

Attention as a Guide for Simultaneous Speech Translation

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

[◇]Fondazione Bruno Kessler

[□]University of Trento

[△]Independent Researcher

{spapi, negri}@fbk.eu, marco.turchi@gmail.com

Abstract

In simultaneous speech translation (SimulST), effective policies that determine when to write partial translations are crucial to reach high output quality with low latency. Towards this objective, we propose EDATT (**E**ncoder-**D**ecoder **A**ttention), an adaptive policy that exploits the attention patterns between audio source and target textual translation to guide an offline-trained ST model during simultaneous inference. EDATT exploits the attention scores modeling the audio-translation relation to decide whether to emit a partial hypothesis or wait for more audio input. This is done under the assumption that, if attention is focused towards the most recently received speech segments, the information they provide can be insufficient to generate the hypothesis (indicating that the system has to wait for additional audio input). Results on en→{de, es} show that EDATT yields better results compared to the SimulST state of the art, with gains respectively up to 7 and 4 BLEU points for the two languages, and with a reduction in computational-aware latency up to 1.4s and 0.7s compared to existing SimulST policies applied to offline-trained models.

1 Introduction

In simultaneous speech translation (SimulST), systems have to generate translations incrementally while concurrently receiving audio input. This requirement poses a significant challenge since the need of generating high-quality outputs has to be balanced with the need to minimize their latency, i.e. the time elapsed (lagging) between when a word is uttered and when it is actually translated by the system.

In direct SimulST systems (Bérard et al., 2016; Weiss et al., 2017),¹ the balance between output

¹In this paper, we focus on direct models that exhibit lower latency and better performance compared to traditional cascade architectures composed of separate automatic speech recognition and machine translation components (Ansari et al., 2020; Anastasopoulos et al., 2021, 2022).

quality and latency is managed by a *decision policy*, which is the strategy for determining, at each time step, whether to emit a partial translation or to wait for additional audio input. Decision policies can be divided into two categories: *fixed* and *adaptive*. Fixed policies are usually based on simple heuristics (Ma et al., 2019), while adaptive policies take into account the actual input content to make the decisions (Zheng et al., 2020). Recent works (Liu et al., 2021b; Zaidi et al., 2021, 2022; Zhang and Feng, 2022) proved the superiority of adaptive policies over fixed ones. However, a major limitation of these policies is that they require training *ad-hoc* and complex SimulST architectures, which results in high computational costs.

Computational costs are also inflated by the common practice of simulating the simultaneous test conditions by providing partial input during training to avoid the quality drops caused by the mismatch between training and test conditions (Ren et al., 2020; Ma et al., 2020b, 2021; Han et al., 2020; Zeng et al., 2021; Liu et al., 2021a; Zaidi et al., 2021, 2022). This practice is independent of the decision policy adopted, and typically requires dedicated trainings for each latency regime. To mitigate this issue, offline-trained ST systems have been employed for simultaneous inference (Liu et al., 2020; Chen et al., 2021; Nguyen et al., 2021) and, along this direction, Papi et al. (2022a) demonstrated that dedicated trainings simulating the inference conditions are not necessary since offline-trained systems outperform those specifically trained for SimulST. The effectiveness of using offline-trained ST models for simultaneous inference has been also confirmed by the last IWSLT 2022 evaluation campaign (Anastasopoulos et al., 2022), where the winning submission to the SimulST task (Polák et al., 2022) is an offline model exploiting the Local Agreement policy by Liu et al. (2020). However, despite its good results, this policy relies on a strategy (the generation of

two consecutive hypotheses prior to starting the emission) that has a significant impact on latency. This raises the need for effective policies that *i)* are adaptive, *ii)* are directly applicable to offline ST models, and *iii)* achieve low latency at low computational costs.

Towards these objectives, we propose EDATT (**E**ncoder-**D**ecoder **A**ttention),² a novel adaptive policy for SimulST that leverages the encoder-decoder attention patterns of an offline-trained ST model to decide when to emit partial translations. In a nutshell, our idea is that the next word of the partial hypothesis at a given time step is safely emitted only if the system does not attend to the most recent audio frames, meaning that the information received up to that time step is sufficient to generate that word. Building on this idea, our contributions are summarized as follows:

- We introduce EDATT, a novel adaptive decision policy for SimulST, which guides offline-trained ST models during simultaneous inference by looking at the attention patterns dynamically computed from the audio input over time;
- We show that EDATT outperforms the Local Agreement policy applied to the same offline ST models at almost all latency regimes, with computational-aware average lagging (AL_CA) reductions up to 1.4s for German and 0.7s for Spanish on MuST-C (Cattoni et al., 2021);
- We show that EDATT also outperforms the state-of-the-art CAAT architecture (Liu et al., 2021b), especially in terms of AL_CA, with gains of up to 7.0 BLEU for German and 4.0 BLEU for Spanish.

2 Background

In terms of architectural choices, Transformer (Vaswani et al., 2017) and its derivatives (Gulati et al., 2020; Chang et al., 2020; Papi et al., 2021; Burchi and Vielzeuf, 2021; Kim et al., 2022; Andrusenko et al., 2022) are the *de-facto* standard both in offline and simultaneous ST (Ansari et al., 2020; Anastasopoulos et al., 2021, 2022).

A generic Transformer model is composed of an encoder, whose role is to map the input speech

²Code, outputs and offline ST models used for our experiments are released under Apache License 2.0 at: <https://github.com/hlt-mt/fbk-fairseq>.

sequence $\mathbf{X} = [x_1, \dots, x_n]$ into an internal representation, and a decoder, whose role is to generate the output textual sequence $\mathbf{Y} = [y_1, \dots, y_m]$ by exploiting the internal representation in an autoregressive manner (Graves, 2013), that is by consuming the previously generated output as additional input when generating the next one.

The encoder and the decoder are composed of a stack of identical blocks, whose components may vary depending on the particular Transformer-based architecture, although they all share the same dot-product attention mechanism (Chan et al., 2016). In general, the attention is a function that maps a query matrix Q and a pair of key-value matrices (K, V) to an output matrix (Bahdanau et al., 2016). The output is obtained as a weighted sum of V , whose weights are computed through a compatibility function between Q and K that, in the case of the scaled dot-product attention used in the original Transformer formulation, is:

$$A(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V$$

where d_k is the dimension of K . The attention A is computed on h heads in parallel, each applying learned linear projections W^Q , W^K , and W^V to the Q , K , and V matrices. These representations are then concatenated and projected using another learned matrix W^O , resulting in the final output:

$$\text{Multihead}(Q, K, V) = \text{Concat}(\text{head}_1, \text{head}_2, \dots, \text{head}_h)W^O$$

where $\text{head}_i = A(QW_i^Q, KW_i^K, VW_i^V)$.

In the encoder layers, Q , K , and V are computed from the same speech input sequence \mathbf{X} , realizing the so-called *self-attention* $A_{\text{self}}(\mathbf{X})$. Differently, in the decoder layer, two types of attention are computed sequentially: *self-attention*, and *encoder-decoder* (or cross) attention. In the encoder-decoder attention, Q comes from the previous decoder layer (or directly from the previously generated output \mathbf{Y} , in the case of the first decoder layer) while K and V come from the output of the encoder, hence the matrix can be expressed as $A_{\text{cross}}(\mathbf{X}, \mathbf{Y})$. In this work, we only exploit the encoder-decoder attention matrix to guide the model during simultaneous inference. Therefore, we use the notation A instead of A_{cross} for simplicity, and henceforth refer to this matrix as the encoder-decoder representation of a specific decoder layer d considering the attention head h .

3 EDATT policy

We propose to exploit the information contained in the encoder-decoder attention matrix of an offline ST model during inference to determine whether to wait for additional audio input or emit a partial translation. The use of attention as the core mechanism of our policy is motivated by related works in machine translation (MT) and language modeling, which prove that attention scores can encode syntactic dependencies (Raganato and Tiedemann, 2018; Htut et al., 2019) and language representations (Lamarre et al., 2022), as well as align source and target tokens (Tang et al., 2018; Zenkel et al., 2019; Garg et al., 2019; Chen et al., 2020). We posit (and demonstrate in Section 5) that this encoder-decoder attention relationship between source audio and target tokens also exists in offline ST models, and can be used to guide them during simultaneous inference.

Our approach builds on the following hypothesis (see Figure 1): at each time step, if the attention is focused towards the end of the input audio sequence (1), the system will probably need more information to correctly produce the current output candidate. On the contrary (2), if the attention concentrates on early audio frames (far enough from the last received ones), the current output candidate can be safely emitted because the early encoded information is sufficient. Accordingly, the model will continue to emit the next token of the partial hypothesis until the above condition is verified, that is until its encoder-decoder attention scores do not focus towards the end of the received speech segment. The rationale is that if the encoder-decoder attention of the predicted token points to the most recent speech information – i.e. attention scores are higher towards the last audio frames received – this information could be incomplete and therefore still insufficient to generate that token.

More formally, at each time step t , EDATT determines whether to emit the next token y_j , given the previously generated tokens $\mathbf{Y}_{j-1} = [y_1, \dots, y_{j-1}]$ and the partial audio input sequence \mathbf{X}_t , by looking at the sum of the last λ encoder-decoder attention weights of the vector $A_j(\mathbf{X}_t, \mathbf{Y}_{j-1})$. Specifically, y_j is emitted if:

$$\sum_{i=t-\lambda}^t A_{i,j}(\mathbf{X}_t, \mathbf{Y}_{j-1}) < \alpha, \quad \alpha \in (0, 1) \quad (1)$$

where α is a hyperparameter that controls the

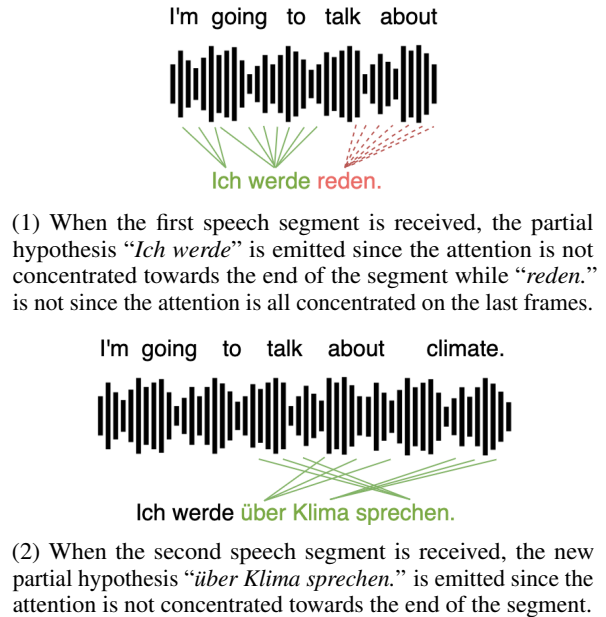


Figure 1: Example of the EDATT policy. Links indicate where the attention weights point to.

quality-latency trade-off: lower values of α increase the latency, as they reduce the possibility to satisfy Equation 1 (i.e. the sum of the last λ encoder-decoder attention weights will likely exceed α), and vice versa. When Equation 1 is satisfied, y_j is emitted and the same process is repeated for y_{j+1} , and so on. The process continues until we reach the token y_{j+w} for which Equation 1 is no longer verified. At that point, the emission is stopped and the total number of tokens emitted at time step t is w .

4 Experimental Settings

4.1 Data

To be comparable with previous works (Ren et al., 2020; Ma et al., 2020b; Zeng et al., 2021; Liu et al., 2021b; Papi et al., 2022a; Zhang and Feng, 2022), we train our models on MuST-C en→{de, es} (Cattoni et al., 2021). The choice of the two target languages is also motivated by their different word ordering: Subject-Object-Verb (SOV) for German and Subject-Verb-Object (SVO) for Spanish. This opens the possibility of validating our approach on target-language word orderings that are respectively different and similar with respect to the English (i.e. SVO) source audio. We also perform data augmentation by applying sequence-level knowledge distillation (Kim and Rush, 2016; Gaido et al., 2021b, 2022a) as in (Liu et al., 2021b; Papi et al., 2022a), for which the transcripts of

MuST-C $en \rightarrow \{de, es\}$ are translated with an MT model (more details can be found in Appendix A) and used together with the gold reference during training. Data statistics are given in Appendix B.

4.2 Architecture and Training Setup

For our experiments, we use the bug-free implementation by Papi et al. (2023) of the Conformer-based encoder-decoder model for ST (Guo et al., 2021). The offline model is made of 12 Conformer encoder layers (Gulati et al., 2020) and 6 Transformer decoder layers ($d_{max} = 6$) with a total of $\sim 115M$ parameters. Each encoder/decoder layer has 8 attention heads ($h_{max} = 8$). The input is represented as 80 audio features extracted every 10ms with sample window of 25 and processed by two 1D convolutional layers with stride 2 to reduce its length by a factor of 4 (Wang et al., 2020). Utterance-level Cepstral Mean and Variance Normalization (CMVN) and SpecAugment (Park et al., 2019) are applied during training. Detailed settings are described in Appendix A.

4.3 Inference and Evaluation

We use the SimulEval tool (Ma et al., 2020a) to simulate simultaneous conditions and evaluate all the models. For our policy, we vary α of Equation 1 in the range $[0.6, 0.4, 0.2, 0.1, 0.05, 0.03]$ and set the size of the speech segment to 800ms. During inference, the features are computed on the fly and CMVN normalization is based on the global mean and variance estimated on the MuST-C training set. All inferences are performed on a single NVIDIA K80 GPU with 12GB memory as in the IWSLT Simultaneous evaluation campaigns.

We use sacreBLEU (Post, 2018)³ to evaluate translation quality, and Average Lagging (Ma et al., 2019) – or AL – to evaluate latency, as in the default SimulEval evaluation setup. As suggested by Ma et al. (2020b), for our comparisons with other approaches we also report computational-aware average lagging (AL_CA), which measures the real elapsed time instead of the ideal one considered by AL, thus giving a more realistic latency measure when the system operates in real time. Its computation is also provided by SimulEval.

4.4 Terms of Comparison

We conduct experimental comparisons with the state-of-the-art architecture for SimulST (CAAT)

and, respectively, the current best (Local Agreement) and the most widely used (Wait-k) policies that can be directly applied to our offline ST systems for simultaneous inference. In detail:

Cross Attention Augmented Transformer (CAAT) – the state-of-the-art architecture for SimulST (Liu et al., 2021b), winner of the IWSLT 2021 SimulST task (Anastasopoulos et al., 2021). Inspired by the Recurrent Neural Network Transducer (Graves, 2012), it is made of three Transformer stacks: the encoder, the predictor, and the joiner. These three elements are jointly trained to optimize translation quality while keeping latency under control. We train and evaluate the CAAT model using the code provided by the authors,⁴ and on the same data used for our offline ST model.

Local Agreement (LA) – the state-of-the-art decision policy introduced by Liu et al. (2020), and used by the winning system at IWSLT 2022 (Anastasopoulos et al., 2022). It consists in generating a partial hypothesis from scratch each time a new speech segment is added, and emitting it – or part of it – if it coincides with one of those generated in the previous l time steps, where l is a hyperparameter. Since Liu et al. (2020) empirically found that considering only the most recent previously generated tokens ($l = 1$) as memory works better, we adopt the same strategy to apply this policy.

Wait-k – the simplest and most widely used decision policy in SimulST (Ren et al., 2020; Ma et al., 2020b; Zeng et al., 2021). It consists in waiting for a fixed number of words (k) before starting to emit the translation, and then proceeding by alternating waiting and writing operations. Since in SimulST the information about the number of words is not explicitly contained in the audio input, a word detection strategy is used to determine this information. Detection strategies can be fixed, when it is assumed that each word has a pre-defined fixed duration, or adaptive, when the information about the number of words is inferred from the audio content. Following Papi et al. (2022a), we adopt a CTC-based adaptive word detection strategy to detect the number of words. In addition, to be comparable with the other approaches, we employ beam search to generate each token.

³BLEU+case.mixed+smooth.exp+tok.13a+version.1.5.1

⁴<https://github.com/danliu2/caat>




Figure 2: Encoder-decoder attention scores on a random sample of the MuST-C en→de dev set, before (a) and after (b) the filtering of the last frame from the attention matrix.

5 Attention Analysis

To validate our hypothesis and study the feasibility of our method, we start by exploring the encoder-decoder attention matrices of the offline trained models. We proceed as follows: first, by visualizing the attention weights, we check for the existence of patterns that could be exploited during simultaneous inference. Then, we analyze the performance of the EDATT policy to discover the best value of λ , the decoder layer d , and the attention head h from which to extract the attention scores that better balance the quality-latency trade-off.

Do attention patterns exist also in ST? To answer this question, we conducted an analysis of the encoder-decoder matrices obtained from the MuST-C en-de dev set. Through the visualization of attention weights, we observed a consistent phenomenon across our two language directions (en→{de, es}): the attention weights concentrate on the last frame, regardless of the input length, as shown in Figure 2a. This behaviour has already been observed in prior works on attention analysis, showing that attention often concentrates on the initial or final token (Clark et al., 2019; Kovaleva et al., 2019; Kobayashi et al., 2020; Ferrando et al., 2022), with up to 97% of attention weights being allocated to these positions. As this might hinder the possibility to effectively visualize attention patterns, similarly to (Vig and Belinkov, 2019), we filtered out the last frame from the attention matrix and then re-normalized it. In this way, as shown in Figure 2b, we obtained a clear pseudo-diagonal pattern compared to the previous unfiltered representation. Such correspondence emerging from the

encoder-decoder attention scores after the removal of the last frame indicates a relationship between the source audio frames and target translation texts that can be exploited by our adaptive attention-based policy during simultaneous inference.




Figure 3: Effect of λ on MuST-C en→{de, es} dev set. We visualize the results with $AL \leq 2.5s$.

What is the optimal value of λ ? To find the best number of frames (λ) on which to apply Equation 1, we analyse the behavior of EDATT by varying α and setting $\lambda \in [2, 4, 6, 8]$.⁵ For this analysis, we extract the attention scores from the 5th decoder layer ($d = 5$) by averaging across the matrices obtained from each attention head ($h = [1, \dots, 8]$) in accordance with the findings of (Garg et al., 2019) about the layer that best represents word alignment.

⁵We do not report the experiments with $\lambda = 1$ since we found that it consistently degrades translation quality. We also experimented with different ways to determine λ , such as using a percentage instead of a fixed number, but none of them yielded significant differences.

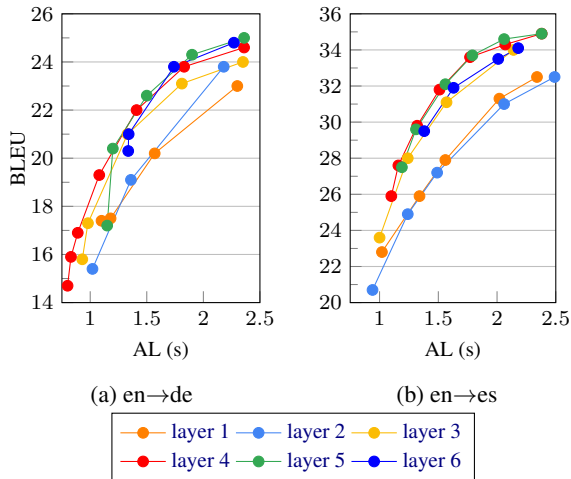


Figure 4: SimulST results on MuST-C dev set $en \rightarrow \{de, es\}$ for each decoder layer d . We visualize the results with $AL \leq 2.5s$.

We perform the analysis on the MuST-C dev set for both language pairs, and present the results in Figure 3. As we can see, as the value of λ increases, the curves shift towards the right, indicating an increase in latency. This means that, consistently across languages, considering too many frames towards the end ($\lambda \geq 6$) affects latency with little effect on quality. Since $\lambda = 2$ yields the lowest latency ($AL \approx 1.2s$) in both languages, and especially in Spanish, we select this value for the following experiments. This outcome is noteworthy as it demonstrates that, at least in our settings, the same optimal value of λ applies to diverse target languages with different word ordering. However, this might not hold for different source and/or target languages, advocating for future explorations as discussed in the Limitations section.

What is the best layer? After determining the optimal value of λ , we proceed to analyze the EDATT performance by varying the decoder layer from which the encoder-decoder attention is extracted. We conduct this study by using $\lambda = 2$, as previously determined to be the optimal value for both languages. In Figure 4, we present the SimulST results (in terms of AL-BLEU curves) for each decoder layer $d = [1, \dots, 6]$.⁶ As we can see, on both languages, Layers 1 and 2 consistently perform worse than the other layers. Also, Layer 3 achieves inferior quality compared to Layers ≥ 4 , especially at medium-high latency ($AL \geq 1.2s$) despite performing better than Layers 1 and 2.

⁶We also tried to make the average of the encoder-decoder attention matrices of each layer but this led to worse results.

Head	en→de			en→es		
	1.2s	1.6s	2s	1.2s	1.6s	2s
Head 1	17.6	19.2	20.5	27.6	30.8	32.1
Head 2	19.0	21.9	23.4	-	31.9	33.9
Head 3	-	22.3	23.9	27.2	29.8	31.1
Head 4	-	21.5	23.3	-	28.4	30.7
Head 5	19.2	22.2	23.8	-	30.9	32.5
Head 6	18.7	21.2	22.7	-	32.0	33.3
Head 7	-	21.9	23.5	-	30.8	32.6
Head 8	19.2	20.7	21.6	-	31.7	33.9
Average	20.3	22.8	24.0	28.6	32.4	34.1

Table 1: BLEU scores on MuST-C dev set $en \rightarrow \{de, es\}$ for each attention head h of Layer 4. Latency (AL) is reported in seconds. “-” means that the BLEU value is not available or calculable. The last row represents the numerical values of Layer 4 curves of Figure 4 obtained by averaging across all 8 heads.

This aligns with the findings of Garg et al. (2019), which observed inferior performance by the first three layers in the alignment task for MT models. Concerning Layer 6, both graphs show that the curves cannot achieve lower latency, starting at around $1.5s$ of AL. This phenomenon is also valid for Layer 5 compared to Layer 4, although being much less pronounced. We also observe that Layer 5 achieves the best performance at higher latency on both languages. However, since Layers 5 and 6 never achieve low latency (AL never approaches $1.2s$), we can conclude that the optimal choice for the simultaneous scenario is Layer 4. This is in line with Lamarre et al. (2022), which indicates the middle layers as the best choice to provide accurate predictions for language representations. As a consequence, we will use $d = 4$ for the subsequent experiments with EDATT.

Would a single attention head encode more useful information? According to prior research examining the usefulness of selecting a single or a set of attention heads to perform natural language processing and translation tasks (Jo and Myaeng, 2020; Behnke and Heafeld, 2020; Gong et al., 2021), we also investigate the behavior of the EDATT policy by varying the attention head h from which the encoder-decoder attention matrix A is extracted. In Table 1,⁷ we present the results obtained from each attention head $h = [1, \dots, 8]$.⁸ Firstly, we observe

⁷A tabular format is used instead of AL-BLEU curves as many parts of the curves are indistinguishable from each other. $AL = 1.2s$ is the first latency measure reported because it is the minimum value spanned by the head-wise curves, and $AL = 2s$ is the last one since increasing latency above this value does not significantly improve translation quality (BLEU).

⁸Since obtaining a specific latency in seconds is not possible with this method, we interpolate the previous and succes-

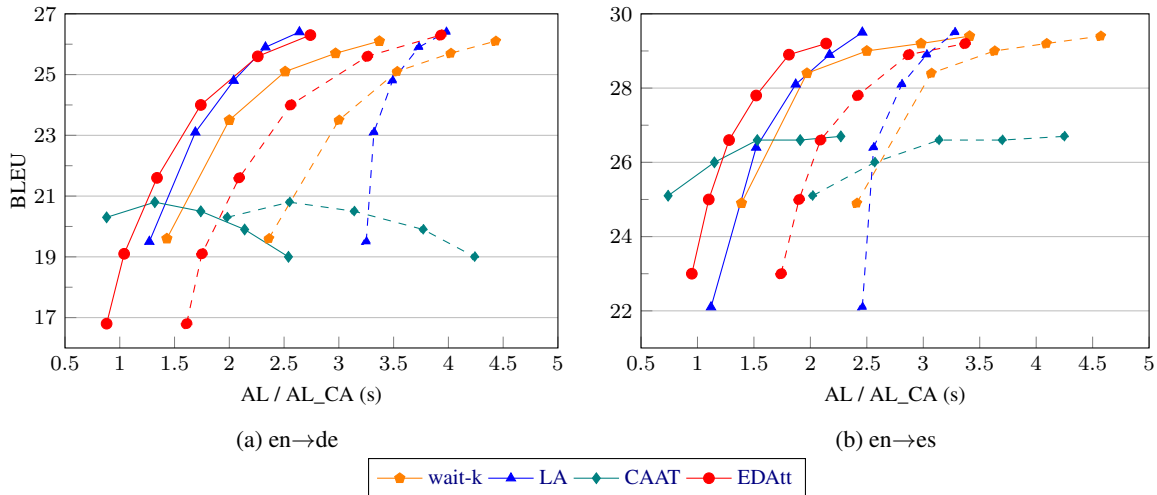


Figure 5: Comparison with the SimulST systems described in Section 4.4 on MuST-C en→{de, es} tst-COMMON. Solid curves represent AL, dashed curves represent AL_CA.

that many heads are unable to achieve low latency, particularly for Spanish. Furthermore, there is no consensus on the optimal head among languages or at different latencies (e.g. Head 6 is the best in Spanish at 1.6s, but it does not achieve lower latency). However, we notice that the average across all heads (last row) has an overall better performance compared to the encoder-decoder matrices extracted from each individual head, and this holds true for both languages. Consequently, we choose to compute the average over the attention heads to apply our EDAtt policy in order to achieve a better quality-latency trade-off for SimulST.

6 Results

6.1 Comparison with Other Approaches

For the comparison of EDAtt with the SimulST systems described in Section 4.4, we report in Figure 5 both AL (solid curves) and AL_CA (dashed curves) as latency measures to give a more realistic evaluation of the performance of the systems in real time, as recommended in (Ma et al., 2020b; Papi et al., 2022a). Results with other metrics, DAL (Cherry and Foster, 2019) and LAAL (Papi et al., 2022b), are provided in Appendix C for completeness. Numeric values for all the plots are presented in Section D. For our policy, we extract the encoder-decoder attention matrix from Layer 4 ($d = 4$), average the weights across heads, and set $\lambda = 2$ as it was found to be the optimal setting on the MuST-C dev set for both languages, as previously discussed in Section 5.

sive points to estimate the BLEU value, when needed.

Quality-latency curves for en→de and en→es show similar trends. The EDAtt policy achieves better overall results compared to the LA and wait-k policies applied to offline ST models. EDAtt consistently outperforms the wait-k policy, with gains ranging from 1.0 to 2.5 BLEU for German and 1.0 to 3 for Spanish, when considering both ideal (AL) and computationally aware (AL_CA) latency measures. Additionally, it is able to achieve lower latency, as the starting point of the wait-k policy is always around 1.5s, while EDAtt starts at 1.0s. In comparison to the LA policy, we observe an AL_CA reduction of up to 1.4s for German and 0.7s for Spanish. Moreover, the computational overhead of EDAtt is consistently lower, 0.9s on average between languages, against 1.3s of LA. Therefore, the computational cost of our policy is 30% lower compared to the LA policy. Additionally, EDAtt outperforms LA at almost every latency, with gains up to 2.0 BLEU for German and 3.0 for Spanish.

Compared with CAAT, when ideal latency is considered (solid curves), we notice that EDAtt achieves higher quality at medium-high latency ($AL \geq 1.2s$), with BLEU gains up to 5.0 points for German and 2.0 for Spanish. When $AL < 1.2s$, instead, there is a decrease in performance with BLEU drops ranging from 1.5 to 4.0 for German and 1.0 to 2.5 for Spanish. However, when considering the realistic computational-aware latency measure AL_CA (dashed curves), we observe that the EDAtt curves are always to the left of those of the CAAT system, indicating that our policy always outperforms it with BLEU gains up to 6.0 points

for German and 2.0 for Spanish.

In light of this, we can conclude that EDATT achieves new state-of-the-art results in terms of computational-aware metrics, while also being superior at medium-high latency when considering the less realistic computational-unaware measure.

6.2 Effects of Accelerated Hardware

To further investigate the computational efficiency of EDATT, we conducted experiments on all the systems described in Section 4.4 using a highly accelerated GPU, an NVIDIA A40 with 48GB memory, during simultaneous inference.

Figure 6 reports the results in terms of quality-latency trade-off. When comparing the curves with the computationally aware ones in Figure 5 (dashed), it can be observed that the LA policy seems to benefit more from the use of expensive accelerated hardware, with a latency reduction of 0.5-1s. However, this reduction is not sufficient to reach a latency lower than 2s with this policy. Considering the other systems, both wait-k and CAAT curves show a slight left shift (by less than 0.5s), similar to EDATT.⁹

In conclusion, our policy proved to be superior even when using accelerated and expensive hardware, further strengthening the previously discussed findings. Moreover, these results indicate that there are no significant differences between the systems when using less or more accelerated GPU hardware and advocate for the wider use of computationally aware metrics in future research.

7 Related Works

The first policy for SimulST was proposed by Ren et al. (2020) and is derived from the wait-k policy (Ma et al., 2019) developed for simultaneous *text-to-text* translation. Most of subsequent studies have also adopted the wait-k policy (Ma et al., 2020b; Han et al., 2020; Chen et al., 2021; Zeng et al., 2021; Karakanta et al., 2021; Nguyen et al., 2021; Papi et al., 2022a). In parallel, several strategies have been developed to directly learn the best policy during training by means of *ad-hoc* architectures (Ma et al., 2021; Liu et al., 2021a,b; Chang and Lee, 2022) and training procedures aimed at

⁹Despite the benefits in terms of quality-latency trade-off, the significantly higher costs of the A40 GPU over the K80 GPU (4.1 vs 0.9 USD/h in Amazon Web Services, <https://aws.amazon.com/it/ec2/pricing/on-demand/>) makes unlikely that such a GPU will soon be of widespread use for simultaneous inference.

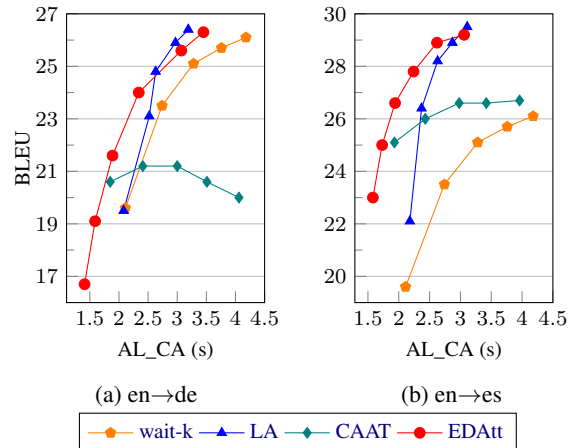


Figure 6: Effect of using NVIDIA A40 GPU on MuST-C en->{de, es} tst-COMMON considering all the systems of Section 4.4. Results are computationally aware.

reducing latency (Liu et al., 2021a,b; Zaidi et al., 2021, 2022; Chang and Lee, 2022; Zhang and Feng, 2022; Omachi et al., 2022). The latter adaptive policies obtained better performance according to the most recent results observed in (Anastasopoulos et al., 2021, 2022). We define our policy as adaptive as well, as it relies on the encoder-decoder attention mechanism, whose dynamics are influenced by the audio input that increases incrementally over time. However, EDATT completely differs from prior works on adaptive policies that exploit attention (Zaidi et al., 2021, 2022; Chang and Lee, 2022; Zhang and Feng, 2022) because is the first policy that does not require influencing the behaviour of the attention weights through dedicated training strategies, therefore being directly applicable to offline-trained ST models. By doing so, we realize *i*) an adaptive policy, *ii*) directly applicable to offline-trained ST models, *iii*) which achieves low latency at low computational costs.

8 Conclusions

After investigating the encoder-decoder attention behavior of offline ST models, we presented EDATT, a novel adaptive decision policy for SimulST that guides an offline ST model to wait or to emit a partial hypothesis by looking at its encoder-decoder attention weights. Comparisons with state-of-the-art SimulST architectures and decision policies reveal that, at lower computational costs, EDATT outperforms the others at almost every latency, with translation quality gains of up to 7.0 BLEU for en->de and 4.0 BLEU for en->es. Moreover, it is also capable of achieving a

computational-aware latency of less than 2s with a reduction of 0.7-1.4s compared to existing decision policies applied to the same offline ST systems.

Acknowledgments

The authors thank Marco Gaido for his valuable support during the paper writing. We acknowledge the support of the PNRR project FAIR - Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU, and of the project “AI@TN” funded by the Autonomous Province of Trento, Italy.

Limitations

Although applicable to any offline ST models, the EDATT policy and its behavior have been analysed on models applying CTC compression. Thus, the audio input undergoes a transformation that does not only reduce its dimension but also compresses it into more meaningful units, similar to words or subwords. As a consequence, the hyper-parameters regarding the number of frames to which apply the policy (λ) can vary and depend on the specific ST model. This would require having a validation set on which to search the best value of λ before directly testing. Moreover, the EDATT policy has been tested on Western European languages and, even if there is no reason suggesting that this cannot be applied (after a proper hyper-parameter search) to other languages, its usage on non-Western European target languages and on a source language different from English has not been verified in this work and is left for future endeavours.

References

Antonios Anastasopoulos, Loïc Barrault, Luisa Benivogli, Marceley Zanon Boito, Ondřej Bojar, Roldano Cattoni, Anna Currey, Georgiana Dinu, Kevin Duh, Maha Elbayad, Clara Emmanuel, Yannick Estève, Marcello Federico, Christian Federmann, Souhir Gahbiche, Hongyu Gong, Roman Grundkiewicz, Barry Haddow, Benjamin Hsu, Dávid Javorský, Věra Kloudová, Surafel Lakew, Xutai Ma, Prashant Mathur, Paul McNamee, Kenton Murray, Maria Nădejde, Satoshi Nakamura, Matteo Negri, Jan Niehues, Xing Niu, John Ortega, Juan Pino, Elizabeth Salesky, Jiatong Shi, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Yogesh Virkar, Alexander Waibel, Changhan Wang, and Shinji Watanabe. 2022. [Findings of the IWSLT 2022 evaluation campaign](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 98–157, Dublin, Ireland (in-person and online).

Antonios Anastasopoulos, Ondřej Bojar, Jacob Bremerman, Roldano Cattoni, Maha Elbayad, Marcello Federico, Xutai Ma, Satoshi Nakamura, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Alex Waibel, Changhan Wang, and Matthew Wiesner. 2021. [Findings of the IWSLT 2021 Evaluation Campaign](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, Online.

Andrei Andrusenko, Rauf Nasretidinov, and Aleksei Romanenko. 2022. [Uconv-conformer: High reduction of input sequence length for end-to-end speech recognition](#). *arXiv preprint arXiv:2208.07657*.

Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, Alexander Waibel, and Changhan Wang. 2020. [FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34, Online.

Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philémon Brakel, and Yoshua Bengio. 2016. [End-to-end attention-based large vocabulary speech recognition](#). In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4945–4949.

Maximiliana Behnke and Kenneth Heafield. 2020. [Losing heads in the lottery: Pruning transformer attention in neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2664–2674, Online.

Maxime Burchi and Valentin Vielzeuf. 2021. [Efficient conformer: Progressive downsampling and grouped attention for automatic speech recognition](#). In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 8–15.

Alexandre Bérard, Olivier Pietquin, Christophe Serivan, and Laurent Besacier. 2016. [Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation](#). In *NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain.

Roldano Cattoni, Mattia Antonino Di Gangi, Luisa Benivogli, Matteo Negri, and Marco Turchi. 2021. [Must-c: A multilingual corpus for end-to-end speech translation](#). *Computer Speech & Language*, 66:101155.

William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. 2016. [Listen, attend and spell: A neural network for large vocabulary conversational speech recognition](#). In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964.

- Chih-Chiang Chang and Hung-Yi Lee. 2022. [Exploring Continuous Integrate-and-Fire for Adaptive Simultaneous Speech Translation](#). In *Proc. Interspeech 2022*, pages 5175–5179.
- Xuankai Chang, Aswin Shanmugam Subramanian, Pengcheng Guo, Shinji Watanabe, Yuya Fujita, and Motoi Omachi. 2020. End-to-end asr with adaptive span self-attention. In *INTERSPEECH*.
- Junkun Chen, Mingbo Ma, Renjie Zheng, and Liang Huang. 2021. [Direct simultaneous speech-to-text translation assisted by synchronized streaming ASR](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4618–4624, Online.
- Yun Chen, Yang Liu, Guanhua Chen, Xin Jiang, and Qun Liu. 2020. [Accurate word alignment induction from neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 566–576, Online.
- Colin Cherry and George Foster. 2019. [Thinking slow about latency evaluation for simultaneous machine translation](#).
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy.
- Mattia A. Di Gangi, Marco Gaido, Matteo Negri, and Marco Turchi. 2020. On Target Segmentation for Direct Speech Translation. In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (AMTA 2020)*, pages 137–150, Virtual.
- Javier Ferrando, Gerard I Gállego, Belen Alastruey, Carlos Escolano, and Marta R Costa-jussà. 2022. Towards opening the black box of neural machine translation: Source and target interpretations of the transformer. *arXiv e-prints*, pages arXiv–2205.
- Marco Gaido, Mauro Cettolo, Matteo Negri, and Marco Turchi. 2021a. [CTC-based compression for direct speech translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 690–696, Online.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2021b. [On Knowledge Distillation for Direct Speech Translation](#). In *Proceedings of CLiC-IT 2020*, Online.
- Marco Gaido, Matteo Negri, and Marco Turchi. 2022a. Direct speech-to-text translation models as students of text-to-text models. *Italian Journal of Computational Linguistics*.
- Marco Gaido, Sara Papi, Dennis Fucci, Giuseppe Fiameni, Matteo Negri, and Marco Turchi. 2022b. [Efficient yet competitive speech translation: FBK@IWSLT2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 177–189, Dublin, Ireland (in-person and online).
- Sarthak Garg, Stephan Peitz, Udhayakumar Nallasamy, and Matthias Paulik. 2019. [Jointly learning to align and translate with transformer models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4453–4462, Hong Kong, China.
- Hongyu Gong, Yun Tang, Juan Pino, and Xian Li. 2021. Pay better attention to attention: Head selection in multilingual and multi-domain sequence modeling. *Advances in Neural Information Processing Systems*, 34:2668–2681.
- Alex Graves. 2012. Sequence transduction with recurrent neural networks. *arXiv preprint arXiv:1211.3711*.
- Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006. Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks. In *Proceedings of the 23rd international conference on Machine learning (ICML)*, pages 369–376, Pittsburgh, Pennsylvania.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented Transformer for Speech Recognition](#). In *Proc. Interspeech 2020*, pages 5036–5040.
- Pengcheng Guo, Florian Boyer, Xuankai Chang, Tomoki Hayashi, Yosuke Higuchi, Hirofumi Inaguma, Naoyuki Kamo, Chenda Li, Daniel Garcia-Romero, Jiatong Shi, Jing Shi, Shinji Watanabe, Kun Wei, Wangyou Zhang, and Yuekai Zhang. 2021. [Recent developments on espnet toolkit boosted by conformer](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5874–5878.
- Hou Jeung Han, Mohd Abbas Zaidi, Sathish Reddy Indurthi, Nikhil Kumar Lakumarapu, Beomseok Lee, and Sangha Kim. 2020. [End-to-end simultaneous translation system for IWSLT2020 using modality agnostic meta-learning](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 62–68, Online.

- Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R Bowman. 2019. Do attention heads in bert track syntactic dependencies? *arXiv preprint arXiv:1911.12246*.
- Jae-young Jo and Sung-Hyon Myaeng. 2020. Roles and utilization of attention heads in transformer-based neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3404–3417, Online.
- Alina Karakanta, Sara Papi, Matteo Negri, and Marco Turchi. 2021. Simultaneous speech translation for live subtitling: from delay to display. In *Proceedings of the 1st Workshop on Automatic Spoken Language Translation in Real-World Settings (ASLTRW)*, pages 35–48, Virtual.
- Sehoon Kim, Amir Gholami, Albert Shaw, Nicholas Lee, Karttikeya Mangalam, Jitendra Malik, Michael W Mahoney, and Kurt Keutzer. 2022. Squeezeformer: An efficient transformer for automatic speech recognition. *arxiv:2206.00888*.
- Yoon Kim and Alexander M. Rush. 2016. Sequence-Level Knowledge Distillation. In *Proc. of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. Attention is not only a weight: Analyzing transformers with vector norms. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of BERT. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4365–4374, Hong Kong, China.
- Mathis Lamarre, Catherine Chen, and Fatma Deniz. 2022. Attention weights accurately predict language representations in the brain. *bioRxiv*.
- Dan Liu, Mengge Du, Xiaoxi Li, Yuchen Hu, and Lirong Dai. 2021a. The USTC-NELSLIP systems for simultaneous speech translation task at IWSLT 2021. In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 30–38, Bangkok, Thailand (online).
- Dan Liu, Mengge Du, Xiaoxi Li, Ya Li, and Enhong Chen. 2021b. Cross attention augmented transducer networks for simultaneous translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 39–55, Online and Punta Cana, Dominican Republic.
- Danni Liu, Gerasimos Spanakis, and Jan Niehues. 2020. Low-Latency Sequence-to-Sequence Speech Recognition and Translation by Partial Hypothesis Selection. In *Proc. Interspeech 2020*, pages 3620–3624.
- Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. 2019. STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3036, Florence, Italy.
- Xutai Ma, Mohammad Javad Dousti, Changan Wang, Jiatao Gu, and Juan Pino. 2020a. SIMULEVAL: An evaluation toolkit for simultaneous translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 144–150, Online.
- Xutai Ma, Juan Pino, and Philipp Koehn. 2020b. SimulMT to SimulST: Adapting simultaneous text translation to end-to-end simultaneous speech translation. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 582–587, Suzhou, China.
- Xutai Ma, Yongqiang Wang, Mohammad Javad Dousti, Philipp Koehn, and Juan Pino. 2021. Streaming simultaneous speech translation with augmented memory transformer. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7523–7527. IEEE.
- Ha Nguyen, Yannick Estève, and Laurent Besacier. 2021. An empirical study of end-to-end simultaneous speech translation decoding strategies. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7528–7532. IEEE.
- Motoi Omachi, Brian Yan, Siddharth Dalmia, Yuya Fujita, and Shinji Watanabe. 2022. Align, write, re-order: Explainable end-to-end speech translation via operation sequence generation. *arXiv preprint arXiv:2211.05967*.
- Sara Papi, Marco Gaido, Matteo Negri, and Andrea Pilzer. 2023. Reproducibility is nothing without correctness: The importance of testing code in nlp. *ArXiv*, abs/2303.16166.
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2021. Speechformer: Reducing information loss in direct speech translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1698–1706, Online and Punta Cana, Dominican Republic.
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2022a. Does simultaneous speech translation need simultaneous models? In *Findings of the*

- Association for Computational Linguistics: EMNLP 2022*, pages 141–153, Abu Dhabi, United Arab Emirates.
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2022b. [Over-generation cannot be rewarded: Length-adaptive average lagging for simultaneous speech translation](#). In *Proceedings of the Third Workshop on Automatic Simultaneous Translation*, pages 12–17, Online.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. [SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition](#). In *Proc. Interspeech 2019*, pages 2613–2617.
- Peter Polák, Ngoc-Quan Pham, Tuan Nam Nguyen, Danni Liu, Carlos Mullov, Jan Niehues, Ondřej Bojar, and Alexander Waibel. 2022. [CUNI-KIT system for simultaneous speech translation task at IWSLT 2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 277–285, Dublin, Ireland (in-person and online).
- Matt Post. 2018. [A Call for Clarity in Reporting BLEU Scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels.
- Alessandro Raganato and Jörg Tiedemann. 2018. [An analysis of encoder representations in transformer-based machine translation](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 287–297, Brussels, Belgium.
- Yi Ren, Jinglin Liu, Xu Tan, Chen Zhang, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2020. [SimulSpeech: End-to-end simultaneous speech to text translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3787–3796, Online.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proc. of 2016 IEEE CVPR*, pages 2818–2826, Las Vegas, Nevada, United States.
- Gongbo Tang, Rico Sennrich, and Joakim Nivre. 2018. [An analysis of attention mechanisms: The case of word sense disambiguation in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 26–35, Brussels, Belgium.
- Jörg Tiedemann. 2016. [OPUS – parallel corpora for everyone](#). In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation: Projects/Products*, Riga, Latvia.
- 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*, volume 30.
- Jesse Vig and Yonatan Belinkov. 2019. [Analyzing the structure of attention in a transformer language model](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 63–76, Florence, Italy.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020. [fairseq s2t: Fast speech-to-text modeling with fairseq](#). In *Proceedings of the 2020 Conference of the Association for Computational Linguistics (ACL): System Demonstrations*.
- Ron J. Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. [Sequence-to-Sequence Models Can Directly Translate Foreign Speech](#). In *Proceedings of Interspeech 2017*, pages 2625–2629, Stockholm, Sweden.
- Mohd Abbas Zaidi, Beomseok Lee, Sangha Kim, and Chanwoo Kim. 2022. [Cross-Modal Decision Regularization for Simultaneous Speech Translation](#). In *Proc. Interspeech 2022*, pages 116–120.
- Mohd Abbas Zaidi, Beomseok Lee, Nikhil Kumar Lakumarapu, Sangha Kim, and Chanwoo Kim. 2021. [Decision attentive regularization to improve simultaneous speech translation systems](#). *arXiv preprint arXiv:2110.15729*.
- Xingshan Zeng, Liangyou Li, and Qun Liu. 2021. [Real-Trans: End-to-end simultaneous speech translation with convolutional weighted-shrinking transformer](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2461–2474, Online.
- Thomas Zenkel, Joern Wuebker, and John DeNero. 2019. [Adding interpretable attention to neural translation models improves word alignment](#). *arXiv preprint arXiv:1901.11359*.
- Shaolei Zhang and Yang Feng. 2022. [Information-transport-based policy for simultaneous translation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 992–1013, Abu Dhabi, United Arab Emirates.
- Baigong Zheng, Kaibo Liu, Renjie Zheng, Mingbo Ma, Hairong Liu, and Liang Huang. 2020. [Simultaneous translation policies: From fixed to adaptive](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2847–2853, Online.

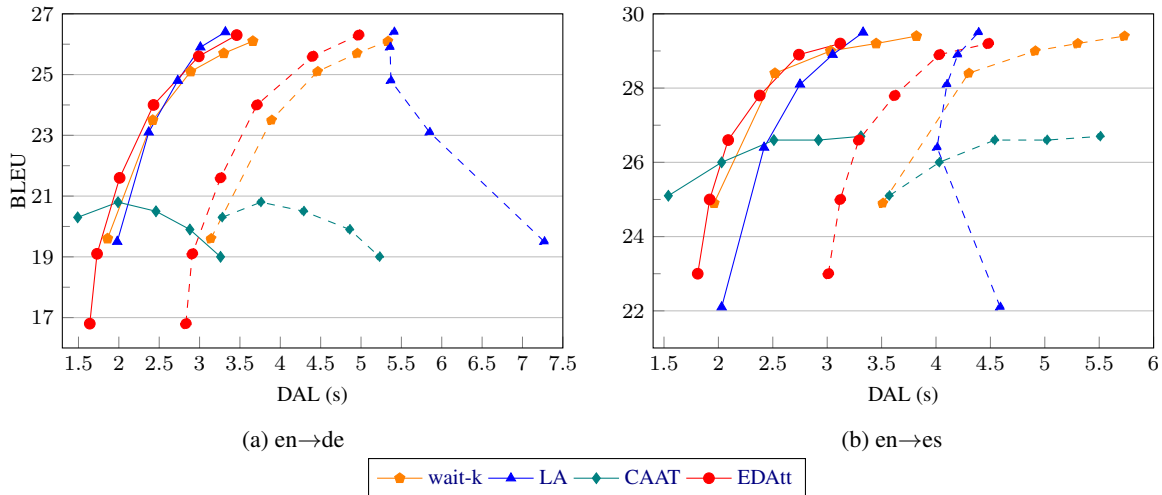


Figure 7: DAL results for the SimulST systems of Section 4.4. Solid curves represent DAL, dashed curves represent DAL_CA.

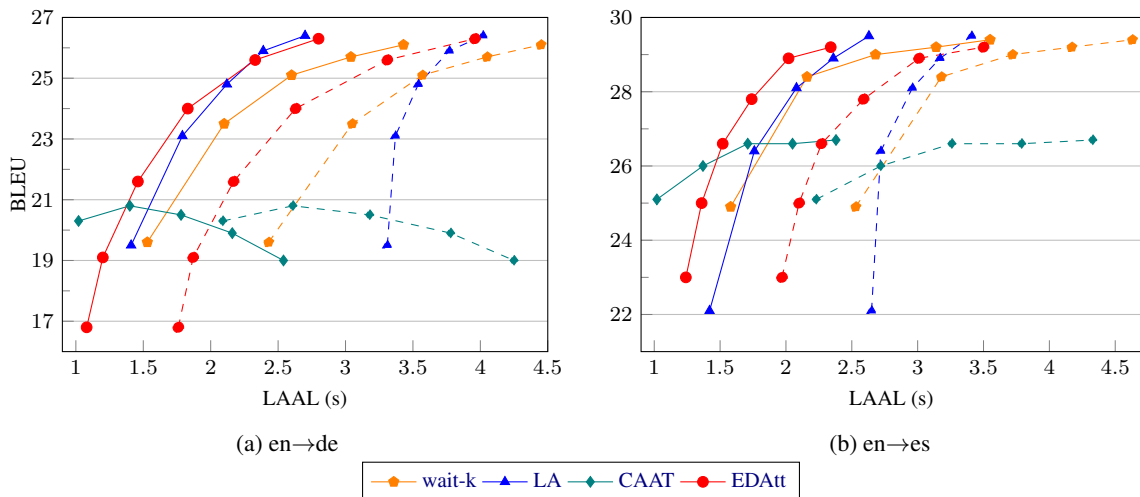


Figure 8: LAAL results for the SimulST systems of Section 4.4. Solid curves represent LAAL, dashed curves represent LAAL_CA.

A Training Settings

We use 512 as embedding size and 2,048 hidden neurons in the feed-forward layers both in the encoder and in the decoder. We set dropout at 0.1 for feed-forward, attention, and convolution layers. Also, in the convolution layer, we set 31 as kernel size for the point- and depth-wise convolutions. The vocabularies are based on Sentence-Piece (Sennrich et al., 2016) with dimension of 8,000 (Di Gangi et al., 2020) for the target side (de, es) and of 5,000 (Wang et al., 2020) for the source side (en). We optimize with Adam (Kingma and Ba, 2015) by using the label-smoothed cross-entropy loss with 0.1 as smoothing factor (Szegedy et al., 2016). We employ Connectionist Temporal Classification – or CTC – (Graves et al., 2006) as

auxiliary loss to avoid pre-training (Gaido et al., 2022b) and also to compress the input audio, reducing RAM consumption and speeding up inference (Gaido et al., 2021a). The learning rate is set to $5 \cdot 10^{-3}$ with Noam scheduler (Vaswani et al., 2017) and warm-up steps of 25k. We stop the training after 15 epochs without loss decrease on the dev set and average 7 checkpoints around the best (best, three preceding, and three succeeding). Trainings are performed on 4 NVIDIA A40 GPUs with 40GB RAM. We set 40k as the maximum number of tokens per mini-batch, 2 as update frequency, and 100,000 as maximum updates (~ 23 hours).

The MT models used for knowledge distillation are trained on OPUS (Tiedemann, 2016) en→{de, es} sections and are plain Transformer architec-

tures with 16 attention heads and 1024 embedding features in the encoder/decoder, resulting in $\sim 212\text{M}$ parameters. We achieve 32.1 and 35.8 BLEU on, respectively, MuST-C tst-COMMON German and Spanish.

B Data Statistics

MuST-C training data (train set) has been filtered: samples containing audio longer than 30s are discarded to reduce GPU computational requests. The total number of samples used during our trainings is shown in Table 2.

split	en→de	en→es
train	225,277*	260,049*
dev	1,423	1,316
tst-COMMON	1,422	1,315

Table 2: Number of samples for each split of MuST-C. * means this number doubled due to the use of KD.

C Main Results with Different Latency Metrics

Apart from AL, two metrics can be adopted to measure latency in simultaneous. The first one is the Differentiable Average Lagging – or DAL – (Cherry and Foster, 2019), a differentiable version of AL, and the Length-Adaptive Average Lagging – or LAAL – (Papi et al., 2022b), which is a modified version of AL that accounts also for the case in which the prediction is longer compared to the reference. Figure 7 and 8 show the results of the systems of Figure 5 by using, respectively, DAL and LAAL considering both computational aware (CA) and unaware metrics for German and Spanish. Numeric values are presented in Section D.

As we can see, the results of Figure 7 and 8 confirm the phenomena found in Section 5, indicating EDATT as the best system among languages and latency values. We observe also that DAL reports higher latency for all systems (it spans from 3 to 7.5s for German and to 5.5s for Spanish), with a counter-intuitive curve for the LA method considering its computational aware version. However, we acknowledge that DAL is less suited than AL/LAAL to evaluate current SimulST systems: in its computation, DAL gives a minimum delay for each emitted word while all the systems considered in our analysis can emit more than one word at once, consequently being improperly penalized in the evaluation.

D Numeric Values for Main Results

Table 3 on the next page.

en-de							
Policy	BLEU	AL	AL_CA	LAAL	LAAL_CA	DAL	DAL_CA
wait-k	19.6	1.43	2.36	1.53	2.43	1.86	3.14
	23.5	2.00	3.00	2.10	3.05	2.42	3.89
	25.1	2.51	3.53	2.60	3.57	2.89	4.46
	25.7	2.97	4.02	3.04	4.05	3.30	4.95
	26.1	3.37	4.43	3.43	4.45	3.66	5.33
LA	19.5	1.27	3.25	1.41	3.31	1.98	7.27
	23.1	1.69	3.32	1.79	3.37	2.37	5.85
	24.8	2.04	3.49	2.12	3.54	2.73	5.37
	25.9	2.33	3.73	2.39	3.77	3.01	5.36
	26.4	2.64	3.98	2.70	4.02	3.32	5.41
CAAT	20.3	0.88	1.98	1.02	2.09	1.49	3.28
	20.8	1.32	2.55	1.40	2.61	1.99	3.76
	20.5	1.74	3.14	1.78	3.18	2.46	4.29
	19.9	2.14	3.77	2.16	3.78	2.88	4.86
	19.0	2.54	4.24	2.54	4.25	3.26	5.23
EDATT	16.8	0.88	1.61	1.08	1.76	1.64	2.83
	19.1	1.04	1.75	1.20	1.87	1.73	2.91
	21.6	1.34	2.09	1.46	2.17	2.01	3.26
	24.0	1.74	2.56	1.83	2.63	2.43	3.71
	25.6	2.26	3.26	2.33	3.31	2.99	4.40
	26.3	2.74	3.93	2.80	3.96	3.46	4.97
en-es							
Policy	BLEU	AL	AL_CA	LAAL	LAAL_CA	DAL	DAL_CA
wait-k	24.9	1.39	2.41	1.58	2.53	1.96	3.51
	28.4	1.97	3.07	2.16	3.18	2.52	4.30
	29.0	2.50	3.63	2.68	3.72	3.03	4.91
	29.2	2.98	4.09	3.14	4.17	3.45	5.30
	29.4	3.41	4.57	3.55	4.63	3.82	5.73
LA	22.1	1.12	2.46	1.42	2.65	2.03	4.59
	26.4	1.52	2.56	1.76	2.72	2.42	4.01
	28.1	1.87	2.81	2.08	2.96	2.75	4.10
	28.9	2.17	3.03	2.36	3.17	3.05	4.20
	29.5	2.46	3.28	2.63	3.41	3.33	4.39
CAAT	25.1	0.74	2.02	1.02	2.23	1.54	3.57
	26.0	1.15	2.57	1.37	2.72	2.03	4.03
	26.6	1.53	3.14	1.71	3.26	2.51	4.54
	26.6	1.91	3.70	2.05	3.79	2.92	5.02
	26.7	2.27	4.25	2.38	4.33	3.31	5.51
EDATT	23.0	0.95	1.74	1.24	1.97	1.81	3.01
	25.0	1.10	1.90	1.36	2.10	1.92	3.12
	26.6	1.28	2.09	1.52	2.27	2.09	3.29
	27.8	1.52	2.42	1.74	2.59	2.38	3.62
	28.9	1.81	2.87	2.02	3.01	2.74	4.03
	29.2	2.14	3.37	2.34	3.50	3.12	4.48

Table 3: Numeric values for the plots presented in Sections 6 and C.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Last section of the paper (no number).
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract and Introduction (Section 1)
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

We will release the code, models, and outputs of the scientific artifacts of Section 3. The use of other scientific artifacts such as datasets is described in Section 4.

- B1. Did you cite the creators of artifacts you used?
Section 4.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
In Section 1, footnote 1.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
We use datasets as is and we build our models from scratch.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
The models we will release are trained on MuST-C for English->German, Spanish as mentioned in Section 4.1.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Appendix A.

C Did you run computational experiments?

They are described in Section 4 and the results are reported in Section 5.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix B.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 5.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

We provide quality-latency graphs (BLEU-AL charts) for one run, as it is usually done in Simultaneous Speech Translation, but we report both ideal and computational-aware latency measures and, for the latter, we also provide results for different hardware (GPU) in Section 5.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

We use FBK-fairseq with its default settings unless stated otherwise, as reported in Appendix B.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.