

Between Flexibility and Consistency: Joint Generation of Captions and Subtitles

Alina Karakanta^{1,2}, Marco Gaido^{1,2}, Matteo Negri¹, Marco Turchi¹

¹ Fondazione Bruno Kessler, Via Sommarive 18, Povo, Trento - Italy

² University of Trento, Italy

`akarakanta,mgaido,negri,turchi}@fbk.eu`

Abstract

Speech translation (ST) has lately received growing interest for the generation of subtitles without the need for an intermediate source language transcription and timing (i.e. captions). However, the joint generation of source captions and target subtitles does not only bring potential output quality advantages when the two decoding processes inform each other, but it is also often required in multilingual scenarios. In this work, we focus on ST models which generate consistent captions-subtitles in terms of structure and lexical content. We further introduce new metrics for evaluating subtitling consistency. Our findings show that joint decoding leads to increased performance and consistency between the generated captions and subtitles while still allowing for sufficient flexibility to produce subtitles conforming to language-specific needs and norms.

1 Introduction

New trends in media localisation call for the rapid generation of subtitles for vast amounts of audiovisual content. Speech translation, and especially direct approaches (Bérard et al., 2016; Bahar et al., 2019), have recently shown promising results with high efficiency because they do not require a transcription (manual or automatic) of the source speech but generate the target language subtitles directly from the audio. However, obtaining the intralingual subtitles (hereafter “captions”) is necessary for a range of applications, while in some settings captions need to be displayed along with the target language subtitles. Such “bilingual subtitles” are useful in multilingual online meetings, in countries with multiple official languages, or for language learners and audiences with different accessibility needs. In those cases, captions and subtitles should not only be consistent with the visual and acoustic dimension of the audiovisual

material but also between each other, for example in the number of blocks (pieces of time-aligned text) they occupy, their length and segmentation. Consistency is vital for user experience, for example in order to elicit the same reaction among multilingual audiences, or to facilitate the quality assurance process in the localisation industry.

Previous work in ST for subtitling has focused on generating interlingual subtitles (Matusov et al., 2019; Karakanta et al., 2020a), a) without considering the necessity of obtaining captions consistent with the target subtitles, and b) without examining whether the joint generation leads to improvements in quality. We hypothesise that knowledge sharing between the tasks of transcription and translation could lead to such improvements. Moreover, joint generation with a single system can avoid the maintenance of two different models, increase efficiency, and in turn speed up the localisation process. Lastly, if joint generation improves consistency, joint models could increase automation in subtitling applications where consistency is a desideratum.

In this work, we address these issues for the first time, by jointly generating both captions and subtitles for the same audio source. We experiment with the following models: 1) **Shared Direct** (Weiss et al., 2017), where the speech encoder is shared between the transcription and the translation decoder, 2) **Two-Stage** (Kano et al., 2017), where the transcript decoder states are passed to the translation decoder, and 3) **Triangle** (Anastasopoulos and Chiang, 2018), which extends the two-stage by adding a second attention mechanism to the translation decoder which attends to encoded speech inputs. We compare these models with the established approaches in ST for subtitling: an independent direct ST model and a cascade (ASR+MT) model. Moreover, we extend the evaluation beyond the usual metrics used to assess transcription and

translation quality (respectively WER and BLEU), by also evaluating the form and consistency of the generated subtitles.

Sperber et al. (2020) introduced several lexical and surface metrics to measure consistency of ST outputs, but they were only applied to standard, non-subtitle, texts. Subtitles, however, are a particular type of text structured in blocks which accompany the action on screen. Therefore, we propose to measure their consistency by taking advantage of this structure and introduce metrics able to reward subtitles that share similar structure and content.

Our contributions can be summarised as follows:

- We employ ST to directly generate both captions and subtitles without the need for human pre-processing (transcription, segmentation).
- We propose new, task-specific metrics to evaluate subtitling consistency, a challenging and understudied problem in subtitling evaluation.
- We show increased performance and consistency between the generated captions and subtitles compared to independent decoding, while preserving adequate conformity to subtitling constraints.

2 Background

2.1 Bilingual subtitles

New life conditions maximised the time spent in front of screens, transforming today's mediascape in a complex puzzle of new actors, voices and workflows. Face-to-face meetings and conferences moved online, opening up participation for global audiences. In these new settings multilinguality and versatility are dominant, and manifested in business meetings with international partners, conferences with world-wide coverage, multilingual classrooms and audiences with mixed accessibility needs. Given these growing needs for providing inclusive and equal access to audiovisual material for a multifaceted audience spectrum, efficiently obtaining high-quality captions and subtitles is becoming increasingly relevant.

Traditionally, displaying subtitles in two languages in parallel (bilingual or dual subtitles) has been common in countries with more than one official languages, such as Belgium and Finland (Gottlieb, 2004). Recently, however, captions along with subtitles have been employed in other countries to attract wider audiences, e.g. in Mainland

China English captions are displayed along with Mandarin subtitles. Interestingly, despite doubling the amount of text that appears on the screen and the high redundancy, it has been shown that bilingual subtitles do not significantly increase users' cognitive load (Liao et al., 2020).

One group which undoubtedly benefits from the parallel presence of captions and subtitles are language learners. Captions have been found to increase learners' L2 vocabulary (Sydorenko, 2010) and improve listening comprehension (Guichon and McLornan, 2008). Subtitles in the learners' native language (L1) are an indispensable tool for comprehension and access to information, especially for beginners. In bilingual subtitles, the captions support learners in understanding the speech and acquiring terminology, while subtitles serve as a dictionary, facilitating bilingual mapping (García, 2017). Consistency is particularly important for bilingual subtitles. Terms should fall in the same block and in similar positions. Moreover, similar length and equal number of lines can prevent long distance saccades, assisting in spotting the necessary information in the two language versions. Several subtitling tools have recently allowed for combining captions and subtitles on the same video (e.g. Dualsub¹) and bilingual subtitles can be obtained for Youtube videos² and TED Talks.³

Another aspect where consistency between captions and subtitles is present is in subtitling templates. A subtitling template is a source language/English version of a subtitle file already segmented and containing timestamps, which is used to directly translate the text in target languages while preserving the same structure (Cintas and Remael, 2007; Georgakopoulou, 2019; Netflix, 2021). This process reduces the cost, turn-around times and effort needed to create a separate timed version for each language, and facilitates quality assurance since errors can be spotted across the same blocks (Nikolić, 2015). These benefits motivated our work towards simultaneously generating two language versions with the maximum consistency, where the caption file can further serve as a template for multilingual localisation. This paper is a first step towards maximising automation for the generation of high-quality multiple language/accessibility subtitle versions.

¹<https://www.dualsub.xyz/>

²<https://www.watch-listen-read.com/>

³<https://amara.org/en/teams/ted/>

2.2 MT and ST for subtitling

Subtitling has long sparked the interest of the Machine Translation (MT) community as a challenging type of translation. Most works employing MT for subtitling stem from the statistical era (Volk et al., 2010; Etchegoyhen et al., 2014) or even before, with example-based approaches (Melero et al., 2006; Armstrong et al., 2006; Nyberg and Mitamura, 1997; Popowich et al., 2000; Piperidis et al., 2005). With the neural era, the interest in automatic approaches to subtitling revived. Neural Machine Translation (NMT) led to higher performance and efficiency and opened new paths and opportunities. Matusov et al. (2019) customised a cascade of ASR and NMT for subtitling, using domain adaptation with fine-tuning and improving subtitle segmentation with a specific segmentation module. Similarly, using cascades, Koponen et al. (2020) explored sentence- and document-level NMT for subtitling and showed productivity gains for some language pairs. However, bypassing the need to train and maintain separate components for transcription, translation and segmentation, direct end-to-end ST systems are now being considered as a valid and potentially more promising alternative (Karakanta et al., 2020a). Indeed, besides the architectural advantages, they come with the promise to avoid error propagation (a well known issue of cascade solutions), reduce latency, and better exploit speech information (e.g. prosody) without loss of information thanks to a less mediated access to the source utterance. To our knowledge, no previous work has yet explored the effectiveness of joint automatic generation of captions and subtitles.

2.3 Joint generation of transcription and translation

The idea of generating transcript and translation has been previously addressed in (Weiss et al., 2017; Anastasopoulos and Chiang, 2018). These papers presented different solutions (e.g. shared decoder and triangle) with the goal of improving translation performance by leveraging both ASR and ST data in direct ST. Later, Sperber et al. (2020) evaluated these methods with the focus of jointly producing consistent source and target language texts. Their underlying intuition is that, since in cascade solutions the translation is derived from the transcript, cascades should achieve higher consistency than direct solutions. Their results, however, showed that triangle models achieve the highest consistency

among the architectures tested (considerably better than that of cascade systems) and have competitive performance in terms of translation quality. Direct independent and shared models, instead, do not achieve the translation quality and consistency of cascades. However, all these previous efforts fall outside the domain of automatic subtitling and ignore the inner structure of the subtitles and their relevance when considering consistency.

3 Methodology

3.1 Models

To study the effectiveness of the different existing ST approaches in the subtitling scenario, we experiment with the following models:

The **Multitask Direct Model (DirMu)** model consists of a single audio encoder and two separate decoders (Weiss et al., 2017): one for generating the source language captions, and the other for the target language subtitles. The weights of the encoder are shared. The model can exploit knowledge sharing between the two tasks, but allows for some degree of flexibility since inference for one task is not directly influenced by the other task.

The **Two-Stage (2ST)** model (Kano et al., 2017) also has two decoders, but the transcription decoder states are passed to the translation decoder. This is the only source of information for the translation decoder as it does not attend to the encoder output.

The **Triangle (Tri)** model (Anastasopoulos and Chiang, 2018) is similar to the two-stage model, but with the addition of an attention mechanism to the translation decoder, which attends to the output embeddings of the encoder. Both 2ST and Tri support coupled inference and joint training.

We compare these models with common solutions for ST. The **Cascade (Cas)** model is a combination of an ASR + NMT components; the ASR transcribes the audio into text in the source language, which is then passed to an NMT system for translation into the target language. The two components are trained separately and can therefore take advantage of richer data for the two tasks. The cascade features full dependence between transcription and translation, which will potentially lead to high consistency.

The **Direct Independent (DirInd)** system consists of two independent direct ST models, one for the transcription (as in the ASR component of the cascade) and one for the translation. It hence lies on the flexibility edge of flexibility-consistency

spectrum compared to the models above.

3.2 Evaluation of subtitling consistency

While some metrics for evaluating transcription-translation consistency have been proposed in (Sperber et al., 2020), these do not capture the peculiarities of subtitling. The goal for consistent captions/subtitles is having the same structure (same number of subtitle blocks) and same content (each pair of caption-subtitle blocks has the same lexical content). Consider the following example:

0:00:50,820, 00:00:53,820

To put the assumptions very clearly:

Enonçons clairement nos hypothèses : le capitalisme,

00:00:53,820, 00:00:57,820

capitalism, after 150 years, has become acceptable,
après 150 ans, est devenu acceptable, au même titre

00:00:58,820, 00:01:00,820

and so has democracy.

que la démocratie.

In the example above, three blocks appear sequentially on the screen based on timestamp information, and each of them contains one line of text in English (caption) and French (subtitle). Since the source utterance is split across the same number of blocks (3), the captions and subtitles have the same structure. However, the captions and subtitles do not have the same lexical content. The first block contains the French words *le capitalisme*, which appear in the second block for the English captions. Similarly, *au même titre* corresponds to the third block in relation to the captions. This is problematic because terms do not appear in the same blocks (e.g. capitalism), and also leads to sub-optimal segmentation, since the French subtitles are not complete semantic units (logical completion occurs after *hypothèses* and *acceptable*).

We hence define the consistency between captions and subtitles based on two aspects: the structural and the lexical consistency. Structural consistency refers to the way subtitles are distributed on a video. In order to be structurally consistent, captions and subtitles for each source utterance should be split across the same number of blocks. This is a prerequisite for bilingual subtitles, since each caption-subtitle pair has the same timestamps. In other words, the captions and subtitles should appear and disappear simultaneously. Therefore, we define **structural consistency** as the percentage of utterances having the same number of blocks between captions and subtitles.

The second aspect of subtitling consistency is lexical consistency. Lexical consistency means that each caption-subtitle pair has the same lexical content. It is particularly important for ensuring synchrony between the content displayed in the captions and subtitles. This facilitates language learning, when terms appear in similar positions, and quality assurance, as it is easier to spot errors in parallel text. We define **lexical consistency** as the percentage of words in each caption-subtitle pair that are aligned to words belonging in the same block. In our example, there are six tokens of the subtitles which are not aligned to captions of the same block: *le capitalisme*, *au même titre*. For obtaining this score, we compute the number of words in each caption aligned to the corresponding subtitle and vice versa. For each caption-subtitle pair, this process results in two lexical consistency scores: $Lex_{caption \rightarrow subtitle}$ and $Lex_{subtitle \rightarrow caption}$, where, in the former, the number of aligned words is normalised by the number of words in the caption, while, in the latter, by the number of words in the subtitle. These two quantities are then averaged into a single value (Lex_{pair}). The corpus-level lexical consistency is obtained by averaging the Lex_{pair} of all caption-subtitle pairs in the test set.

4 Experimental setting

4.1 Data

For our experiments we use MuST-Cinema (Karakanta et al., 2020b), an ST corpus compiled from subtitles of TED talks. For a sound comparison with Karakanta et al. (2020a), we conduct the experiments on 2 language pairs, English→French and English→German. The breaks between subtitles are marked with special symbols, <eob> for breaks between blocks of subtitles and <eol> for new lines inside the same block. The training data contain 408 and 492 hours of pre-segmented audio (229K and 275K sentences) for German and French respectively. For tuning and evaluation we use the official development and test sets. We expect the captions and subtitles of TED Talks to have high consistency, since the captions serve as the basis for translating the speech in target subtitles.

The text data is segmented into sub-words with Sentencepiece (Kudo and Richardson, 2018) with the unigram setting. In line with recent works in ST, we found that a small vocabulary size is beneficial for the performance of ST models. Therefore,

we set a shared vocabulary of 1024 for all models except the MT component of the cascade, where vocabulary size is set to 24k. The special symbols <eob> and <eol> are kept as a single token.

For the audio input, we use 40-dimensional log Mel filterbank speech features. The ASR encoder was pretrained on the IWSLT 2020 data, i.e. Europarl-ST (Iranzo-Sánchez et al., 2020), Librispeech (Panayotov et al., 2015), How2 (Sanabria et al., 2018), Mozilla Common-Voice,⁴ MuST-C (Cattoni et al., 2020), and the ST TED corpus.⁵

4.2 Model training

The ASR and ST models are trained using the same settings. The architecture used is S-Transformer, (Di Gangi et al., 2019), an ST adaptation of Transformer, which has been shown to achieve high performance on different speech translation benchmarks. Following state-of-the-art systems (Potapczyk and Przybysz, 2020; Gaido et al., 2020), we do not add 2D self-attentions. The size of the encoder is set to 11 layers, and to 4 layers for the decoder. The ASR model used to pretrain the encoder, instead, has 8 encoder and 6 decoder layers. The additional 3 encoder layers are initialised randomly, similarly to the adaptation layer proposed by Bahar et al. (2019). As distance penalty, we choose the logarithmic distance penalty. We optimise using Adam (Kingma and Ba, 2015) (betas 0.9, 0.98), 4000 warm-up steps with initial learning rate of 0.0003, and learning rate decay with the inverse square root of the iteration. We apply label smoothing of 0.1, and dropout (Srivastava et al., 2014) is set to 0.2. We further use SpecAugment (Park et al., 2019), a technique for online data augmentation, with augment rate of 0.5. Training is completed when the validation perplexity does not improve for 3 consecutive epochs.

The MT component is based on the Transformer architecture (big) (Vaswani et al., 2017) with similar settings to the original paper. Since the ASR component outputs punctuation, no other pre-processing (except for BPE) is applied to the training data. In order to ensure a fair comparison with the direct and joint models, the MT component is trained only on MuST-Cinema data.

All experiments are run with the fairseq toolkit (Ott et al., 2019). Training is performed on

⁴<https://voice.mozilla.org/>

⁵http://iwslt.org/doku.php?id=offline_speech_translation

two K80 GPUs with 11 GB memory and models converged in about five days. Our implementation of the DirMu, Tri and 2ST models is publicly available at: https://github.com/mgaido91/FBK-fairseq-ST/tree/acl_2021

4.3 Evaluation

We evaluate three aspects of the automatically generated captions and subtitles: 1) quality, 2) form, and 3) consistency. For **quality** of transcription we compute WER on unpunctuated, lowercased output, while for quality of translation we use SacreBLEU (Post, 2018).⁶ We report scores computed at the level of utterances, where the output sentences contain subtitle breaks. A break symbol is considered as another token contributing to the score.

For evaluating the **form** of the subtitles, we focus on the conformity to the subtitling constraints of length and reading speed, as well as proper segmentation, as proposed in (Karakanta et al., 2019). We compute the percentage of subtitles conforming to a maximum length of 42 characters/line and a maximum reading speed of 21 characters/second.⁷ The plausibility of segmentation is evaluated based on syntactic properties. Subtitle breaks should be placed in such a way that keeps syntactic and semantic units together. For example, an adjective should not be separated from the noun it describes. We consider as plausible only those breaks following punctuation marks or those between a content word (chunk) and a function word (chink). We obtain Universal Dependencies⁸ PoS-tags using the Stanza toolkit (Qi et al., 2020) and calculate the percentage of break symbols falling either in the punctuation or the content-function groups as plausible segmentation.

Lastly, we evaluate structural and lexical **consistency** between the generated captions and corresponding subtitles, as described in Section 3.2. Word alignments are obtained using fast_align (Dyer et al., 2013) on the concatenation of MuST-Cinema training data and the system outputs. Text is tokenised using Moses tokeniser and the consistency percentage is computed on tokenised text.

En→Fr	WER	SacreBLEU	Length	Read_speed	Segment.	Struc.	Lex.
Cas	19.69	26.9	.94 / .93	.85 / .70	.86 / .82	.98	.99
DirInd	19.69	24.0	.94 / .94	.85 / .73	.86 / .84	.75	.86
DirMu	17.73	25.2	.95 / .93	.85 / .73	.87 / .80	.77	.87
2ST	19.05	25.6	.95 / .94	.85 / .71	.87 / .82	.83	.84
Tri	18.93	25.3	.93 / .91	.85 / .72	.87 / .82	.82	.92
<hr/>							
En→De							
Cas	18.52	19.9	.94 / .90	.62 / .58	.86 / .76	.95	.96
DirInd	18.52	18.1	.94 / .92	.62 / .59	.86 / .78	.73	.86
DirMu	16.95	18.7	.95 / .92	.62 / .59	.87 / .73	.75	.82
2ST	18.93	<i>19.6</i>	.94 / .92	.62 / .60	.86 / .76	.82	.81
Tri	19.10	<i>19.8</i>	.93 / .92	.62 / .60	.87 / .76	.78	.91

Table 1: Results for quality (WER and BLEU), subtitling conformity (Length, Reading speed and Segmentation), and subtitling consistency (Structural and Lexical) for model outputs for French and German. Conformity scores are reported for captions / subtitles. **Bold** denotes the best score. Results that are not statistically significant – according to pairwise bootstrap resampling (Koehn, 2004), $p < 0.05$ – than the best score are reported in *italics*.

5 Results

5.1 Transcription/Translation quality

We first examine the quality of the systems’ outputs. The first two columns of Table 1 show the WER and SacreBLEU score for the examined models.

In terms of **transcription quality**, `DirMu` (Multitask Direct – see Section 3.1) obtains the lowest WER for both languages (17.73 for French and 16.95 for German). As far as the rest of the models are concerned, there is a different tendency for French and German. `Tri` (Triangle) and `2ST` (Two-Stage) perform equally better than the `Cas/DirInd` for French, while the `Cas/DirInd` have higher transcription quality than `Tri` and `2ST` for German. An explanation for this incongruity is that these two models perform coupled inference, therefore the benefit of the joint decoding for the transcription can be related to similarities in terms of vocabulary between the two languages. Since French has a higher vocabulary similarity to English, with many words in TED Talks being cognates (e.g. specialised terminology), it is possible that joint decoding favours the transcription for French but not for German.

When it comes to **translation quality**, `Cas` outperforms all other models for French with 26.9 BLEU points, while the differences are not statistically significant among `DirMu`, `2ST` and `Tri`. For German, however, `Cas`, `2ST` and `Tri` perform on par. The model obtaining the lowest scores is `DirInd`. This finding confirms our hypothesis that

⁶BLEU+c.mixed+.1+s.exp+tok.13a+v.1.4.3

⁷In line with the TED Talk guidelines: <https://www.ted.com/participate/translate/guidelines>

⁸<https://universaldependencies.org/>

joint decoding, despite being more complex, improves translation quality thanks to the knowledge shared between the two tasks at decoding time.

In comparison to previous works, our transcription results are contrary to Sperber et al. (2020), who obtained the lowest WER with the cascade and direct independent models. However, for translation quality our best models are `Cas`, `2ST` and `Tri`, as in previous work. Moreover, in line with Anastasopoulos and Chiang (2018), the gains for `Tri` are higher for translation than for transcription. Comparing the BLEU score of our `DirInd` models to the models in (Karakanta et al., 2020a), we found that our models achieve higher performance with 20.07 BLEU compared to 18.76 for French and 13.55 compared to 11.92 for German.

All in all, we found that coupled-inference, supported by `Cas`, `2ST` and `Tri`, improves translation but not transcription quality. On the contrary, multitasking as in `DirMu` is beneficial for transcription, possibly because of a reinforcement of the speech encoder. However, could this improvement come at the expense of conformity to the subtitling constraints?

5.2 Subtitling conformity

Columns 3-5 of Table 1 show the percentage of captions/subtitles conforming to the length, reading speed and segmentation constraints discussed in Section 4.3. We observe that joint decoding does not lead to significant losses in conformity. Specifically, the captions generated by `DirMu` have the highest conformity in terms of length (95%), reading speed (85% and 62%) and segmentation quality (87%). Moreover, the high conformity score for

DirMu correlates with the low WER, showing that quality goes hand-in-hand with conformity.

For the conformity of the target language subtitles, instead, the picture is different. Even though the differences are not large, Cas has lower conformity to length (93% and 90%) and reading speed (70% and 58%). The segmentation scores show that, despite their high translation quality, the systems featuring coupled inference (Cas, Tri and 2ST) are constrained by the structure of the captions and segment subtitles in positions which are not optimal for the target language norms (82% and 76%). DirInd, on the contrary, has higher conformity compared to the other models (94% and 92% for length, 73% and 59% for reading speed), as well as segmentation quality (84% and 78%). DirInd is left to determine the most plausible segmentation for the target language without being bound by consistency constraints from the source. The lowest segmentation quality of subtitles is achieved by DirMu (80% and 73%).

We can conclude that the quality improvements of coupled inference and multi-tasking come with a slight compromise of subtitling conformity, as a result of loss of flexibility in decoding.

5.3 Subtitling consistency

The last two columns of Table 1 present the results for the subtitling consistency.

In terms of **Structural consistency** (Struc.), the model achieving the highest scores is Cas, with 98% and 95% of the utterances being distributed along the same number of blocks. As expected, the lowest structural consistency is achieved by DirInd (75% and 73%), which determines independently the positions of the block symbols. Among the joint models, Tri outputs captions and subtitles with higher consistency than DirMu, but both are outperformed by 2ST (83% and 82%). Our hypothesis is that by attending only the caption decoder, 2ST behaves similarly to the cascade, and the translation decoder better replicates the block structure. We noted that the reference captions and subtitles have lower consistency (92% both for French and German) than the cascade. This shows that the cascade copies the same <eob> tokens and achieves extreme structural consistency, which is a desideratum for our study case but may be harmful in other scenarios, since it leads to lower conformity (see Section 5.2). Indeed, in scenarios where consistency is not a key, subtitlers should have the

flexibility to adjust subtitling segmentation to suit the needs of their target languages (Oziemblewska and Szarkowska, 2020).

The **Lexical consistency** (Lex.) results show that Cas is the model with the highest content overlap in parallel caption-subtitle blocks with 99% and 96% of the words being aligned to the same block. As with the structural consistency, the lexical consistency of the cascade is higher than the references (95% for French, 86% for German). The direct model with the highest lexical consistency is Tri (92% and 91%). Interestingly, despite its high structural consistency, 2ST does not distribute the content consistently in the parallel blocks, achieving the lowest conformity (81%). The DirMu also achieves lower consistency than DirInd for German (82% compared to 86%) but not for French (87% compared to 86%). It is worth noting that lexical consistency is generally lower for German than for French. Indeed, a 100% lexical consistency between subtitles in languages with different word order is not always feasible or even appropriate. For example, the main verb in an English subordinate clause appears in the second position while in German at the end of the sentence. In order to adhere to grammatical rules, words in subtitles of different languages often have inter-block reordering. Therefore, the balance between flexibility and consistency is manifested here as a compromise between grammaticality and preservation of the same lexical content on each pair of subtitles.

To sum up, the results of structural consistency show that the models are able to preserve the block structure between captions and subtitles in more than 75% of the utterances. In addition, the high lexical consistency shows that the block symbols are not inserted randomly, but placed in a way that preserves the same lexical content in the parallel blocks.

All in all, our results show that the evaluation of captions and subtitles is a multifaceted process that needs to be addressed from multiple aspects: quality, conformity and consistency. Missing one of the three can lead to wrong conclusions. For instance, only considering quality and consistency could lead to disregard the importance of conformity and consider independent solutions an obsolete technology. Secondly, among the Direct architectures, the use of techniques that allow linking the generation process of captions and subtitles helps to achieve overall better quality and consistency than inde-

pendent decoding, with a slight discount in conformity, especially for the target subtitles. Between the `DirMu`, `2ST` and `Tri`, there is not a model that outperforms all the others in all the metrics, so the choice mainly depends on the application scenario. Lastly, comparing the `Cascade` and the `Direct`, the `Cascade` seems to be the best choice, but recent advancements in `Direct` approaches result in competitive solutions with increased efficiency of maintaining one model for both tasks.

6 Analysis

6.1 Evaluation of Lexical Consistency

In this section, we test the reliability of the lexical consistency metric. The metric depends on the successful word alignment, which, especially for low quality text, might be sub-optimal. We therefore manually count the number of words in the subtitles which do not appear in the corresponding captions. The task is performed on the first 347 sentences of the output of `DirMu` for French and German. We then estimate the mean absolute error between the consistency metric computed using the manual and the automatic alignments. As an additional step, we compute how often the automatic and the manual annotations agree in their judgement of consistent/non-consistent content in each block.

The mean absolute error between the manually and the automatically computed score is .08 for French and .11 for German. The metric may not be able to account for very small score differences between systems, however, when inspecting the differences between manual and automatic annotation we noticed that most errors appear in very low quality outputs or where lexical content was missing, and lead to a misalignment of only a few words. These cases were in fact challenging even for the human annotator. Instead, the agreement in the consistent/not-consistent judgement is high, with .85 for French and .75 for German. Considering the difficulty of aligning sentences belonging to languages with different word ordering, and the lower quality of German outputs, it is not surprising that the word aligner from English to German affects more our metric. However, these results show that the real impact is moderate and the metric is consistent with the human judgements in the majority of cases.

6.2 Does structural consistency extend to line breaks?

But what happens with the line breaks? Does a one-line caption correspond to a one-line subtitle in the output of our models? Having the same number of lines between caption and subtitle blocks is a more challenging scenario, since the subtitles tend to expand because of different length ratios between languages and translation strategies such as explicitation. For instance, for the target languages considered in this work (French and German) the length of the target subtitles when subtitling from English has been reported to be 5%-35% higher.⁹ If structural consistency is enforced to line breaks, it may compromise either the quality of the translation or the conformity to the subtitling constraints. In case of a one-liner caption, important information may be not rendered in the corresponding subtitle in order to match a shorter length of the caption, or the length constraint will be violated since the longer subtitle will not be adequately segmented in two lines. In order to ensure that our models do not push the structural consistency to an extreme, we compute the percentage of caption-subtitle blocks having the same number of lines.

	Cas	DirInd	DirMu	2ST	Tri	Ref
Fr	67%	49%	54%	59%	57%	67%
De	66%	47%	55%	53%	51%	66%

Table 2: Percentage of subtitle blocks containing the same number of lines for French and German outputs.

Table 2 confirms that caption and subtitle blocks do not always have the same number of lines, since only 67% and 66% of blocks in the caption/subtitle references have the same number of lines. When it comes to the models, the `Cascade` exactly matches the percentages of the references, while the `Direct` models have even lower percentage of equal number of lines. Among the `Direct` models, again the `DirInd` shows the lowest similarity. We observed that more line breaks were present in the target subtitles, which ensures length conformity, since the target subtitles expand (source-target character ratio of 0.91 for French and 0.93 for German).

⁹<https://www.andiamo.co.uk/resources/expansion-and-contraction-factors/>
<http://www.aranchodoc.com/wp-content/uploads/2017/07/Text-Expansion-Contraction-Chart3.png>
<https://www.kwintessential.co.uk/resources/expansion-retraction>

Therefore, the fact that structural consistency allows for flexibility in relation to the number and position of line breaks is key to achieving high quality and conformity.

7 Conclusions

In this work we explored joint generation of captions and subtitles as a way to increase efficiency and consistency in scenarios where this property is a desideratum. To this aim, we proposed metrics for evaluating subtitling consistency, tailored to the structural peculiarities of this type of translation. We found that coupled inference, either by models supporting end-to-end training (2ST, Tri) or not (Cas), leads to quality and consistency improvements, but with a slight degradation of the conformity to target subtitle constraints. The final architectural choice depends on the flexibility versus conformity requirements of the application scenario.

The findings of this work have provided initial insights related to the joint generation of captions and subtitles. One future research direction is towards improving the quality of generation by using more recent, higher-performing ST architectures. For example, Liu et al. (2020) extended the notion of the dual decoder by adding an interactive attention mechanism which allows the two decoders to exchange information and learn from each other, while synchronously generating transcription and translation. Le et al. (2020) proposed two variants of the dual decoder of Liu et al. (2020), the cross and parallel dual decoder, and experimented with multilingual ST. While neither of these works reported results on consistency, we expect that they are relevant to our scenario and have the potential of jointly generating multiple language/accessibility versions with high consistency. Moving beyond generic architectures, in the future we are planning to experiment with tailored architectures for improving consistency between automatically generated captions and subtitles. One important insight emerging from this work is that different degrees of conformity are required, or even appropriate, depending on the application scenario and languages involved. Given these challenges, we are aiming at developing approaches which allow for tuning the output to the desired degree of conformity, whether lexical, structural or both. We hope that this work will contribute to the line of research efforts towards improving

efficiency and quality of automatically generated captions and subtitles.

References

- 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. Association for Computational Linguistics.
- Stephen Armstrong, Colm Caffrey, and Marian Flanagan. 2006. [Translating DVD Subtitles from English-German and English-Japanese Using Example-Based Machine Translation](#). In *MuTra 2006—Audio-visual Translation Scenarios: Conference Proceedings*.
- Parnia Bahar, Tobias Bieschke, and Hermann Ney. 2019. A Comparative Study on End-to-end Speech to Text Translation. In *Proceedings of International Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 792–799, Sentosa, Singapore.
- Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation. In *NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain.
- Roldano Cattoni, Mattia A. Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2020. MuST-C: A Multilingual Corpus for end-to-end Speech Translation. *Computer Speech & Language Journal*. Doi: <https://doi.org/10.1016/j.csl.2020.101155>.
- Jorge Diaz Cintas and Aline Remael. 2007. *Audiovisual Translation: Subtitling*. Translation practices explained. Routledge.
- Mattia Antonino Di Gangi, Matteo Negri, and Marco Turchi. 2019. Adapting Transformer to End-to-end Spoken Language Translation. In *INTERSPEECH*, pages 1133–1137.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. [A simple, fast, and effective reparameterization of IBM model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- 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, 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. Association for Computational Linguistics.
- Boni García. 2017. [Bilingual subtitles for second-language acquisition and application to engineering education as learning pills](#). *Computer Applications in Engineering Education*, 25(3):468–479.
- Panayota Georgakopoulou. 2019. Template files: The Holy Grail of subtitling. *Journal of Audiovisual Translation*, 2(2):137–160.
- Henrik Gottlieb. 2004. Subtitles and international anglicization. *Nordic Journal of English Studies*, 3:219–230.
- Nicolas Guichon and Sinead McLornan. 2008. The effects of multimodality on 12 learners: Implications for call resource design. *System*, pages 85–93.
- Javier Iranzo-Sánchez, Joan Albert 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.
- Takatomo Kano, Sakriani Sakti, and Satoshi Nakamura. 2017. [Structured-based Curriculum Learning for End-to-end English-Japanese Speech Translation](#). In *Annual Conference of the International Speech Communication Association (InterSpeech)*, pages 2630–2634.
- Alina Karakanta, Matteo Negri, and Marco Turchi. 2019. [Are Subtitling Corpora really Subtitle-like?](#) In *Sixth Italian Conference on Computational Linguistics, CLIC-It*.
- 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. Association for Computational Linguistics.
- 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. European Language Resources Association.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference for Learning Representations*.
- 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. Association for Computational Linguistics.
- 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: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Hang Le, Juan Pino, Changhan Wang, Jiatao Gu, Didier Schwab, and Laurent Besacier. 2020. [Dual-decoder transformer for joint automatic speech recognition and multilingual speech translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3520–3533, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Sixin Liao, Jan-Louis Kruger, and Stephen Doherty. 2020. [The impact of monolingual and bilingual subtitles on visual attention, cognitive load, and comprehension](#). *The Journal of Specialised Translation*.
- Yuchen Liu, Jiajun Zhang, Hao Xiong, Long Zhou, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2020. [Synchronous Speech Recognition and Speech-to-Text Translation with Interactive Decoding](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*. Association for the Advancement of Artificial Intelligence (www.aaai.org).
- 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. Association for Computational Linguistics.
- Maite Melero, Antoni Oliver, and Toni Badia. 2006. [Automatic Multilingual Subtitling in the eTITLE Project](#). In *Proceedings of ASLIB Translating and the Computer 28*.
- Netflix. 2021. Subtitle template timed text style guide. <https://partnerhelp.netflixstudios.com/hc/en-us/articles/219375728-English-Template-Timed-Text-Style-Guide>. Last accessed: 02/05/2021.
- Kristijan Nikolić. 2015. [The pros and cons of using templates in subtitling](#). In Rocío Baños Piñero and Jorge Díaz Cintas, editors, *Audiovisual Translation*

- in a *Global Context: Mapping an Ever-changing Landscape*, pages 192–202. Palgrave Macmillan UK, London.
- Eric Nyberg and Teruko Mitamura. 1997. [A Real-Time MT System for Translating Broadcast Captions](#). In *Proceedings of the Sixth Machine Translation Summit*, pages 51–57.
- 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 Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Magdalena Oziemblewska and Agnieszka Szarkowska. 2020. [The quality of templates in subtitling. A survey on current market practices and changing subtitler competences](#). *Perspectives*, 0(0):1–22.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An ASR corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, South Brisbane, Queensland, Australia.
- 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](#). *Interspeech 2019*.
- Stelios Piperidis, Iason Demiros, and Prokopis Prokopidis. 2005. [Infrastructure for a Multilingual Subtitle Generation System](#). In *9th International Symposium on Social Communication*, pages 24–28.
- Fred Popowich, Paul McFetridge, Davide Turcato, and Janine Toole. 2000. [Machine Translation of Closed Captions](#). *Machine Translation*, pages 311–341.
- 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. Association for Computational Linguistics.
- Tomasz Potapczyk and Pawel Przybysz. 2020. [SRPOL’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.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, 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).
- Matthias Sperber, Hendra Setiawan, Christian Gollan, Udhyakumar Nallasamy, and Matthias Paulik. 2020. [Consistent transcription and translation of speech](#). *Transactions of the Association for Computational Linguistics*, 8:695–709.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: a simple way to prevent neural networks from overfitting](#). *Journal of machine learning research*, 15(1):1929–1958.
- Tetyana Sydorenko. 2010. [Modality of input and vocabulary acquisition](#). *Lang Learn Technol*, pages 50–73.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, pages 6000–6010.
- 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.
- 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.