

# Dodging the Data Bottleneck: Automatic Subtitling with Automatically Segmented ST Corpora

Sara Papi<sup>♣,◇</sup>, Alina Karakanta<sup>♣,◇</sup>, Matteo Negri<sup>♣</sup>, Marco Turchi<sup>♣\*</sup>

<sup>♣</sup>Fondazione Bruno Kessler

<sup>◇</sup>University of Trento

<sup>♣</sup>Zoom Video Communications

{spapi, akarakanta, negri}@fbk.eu, marco.turchi@zoom.us

## Abstract

Speech translation for subtitling (SubST) is the task of automatically translating speech data into well-formed subtitles by inserting subtitle breaks compliant to specific displaying guidelines. Similar to speech translation (ST), model training requires parallel data comprising audio inputs paired with their textual translations. In SubST, however, the text has to be also annotated with subtitle breaks. So far, this requirement has represented a bottleneck for system development, as confirmed by the dearth of publicly available SubST corpora. To fill this gap, we propose a method to convert existing ST corpora into SubST resources without human intervention. We build a segmenter model that automatically segments texts into proper subtitles by exploiting audio and text in a multimodal fashion, achieving high segmentation quality in zero-shot conditions. Comparative experiments with SubST systems respectively trained on manual and automatic segmentations result in similar performance, showing the effectiveness of our approach.

## 1 Introduction

Massive amounts of audiovisual content are available online, and this abundance is accelerating with the spread of online communication during the COVID-19 pandemic. The increased production of pre-recorded lectures, presentations, tutorials and other audiovisual products raises an unprecedented demand for subtitles in order to facilitate comprehension and inclusion of people without access to the source language speech. To keep up with such a demand, automatic solutions are seen as a useful support to the limited human workforce of trained professional subtitlers available worldwide (Tardel, 2020). Attempts to automatise subtitling have focused on Machine Translation for translating human- or automatically-generated source

language subtitles (Volk et al., 2010; Etchegoyhen et al., 2014; Matusov et al., 2019; Koponen et al., 2020). Recently, direct ST systems (Bérard et al., 2016; Weiss et al., 2017) have been shown to achieve high performance while generating the translation in the target language without intermediate transcription steps. For automatic subtitling, Karakanta et al. (2020a) suggested that, by directly generating target language subtitles from the audio (i.e. predicting subtitle breaks together with the translation), the model can improve subtitle segmentation by exploiting additional information like pauses and prosody. However, the scarcity of SubST corpora makes it hard to build competitive systems for automatic subtitling, especially if no corpus is available for specific languages/domains.

One solution to the SubST data bottleneck could be leveraging ST corpora by inserting subtitle breaks on their target side. Automatic segmentation of a text into subtitles is normally implemented with rule-based approaches and heuristics, e.g. a break is inserted before a certain length limit is reached. More involved algorithms (SVM, CRF, seq2seq) predict breaks using a segmenter model trained on subtitling data for a particular language (Álvarez et al., 2016, 2017; Karakanta et al., 2020c). Still, the performance of these models relies on high-quality segmentation annotations for each language, which web-crawled subtitling corpora like OpenSubtitles (Lison et al., 2018) rarely contain.

In this work, we address the scarcity of SubST corpora by developing a multimodal segmenter<sup>1</sup> able to automatically annotate existing ST corpora with subtitle breaks in a zero-shot fashion. Specifically, our segmenter exploits, for the first time in this scenario, the source language audio (here: en) and segmented target text already available in a few languages (here: de, en, fr, it). Its key strength is the ability to segment not only target languages

\* Work done when working at FBK.

<sup>1</sup>Code and model available at <https://github.com/hlt-mt/FBK-fairseq>.

for which high-quality segmented data is available but also unseen languages having some degree of similarity with those covered by the original ST resource(s). This opens up the possibility to automatically obtain synthetic SubST training data for previously not available languages. Along this direction, our zero-shot segmentation results on two unseen languages (es, nl) show that training a SubST system on automatically-segmented data leads to comparable performance compared to using a gold, manually-segmented corpus.

## 2 Methodology

Our method to leverage ST corpora for SubST can be summarized as follows: *i)* we train different segmenters on available human-segmented subtitling data to select the best performing one; *ii)* we run the selected segmenter in a *zero-shot* fashion (i.e. without fine-tuning or adaptation) to insert subtitle breaks in unsegmented text data of *unseen* languages; *iii)* then, the automatically annotated texts are paired with the corresponding audio to obtain a synthetic parallel SubST corpus; *iv)* finally, a SubST model is trained on the synthetic corpus.

We test our method on two language pairs (en-es, en-nl) by comparing the results of SubST models trained on synthetic data with those of identical models trained on original gold data.

### 2.1 Segmenter

We adopt the general segmentation approach of (Karakanta et al., 2020b) where a sequence to sequence *Textual segmenter*, trained on pairs of unsegmented-segmented text, takes unsegmented text as input and inserts subtitle breaks.

To improve segmentation quality, we extend this approach in two ways. Our first extension is multimodal training. Since speech phenomena, such as pauses and silences, can strongly influence the structure of the subtitles (Carroll and Ivarsson, 1998), we expect that information from the speech modality could improve segmentation. To explore this hypothesis, we extend the textual segmenter with a multimodal architecture (Sulubacak et al., 2020), which receives input from different modalities: in our case, audio and text.<sup>2</sup> Our *Multimodal segmenter* is built using an architecture with two encoders: one for the text (with the same structure

<sup>2</sup>Images and videos with subtitling material are often protected by copyright and thus not publicly available. Improving the segmenter with data from the visual modality is thus left to future work depending on the availability of such resources.

as the textual segmenter) and one for the audio. We combine the encoder states obtained by the two encoders using parallel cross-attention (Bawden et al., 2018),<sup>3</sup> as it proved to be effective both in speech and machine translation (Kim et al., 2019; Gaido et al., 2020). Parallel attention (Figure 1) is computed by attending at the same intermediate representation (the decoder self-attention); then, the audio encoder cross-attention and the text encoder cross-attention are summed together and fed to the feed-forward layer.

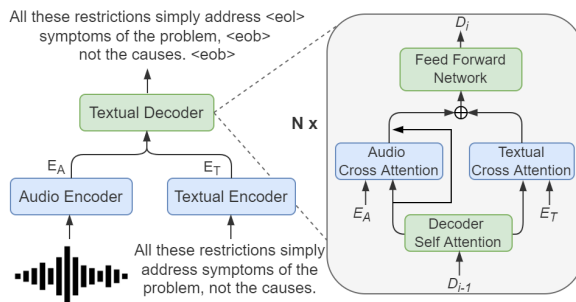


Figure 1: Parallel Multimodal segmenter architecture.

Since subtitling constraints are the same across several languages, our second extension is to learn segmentation multilingually. To this aim, we follow standard approaches used in MT and ST, respectively (see Appendix A for more details): for the textual segmenter we combine samples from multiple languages in the same training step (Ott et al., 2018); for the multimodal segmenter we add a prefix language token to the target text (Inaguma et al., 2019). As in MT (Ha et al., 2016), multilingual training in ST has been shown to enhance performance (Wang et al., 2020) while maintaining only one model for multiple languages.

### 2.2 Data, Baselines and Evaluation

**Data.** To train our textual and multimodal segmenters, we use en→{de/fr/it} sections of MuST-Cinema (Karakanta et al., 2020b), the only publicly available SubST dataset. More details about the data selection are provided in Appendix A. To test the segmenters in zero-shot conditions (Section 4) and train our SubST models (Section 5), we select two target languages also contained in MuST-Cinema:<sup>4</sup> Dutch (an SOV – Subject-Verb-Object – language) and Spanish (SVO). Using the corpus

<sup>3</sup>We also tested sequential cross-attention (Zhang et al., 2018) but do not report results since they are slightly worse compared to parallel cross-attention.

<sup>4</sup>Though present in MuST-Cinema, es and nl data are only used for testing purposes so as to simulate the zero-shot con-

notation, subtitle breaks are defined as: *block break* <eob>, which marks the end of the current subtitle displayed on screen, and *line break* <eol>, which splits consecutive lines inside the same block.

**Baselines.** We compare the performance of the segmenters with two baselines. One is a rule-based method (*Count Chars*) where a break is inserted before a 42-character limit. This is the simplest method to always produce length-conforming subtitles and serves as a lower bound for segmentation performance. Our second baseline (*Supervised*) is a neural textual segmenter trained on OpenSubtitles, the largest collection of publicly available textual subtitling data, for the respective language (es, nl). Although OpenSubtitles is available for a variety of languages, it has some limitations: it does not contain audio, the subtitle and segmentation quality varies since subtitles are often machine-translated or created by non-professionals, and line breaks were lost when pre-processing the subtitles to create the corpus. These limitations may have a detrimental effect on the quality of segmenters trained on this data (Karakanta et al., 2019). Complete details on experimental settings are presented in Appendix A.

**Evaluation.** To evaluate both the quality of the SubST output and the accuracy of our segmenters, we resort to reference-based evaluation. For translation quality of the SubST output we use BLEU (Post, 2018)<sup>5</sup>, computed on the text from which the subtitle breaks are removed. For segmentation accuracy, we use *Sigma* (Karakanta et al., 2022), a novel subtitle segmentation metric based on BLEU. Sigma is the ratio of the segmentation achieved for a given text to the best segmentation that could be achieved. Contrary to other standard segmentation metrics, such as F1, it can be computed when the output text is different than the reference text. To ensure that the system does not over- or under-generate subtitle breaks, we additionally report *Break coverage* computed as follows:

$$Coverage(\%) = \left( \frac{\#\langle break \rangle_{pred}}{\#\langle break \rangle_{ref}} \cdot 100 \right) - 100$$

where <break> corresponds to either <eol> or <eob>. EOL and EOB coverage obtains negative values when the segmenter inserts less breaks than required or positive values when it inserts more. Lastly, we use *length conformity* (or characters

ditions required to select the best segmenter and evaluate our SubST systems.

<sup>5</sup>BLEU+c.mixed+#.1+s.exp+tok.13a+v.1.5.1

per line – CPL), corresponding to the percentage of subtitles not exceeding the allowed maximum length of 42 CPL, as per TED guidelines.<sup>6</sup>

### 3 Segmentation on *seen* languages

We train the mono/multi-lingual versions of our *Textual/Multimodal* segmenters for the four languages (de, en, fr, it) and measure their performance in terms of Sigma and CPL. The results are shown in Table 1.

Looking at the Sigma values, both the *Textual* and the *Multimodal* segmenter perform better than the rule-based baseline, despite a small drop in CPL. The *Multimodal* segmenter always outperforms the *Textual* one by 2 Sigma points on average and inserts break symbols more accurately. Moreover, it benefits from multilingual training on all languages. In contrast, overall subtitle conformity is higher for the *Textual* segmenter in 3 out of 4 languages, where its CPL scores are 1.2-2.6 percentage points above those obtained by the *Multimodal* one. In addition, except for one case (German), higher CPL values are obtained with monolingual training.

### 4 Zero-shot segmentation

Aiming to build a SubST model for unseen languages (es, and nl), we first select the best segmenter for generating synthetic en→es/nl data. As shown in Table 2, all the models that receive only text as input (*Count Chars*, *Supervised* and *Textual*) achieve low segmentation performance, with Sigma ranging between 63-75. The zero-shot *Textual* segmenter achieves higher segmentation quality compared to the *Count Chars* and *Supervised* baselines by 10 points. However, its main drawback is the inability to copy the actual text, as shown by the BLEU values of 61 for nl and 69 for es. In this respect, the baselines perform much better. Despite being trained on subtitling data for the particular language, the low segmentation performance of *Supervised* can be attributed to the different domain compared to the MuST-Cinema test set. For example, MuST-Cinema mainly contains long sentences with multiple breaks, while in OpenSubtitles we rarely come across sentences with more than three breaks. Moreover, both *Supervised* and *Textual* generate subtitles conforming to the CPL constraint in only 70% of the cases, despite

<sup>6</sup><https://www.ted.com/participate/translate/subtitling-tips>

Segmenter	Training	English		French		German		Italian	
		Sigma	CPL	Sigma	CPL	Sigma	CPL	Sigma	CPL
Count Chars		63.71	100%	62.87	100%	62.34	100%	61.49	100%
Textual	mono	84.87	<b>96.6%</b>	83.68	<b>96.7%</b>	83.62	90.9%	82.22	90.0%
	multi	85.98	88.5%	84.56	94.3%	84.02	90.9%	83.04	<b>91.2%</b>
Multimodal	mono	85.76	94.8%	84.25	93.9%	84.22	<b>91.4%</b>	82.62	89.9%
	multi	<b>87.44</b>	95.0%	<b>86.49</b>	94.1%	<b>86.4</b>	89.9%	<b>85.33</b>	90.0%

Table 1: Segmentation results on *seen* languages.

having received only length-conforming subtitles as training data. The negative values of EOL and EOB coverage show that all textual methods under-generate subtitle breaks. From these results we can conclude that zero-shot segmentation does not perform satisfactorily with textual input only.

Dutch					
Segmenter	BLEU	Sigma	CPL	EOL	EOB
Count Chars	<b>100</b>	63.2	<b>100%</b>	<b>-21.2%</b>	-7.1%
Supervised	89.5	64.4	71.2%	-31.4%	-51.3%
Textual	61.3	74.4	77.8%	-23.4%	-9.9%
Multimodal	99.9	<b>80.3</b>	91.4%	-27.2%	<b>+0.4%</b>
Spanish					
Segmenter	BLEU	Sigma	CPL	EOL	EOB
Count Chars	<b>100</b>	63.2	<b>100%</b>	-24.6%	<b>-4.4%</b>
Supervised	92.6	64.1	71.2%	-32.3%	-45.4%
Textual	69.6	75.8	70.1%	-47.6%	-19.3%
Multimodal	99.6	<b>78.7</b>	91.8%	<b>-22.4%</b>	+4.7%

Table 2: Segmentation results on *unseen* languages.

In comparison, the *Multimodal* segmenter performs significantly better. It reaches an absolute gain of 6.1 Sigma points for nl and 2.9 for es compared to *Textual*. Moreover, contrary to *Textual* and *Supervised*, the *Multimodal* model learnt to perfectly copy the text, as shown by the high BLEU scores (up to 99.9 on nl), close to the maximum score of a method – *Count Chars* – that by design does not change the original text. The CPL results are in agreement with BLEU: for both languages, the *Multimodal* model respects the length constraint in more than 91% of the subtitles. Strikingly, even if the two target languages were never seen by the model, these results are similar to those obtained on seen languages (see Table 1). Unlike the rest of the models, *Multimodal* is the only model which does not under-generate  $\langle eob \rangle$ . This is in line with the results of Karakanta et al. (2020a), who showed that exploiting the audio in ST is beneficial for inserting subtitle breaks ( $\langle eob \rangle$ , for instance, typically corresponds to longer speech pauses). The results are more discordant for the EOL Coverage. On es, *Multimodal* shows a lower tendency to under-generate, while on nl both mod-

els fail to insert at least the 23.4% of  $\langle eol \rangle$ . We assume this phenomenon is caused by the lower frequency of  $\langle eol \rangle$  in the corpus, since a subtitle can be composed of only one line, as well as by the higher difficulty in placing the break for which the system cannot resort to speech clues (e.g. pauses).

**Ablation.** To test the effectiveness of the *Multimodal* model also in the absence of similar languages in the training set, we train it on a limited set of Latin languages (Italian and French) and test it on Dutch, which is a Germanic language.

The results (*fr*, *it only*) are shown in Table 3. Even if trained on only two languages from a different language group, the *fr*, *it only Multimodal* model shows competitive results. In terms of segmentation, there is only a slight degradation of 3 Sigma points compared to the full multilingual *Multimodal* model and a 3.6% drop in CPL conformity, which could be attributed to a lower EOL coverage. However, it is still significantly better in terms of Sigma, CPL conformity and EOB coverage compared to all the other segmenters (*Count Chars*, *Supervised*, and *Textual*). In terms of changes to the text, as show by BLEU, it is on par with *Supervised*, a model trained only on Dutch subtitles, and better than the *Textual* by 25 BLEU points. The presence of related languages seems to help the model better copy the text, since the main drop compared to the full *Multimodal* model is in terms of BLEU. Overall, we can conclude that the presence of related languages in the training set can enhance the performance, but the segmentation accuracy and conformity are only minimally affected. The results obtained by the *fr*, *it only Multimodal* confirm the ability and superiority of this model in segmenting texts on unseen languages also belonging to different language groups.

**Limitations.** So far, our results indicate the effectiveness of *Multimodal* segmentation to automatically turn existing ST corpora into SubST training data. In addition, at least for the Western European languages considered in our experiments, our

Segmenter	BLEU	Sigma	CPL	EOL	EOB
Count Chars	<b>100</b>	63.2	<b>100%</b>	<b>-21.2%</b>	-7.1%
Supervised	89.5	64.4	71.2%	-31.4%	-51.3%
Textual	61.3	74.4	77.8%	-23.4%	-9.9%
Multimodal	99.9	<b>80.3</b>	91.4%	-27.2%	<b>+0.4%</b>
-fr, it only	88.9	77.0	87.8%	-34.8%	<b>-0.4%</b>

Table 3: Ablation results on MuST-Cinema amara en→nl. All but the last line are from Table 2.

approach can be successfully applied in zero-shot settings involving languages not present in the training data which also belong to different language groups. Not being possible to verify due to the lack of suitable benchmarks, the possibility of porting our approach to scenarios involving different alphabets is not verified in this work. This would require, at least, a vocabulary adaptation which represents a well-known problem in multilingual approaches to MT/ST (Garcia et al., 2021). Nevertheless, even in the worst case in which some degree of similarity across languages is required for zero-shot automatic segmentation, we believe that these results indicate a viable path towards overcoming the scarcity of SubST resources. In the next section, we will test this hypothesis.

## 5 SubST with Synthetic Data

Since our *Multimodal* segmenter achieves the best performance overall, we use it to automatically generate the synthetic counterpart of the en→{es, nl} sections of MuST-Cinema. The resulting data is respectively used to train two SubST systems. The goal is to achieve comparable performance to that of similar models trained on manually segmented subtitles. For this purpose, using the same architecture, we also train two systems on the original manual segmentations of MuST-Cinema.

Dutch					
Data	BLEU	Sigma	CPL	EOL	EOB
Original	<b>25.3*</b>	<b>81.58</b>	91.2%	-36.8%	+8.0%
Synthetic	24.3*	75.52	<b>94.7%</b>	<b>-20.4%</b>	<b>+4.8%</b>
Spanish					
Data	BLEU	Sigma	CPL	EOL	EOB
Original	<b>30.7*</b>	<b>79.21</b>	<b>96.7%</b>	<b>-10.0%</b>	+10.9%
Synthetic	<b>30.7*</b>	77.84	94.2%	-21.5%	<b>+9.9%</b>

Table 4: Results of the SubST systems. The \* stands for statistically **not** significant results according to the bootstrap resampling test (Koehn, 2004).

As shown in Table 4, the SubST system trained on our automatically segmented data (*Synthetic*) shows comparable performance with the system

trained on the original segmentation (*Original*). The BLEU between the two models is identical for es, while for nl the difference is not significant. On the contrary, the Sigma for the system trained on manual segmentations is higher than for the synthetic ones by 6 points for nl but less than 2 for es. These results highlight that the breaks introduced by a non-perfect automatic segmentation influence the way the subtitle breaks are placed in the translation but not necessarily the translation itself. For the length constraint, both systems obtain high CPL conformity, with the *Synthetic* model scoring 3.5% more on nl and 2.5% less on es. This is related to the number of <eol> and <eob> inserted by the system: the more subtitle breaks are present, the more fine-grained is the segmentation, leading to higher conformity. Indeed, CPL is higher when the Break Coverage is high.

**Manual Analysis.** Upon examination of the segmentation patterns of the two en→es systems,<sup>7</sup> we did not identify particular differences. Specifically, the inserted <eob> tags follow punctuation marks in 76% of the cases for both models and are followed by prepositions and conjunctions in 32% and 29% for *Original* and *Synthetic* respectively. Similar patterns between outputs were observed for <eol> too, which is followed by a comma in the majority of cases and by the same function words as <eob>. These results suggest that systems trained on automatically segmented data are able to reproduce similar segmentation patterns to those trained on original data without showing a significant degradation in the translation.

## 6 Conclusions

We presented an automatic segmenter able to turn existing ST corpora into SubST training data. Through comparative experiments on two language pairs in zero-shot conditions, we showed that SubST systems trained on this synthetic data are competitive with those built on human-annotated subtitling corpora. Building on these positive results, and conditioned to the availability of suitable benchmarks, verifying the portability of the approach to a larger set of languages and domains is our priority for future work.

<sup>7</sup>We were unable to replicate the analysis on nl as we do not have the required linguistic competences.

## References

- Aitor Álvarez, Marina Balenciaga, Arantza del Pozo, Haritz Arzelus, Anna Matamala, and Carlos-D. Martínez-Hinarejos. 2016. [Impact of automatic segmentation on the quality, productivity and self-reported post-editing effort of intralingual subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3049–3053, Portorož, Slovenia.
- Aitor Álvarez, Carlos-D. Martínez-Hinarejos, Haritz Arzelus, Marina Balenciaga, and Arantza del Pozo. 2017. [Improving the automatic segmentation of subtitles through conditional random field](#). *Speech Communication*, 88:83–95.
- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2018. [Evaluating discourse phenomena in neural machine translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the ACL: Human Language Technologies*, pages 1304–1313, New Orleans, Louisiana.
- 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 *Proceedings of the NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain.
- Mary Carroll and Jan Ivarsson. 1998. *Code of Good Subtitling Practice*. Simrishamn: TransEdit.
- Thierry Etchegoyhen, Lindsay Bywood, Mark Fishel, Panayota Georgakopoulou, Jie Jiang, Gerard van Loenhout, Arantza del Pozo, Mirjam Sepesy Maučec, Anja Turner, and Martin Volk. 2014. [Machine Translation for Subtitling: A Large-Scale Evaluation](#). In *Proceedings of the 9th International Conference on Language Resources and Evaluation (LREC)*, pages 46–53.
- Marco Gaido, Mauro Cettolo, Matteo Negri, and Marco Turchi. 2021. [CTC-based compression for direct speech translation](#). In *Proceedings of the 16th Conference of the European Chapter of the ACL: Main Volume*, pages 690–696, Online.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, Mauro Cettolo, and Marco Turchi. 2020. [Contextualized Translation of Automatically Segmented Speech](#). In *Proceedings of Interspeech 2020*, pages 1471–1475.
- Xavier Garcia, Noah Constant, Ankur Parikh, and Orhan Firat. 2021. [Towards continual learning for multilingual machine translation via vocabulary substitution](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1184–1192, Online.
- Alex Graves, Santiago Fernández, Faustino 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 '06*, page 369–376, New York, NY, USA.
- Thanh-Le Ha, Jan Niehues, and Alexander Waibel. 2016. [Toward multilingual neural machine translation with universal encoder and decoder](#). *Institute for Anthropomatics and Robotics*, 2(10.12):16.
- Hirofumi Inaguma, Kevin Duh, Tatsuya Kawahara, and Shinji Watanabe. 2019. [Multilingual end-to-end speech translation](#). In *Proceedings of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 570–577. IEEE.
- Alina Karakanta, François Buet, Mauro Cettolo, and François Yvon. 2022. [Evaluating subtitle segmentation for end-to-end generation systems](#). In *Proceedings of LREC*, pages 3069–3078, Marseilles, France.
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2019. [Are subtitling corpora really subtitle-like?](#) In *Proceedings of the Sixth Italian Conference on Computational Linguistics (CLiC-It)*, Bari, Italy.
- 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 *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3727–3734, Marseille, France.
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2020c. [Point Break: Surfing Heterogeneous Data for Subtitle Segmentation](#). In *Proceedings of the Seventh Italian Conference on Computational Linguistics (CLiC-It)*, Bologna, Italy.
- Yunsu Kim, Duc Thanh Tran, and Hermann Ney. 2019. [When and why is document-level context useful in neural machine translation?](#) In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 24–34, Hong Kong, China.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain.
- 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 (EAMT 2020)*, pages 115–124.
- 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 (Demonstrations)*, pages 66–71, Brussels, Belgium.
- Pierre Lison, Jörg Tiedemann, and Milen Kouylekov. 2018. [OpenSubtitles2018: Statistical rescore of sentence alignments in large, noisy parallel corpora](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan.
- 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.
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the ACL (Demonstrations)*, pages 48–53, Minneapolis, Minnesota.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium.
- 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*, pages 84–91, Bangkok, Thailand (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.
- Umut Sulubacak, Ozan Caglayan, Stig-Arne Grönroos, Aku Rouhe, Desmond Elliott, Lucia Specia, and Jörg Tiedemann. 2020. Multimodal machine translation through visuals and speech. *Machine Translation*, 34(2):97–147.
- Anke Tardel. 2020. [Effort in Semi-Automatized Subtitling Processes: Speech Recognition and Experience during Transcription](#). *Journal of Audiovisual Translation*, 3(2):79–102.
- 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.
- Martin Volk, Rico Sennrich, Christian Hardmeier, and Frida Tidström. 2010. Machine Translation of TV Subtitles for Large Scale Production. In *Proceedings of the Second Joint EM+/CNGL Workshop "Bringing MT to the User: Research on Integrating MT in the Translation Industry (JEC'10)*, pages 53–62, Denver, CO.
- 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 1st Conference of the Asia-Pacific Chapter of the ACL and the 10th IJCNLP (Demonstrations)*, pages 33–39, Suzhou, China.
- 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.
- Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 2018. [Improving the transformer translation model with document-level context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 533–542, Brussels, Belgium.

## A Experimental Settings

### A.1 Data Selection

For the initial experiments aimed to train textual and multimodal segmenters and to select the best one (step 1 of the process described in Section 2), we use three sections of MuST-Cinema (Karakanta et al., 2020b), the only to date publicly available SubST dataset,<sup>8</sup> namely French, German, and Italian. Each section contains paired audio utterances, English transcripts, and translations in the corresponding language, where both sides of the text are built from subtitles created by humans. For French (275K sentences), German (229K sentences) and Italian (253K sentences), we collect the segmented translations of the corresponding MuST-Cinema sections. For English, we concatenate the segmented transcripts of the previous three sections (757K sentences). For each language (de, en, fr, it), the training data for the segmenter consists of unsegmented texts and, in the case of the multimodal segmenter, also audio as the source input, and of segmented texts (subtitles) as the target. Using the corpus notation, subtitle breaks are defined as: *block break* <eob>, which marks the end of the current subtitle displayed on screen, and *line break* <eol>, which splits consecutive lines inside the same block. For unsegmented texts, <eob> and <eol> are removed.

### A.2 Systems

We use the Adam optimizer and inverse square-root learning rate (lr) scheduler for all the trainings.

<sup>8</sup><https://ict.fbk.eu/must-cinema/> - License: CC BY-NC-ND 4.0

The *Textual* segmenter is a Transformer-based (Vaswani et al., 2017) architecture consisting of 3 encoder layers and 3 decoder layers. We set the hyper-parameters as in the fairseq (Ott et al., 2019) multilingual translation task, both for the mono- and multilingual textual segmenters. For the multilingual model, a mini-batch for each language direction (here: 4) is built and the model weights are updated after each mini-batch, a mechanism already present in fairseq Multilingual Machine Translation (Ott et al., 2019).

The *Multimodal* segmenter is an extension of the textual segmenter encoder-decoder structure with an additional speech encoder composed of 12 Transformer encoder layers as in the original speech-to-text task (Wang et al., 2020) but with the addition of a CTC (Graves et al., 2006) module to avoid the speech encoder pre-training (Gaido et al., 2021). The training of the multilingual models is performed by pre-pending the language token (en, de, fr, it) to the target sentence (Inaguma et al., 2019), a mechanism already present in the Fairseq Speech-to-text library (Wang et al., 2020). The encoder and decoder embeddings are shared. We select the hyper-parameters of the original implementation,<sup>9</sup> except for a higher learning rate of  $1 \cdot 10^{-3}$ , since pre-training was skipped. The vocabulary is generated using SentencePiece (Kudo and Richardson, 2018), setting the size to 10k unigrams both for the mono- and multilingual segmenters.

For the *Supervised* baseline using OpenSubtitles data, we follow the data selection process for the highest-performing segmenter in (Karakanta et al., 2020c) (*OpenSubs-42*). We first filter sentences with subtitles of maximum 42 characters. Since line breaks are not present in OpenSubtitles, we substitute `<eob>` symbols with `<eol>` with a probability of 0.25, paying attention not to insert two consecutive `<eol>`. This proportion reflects the `<eol>/<eob>` distribution featured by the MuST-Cinema training set. We noted that almost 90% of the sentences filtered contain only one subtitle. This is not very informative for the segmenter, since the only operation required is inserting one `<eob>` at the end of the sentence. For this reason, we further select only sentences with at least two subtitles (or two subtitle lines). This results in 2,956,207 sentences for es and 683,382 sentences for nl. We then add the same number of

sentences containing only one subtitle. After this process, we obtain 5,912,414 sentences for es and 1,366,764 sentences for nl. The supervised baseline is trained with the same settings as the textual monolingual segmenter.

For the *Count Chars* baseline, a break is inserted before reaching the 42-character limit, as per TED guidelines. If the 42-character limit is reached in the middle of a word, the break is inserted before this word. This method will always obtain a 100% conformity to the length constraint. As with the data filtering process, `<eol>` is inserted with probability of 0.25.

For the SubST models discussed in Section 5, we use the speech-to-text task *small* architecture of fairseq with the additional CTC module as in (Papi et al., 2021).

We use 4 GPUs K80 for training all the architectures: it takes around 1 day for the textual-only and around 1 week for the multimodal segmenters and the SubST models. All results are obtained by averaging 7 checkpoints (best, three preceding and three succeeding checkpoints).

<sup>9</sup>[https://github.com/pytorch/fairseq/blob/main/examples/speech\\_to\\_text/docs/mustc\\_example.md](https://github.com/pytorch/fairseq/blob/main/examples/speech_to_text/docs/mustc_example.md)