corpora with label-smoothed cross-entropy (Szegedy et al., 2016). In all the three steps, we use SpecAugment (Park et al., 2019) and time stretch (Nguyen et al., 2020) as data augmentation techniques.

The ASR model is similar to ST models, but we use 8 encoder layers and 6 decoder layers. For MT, instead, the Transformer attentions has 16 heads and hidden-layer features are two times those of ST and ASR models.

We experimented with translation from English speech into two target languages: German and Italian. To train the MT model used for knowledge distillation, we employed the WMT 2019 datasets (Barrault et al., 2019) and the 2018 release of OpenSubtitles (Lison and Tiedemann, 2016) for English-German and OPUS (Tiedemann, 2016) for English-Italian. All the data were cleaned with Modern MT (Bertoldi et al., 2017). The ASR model, whose encoder was used to initialize that of the ST model, was trained on TED-LIUM 3 (Hernandez et al., 2018), Librispeech (Panayotov et al., 2015), Mozilla Common Voice², How2 (Sanabria et al., 2018), and the audio-transcript pairs of the ST corpora. The ST corpora were MuST-C (Cattani et al., 2021) and Europarl-ST (Iranzo-Sánchez et al., 2020) for both target languages. We filtered out samples with input audio longer than 20s to avoid out-of-memory errors. The text is encoded

²<https://voice.mozilla.org/>

using BPE (Sennrich et al., 2016) with 8,000 merge rules (Di Gangi et al., 2020).

4 Results

We compute ST results in terms of BLEU (Papineni et al., 2002) and TER (Snover et al., 2006) on the MuST-C and Europarl-ST test sets for en-de and en-it. In MuST-C the two test sets are identical with regard to the audio, while in Europarl-ST they contain different recordings.

As shown in Table 3, fixed-length segmentation always outperforms the best VAD, both in terms of BLEU and TER. This may be surprising but it confirms previous findings in ASR (Sinclair et al., 2014): also in ST, VAD is more costly and less effective than a naive fixed-length segmentation. Besides, it suggests that the resulting segments’ length is more important than the precision of the split times. This observation motivates the definition of our proposed techniques. Compared to fixed-length segmentation, the SRPOL-like method provides better results for en-it, but worse for en-de, indicating that the syntactic properties of the source and target languages are an important factor for audio segmentation (see §5).

Our proposed method (*Pause in 17-20s* in Table 3) outperforms the others on all test sets but Europarl en-it, in which SRPOL-like has a higher BLEU (but worse TER). The version with forced splits on 550ms pauses is inferior to the version without forced splits on the German test sets, but it is on par for MuST-C en-it and superior on Europarl en-it, on which it is the best segmentation overall by a large margin. Moreover, its scores are always better than the ones obtained by the SRPOL-like approach, although the length of the produced segments is similar. These results suggest that, although the best version depends on the syntax and the word order of the source and target languages, our method can always outperform the others in terms of both BLEU and TER. Noticeably, it does not introduce latency, since it does not require the




Figure 2: Z-score normalized output lengths (number of words) according to the input segments length.

(a) Hallucinations with non-speech audio	
Audio	<i>Music and applause.</i>
4s seg-ments	[Chinesisch] [Hawaiianischer Gesang] // Chris Anderson: Du bist ein Idiot. // Nicole: Nein. <i>[Chinese] [Hawaiian song] // Chris Anderson: You are an idiot. // Nicole: No.</i>
(b) Hallucinations with sub-sentential utterances	
Audio	Now, chimpanzees are well-known for their aggression. // (Laughter) // But unfortunately, we have made too much of an emphasis of this aspect (...)
Reference	Schimpansen sind bekannt für ihre Aggressivität. // (Lachen) // Aber unglücklicherweise haben wir diesen Aspekt überbetont (...)
4s seg-ments	Publikum: Nein. Schimpansen sind bekannt. // Ich bin für ihre Aggression gegangen. // Aber leider haben wir zu viel Coca-Cola gemacht. // Das ist eine wichtige Betonung dieses Aspekts (...) <i>Audience: No.</i> Chimpanzees are known. // I went for their aggression. // But unfortunately we made too much Coca-Cola. // This is an important emphasis of this aspect (...)
20s seg-ments	Schimpansen sind bekannt für ihre Entwicklung. // Aber leider haben wir zu viel Schwerpunkt auf diesem Aspekt (...) <i>Chimpanzees are known for their development. // But unfortunately, we have expressed too much emphasis on this aspect (...)</i>
(c) Hallucinations and bad translation with sub-sentential utterances	
Audio	(...) where the volunteers supplement a highly skilled career staff, you have to get to the fire scene pretty early to get in on any action.
Reference	(...) in der Freiwillige eine hochqualifizierte Berufsfeuerwehr unterstützten, muss man ziemlich früh an der Brandstelle sein, um mitmischen zu können.
4s seg-ments	(...) wo die Bombenangriffe auf dem Markt waren. // Man muss bis zu 1.000 Angestellte in die USA, nach Nordeuropa kommen. <i>(...) where the bombings were on the market. // You have to come up to 1,000 employees in the USA, to Northern Europe.</i>
20s seg-ments	(...) in der die Freiwilligen ein hochqualifiziertes Karriere-Team ergänzen, muss man ziemlich früh an die Feuerszene kommen, um in irgendeiner Aktion zu gelangen. <i>(...) where the volunteers complement a highly qualified career team, you have to get to the fire scene pretty early in order to get into any action.</i>
(d) Final portions of long segment ignored	
Audio	But still it was a real footrace against the other volunteers to get to the captain in charge to find out what our assignments would be. // When I found the captain, (...)
Reference	Aber es war immer noch ein Wettrennen gegen die anderen Freiwilligen um den verantwortlichen Hauptmann zu erreichen und herauszufinden was unsere Aufgaben sein würden. // Als ich den Hauptmann fand (...)
22s seg-ments	(...) Es war immer noch ein echtes Fussrennen gegen die anderen Freiwilligen. // Als ich den Kapitän fand, (...) <i>(...) It was still a real footrace against the other volunteers. // When I found the captain, (...)</i>

Table 4: Translations affected by errors caused by too short – (a), (b), (c) – or too long – (d) – segments. The symbol “//” refers to a break between two segments. The breaks might be located in different positions in the different segmentations. Over-generated – in examples (a), (b), (c) – and missing – in (d) – content is marked in **bold** respectively in system’s outputs and in the reference.

full audio to be available for splitting it, as the SR-POL-like technique does. In particular, averaged on the two domains, our best results (respectively with and without *forced splits*) reduce the gap with the manual segmentation by 54.71% (en-de) and 30.95% (en-it) compared to VAD-based segmentation.

5 Analysis

A first interesting consideration regards the length of the produced translations. In particular, we analyze the case of fixed-length segmentation (see Fig. 2): in presence of short input segments the output is longer, while it gets shorter in case of seg-

ments longer than 20s. To understand this behavior, we performed a manual inspection of the German translations produced by fixed-length segmentation with 4s, 20s and 22s.

The analysis revealed two main types of errors: overly long (*hallucinations* (Lee et al., 2018)) and overly short outputs. The first type of error occurs when the system is fed with small, sub-sentential segments. In this case, trying to generate well-formed sentences, the system “completes” the translation with text that has no correspondence with the input utterance. The second type of error occurs when the system is fed with segments that exceed the maximum length observed in the training data. In this case, part of the input (even complete clauses, typically towards the end of the utterance) is not realized in the final translation.

Table 4 provides examples of all these phenomena. The first three examples show cases of hallucinations in short (4s) segments, while the last one shows an incomplete translation of a long (22s) segment. In particular:

(a) shows the generation of text not related to the source when the audio contains only noise or silence (e.g. at the beginning of a TED talk recording).

(b) presents the addition of non-existing content in the translation of a sub-sentential segment.

(c) is related to a sub-sentential utterance as well, but in this case the output of the system is affected by both hallucinations and poor translation quality due to the lack of enough context.

(d) reports a segment whose last portion ignored.

The length of the generated outputs also helps understanding the different results obtained by the variants of our method on the two target languages. Indeed, the introduction of forced splits (+ *force split*) produces audio segments that are much shorter ($\sim 8s$ vs $\sim 17s$) and hence, according to the previous consideration, the resulting translation is overall longer. For German, the difference in terms of output length is high ($> 8.5\%$), while for Italian it is much lower (4.33% on MuST-C and 2.49% on Europarl-ST). So, the German results are penalized by the additional hallucinations, while, for Italian translations, the beneficial separation of clauses delimited by terminal juncture dominates.

This different behavior relates to the different syntax of the source and target languages. Indeed, translating from English (an SVO language) into German (an SOV language) requires long-range re-

orderings (Gojun and Fraser, 2012), which can also span over sub-clauses. The Italian phrase structure, instead, is more similar to English. This is confirmed by the shifts counted in TER computation, which are 20% more in German than in Italian. Moreover, in Italian their number does not change between our method with and without forced splits, while in German the version with forced splits has 5-10% more shifts.

6 Conclusions

We studied different segmentation techniques for direct ST. Despite its wide adoption, VAD-based segmentation resulted to be underperforming. We showed that audio segments’ length is a crucial factor to obtain good translations and that the best segmentation approach depends on the structural similarity between the source and target languages. In particular, we demonstrated that the resulting segments should be neither longer than the maximum length of the training samples nor too short (especially when the target language has a different structure). Inspired by these findings, we proposed two variants of a hybrid method that significantly improve on different test sets and languages over the VAD baseline and the other techniques presented in literature. Our approach was designed to be also applicable to audio streams and to allow controlling latency, hence being suitable even for online use cases.

References

- 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 Chanchang Wang. 2020. **FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN**. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34, Online. Association for Computational Linguistics.
- Parnia Bahar, Patrick Wilken, Tamer Alkhouli, Andreas Guta, Pavel Golik, Evgeny Matusov, and Christian Herold. 2020. Start-Before-End and End-to-End: Neural Speech Translation by AppTek and RWTH Aachen University. In *Proceedings of 17th International Workshop on Spoken Language Translation (IWSLT)*, Virtual.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Gra-

- ham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 conference on machine translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Nicola Bertoldi, Roldano Cattoni, Mauro Cettolo, Amin Farajian, Marcello Federico, Davide Caroselli, Luca Mastrostefano, Andrea Rossi, Marco Trombetti, Ulrich Germann, and David Madl. 2017. MMT: New Open Source MT for the Translation Industry. In *Proceedings of the 20th Annual Conference of the European Association for Machine Translation (EAMT)*, pages 86–91, Prague, Czech Republic.
- Roldano Cattoni, Mattia A. Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. [MuST-C: A multilingual corpus for end-to-end speech translation](#). *Computer Speech & Language*, 66:101155.
- Mattia A. Di Gangi, Marco Gaido, Matteo Negri, and Marco Turchi. 2020. [On target segmentation for direct speech translation](#). In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 137–150, Virtual. Association for Machine Translation in the Americas.
- Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2019. [Adapting Transformer to End-to-End Spoken Language Translation](#). In *Proceedings of Interspeech 2019*, pages 1133–1137.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2020. [On Knowledge Distillation for Direct Speech Translation](#). In *Proceedings of the Seventh Italian Conference on Computational Linguistics, CLiC-it 2020, Bologna, Italy, March 1-3, 2021*, volume 2769 of *CEUR Workshop Proceedings*.
- Anita Gojun and Alexander Fraser. 2012. Determining the placement of German verbs in English-to-German SMT. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 726–735, Avignon, France.
- 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 *Proceedings of the Speech and Computer - 20th International Conference (SPECOM)*, pages 198–208, Leipzig, Germany. Springer International Publishing.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. In *Proceedings of NIPS Deep Learning and Representation Learning Workshop*, Montréal, Canada.
- Javier Iranzo-Sánchez, Joan A. Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Giménez. Adrià, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. Europarl-ST: A Multilingual Corpus For Speech Translation Of Parliamentary Debates. In *Proceedings of 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233, Barcelona, Spain.
- 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 *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7180–7184, Brighton, UK.
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2020. [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.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of 3rd International Conference on Learning Representations (ICLR)*, San Diego, California.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fannjiang, and David Sussillo. 2018. Hallucinations in Neural Machine Translation. In *Proceedings of NIPS 2018 Interpretability and Robustness for Audio, Speech and Language Workshop*, Montréal, Canada.
- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles. In *Proceedings of the Tenth International Language Resources and Evaluation Conference (LREC)*, pages 923–929, Portoroz, Slovenia.
- Yuchen Liu, Hao Xiong, Jiajun Zhang, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2019. [End-to-End Speech Translation with Knowledge Distillation](#). In *Proceedings of Interspeech 2019*, pages 1128–1132, Graz, Austria.
- Evgeny Matusov, Arne Mauser, and Hermann Ney. 2006. Automatic Sentence Segmentation and Punctuation Prediction for Spoken Language Translation. In *Proceedings of International Workshop on Spoken Language Translation (IWSLT) 2006*, pages 158–165, Kyoto, Japan.
- Sylvain Meignier and Teva Merlin. 2010. LIUM SpkDiarization: An Open Source Toolkit For Diarization. In *Proceedings of the CMU SPUD Workshop*, Dallas, Texas.
- Thai-Son Nguyen, Sebastian Stueker, Jan Niehues, and Alex Waibel. 2020. Improving Sequence-to-sequence Speech Recognition Training with On-the-fly Data Augmentation. In *Proceedings of the 2020*

- International Conference on Acoustics, Speech, and Signal Processing – IEEE-ICASSP-2020*, Barcelona, Spain.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: An ASR corpus based on public domain audio books. In *Proceedings of 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, South Brisbane, Queensland, Australia.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania.
- 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 *Proceedings of Interspeech 2019*, pages 2613–2617, Graz, Austria.
- Tomasz Potapczyk and Pawel Przybysz. 2020. [SR-POL’s system for the IWSLT 2020 end-to-end speech translation task](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 89–94, Online. Association for Computational Linguistics.
- 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 *Proceedings of Visually Grounded Interaction and Language (ViGIL)*, Montréal, Canada. Neural Information Processing Society (NeurIPS).
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural Machine Translation of Rare Words with Subword Units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Mark Sinclair, Peter Bell, Alexandra Birch, and Ferguson McInnes. 2014. A semi-Markov model for speech segmentation with an utterance-break prior. In *Proceedings of Interspeech 2014*, pages 2351–2355, Singapore.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A Study of Translation Edit Rate with Targeted Human Annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas*, pages 223–231, Cambridge.
- Jongseo Sohn, Nam Soo Kim, and Wonyong Sung. 1999. A statistical model-based voice activity detection. *IEEE Signal Processing Letters*, 6(1):1–3.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, Las Vegas, Nevada, United States.
- Jörg Tiedemann. 2016. OPUS – Parallel Corpora for Everyone. *Baltic Journal of Modern Computing*, page 384. Special Issue: Proceedings of the 19th Annual Conference of the European Association of Machine Translation (EAMT).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of Advances in Neural Information Processing Systems 30 (NIPS)*, pages 5998–6008, Long Beach, California.