



Contextualized Translation of Automatically Segmented Speech

Marco Gaido^{1,2}, Mattia A. Di Gangi^{1,2}, Matteo Negri¹, Mauro Cettolo¹, Marco Turchi¹

¹Fondazione Bruno Kessler, Italy

²University of Trento, Italy

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

Abstract

Direct speech-to-text translation (ST) models are usually trained on corpora segmented at sentence level, but at inference time they are commonly fed with audio split by a voice activity detector (VAD). Since VAD segmentation is not syntax-informed, the resulting segments do not necessarily correspond to well-formed sentences uttered by the speaker but, most likely, to fragments of one or more sentences. This segmentation mismatch degrades considerably the quality of ST models' output. So far, researchers have focused on improving audio segmentation towards producing sentence-like splits. In this paper, instead, we address the issue in the model, making it more robust to a different, potentially sub-optimal segmentation. To this aim, we train our models on randomly segmented data and compare two approaches: fine-tuning and adding the previous segment as context. We show that our context-aware solution is more robust to VAD-segmented input, outperforming a strong base model and the fine-tuning on different VAD segmentations of an English-German test set by up to 4.25 BLEU points.

Index Terms: speech translation, VAD, context, segmentation

1. Introduction

Speech-to-text translation (ST) has been traditionally addressed by pipeline approaches involving several components [1]. The most important blocks are the speech recognition (ASR), which converts the input audio into its transcript, and the neural machine translation (NMT), which translates the transcript into the target language. Direct ST models [2, 3] recently gained attention as an alternative approach, thanks to their appealing promises to overcome some of the pipeline systems' problems, such as error propagation and loss of information present in the audio (prosody in particular).

Both pipeline and direct solutions, however, can be significantly affected by mismatches in the segmentation of the input between training and test data. On one side, the two solutions involve the use of training data segmented at sentence level. This, for instance, holds for the parallel corpora normally used to train the NMT component of the pipeline approach, as well as for all the available ST corpora used for direct ST training. On the other side, at inference time both solutions will be exposed to data segmented according to criteria that look at properties of the audio input rather than at linguistic notions like sentence well-formedness. The most widespread approach consists in fact in using voice activity detection (VAD) to split the audio stream into chunks, which are input to the ST system. In particular, VAD systems determine whether a given short (usually 10-30 ms) audio segment actually contains speech, and this information is used in the context of ST for two purposes: *i*) dividing the audio stream into segments containing uninterrupted speech; *ii*) filtering out audio segments containing other sounds.

Since VAD is solely based on the alternation between hu-

man voice, silences and other sounds, the resulting splits might not correspond to well-formed sentences but to fragments of one or more sentences. The impact of feeding an ST model trained on "clean" data with sub-optimal, not linguistically-motivated segmentations varies according to the characteristics of the VAD employed and its settings. Very aggressive settings reduce the generation of long (cross-sentential) segments, which are difficult to handle by neural models that are typically very sensitive to input length. On the downside, they produce short (sub-sentential) segments that might not provide enough context for proper translation. To address this problem, pipeline systems include an additional component that re-segments the ASR output to provide the NMT with well-formed sentences [4, 5, 6]. Since this solution is not possible for direct ST, where the two steps are not decoupled, researchers have worked on alternative audio segmentation techniques. In the 2019 IWSLT offline ST task [1], for instance, the best direct ST system [7] had one of its key features in the segmentation method.

Instead of working on the segmentation algorithm, in this paper we aim to make our direct ST models more robust to VAD-segmented data. To train them on a data distribution more similar to the one fed at inference time, we generate an artificial dataset by randomly re-segmenting clean (i.e. sentence-based) ST data. Then, we experiment with two approaches: *i*) fine-tuning on the new dataset; *ii*) improving our direct ST model with the capability to look back and attend to the preceding segment as contextual information. Our experiments show that the proposed context-based solution effectively handles the segmentation of different VAD systems and configurations, reducing the drop in translation quality caused by segmentation mismatches in the training and test data by up to 55%.

2. Context-aware ST

The idea of exploiting contextual information to improve translation has been successfully applied in NMT [8, 9, 10, 11]. In our use case, unlike [8], we are interested only in modeling short-range cross-segment dependencies to cope with the sub-optimal breaks introduced by VAD segmentation. We hence consider as context only the segment immediately preceding the one to be translated, leaving out of our study hierarchical approaches modeling the whole document as context. Moreover, while in document-level NMT the best approach is to use the source side of the sentence(s) as contextual information, in the ST scenario it is not trivial to understand which side is best. On one hand, audio source avoids the error propagation and exposure bias introduced by using as context the translations generated at inference time. On the other, these problems are balanced by the easiness of extracting information from text rather than from audio [12]. In this work, we study both options.

To integrate context information into the model, we explore the two solutions that gave the best results for NMT [11]. They

respectively use sequential [9] and parallel [10] decoders. We also experimented with the integration of context information in the encoder [9], but the trainings were either very unstable (when using audio as context) or ineffective, eventually leading to worse results. For this reason, we do not consider this type of integration in the rest of the paper. Finally, supported by previous findings [11], we neither investigate the concatenation of the context with the current input [13], nor the combination of encoded representations of the two [14].

Our base model is an adaptation of Transformer [15]: its encoder is enhanced to take into account the characteristics of speech input by means of two 2D convolutional layers and a logarithmic distance penalty in its self-attention layers [16]. Both the sequential and the parallel decoder use a multi-encoder approach, with an additional encoder dedicated to the context information. However, they differ in the way this information is integrated into the base model. The context encoder is composed of Transformer encoder layers, but its input depends on the modality of the segment used as context, i.e. text or audio. When we use the generated translations as context, its tokens are converted into vectors with *word embeddings* (namely, we re-use the decoder embeddings), summed with *positional encoding* and then provided to the encoder Transformer layers. When we use the audio as context, the input audio features are first processed by the encoder of the base model and then passed to the context encoder [12].

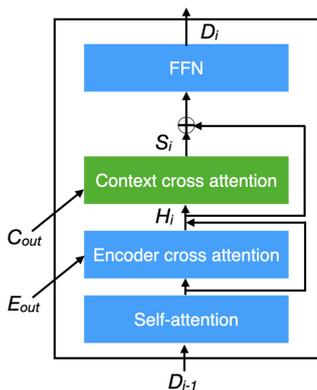


Figure 1: *Sequential context integration.*

Sequential (Figure 1). In each decoder Transformer layer, an additional multi-head cross-attention sub-layer is introduced. It queries the output C_{out} of the context encoder using the output H_i of the i -th encoder cross-attention sub-layer. The result S_i of this operation is combined with H_i using a position-wise gating mechanism, before being fed to the feed-forward network FFN_i . Hence, the output of the i -th decoder layer D_i is:

$$\lambda_i = \sigma(W_{hi}H_i + W_{si}S_i) \quad (1)$$

$$D_i = FFN_i(\lambda_i H_i + (1 - \lambda_i)S_i) \quad (2)$$

Parallel (Figure 2). In each decoder Transformer layer, the output of the self-attention sub-layer is used as query for both the encoder cross-attention and the context cross-attention defined in the same way as in the previous case. The outputs of these two sub-layers are then combined using the position-wise gating mechanism described in Eq.(2).

To avoid over-relying on the context, we add a regularization on the context gate. Our regularization is slightly different from the one proposed by [17]: we always penalize the context

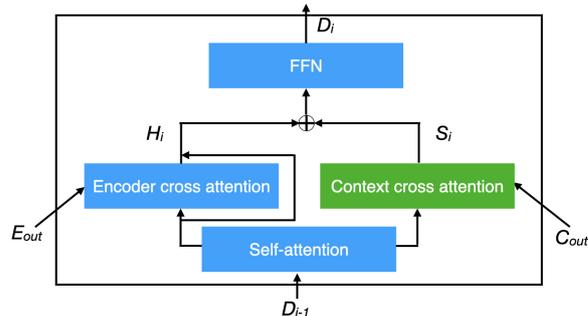


Figure 2: *Parallel context integration.*

information, so that the model will use it only when it is strictly needed. With the regularization factor, the resulting loss is:

$$\mathcal{L}' = \mathcal{L} + \alpha \sum_{i=0}^{N_d} (1 - \lambda_i) \quad (3)$$

3. Experimental settings

3.1. Clean and artificial data

Our base models are trained and evaluated on English-German data drawn from the MuST-C corpus, the largest ST dataset currently available [18]. MuST-C comprises 234K samples (corresponding to about 408 hours of speech) divided into training (229K), validation (1.5K) and test sets (3.5K).

To cope with segmentation mismatches between the clean data used for training and the VAD-processed ones handled at inference time, we generate an automatic re-segmentation of the MuST-C training and validation set. The re-segmentation starts by picking a random (with uniform distribution) *split word* for each sample in the original English transcripts. Each fragment spanning from a *split word* to the word before the next *split word* becomes a segment of the new training set and the preceding fragment becomes its context. We extract the audio corresponding to each resulting transcript by leveraging word alignments computed with Gentle.¹ Then, we retrieve the corresponding translations using word alignments generated with fast_align [19]. In case of missing alignments (either with the audio or with the translation), the sample is discarded. The resulting training dataset contains 225K samples (4K less than the original), while the validation set size is almost unchanged.

A manual check on a sample of the produced aligned segments revealed that about 96% of them are acceptable. The most frequently observed issue is that some translations contain 1-2 words more than the optimum, mostly due to the lack of some word alignments and to word-reordering. This leads to the presence of overlapping words between the context and the target German references in 25% of the samples. In early experiments, this caused model instability at inference time because models learnt to copy the final context words, up to producing nonsensical sequences of repeated tokens. We solved the issue by filtering out the overlapping words from the context.

3.2. VAD and segmentation

As we want our systems to be robust to different VAD outputs, we test our models on two different open source VAD tools:

¹<https://github.com/lowerquality/gentle/>

LIUM [20] and WebRTC’s VAD.² For WebRTC we tested all the possible configurations, varying the *frame size* (allowed values are 10ms, 20ms and 30ms) and the *aggressiveness* (ranging from 0 to 3, extremes included). We discarded those producing either too long (> 60s) or too many segments (> 5, 100, i.e. twice the segments of the original sentence-based segmentation of the MuST-C test set). In this way, we ended up with three configurations, whose characteristics are described in Table 1.

Overall, the segments produced by WebRTC have much higher variance in their length (ranging from 0.40s to 58.62s) compared to LIUM (from 2.50s to 18.63s) and are significantly more (> 3, 500 vs 2, 725). As anticipated in §1, this can affect the final performance of neural ST models, for which handling very long/short segments is difficult. However, from a qualitative standpoint, a manual inspection of 50 samples showed that the split times selected by LIUM are less accurate than those selected by WebRTC: while the former often splits fluent speech, the latter always selects positions in which the speaker is silent.

Table 1: *Statistics for different segmentations of the MuST-C test set. “Man.” refers to the original sentence-based segmentation.*

System Frame size Aggress.	Man.	LIUM	WebRTC		
			30ms 3	20ms 2	20ms 3
% filt. audio	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

3.3. Training settings

All our models are optimized with label smoothed cross entropy [21] using the Adam optimizer [22] with a learning rate starting from $3 \cdot 10^{-4}$, increasing linearly up to $5 \cdot 10^{-4}$ in the first 5, 000 steps and then decaying with inverse square root policy. The overall batch size was 512 (*audio, translation*) pairs. We used the *BIG* configuration from [16] regarding all layers’ hidden sizes. The number of context encoder layers N_e is set to 1, as [9] shows that this leads to the best results. Since [11] has demonstrated that poorly regularized systems can lead to ambiguous results when integrating context, we used 0.2 dropout and *SpecAugment* [23] to prevent this issue.

We performed preliminary experiments on a baseline model (*BASE_MUSTC*) with 8 encoder layers N_e and 6 decoder layers N_d trained on the MuST-C En-De training set. Since models using the generated translations as context are affected by exposure bias, we wanted to test our solution also in more realistic conditions, with a stronger baseline model trained in rich data conditions. This model (*BASE_ALL*) was trained with N_e set to 11 and N_d to 4, on all the data available for the IWSLT 2020 evaluation campaign,³ with knowledge distillation from an MT model and synthetic data generated translating the transcripts of ASR corpora. Its training involves a pre-training on the synthetic data, a fine-tuning on the data having ground-truth translations and a second fine-tuning using label-smoothed cross entropy instead of knowledge distillation [24].

All the context-aware models are initialized with the corresponding baseline model trained on sentence-segmented data.

²<https://webrtc.org/>. We use the open-source Python interface <https://github.com/wiseman/py-webrtcvad>

³http://iwslt.org/doku.php?id=offline_speech_translation

We experimented with freezing all the pre-trained parameters as in [9], but freezing the decoder weights turned out to be harmful. If frozen, decoder’s layers are not able to adapt to the new inputs (with different segmentation) and this slows down convergence and leads to worse results. We hence freeze only the encoder. Our code is based on fairseq [25] and is available at <https://github.com/mgaido91/FBK-fairseq-ST>.

Textual data were pre-processed with tokenization and punctuation normalization performed using Moses [26], and were segmented with 8, 000 BPE merge rules [27]. For the audio, we applied 40 Mel filters with window size of 25ms and stride of 10ms, performing speaker normalization with XNMT [28]. To avoid out-of-memory errors, we excluded from the training set the audio segments longer than 20 seconds.

In all cases, evaluation is performed on the best model according to the loss on the validation set. The metrics used are BLEU [29] and TER [30], computed against the reference translations in the MuST-C En-De test set.

4. Results

We performed preliminary experiments with *BASE_MUSTC* (scoring 21.08 BLEU on the original MuST-C En-De test set) to compare the context integration techniques and select the most suitable one for ST. We then compared the fine-tuning with the context-aware models using the stronger baseline model *BASE_ALL* (scoring 27.55 BLEU on the original test set).

Table 2: *Evaluation results on the VAD-segmented test set. Notes: SRC=audio as context; TGT=generated translation as context; SEQ=sequential; PAR=parallel.*

	LIUM	WebRTC		
		3, 30ms	2, 20ms	3, 20ms
BASE_MUSTC	17.32	17.82	17.75	16.31
SRC SEQ	19.08	18.81	18.00	17.42
SRC PAR	19.25	18.90	18.25	17.30
TGT SEQ	19.57	19.21	18.81	17.60
TGT PAR	20.01	18.98	18.82	17.32

4.1. Context information and integration

Table 2 shows that all the tested approaches outperform the baseline on VAD-segmented data with a margin that ranges from 0.25 to 2.69 BLEU points. This indicates that the context is useful to mitigate the effect of VAD-based segmentation. On LIUM, our models achieve the highest score (*TGT PAR*, 20.01 BLEU) and the largest gain over the baseline; on WebRTC the improvements are significant but smaller. We argue that the reason lies in the different characteristics of the two tools. The split positions selected by LIUM do not always correspond to actual pauses in the audio, which prevents the baseline model from disposing of all the information necessary for translation. This information, instead, is available to the context-aware models as they can access the previous segment. WebRTC, instead, produces very long/short segments, whose effect on context-aware models is limited: the contribution of adding the previous segment is low both in case of very long segments, as only the first part is influenced by it, and in case of very short ones, as having a short segment as context means adding little information. We also experimented with including manually-segmented data, but it was not beneficial for any of our models.

Table 3: Comparison between base model, fine-tuning and context-aware models.

	LIUM		WebRTC					
	BLEU (↑)	TER (↓)	AGG=3, FS=30ms		AGG=2, FS=20ms		AGG=3, FS=20ms	
	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)	BLEU (↑)	TER (↓)
BASE.ALL	19.66	76.57	22.07	67.08	21.98	66.83	19.59	72.62
FINE-TUNE	22.48	64.21	23.48	60.03	23.40	61.54	21.35	63.90
TGT SEQ	23.18	58.60	22.85	58.49	22.59	59.79	21.11	60.51
+ REG	23.88	58.81	23.61	58.57	23.15	60.36	21.88	60.97
TGT PAR	23.77	59.02	23.34	58.94	22.91	60.09	21.75	60.77
+ REG	23.91	58.95	23.51	58.64	23.40	59.95	22.03	60.83

Looking at the context modality (text vs audio), we observe that supplying the previously generated translation (TGT*) yields higher BLEU scores than supplying its corresponding audio (SRC*) with both the integration types (*SEQ and *PAR). This suggests that the audio representation produced by current ST models is less suitable than text to extract useful content information to support translation. In light of these observations, we decided to proceed with *TGT SEQ* and *TGT PAR* in the following experiments with the stronger *BASE.ALL* model.

4.2. Context vs fine-tuning

In this section, we compare the performance of the fine-tuning and the context-aware solutions. In this way, we can disentangle the benefits produced by the context and those due to the use of artificial training data.

The results in Table 3 show that: *i*) fine-tuning on the artificial data produces significant gains over *BASE.ALL* (respectively, 2.82 BLEU points on LIUM and from 1.41 to 1.76 on WebRTC), and *ii*) *TGT PAR* outperforms *TGT SEQ* on all datasets (by 0.32 to 0.64). *TGT PAR* without regularization is superior to the fine-tuning when the VAD splits very aggressively (21.75 vs 21.35 on WebRTC 3, 20ms) or in non-pause positions (23.77 vs 22.48 on LIUM). On the other VAD configurations, the results are close, but inferior to the fine-tuning. Our intuition is that this behavior is caused by the noise added by the context-attention when the context is not needed. This is confirmed by the results obtained adding the context-gate regularization presented in Eq. (3) (*TGT PAR+REG* and *TGT SEQ+REG*). The regularization allows our best context-aware model (*TGT PAR+REG*) to outperform the fine-tuned model on 3 out of 4 VAD configurations tested (in one case BLEU is on par) and improves both integration types. *TGT SEQ* benefits more from it, closing the gap with *TGT PAR*. The value of the hyperparameter α was chosen among 0.01, 0.02, 0.04 and 0.08: we set it to 0.04 as it provided the best loss on the validation set.

The difference between context-aware models and fine-tuning is even more evident if we consider the TER metric (the lower the better). In this case, *TGT SEQ* obtains the best scores in every setting, but the results of all context-aware models are close and are 2 to 6 points better than those obtained with fine-tuning. We also noticed that 1-, 2-, 3- and 4-gram BLEU scores are always significantly higher for the context-aware solutions than for the fine-tuning, even when the overall BLEU scores are similar. The reason lies in the brevity penalty, as the context-aware models produce shorter translations. Interestingly, the best result (23.91 BLEU) is obtained by exploiting the context in one of the hardest segmentations for the base model (19.66 BLEU). This is coherent with the behavior observed in §4.1.

5. Analysis

We performed a manual analysis of the translations produced by the baseline and by our best context-aware model (*TGT PAR + REG*) on the LIUM-segmented test set. The goal was to check whether the gains are actually due to the use of contextual information and to understand how this information is exploited. We noticed three main issues solved by the context-aware approach. They are all related to the presence of sub-sentential fragments located at the beginning or the end of a segment. First, these fragments are often ignored by the baseline model. Being trained only on well-formed sentences from the clean MuST-C corpus, this model seems unable to handle segments reflecting truncated sentences and, instead of returning partial translations, it opts for ignoring part of the input audio. Second, the base model produces *hallucinations* [31] trying to translate a sub-sentential fragment into a well-formed target sentence. Our models, instead, produce the translation corresponding to the incomplete fragment. Third, the baseline model translates the sub-sentential fragment and the adjacent sentence in the same segment into one single output sentence, mixing them. In contrast, our models are able to translate them separately.

6. Conclusions

We studied how to make ST models trained on data segmented at sentence-level robust to VAD-segmented audio supplied at inference time. To this aim, we explored different approaches to integrate contextual information provided by the segment preceding the one to be translated. Our experiments show that adopting a context-aware architecture, combined with training on artificial data generated with random segmentation, is beneficial to improve final translation quality. We also demonstrate that, compared to the best automatic segmentation (22.07 BLEU), context-aware models achieve results that are similar in the worst case (22.03) and significantly better in the best case (23.91). In this case, our context-based approach allows to reduce by 55% the performance gap of the base model (19.66) with respect to optimal (i.e. sentence-level) manual segmentation (27.55). All in all, this suggests that working on models' robustness to sub-optimal VAD segmentation is at least as promising as improving the segmentation itself.

7. Acknowledgements

This work is part of the "End-to-end Spoken Language Translation in Rich Data Conditions" project,⁴ which is financially supported by an Amazon AWS ML Grant.

⁴<https://ict.fbk.eu/units-hlt-mt-e2eslt/>

8. References

- [1] J. Niehues, R. Cattoni, S. Stucker *et al.*, “The IWSLT 2019 Evaluation Campaign,” in *Proc. of 16th International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong, Nov. 2019.
- [2] A. Bérard, O. Pietquin, C. Servan, and L. Besacier, “Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation,” in *Proc. of NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain, Dec. 2016.
- [3] R. J. Weiss, J. Chorowski, N. Jaitly, Y. Wu, and Z. Chen, “Sequence-to-Sequence Models Can Directly Translate Foreign Speech,” in *Proc. of Interspeech 2017*, Stockholm, Sweden, Aug. 2017, pp. 2625–2629.
- [4] E. Matusov, A. Mauser, and H. Ney, “Automatic Sentence Segmentation and Punctuation Prediction for Spoken Language Translation,” in *Proc. of 2006 International Workshop on Spoken Language Translation (IWSLT)*, Kyoto, Japan, Nov. 2006, pp. 158–165.
- [5] Y. Oda, G. Neubig, S. Sakti, T. Toda, and S. Nakamura, “Optimizing Segmentation Strategies for Simultaneous Speech Translation,” in *Proc. of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL)*, Baltimore, Maryland, Jun. 2014, pp. 551–556.
- [6] E. Cho, J. Niehues, and A. Waibel, “NMT-Based Segmentation and Punctuation Insertion for Real-Time Spoken Language Translation,” in *Proc. of Interspeech 2017*, Stockholm, Sweden, Aug. 2017, pp. 2645–2649.
- [7] T. Potapczyk, P. Przybysz, M. Chochowski, and A. Szumaczk, “Samsung’s System for the IWSLT 2019 End-to-End Speech Translation Task,” in *Proc. of 16th International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong, Nov. 2019.
- [8] L. Wang, Z. Tu, A. Way, and Q. Liu, “Exploiting Cross-Sentence Context for Neural Machine Translation,” in *Proc. of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Copenhagen, Denmark, Sep. 2017, pp. 2826–2831.
- [9] J. Zhang, H. Luan, M. Sun, F. Zhai, J. Xu, M. Zhang, and Y. Liu, “Improving the Transformer Translation Model with Document-Level Context,” in *Proc. of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, Oct.-Nov. 2018, pp. 533–542.
- [10] R. Bawden, R. Sennrich, A. Birch, and B. Haddow, “Evaluating Discourse Phenomena in Neural Machine Translation,” in *Proc. of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, New Orleans, Louisiana, Jun. 2018, pp. 1304–1313.
- [11] Y. Kim, D. T. Tran, and H. Ney, “When and Why is Document-level Context Useful in Neural Machine Translation?” in *Proc. of the 4th Workshop on Discourse in Machine Translation (DiscoMT)*, Hong Kong, China, Nov. 2019, pp. 24–34.
- [12] M. A. Di Gangi, V. Nguyen, M. Negri, and M. Turchi, “Instance-based Model Adaptation for Direct Speech Translation,” in *Proc. of 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Barcelona, Spain, May 2020, pp. 7914–7918.
- [13] R. Agrawal, M. Turchi, and M. Negri, “Contextual Handling in Neural Machine Translation: Look Behind, Ahead and on Both Sides,” in *Proc. of the 21st Annual Conference of the European Association for Machine Translation (EAMT)*, Alacant, Spain, May 2018, pp. 11–20.
- [14] E. Voita, P. Serdyukov, R. Sennrich, and I. Titov, “Context-Aware Neural Machine Translation Learns Anaphora Resolution,” in *Proc. of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, Melbourne, Australia, Jul. 2018, pp. 1264–1274.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is All You Need,” in *Proc. of Advances in Neural Information Processing Systems 30 (NIPS)*, Long Beach, California, Dec. 2017, pp. 5998–6008.
- [16] M. A. Di Gangi, M. Negri, and M. Turchi, “Adapting Transformer to End-to-End Spoken Language Translation,” in *Proc. of Interspeech 2019*, Graz, Austria, Sep. 2019, pp. 1133–1137.
- [17] X. Li, L. Liu, R. Wang, G. Huang, and M. Meng, “Regularized context gates on transformer for machine translation,” in *Proc. of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, Online, Jul. 2020, pp. 8555–8562.
- [18] M. A. Di Gangi, R. Cattoni, L. Bentivogli *et al.*, “MuST-C: a Multilingual Speech Translation Corpus,” in *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, Minneapolis, Minnesota, Jun. 2019, p. 2012–2017.
- [19] C. Dyer, V. Chahuneau, and N. A. Smith, “A Simple, Fast, and Effective Reparameterization of IBM Model 2,” in *Proc. of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, Atlanta, Georgia, Jun. 2013, pp. 644–648.
- [20] S. Meignier and T. Merlin, “LIUM SpkDiarization: An Open Source Toolkit For Diarization,” in *Proc. of the CMU SPUD Workshop*, Dallas, Texas, Mar. 2010.
- [21] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the Inception Architecture for Computer Vision,” in *Proc. of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, Nevada, Jun. 2016, pp. 2818–2826.
- [22] D. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *Proc. of 3rd International Conference on Learning Representations (ICLR)*, San Diego, California, May 2015.
- [23] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. of Interspeech 2019*, Graz, Austria, Sep. 2019, pp. 2613–2617.
- [24] M. Gaido, M. A. Di Gangi, M. Negri, and M. Turchi, “End-to-End Speech-Translation with Knowledge Distillation: FBK@IWSLT2020,” in *Proc. of the 17th International Conference on Spoken Language Translation (IWSLT)*, Online, Jul. 2020, pp. 80–88.
- [25] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, “fairseq: A Fast, Extensible Toolkit for Sequence Modeling,” in *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, Minneapolis, Minnesota, Jun. 2019, pp. 48–53.
- [26] P. Koehn, H. Hoang, A. Birch *et al.*, “Moses: Open Source Toolkit for Statistical Machine Translation,” in *Proc. of the 45th Annual Meeting of the Association for Computational Linguistics (ACL)*, Prague, Czech Republic, Jun. 2007, pp. 177–180.
- [27] R. Sennrich, B. Haddow, and A. Birch, “Neural Machine Translation of Rare Words with Subword Units,” in *Proc. of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, Berlin, Germany, Aug. 2016, pp. 1715–1725.
- [28] G. Neubig, M. Sperber, X. Wang *et al.*, “XNMT: The eXtensible Neural Machine Translation Toolkit,” in *Proc. of the 13th Conference of the Association for Machine Translation in the Americas (AMTA)*, Boston, Massachusetts, Mar. 2018, pp. 185–192.
- [29] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: a Method for Automatic Evaluation of Machine Translation,” in *Proc. of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, Philadelphia, Pennsylvania, Jul. 2002, pp. 311–318.
- [30] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul, “A Study of Translation Edit Rate with Targeted Human Annotation,” in *Proc. of the 7th Conference of the Association for Machine Translation in the Americas (AMTA)*, Cambridge, Massachusetts, Aug. 2006, pp. 223–231.
- [31] K. Lee, O. Firat, A. Agarwal, C. Fannjiang, and D. Sussillo, “Hallucinations in Neural Machine Translation,” in *Proc. of NIPS 2018 Interpretability and Robustness for Audio, Speech and Language Workshop*, Montréal, Canada, Dec. 2018.