

Cascade *versus* Direct Speech Translation: Do the Differences Still Make a Difference?

Luisa Bentivogli¹, Mauro Cettolo¹, Marco Gaido^{1,2},
Alina Karakanta^{1,2}, Alberto Martinelli^{2*}, Matteo Negri¹, Marco Turchi¹

¹Fondazione Bruno Kessler

²University of Trento

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

Abstract

Five years after the first published proofs of concept, *direct* approaches to speech translation (ST) are now competing with traditional *cascade* solutions. In light of this steady progress, can we claim that the performance gap between the two is closed? Starting from this question, we present a systematic comparison between state-of-the-art systems representative of the two paradigms. Focusing on three language directions (English–German/Italian/Spanish), we conduct automatic and manual evaluations, exploiting high-quality professional post-edits and annotations. Our multi-faceted analysis on one of the few publicly available ST benchmarks attests for the first time that: *i*) the gap between the two paradigms is now closed, and *ii*) the subtle differences observed in their behavior are not sufficient for humans neither to distinguish them nor to prefer one over the other.

1 Introduction

Speech translation (ST) is the task of automatically translating a speech signal in a given language into a text in another language. Research on ST dates back to the late eighties and its evolution followed the development of the closely related fields of speech recognition (ASR) and machine translation (MT) that, since the very beginning, provided the main pillars for building the so-called *cascade* architectures. With the advent of deep learning, the neural networks widely used in ASR and MT have been adapted to develop a new *direct* ST paradigm. This approach aims to overcome known limitations of the cascade one (e.g. architectural complexity, error propagation) with a single encoder-decoder architecture that directly translates the source signal bypassing intermediate representations.

* The work of Alberto Martinelli was carried out during an internship at Fondazione Bruno Kessler.

Until now, the consolidated underlying technologies and the richness of available data have upheld the supremacy of cascade solutions in industrial applications. However, architectural simplicity, reduced information loss and error propagation are the ace up the sleeve of the direct approach, which has rapidly gained popularity within the research community in spite of the critical bottleneck represented by data paucity.

Within a few years after the first proofs of concept (Bérard et al., 2016; Weiss et al., 2017), the performance gap between the two paradigms has gradually decreased. This trend is mirrored by the findings of the International Workshop on Spoken Language Translation (IWSLT),¹ a yearly evaluation campaign where direct systems made their first appearance in 2018. On English-German, for instance, the BLEU difference between the best cascade and direct models dropped from 7.4 points in 2018 (Niehues et al., 2018) to 1.6 points in 2019 (Niehues et al., 2019b). In 2020, participants were allowed to choose between processing a pre-segmented version of the test set or the one produced by their own segmentation algorithm. As reported in (Ansari et al., 2020), the distance between the two paradigms further decreased to 1.0 BLEU point in the first condition and, for the first time, it was slightly in favor of the best direct model in the second condition, with a small but nonetheless meaningful 0.24 difference.

So, quoting Ansari et al. (2020), is the cascade solution still the dominant technology in ST? Has the direct approach closed the huge initial performance gap? Are there systematic differences in the outputs of the two technologies? Are they distinguishable? Answering these questions is more than running an evaluation exercise. It implies pushing research towards a deeper investigation of direct

¹<http://iwslt.org>

ST, finding a path towards its wider adoption in industrial settings and motivating higher engagement in data exploitation and resource creation to train the data-hungry end-to-end neural systems.

For all these reasons, while Ansari et al. (2020) were cautious in drawing firm conclusions, in this paper we delve deeper into the problem with the first thorough comparison between the two paradigms. Working on three language directions (en-de/es/it), we train state-of-the-art cascade and direct models (§3), running them on test data drawn from the MuST-C corpus (Cattoni et al., 2020).

Systems’ behavior is analysed from different perspectives, by exploiting high-quality post-edits and annotations by professionals. After discussing overall systems’ performance (§4), we move to more fine-grained automatic and manual analyses covering two main aspects: the relation between systems’ performance and specific characteristics of the input audio (§5), and the possible differences in terms of lexical, morphological and word ordering errors (§6). We finally explore whether, due to latent characteristics overlooked by all previous investigations, the output of cascade and direct systems can be distinguished either by a human or by an automatic classifier (§7). Together with a comparative study attesting the parity of the two paradigms on our test data, another contribution of this paper is the release of the manual post-edits that rendered our investigation possible. The data is available at: <https://ict.fbk.eu/mustc-post-edits>.

2 Background

Cascade ST. By concatenating ASR and MT components (Stentiford and Steer, 1988; Waibel et al., 1991), cascade ST architectures represent an intuitive solution to achieve reasonable performance and high adaptability across languages and domains. At the same time, however, they suffer from well-known problems related to the concatenation of multiple systems. First, they require *ad-hoc* training and maintenance procedures for the ASR and MT modules; second, they suffer from error propagation and from the loss of speech information (e.g. prosody) that might be useful to improve final translations. Research has focused on mitigating error propagation by: *i*) feeding the MT system with ASR data structures (e.g. ASR n-best, lattices or confusion networks) which are more informative than the 1-best output (Lavie et al., 1996; Matusov et al., 2005; Bertoldi and Federico, 2005;

Beck et al., 2019; Sperber et al., 2019), and *ii*) making the MT robust to ASR errors, for instance by training it on parallel data incorporating real or emulated ASR errors as in (Peitz et al., 2012; Ruiz et al., 2015; Sperber et al., 2017; Cheng et al., 2019; Di Gangi et al., 2019a). Although the former solutions are effective to some extent, state-of-the-art cascade architectures (Pham et al., 2019; Bahar et al., 2020) prefer the latter, as they are simpler to implement and maintain.

Direct ST. To overcome the limitations of cascade models, Bérard et al. (2016) and Weiss et al. (2017) proposed the first direct solutions bypassing intermediate representations by means of encoder-decoder architectures based on recurrent neural networks. Currently, more effective solutions (Potapczyk and Przybysz, 2020; Bahar et al., 2020; Gaido et al., 2020) rely on ST-oriented adaptations of Transformer (Vaswani et al., 2017) integrating the encoder with: *i*) convolutional layers to reduce input length, and *ii*) penalties biasing attention to local context in the encoder self-attention layers (Povey et al., 2018; Sperber et al., 2018; Di Gangi et al., 2019b). Though effective, these architectures have to confront with training data paucity, a critical bottleneck for neural solutions. The problem has been mainly tackled with data augmentation and knowledge transfer techniques. Data augmentation consists in producing artificial training corpora by altering existing datasets or by generating (*audio, translation*) pairs through speech synthesis or MT (Bahar et al., 2019b; Nguyen et al., 2020; Ko et al., 2015; Jia et al., 2019). Knowledge transfer (Gutstein et al., 2008) consists in passing (here to ST) the knowledge learnt by a neural network trained on closely related tasks (here, ASR and MT). Existing ASR models have been used for encoder pre-training (Bérard et al., 2018; Bansal et al., 2019; Bahar et al., 2019a) and multi-task learning (Weiss et al., 2017; Anastasopoulos and Chiang, 2018; Indurthi et al., 2020). Existing neural MT models have been used for decoder pre-training (Bahar et al., 2019a; Inaguma et al., 2020), joint learning (Indurthi et al., 2020; Liu et al., 2020) and knowledge distillation (Liu et al., 2019).

Previous comparisons. Most of the works on direct ST also evaluate the proposed solutions against a cascade counterpart. The conclusions, however, are discordant. Looking at recent works, Pino et al. (2019) show similar scores, Indurthi et al. (2020) report higher results for their direct model, while

Inaguma et al. (2020) end up with the opposite finding. The main problems of these comparisons are that: *i*) not all the architectures are equally optimized, *ii*) for the sake of fairness in terms of training data, cascade systems are restricted to unrealistic settings with small training corpora that penalize their performance, and *iii*) evaluation always relies only on automatic metrics computed on single references. The IWSLT campaigns (Niehues et al., 2019a; Ansari et al., 2020) set up a shared evaluation framework where systems built on a large set of training data are optimized to achieve the best performance, independently from the underlying architecture. In the last round, direct models approached, and in one case (Potapczyk and Przybysz, 2020) outperformed, the cascade ones. However, the evaluation was run only on one language pair, by solely relying on automatic metrics and single references. In this paper, we overcome these limitations by comparing the two paradigms on three language pairs, using different metrics, multiple references (including professional post-edits) as well as fine-grained automatic and manual analysis procedures.

3 Experimental Setting

3.1 ST Systems

To maximize the cross-language comparability of our analyses, we built the cascade and direct ST systems for en–de/es/it with the same core technology, based on Transformer. Their good quality is attested by the comparison with the winning system at the IWSLT-20 offline ST task (Bahar et al., 2020),² which consists of an ensemble of two cascade models scoring 28.8 BLEU on the en-de portion of the MuST-C *Common* test set. On the same data, our cascade and direct models achieve similar BLEU scores, respectively 28.9 and 29.1 (see Table 1).³ On en-es and en-it, identical architectures perform similarly or better (up to 32.9 BLEU on en-es). Although BLEU scores are not strictly comparable across languages, we can safely consider all our models as state-of-the-art.

For the sake of reproducibility, we provide complete details about data, architectures and training setup in Appendix A.

²In the pre-segmented data condition (Ansari et al., 2020).

³Also the ASR performance of our cascade solution (10.2 WER on MuST-C *Common*) is in line with the results obtained by Bahar et al. (2020) for their best ASR model.

3.2 Evaluation Methodology

Data. Our evaluation data is drawn from the TED-based MuST-C corpus (Cattoni et al., 2020), the largest freely available multilingual corpus for ST. It covers 14 language directions, with English audio segments automatically aligned with their corresponding manual transcriptions and translations. The en–de/es/it MuST-C *Common* test sets contain the same 27 TED talks, for a total of around 2,500 segments largely overlapping across languages.⁴ For all the three language pairs, we selected subsets of MuST-C *Common* containing the same English audio portions from each talk, in order to obtain representative groups of contiguous segments that are comparable across languages. Furthermore, to ensure high data quality, we manually checked the selected samples and kept only those segments for which the *audio-transcript-translation* alignment was correct. Each of the three resulting test sets – henceforth *PE-sets* – is composed of 550 segments, corresponding to about 10,000 English source words.

Post-editing. A key element of our multi-faceted analysis is human post-editing (PE), which consists in manually correcting systems’ output according to the input (the source audio in our case). In PE-based evaluation, the original output is compared against its post-edited version using distance-based metrics like TER (Snover et al., 2006). This allows for counting only the true errors made by a system, without penalising differences due to linguistic variation as it happens when exploiting independent references. This makes PE-based evaluation one of the most prominent methodologies used for translation quality assessment (Snover et al., 2006, 2009; Denkowski and Lavie, 2010; Cettolo et al., 2013; Bojar et al., 2015; Graham et al., 2016; Bentivogli et al., 2018b).

To collect the post-edits for our study, we strictly followed the methodology of the IWSLT 2013-2017 evaluation campaigns (Cettolo et al., 2013), which offered us a consolidated framework and best practices to draw upon. Our cascade and direct systems were both run on the PE-sets to be post-edited. To guarantee high quality post-edits, for each language we hired two professional translators with experience in subtitling and post-editing. Moreover, in order to cope with translators’ vari-

⁴MuST-C *Common* segments can vary across languages due to the automatic procedures of segmentation, audio-text alignment and filtering that were applied to the talks.

ability (i.e. more/less aggressive editing strategies), the outputs of the two ST systems were randomly assigned ensuring that each translator worked on all the 550 segments, post-editing an equal number of outputs from both systems. The task was performed with a CAT tool⁵ that displays the manual transcript of the audio together with the ST output to be edited. However, since ST systems take as input an audio signal, we also provided translators with the audio file of each segment, asking them to post-edit strictly according to it.⁶ For each language pair, the final PE-set used in our study consists of the 550 MuST-C original *audio-transcript-translation* triplets plus two additional sets of reference translations, i.e. the post-edited versions of the two systems’ outputs.

Analyses. The collected post-edits are exploited to assess overall systems’ performance (§4) as well as to carry out deeper quantitative and qualitative analyses aimed to shed light on possible systematic differences in systems’ behavior (§5.1 and §6.1). Focusing on specific aspects of the ST problem, the inquiry is also performed by means of manual annotation of systems’ outputs (§5.2, §6.2 and §7.1). Due to the linguistic nature of this task, centred on fine-grained aspects requiring a variety of skills in both evaluation and ST technology, for such analyses we relied on three researchers in translation technology – one per language pair – with a strong background in linguistics, excellent knowledge of the addressed languages (C2 or native), as well as strong expertise in systems’ evaluation.

4 Overall Systems’ Performance

We compute overall performance results both on the PE-sets and on the MuST-C *Common* test sets. Our primary evaluation is based on the collected post-edits. We consider two TER-based⁷ metrics: *i*) human-targeted TER (HTER) computed between the automatic translation and its human post-edited version, and *ii*) multi-reference TER (mTER) computed against the closest reference among the three available ones (two post-edits and the official reference from MuST-C). The latter metric better accounts for post-editors’ variability, making the evaluation more reliable and informative. For the sake of completeness, in Table 1 we also report Sacre-

⁵www.matecat.com

⁶The *ad-hoc* ST PE guidelines given to translators are included in Appendix B.

⁷www.cs.umd.edu/~snover/tercom

BLEU⁸ (Post, 2018) and TER scores computed only on the official MuST-C *Common* references.

		PE Set				M. Common	
		HTER	mTER	BLEU	TER	BLEU	TER
de	C	28.65	24.41	28.96	53.23	28.86	53.93
	D	30.22	25.60	28.46	52.56	29.05	52.77*
es	C	29.96	25.30	34.05*	50.75	32.93*	53.21*
	D	28.19*	24.02*	32.17	51.08	31.98	54.00
it	C	25.69	23.29	30.04*	54.01	28.56	56.29
	D	26.14	23.26	28.81	54.06	28.56	55.35*

Table 1: Performance of (C)ascade and (D)irect systems on the PE-sets and MuST-C *Common* test sets. Statistically significant differences (*) are computed with Paired Bootstrap Resampling (Koehn, 2004).

A bird’s-eye view of the results shows that, in more than half of the cases, performance differences between cascade and direct systems are not statistically significant. When they are, the raw count of wins for the two approaches is the same (4), attesting their substantial parity.

Looking at our primary metrics (HTER and mTER), systems are on par on en-it and en-de, while for en-es the direct approach significantly outperforms the cascade one. This difference, however, does not emerge with the other metrics. Indeed, BLEU and TER scores computed against the official references are less coherent across metrics and test sets. For instance, on the en-it PE-set the cascade system significantly outperforms the direct one in terms of BLEU score, while TER shows the opposite on MuST-C *Common*. Interestingly, the scores obtained using independent references can also disagree with those computed with post-edits. This is the case of en-es, where significant HTER and mTER reductions attest the superiority of the direct system, while most BLEU and TER scores are still in favor of the cascade.

On the one hand, primary evaluation scores suggest that the rapidly advancing direct technology has eventually reached the traditional cascaded approach. On the other, the highlighted incongruities confirm widespread concerns about the reliability of fully automatic metrics – based on independent references – to properly evaluate neural systems (Way, 2018). This calls for deeper quantitative and qualitative analyses. Those presented in the next sections investigate performance differences focusing on two main aspects: the impact of specific input audio properties (§5), and the linguistic errors made by the systems (§6).

⁸BLEU+c.mixed+#.1+s.exp+tok.13a+v.1.4.3

5 ST Quality and Audio Properties

5.1 Automatic Analysis

The two ST approaches handle the input audio differently: the cascade one by means of a dedicated ASR component that produces intermediate transcripts; the direct one by extracting all the relevant information to translate in an end-to-end fashion. Is it therefore possible that some audio properties have different impact on their results? Overall performance being equal, answering this question would help to understand if one approach is preferable over the other under specific audio conditions.

Among other possible factors (e.g. noise, recording conditions, overlapping speakers) we tried to shed light on this aspect by focusing on two common factors: audio duration and speech rate. To this aim, we grouped the sentences in the PE-set according to the sentence-wise HTER percentage difference – i.e. the difference between the cascade and direct HTER scores divided by their average.

The threshold for considering performance differences as significant was set to 10%. The resulting groups contain sentences where: *i*) cascade is significantly better than direct, *ii*) direct is significantly better than cascade, *iii*) the difference between the two is not significant, and *iv*) both systems have HTER=0. For each group, we calculated the average audio duration and the corresponding speech rate in terms of phonemes⁹ per second.

Results are shown in Table 2, where – for the sake of completeness – also the length of the reference audio transcript is given, together with the average HTER of the systems.

As we can see, results are coherent across languages: audio duration and speech rate averages do not differ, neither when one system performs significantly better than the other, nor when the HTER differences are not significant. We can hence conclude that, if audio duration and speech rate have any influence on systems’ performance, our analysis does not highlight specific conditions that are more favorable to one approach than to the other. Both are equally robust with respect to the audio properties here considered.

5.2 Manual Analysis

Handling the input audio differently, the two approaches have inherent strengths and weaknesses.

⁹Obtained by processing the transcripts with eSpeak (espeak.sourceforge.net).

		#sentences	audio duration (seconds)	speech rate (phonemes/s)	ref transcr length (#words)	C HTER	D HTER
de	C better	240	6.15	14.43	19.75	16.30	40.53
	D better	191	6.00	14.52	18.88	44.85	17.89
	No Diff	45	6.68	14.31	22.07	40.74	40.18
	HTER 0	74	2.71	15.53	9.64	0	0
es	C better	215	5.92	12.20	19.52	16.09	38.76
	D better	234	6.28	12.09	20.39	46.45	21.14
	No Diff	54	6.47	12.01	20.26	40.22	40.37
	HTER 0	47	3.09	13.14	10.23	0	0
it	C better	231	6.03	12.31	19.41	14.82	36.40
	D better	212	6.06	12.21	19.33	37.65	15.80
	No Diff	55	6.93	11.94	21.73	35.39	35.37
	HTER 0	52	2.96	12.68	10.33	0	0

Table 2: Comparison of (C)ascade and (D)irect performance based on different audio properties.

In particular, although suffering from the well-known scarcity of sizeable training corpora, direct solutions come with the promise (Sperber and Paulik, 2020) of: *i*) higher robustness to error propagation, and *ii*) reduced loss of speech information (e.g. prosody). Our next qualitative analysis tries to delve into these aspects by looking at audio understanding and prosody issues.

Audio understanding. Errors due to wrong audio understanding are easy to identify for cascade systems – since they are evident in the intermediate ASR transcripts – but harder to spot for direct systems, whose internal representations are by far less accessible. In this case, errors can still be identified in mistranslations corresponding to words which are phonetically similar to parts of the input audio – e.g. *nice voice* mistranslated in German as *nette Jungen (nice boys)*. To spot such errors, our annotators carefully inspected the PE-set by comparing the audio, the reference transcripts and systems’ output translations for both the cascade and direct models, as well as the ASR transcripts for the cascade one. Some interesting examples of the identified errors are reported in Table 3.

AUDIO	to the er- euh [disfluency] Egyptian government
C	der eruptiven [Eng. “eruptive”] Regierung ...
D	an die Regierung Ägyptens
AUDIO	dominated by big, scary guys,...
C	dominados por <i>grandes tipos aterradores</i>
D	dominados por los chicos de Big Kerry
AUDIO	I think, like her ,...
C	Penso che, come qui [Eng. “here”], ...
D	Penso che, come i capelli [Eng. “hair”], ...

Table 3: Examples of audio understanding errors.

As shown in Table 4, audio understanding errors are quite common for both systems in all language pairs. However, both the number of errors and the number of sentences they affect is significantly lower for the direct one. We observed that this is the case especially for “more difficult” sentences, such as sentences with poor audio quality and overlapping or disfluent speech.

Though far from being conclusive (we acknowledge that, due to the “opacity” of direct models, their error counts might be slightly underestimated), this analysis seems to confirm the theoretical advantages of direct ST. This finding advocates for more thorough future investigations on neural networks’ interpretability, targeting its empirical verification on larger and diverse benchmarks.

	Both	C	D	C_{tot}	D_{tot}	C_{sent}	D_{sent}
de	51	96	52	147	103	117	91
es	82	108	66	190	148	150	127
it	87	82	69	169	156	143	138

Table 4: Audio understanding errors in the PE-set and number of sentences containing at least one such error.

Prosody. Prosody is central to disambiguating utterances, as it reflects language elements which may not be encoded by grammar and vocabulary choices. While prosody is directly encoded by the direct system, it is lost in the unpunctuated input received by the MT component of a cascade. Besides few interrogative sentences, our annotators were able to isolate only a handful of utterances whose prosodic markers result in different interpretations by the two models. Concerning interrogatives, both systems managed to translate them correctly in most cases (24 for cascade and 25 for direct out of 31). This is not surprising given the syntactic structure of English questions, which is explicit and does not rely solely on prosody (e.g. compared to Italian). In all other cases (examples in Table 5), the direct model’s higher sensitivity to prosody seems to give it an edge on cascade in disambiguating and correctly rendering the utterance meaning. Also this finding calls for future inquiries aimed to check the regularity of these differences on larger datasets.

6 Linguistic Errors

6.1 Automatic Analysis

For this analysis, we rely on the publicly available tool¹⁰ used by Bentivogli et al. (2018a) to analyse

¹⁰wit3.fbk.eu/2016-02, details in Appendix C.

src nation states — governments doing the attacks
C Regierungen der Nationalstaaten [governments of nation states]
D Nationen, Regierungen [nations, governments]
src like the one we saw before, moving
C como el que vimos antes de moverse [like the one we saw before moving]
D como el que hemos visto antes, moviéndose [like the one we saw before, moving]
src Photos like this: construction going on
C Foto come questa costruzione [Photos like this construction]
D Foto come queste: costruzione [Photos like these: construction]

Table 5: The two approaches dealing with prosody.

	en-de			en-es			en-it		
	C	D	$\Delta\%$	C	D	$\Delta\%$	C	D	$\Delta\%$
L	2481	2560	+3.2	2674	2497	-6.6	2264	2264	0.0
M	468	536	+14.5	535	494	-7.7	433	470	+8.6
R	398	476	+19.6	308	290	-5.8	230	226	-1.7
	3347	3572	+6.7	3517	3281	-6.7	2927	2960	+1.1

Table 6: Distribution of (L)exical, (M)orphological and (R)ordering errors. Absolute numbers are presented together with the percentage of reduction/increase of the (D)irect system with respect to the (C)ascade ($\Delta\%$).

what linguistic phenomena are best modeled by MT systems. The tool exploits manual post-edits and HTER-based computations to detect and classify translation errors according to three linguistic categories: lexicon, morphology and word order. Table 6 presents their distribution.

As expected from the HTER scores in Table 1, results vary across language pairs. On en-it, systems show pretty much the same number of errors, with a slight percentage gain (+1.1) in favor of the cascade. For the other two pairs, differences are more marked and opposite, with an overall error reduction for the direct system on en-es (-6.7) and in favor of the cascade on en-de (+6.7).

Looking at the distribution of errors across categories, while for en-es the direct system is always better and the percentage reduction is homogeneously distributed, for en-de the better performance of the cascade is concentrated in the morphology and word order categories. Since English and German are the most different languages in terms of morphology and word order, this result suggests that cascade systems still have an edge on the direct ones in their ability to handle morphology and word reordering. This is further supported by en-it: the only difference, in favor of the cascade, is indeed observed in the morphology category.

6.2 Manual Analysis

Since lexical errors represent by far the most frequent category for both approaches in all language pairs, we complement the automatic analysis with a more fine-grained manual inspection, further distinguishing among lexical errors due to missing words, extra words, or wrong lexical choice.¹¹

The analysis was carried out on subsets of the PE-set, created in such a way to be suitable for manual annotation. Namely, we removed sentences for which the output of the two systems is: *i*) identical, *ii*) judged correct by post-editors (HTER=0), or *iii*) too poor to be reliably annotated for errors (HTER>40%). The resulting sets contain 207 sentences for en-de, 238 for en-es, and 285 for en-it.

This analysis reveals that, for all language pairs, wrong lexical choice is the most frequent error type (~65% of lexical errors on average) followed by missing words (~30%), and extra words (~5%).

While errors due to lexical choice and superfluous words vary across languages, we observe a systematic behavior with respect to missing words (words that are present in the audio but are not translated). As we can see in Table 7, direct systems lose more information from the source input than their cascade counterparts, in terms of both single words and contiguous word sequences. It is particularly interesting to notice that also for en-es – where the direct system is significantly stronger than the cascade – the issue is still evident, although to a lesser extent. Table 8 collects examples of the encountered lexical phenomena.

	single words		word sequences		total # words		
	C	D	C	D	C	D	Δ%
de	25	34	6	10	42	58	+38.10
es	26	40	10	11	59	68	+15.25
it	53	83	14	18	96	128	+33.33

Table 7: Missing words for (C)ascade and (D)irect systems. Absolute numbers vary across languages as they reflect the different size of the annotated subsets.

Finally, we report that a non-negligible amount of missing words (between 10% and 20%) is represented by discourse markers, i.e. words or phrases used to connect and manage what is being said (e.g. “you know”, “well”, “now”). Although this is

¹¹Various error taxonomies covering different levels of granularity have been developed, and the distinction between these types of lexical errors is widely adopted, including the DQF-MQM framework – <https://info.taus.net/dqf-mqm-error-typology-templ>

AUDIO	“That’s fine”, says George,
C	“Das ist in Ordnung.” [–] George,
D	“Das ist in Ordnung, [–] George,”
AUDIO	Well after two years, ...
C	<i>Bueno</i> , después de dos años, ...
D	[–] Después de dos años, ...
AUDIO	My wife and kids and I, moved to ...
C	Io e mia moglie <i>e i miei figli</i> ci siamo trasferiti...
D	Io e mia moglie [–] ci siamo trasferiti...

Table 8: Examples of missing words.

a frequent phenomenon in speech, not translating discourse markers cannot be properly considered as an error, since markers *i*) do not carry semantic information, and *ii*) can be intentionally dropped in some use cases, such as in subtitling.

7 Classifiers’ Verdict

So far, our inquiry has been entirely driven by pre-defined assumptions (the importance of certain audio properties) and linguistic criteria (the focus on specific error types). This *top-down* approach, however, might fail to disclose important differences, which were not specifically sought after when analysing the two paradigms. This consideration motivates the adoption of the complementary *bottom-up* approach that concludes our comparative study by answering the question: is the output of cascade and direct systems distinguishable? Understanding if and why discriminating between the two is possible would not only suggest new issues to look at. It would also highlight possible output regularities that, despite the similar overall performance, make one paradigm preferable over the other in specific application scenarios. To this aim, we set up a classification experiment, comparing the ability of humans to correctly identify the output of the two systems with the performance of an automatic text classifier.

7.1 Human Classification

After getting acquainted with systems’ output through the previous manual analyses, our assessors were instructed to perform a classification task. The classification had to be performed on 10 blocks of items comprising a set of unseen English contiguous sentences (gold transcripts) from the MuST-C *Common* test set, and two sets of anonymized translations, one produced by the cascade and one by the direct model. For each block, the assessors had to assign each set of translations to the correct system, or label them as indistinguishable. To investigate whether more context helps in the assign-

ment, we set up two experiments with respectively 10 and 20 contiguous sentences per block.

# of sentences	en-de		en-es		en-it	
	10	20	10	20	10	20
Correct	7	6	4	4	4	3
Wrong	2	2	2	3	1	2
Indistinguishable	1	2	4	3	5	5
Total # of blocks	10	10	10	10	10	10

Table 9: Results of human classification.

The results in Table 9 show that en-es and en-it systems are not distinguishable, since only a maximum of 4 blocks out of 10 were correctly classified, while most en-de blocks were correctly classified. According to the en-de assessor, this is due to the fact that the structure of the sentences generated by the direct system is very similar to that of the corresponding English sources. This characteristic stands out in German, which differs from English in terms of word order more than Italian and Spanish. This type of behavior does not necessarily imply the presence of errors but, like a fingerprint, makes the en-de direct system more recognizable by a human. Furthermore, being sub-optimal for German, this structure can cause preferential edits by the post-editors, which would be in line with the concentration of errors in the word order category observed in Table 6 (+19.6%).

Assessing the importance of context, the ability of humans to distinguish the systems does not improve when passing from 10 to 20 sentences per block. This suggests that the behavioral differences between cascade and direct systems are so subtle that, on larger samples, they mix up and balance making their fingerprints less traceable.

7.2 Automatic Classification

As a complement to the human classification experiment, we check whether an automatic tool is able to accomplish a similar task. Our classifier combines n -gram language models with the Naive Bayes algorithm, as proposed in (Peng and Schuurmans, 2003). We trained two 5-gram models, respectively using translations by the cascade and the direct systems. At classification time, given a translated text, the classifier computes the perplexity of the two models and assigns the *cascade* or *direct* label based on the model with the lowest perplexity. Also these experiments were carried out on the MuST-C *Common* set. The classifier was tested via k -fold cross-validation, for different values of k – i.e. different sizes of text to classify.

As shown in Figure 1, contrary to humans, the more data the classifier receives, the higher its accuracy in discriminating between systems. Already at a size of 20 sentences, accuracy is always $\sim 80\%$. This suggests that systems have their own “language”, a fluency-related fingerprint.

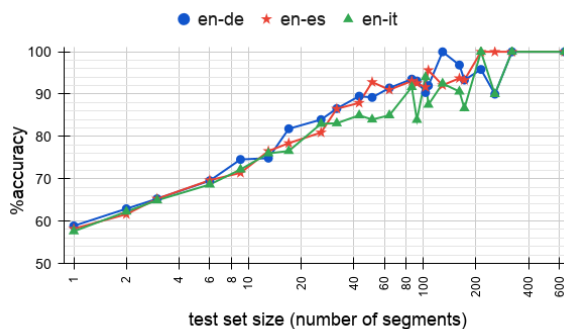


Figure 1: Results of automatic classification for different sizes of system output blocks (1-600 sentences).

To check this finding, we measured outputs’ *lexical diversity* in terms of moving average Type-Token Ratio – maTTR (Covington and McFall, 2010) – and with the Measure of Textual Lexical Diversity (MTLD) by McCarthy and Jarvis (2010).

Table 10 shows that the cascade output exhibits higher lexical diversity on all languages, with smaller differences on en-de and en-es compared to en-it. A plausible conclusion is that the cascade produces richer output, whose variety does not necessarily result in better translations nor is appreciated by humans. Indeed, annotators were able to correctly distinguish the output only for en-de, where lexical diversity is similar (see §7.1).

	en-de		en-es		en-it	
	maTTR	MTLD	maTTR	MTLD	maTTR	MTLD
R	73.11	97.02	69.81	77.19	74.50	109.79
C	71.84	83.64	68.42	67.68	73.20	97.82
D	71.45	83.27	67.99	65.59	72.60	90.78

Table 10: Lexical diversity of the human (R)eference, (C)ascade and (D)irect outputs.

8 Conclusion and Final Remarks

There is a time when the possible transition from consolidated technological frameworks to new emerging paradigms depends on answering fundamental questions about their potential, strengths and weaknesses. A time when technology developers are faced with the choice of where to direct their future investments. Five years after its appearance

on the scene, the direct approach to ST confronts the community with similar questions in relation to the traditional cascade paradigm that it aims to overtake. Our investigation showed that, in spite of the known data paucity conditions still penalizing the direct approach, the two technologies now perform substantially on par. Subtle differences in their behavior exist: overall performance being equal, the cascade still seems to have an edge in terms of morphology, word ordering and lexical diversity, which is balanced by the advantages of direct models in audio understanding and in capturing prosody. However, they do not seem sufficient and consistent enough across languages to make the output of the two approaches easily distinguishable, nor to make one model preferable to the other. Back to our title, they no longer make a difference.

We are aware that the generalizability of these results depends on several factors such as the considered languages, systems and benchmarks, as well as the human workforce deployed for the inquiry. Here, with the help of professionals, we proposed multi-faceted quantitative and qualitative analyses, run on the output of state-of-the-art systems on three language pairs – though, by now, covering only the most-explored and data-favorable condition, which has English as source. Although our findings hold for a specific scenario, in which free data were at our disposal (and to which we contribute back by releasing high-quality post-edits), they might not be generalizable to other (e.g. difficult, distant) languages and other (e.g. highly specialized) domains. Nevertheless, we present them as a timely contribution towards answering a burning question within the ST community.

Acknowledgements

The creation of the post-edits used in this work was funded by the European Association for Machine Translation (EAMT) through its 2020 Sponsorship of Activities programme. The computational costs were covered by the “End-to-end Spoken Language Translation in Rich Data Conditions” project,¹² which was financially supported by an Amazon AWS ML Grant.

¹²<https://ict.fbk.eu/units-hlt-mt-e2eslt/>

References

- Antonios Anastasopoulos and David Chiang. 2018. Tied Multitask Learning for Neural Speech Translation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 82–91, New Orleans, US-LA.
- 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 International Conference on Spoken Language Translation (IWSLT)*, Virtual Event.
- Amittai Axelrod, Xiaodong He, and Jianfeng Gao. 2011. Domain Adaptation via Pseudo In-Domain Data Selection. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, Edinburgh, Scotland, UK.
- Parnia Bahar, Tobias Bieschke, and Hermann Ney. 2019a. A Comparative Study on End-to-end Speech to Text Translation. In *Proceedings of the International Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 792–799, Sentosa, Singapore.
- Parnia Bahar, Patrick Wilken, Tamer Alkhoul, Andreas Guta, Pavel Golik, Evgeny Matusov, and Christian Herold. 2020. Start-Before-End and End-to-End: Neural Speech Translation by AppTek and RWTH Aachen University. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Virtual Event.
- Parnia Bahar, Albert Zeyer, Ralf Schlüter, and Hermann Ney. 2019b. On Using SpecAugment for End-to-End Speech Translation. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Sameer Bansal, Herman Kamper, Karen Livescu, Adam Lopez, and Sharon Goldwater. 2019. Pre-training on High-resource Speech Recognition Improves Low-resource Speech-to-text Translation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, Minneapolis, US-MN.
- Daniel Beck, Trevor Cohn, and Gholamreza Haffari. 2019. Neural speech translation using lattice transformations and graph networks. In *Proceedings of the EMNLP Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-13)*, pages 26–31, Hong Kong.
- Luisa Bentivogli, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2018a. Neural versus phrase-based MT quality: an in-depth analysis on

- English–German and English–French. *Computer Speech and Language*, 49:52 – 70.
- Luisa Bentivogli, Mauro Cettolo, Marcello Federico, and Christian Federmann. 2018b. Machine Translation Human Evaluation: an investigation of evaluation based on Post-Editing and its relation with Direct Assessment. In *Proceedings of the International Conference on Spoken Language Translation (IWSLT)*, Bruges, Belgium.
- Alexandre Bérard, Laurent Besacier, Ali Can Kocabiyikoglu, and Olivier Pietquin. 2018. End-to-End Automatic Speech Translation of Audiobooks. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6224–6228, Calgary, Canada.
- Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation. In *Proceedings of the NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain.
- Nicola Bertoldi and Marcello Federico. 2005. A new decoder for spoken language translation based on confusion networks. In *Proceedings of the IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 86–91, San Juan, Puerto Rico.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. Findings of the 2015 Workshop on Statistical Machine Translation. In *Proceedings of the Workshop on Statistical Machine Translation (WMT)*, Lisbon, Portugal.
- Roldano Cattoni, Mattia A. Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2020. MuST-C: A Multilingual Corpus for end-to-end Speech Translation. *Computer Speech & Language Journal*. Doi: <https://doi.org/10.1016/j.csl.2020.101155>.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. 2013. Report on the 10th IWSLT Evaluation Campaign. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Heidelberg, Germany.
- Qiao Cheng, Meiyuan Fang, Yaqian Han, Jin Huang, and Yitao Duan. 2019. Breaking the data barrier: Towards robust speech translation via adversarial stability training. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Michael A. Covington and Joe D. McFall. 2010. Cutting the Gordian Knot: The Moving-Average Type–Token Ratio (MATTR). *Journal of Quantitative Linguistics*, 17(2):94–100.
- Michael Denkowski and Alon Lavie. 2010. Choosing the right evaluation for machine translation: An examination of annotator and automatic metric performance on human judgment tasks. In *Proceedings of the Conference of the Association of Machine Translation in the Americas (AMTA)*, Denver, US-CO.
- Mattia A. Di Gangi, Robert Enyedi, Alessandra Brusadin, and Marcello Federico. 2019a. Robust Neural Machine Translation for Clean and Noisy Speech Transcripts. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2019b. Adapting Transformer to End-to-end Spoken Language Translation. In *Proceedings of the Conference of the International Speech Communication Association (INTERSPEECH)*, Graz, Austria.
- Mattia Antonino Di Gangi, Matteo Negri, and Marco Turchi. 2019c. One-to-many multilingual end-to-end speech translation. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU), 14-18 December 2019, Sentosa, Singapore*.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2020. End-to-End Speech-Translation with Knowledge Distillation: FBK@IWSLT2020. In *Proceedings of the International Conference on Spoken Language Translation (IWSLT)*, Virtual Event.
- Yvette Graham, Timothy Baldwin, Meghan Dowling, Maria Eskevich, Teresa Lynn, and Lamia Tounsi. 2016. Is all that Glitters in Machine Translation Quality Estimation really Gold? In *Proceedings of the International Conference on Computational Linguistics (COLING)*, pages 3124–3134, Osaka, Japan.
- Steven Gutstein, Olac Fuentes, and Eric Freudenthal. 2008. Knowledge Transfer in Deep Convolutional Neural Nets. *International Journal on Artificial Intelligence Tools*, 17(03):555–567.
- François Hernandez, Vincent Nguyen, Sahar Ghanay, Natalia Tomashenko, and Yannick Estève. 2018. TED-LIUM 3: Twice as Much Data and Corpus Repartition for Experiments on Speaker Adaptation. In *Proceedings of the Speech and Computer - 20th International Conference (SPECOM)*, pages 198–208, Leipzig, Germany. Springer International Publishing.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. In *Proceedings of NIPS Deep Learning and Representation Learning Workshop*, Montréal, Canada.
- Hirofumi Inaguma, Shun Kiyono, Kevin Duh, Shigeki Karita, Nelson Yalta, Tomoki Hayashi, and Shinji Watanabe. 2020. ESPnet-ST: All-in-One Speech Translation Toolkit. In *Proceedings of the Annual Meeting of the Association for Computational*

- Linguistics (ACL): System Demonstrations*, Virtual Event.
- Sathish R. Indurthi, Houjeung Han, Nikhil K. Lakumarapu, Beomseok Lee, Insoo Chung, Sangha Kim, and Chanwoo Kim. 2020. End-end Speech-to-Text Translation with Modality Agnostic Meta-Learning. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7904–7908, Barcelona, Spain.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Giménez. Adrià, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. **Europarl-ST: A Multilingual Corpus For Speech Translation Of Parliamentary Debates**. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233, Barcelona, Spain.
- Ye Jia, Melvin Johnson, Wolfgang Macherey, Ron J. Weiss, Yuan Cao, Chung-Cheng Chiu, Naveen Ari, Stella Laurenzo, and Yonghui Wu. 2019. Leveraging Weakly Supervised Data to Improve End-to-End Speech-to-Text Translation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7180–7184, Brighton, UK.
- Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. 2015. Audio Augmentation for Speech Recognition. In *Proceedings of the Conference of the International Speech Communication Association (INTERSPEECH)*, pages 3586–3589, Dresden, Germany.
- Philipp Koehn. 2004. Statistical Significance Tests for Machine Translation Evaluation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Barcelona, Spain.
- Alon Lavie, Donna Gates, Marsal Gavaldà, Laura Mayfield Tomokiyo, Alex Waibel, and Lori Levin. 1996. Multi-lingual translation of spontaneously spoken language in a limited domain. In *Proceedings of the International Conference on Computational Linguistics (COLING)*, Copenhagen, Denmark.
- Yuchen Liu, Hao Xiong, Jiajun Zhang, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2019. **End-to-End Speech Translation with Knowledge Distillation**. In *Proceedings of the Conference of the International Speech Communication Association (INTERSPEECH)*, pages 1128–1132, Graz, Austria.
- Yuchen Liu, Junnan Zhu, Jiajun Zhang, and Chengqing Zong. 2020. **Bridging the Modality Gap for Speech-to-Text Translation**. arxiv.org/pdf/2010.14920.pdf.
- Evgeny Matusov, Hermann Ney, and Ralph Schluter. 2005. Phrase-based translation of speech recognizer word lattices using loglinear model combination. In *Proceedings of the IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 110–115, San Juan, Puerto Rico.
- Philip M. McCarthy and Scott Jarvis. 2010. **MTLD, vocd-D, and HD-D: A validation study of sophisticated approaches to lexical diversity assessment**. *Behavior Research Methods*, 42(2):381–392.
- Graham Neubig, Matthias Sperber, Xinyi Wang, Matthieu Felix, Austin Matthews, Sarguna Padmanabhan, Ye Qi, Devendra Singh Sachan, Philip Arthur, Pierre Godard, John Hewitt, Rachid Riad, and Liming Wang. 2018. XNMT: The extensible neural machine translation toolkit. In *Conference of the Association for Machine Translation in the Americas (AMTA) Open Source Software Showcase*, Boston.
- Thai-Son Nguyen, Sebastian Stueker, Jan Niehues, and Alex Waibel. 2020. Improving Sequence-to-sequence Speech Recognition Training with On-the-fly Data Augmentation. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Barcelona, Spain.
- Jan Niehues, Roldano Cattoni, Sebastian Stueker, Matteo Negri, Marco Turchi, et al. 2019a. The IWSLT 2019 Evaluation Campaign. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Jan Niehues, Roldano Cattoni, Sebastian Stueker, Mauro Cettolo, Marco Turchi, and Marcello Federico. 2018. The IWSLT 2018 Evaluation Campaign. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Bruges, Belgium.
- Jan Niehues, Roldano Cattoni, Sebastian Stueker, Matteo Negri, Marco Turchi, Thanh-Le ha, Elizabeth Salesky, Ramon Sanabria, Loic Barrault, Lucia Specia, and Marcello Federico. 2019b. The IWSLT 2019 Evaluation Campaign. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. **fairseq: A fast, extensible toolkit for sequence modeling**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an ASR Corpus Based on Public Domain Audio Books. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, Brisbane, Australia.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. **SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition**. In *Proceedings of the Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2613–2617, Graz, Austria.

- Stephan Peitz, Simon Wiesler, Markus Nußbaum-Thom, and Hermann Ney. 2012. Spoken language translation using automatically transcribed text in training. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Fuchun Peng and Dale Schuurmans. 2003. [Combining Naive Bayes and n-Gram Language Models for Text Classification](#). In *Proceedings of the European Conference on Information Retrieval (ECIR)*, pages 335–350, Pisa, Italy.
- Ngoc-Quan Pham, Thai-Son Nguyen, Thanh-Le Ha, Juan Hussain, Felix Schneider, Jan Niehues, Sebastian Stüker, and Alexander Waibel. 2019. The IWSLT 2019 KIT Speech Translation System. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Juan Pino, Liezl Puzon, Jiatao Gu, Xutai Ma, Arya D. McCarthy, and Deepak Gopinath. 2019. [Harnessing Indirect Training Data for End-to-End Automatic Speech Translation: Tricks of the Trade](#). In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Hong Kong.
- Matt Post. 2018. [A Call for Clarity in Reporting BLEU Scores](#). In *Proceedings of the Conference on Machine Translation (WMT)*, pages 186–191, Brussels, Belgium.
- Tomasz Potapczyk and Pawel Przybylski. 2020. SRPOL’s System for the IWSLT 2020 End-to-End Speech Translation Task. In *Proceedings of the International Conference on Spoken Language Translation (IWSLT)*, Virtual Event.
- Daniel Povey, Hossein Hadian, Pegah Ghahremani, Ke Li, and Sanjeev Khudanpur. 2018. A Time-Restricted Self-Attention Layer for ASR. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5874–5878, Calgary, Canada.
- Nicholas Ruiz, Qin Gao, William Lewis, and Marcello Federico. 2015. Adapting machine translation models toward misrecognized speech with text-to-speech pronunciation rules and acoustic confusability. In *Proceedings of the Conference of the International Speech Communication Association (INTERSPEECH)*, Dresden, Germany.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. [How2: A Large-scale Dataset For Multimodal Language Understanding](#). In *Proceedings of Visually Grounded Interaction and Language (ViGIL)*, Montréal, Canada. Neural Information Processing Society (NeurIPS).
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural Machine Translation of Rare Words with Subword Units. *arXiv preprint arXiv:1508.07909*.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A Study of Translation Edit Rate with Targeted Human Annotation. In *Proceedings of the Conference of the Association for Machine Translation of the Americas (AMTA)*, pages 223–231, Cambridge, US-MA.
- Matthew Snover, Nitin Madnani, Bonnie J Dorr, and Richard Schwartz. 2009. Fluency, Adequacy, or HTER?: exploring different human judgments with a tunable MT metric. In *Proceedings of the Workshop on Statistical Machine Translation (WMT)*, pages 259–268, Athens, Greece.
- Matthias Sperber, Graham Neubig, Ngoc-Quan Pham, and Alex Waibel. 2019. Self-attentional models for lattice inputs. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1185–1197, Florence, Italy.
- Matthias Sperber, Jan Niehues, Graham Neubig, Sebastian Stüker, and Alex Waibel. 2018. [Self-Attentional Acoustic Models](#). In *Proceedings of the Conference of the International Speech Communication Association (INTERSPEECH)*, pages 3723–3727, Hyderabad, India.
- Matthias Sperber, Jan Niehues, and Alex Waibel. 2017. Toward Robust Neural Machine Translation for Noisy Input Sequences. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*, Tokyo, Japan.
- Matthias Sperber and Matthias Paulik. 2020. [Speech translation and the end-to-end promise: Taking stock of where we are](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 7409–7421, Virtual Event.
- Fred W.M. Stentiford and Martin G. Steer. 1988. Machine translation of speech. *British Telecom Technology Journal*, 6(2):116–122.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, Las Vegas, Nevada, United States.
- Jörg Tiedemann. 2016. Opus – parallel corpora for everyone. *Baltic Journal of Modern Computing*, page 384. Special Issue: Proceedings of the 19th Annual Conference of the European Association of Machine Translation (EAMT).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of Advances in Neural Information Processing Systems 30 (NIPS)*, pages 5998–6008, Long Beach, US-CA.
- Alex Waibel, Ajay N Jain, Arthur E McNair, Hiroaki Saito, Alexander G Hauptmann, and Joe Tebelskis. 1991. Janus: a speech-to-speech translation system

using connectionist and symbolic processing strategies. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 793–796, Toronto, Canada.

Andy Way. 2018. Quality Expectations of Machine Translation. In S. Castilho, J. Moorkens, F. Gaspari, and S. Doherty, editors, *Translation quality assessment: From Principles to Practice*, pages 159–178. Springer.

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 the Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2625–2629, Stockholm, Sweden.

A Systems’ Description

In this section we describe the ST models created for our study (see Section 3.1). All the details about the different trainings are given below, while the validation set was common to all trainings, since we used the MuST-C *dev set*.

The source code for the ASR and the direct ST models is available at: <https://github.com/mgaido91/FBK-fairseq-ST>.

The source code for the MT component of the cascade model can be found at: <https://github.com/modernmt/modernmt>.

A.1 Cascade approach

The Cascade system is composed of a pipeline of automatic speech recognition (ASR) and machine translation (MT) models.

The ASR model is a slightly revisited version (Gaido et al., 2020) of the S-Transformer (Di Gangi et al., 2019b), where the two 2D self-attention layers are replaced with two Transformer encoder layers (for a total of 8 layers), while the decoder is the same (with 6 layers). Hence, the model processes the input with two 3x3 2D CNNs (having 64 filters), whose output is first projected into a higher-dimensional space and then summed with positional embeddings before being fed to the Transformer encoder layers; Transformer encoder layers use logarithmic distance penalty. The attention mechanism consists of 8 attention heads. The dimensionality of input and output is 512, while the inner-layers have dimensionality 2048. The resulting number of parameters is 63M.

The ASR model was trained with the goal of achieving state-of-the-art performance. To this aim, we relied on two data augmentation techniques

that were shown to yield competitive models at the IWSLT-2020 evaluation campaign (Ansari et al., 2020), namely: *i*) SpecAugment (Park et al., 2019) applied with probability 0.5 by masking two bands on the frequency axis (with 13 as maximum mask length) and two on the time axis (with 20 as maximum mask length), and *ii*) time stretch (Nguyen et al., 2020) with probability of 0.3 and stretching factor sampled uniformly for each utterance between 0.8 and 1.25. The ASR model was trained on 1.25M utterance-transcript pairs coming from the ASR corpora Librispeech (Panayotov et al., 2015), Mozilla Common Voice,¹³ How2 (Sanabria et al., 2018), TEDLIUM-v3 (Hernandez et al., 2018), as well as the ST corpora Europarl-ST (Iranzo-Sánchez et al., 2020) and MuST-C (Cattoni et al., 2020).¹⁴ We filtered out all pairs whose utterance was longer than 20 seconds. The audio input was preprocessed with XNMT¹⁵ (Neubig et al., 2018) to extract 40 features per time frame (with 25ms windows and 10ms sliding) and per-speaker normalization was applied. The text was preprocessed by normalizing punctuation and de-escaping special characters, and was tokenized with Moses.¹⁶ Then it was encoded with a BPE (Sennrich et al., 2015) code learnt on the OPUS data¹⁷ using 8k merge rules.

The MT component is built on the ModernMT framework¹⁸ which features machine translation implementing the Transformer architecture. We trained either a Base (en-it) or a Big (en-{de,es}) Transformer model (Vaswani et al., 2017) with 6 blocks in the encoder and 6 in the decoder, 512/1024 as input size, the same as output size, 2048/4096 as inner dimension and 8/16 attention heads. The total number of parameters is about 61M for the Base model, 210M for the Big models.

As regards pre-processing, for all the three language directions we used the internal ModernMT procedures.

In training, models are optimized with Adam using $\beta_1=0.9$, $\beta_2=0.98$; the learning rate is linearly increased during the warmup (8k iterations) up to

¹³<https://voice.mozilla.org/>

¹⁴For English-German, the ST corpora include also the Speech-Translation TED corpus provided in the IWSLT offline-speech-translation task: http://iwslt.org/doku.php?id=offline_speech_translation

¹⁵<https://github.com/neulab/xnmt>

¹⁶<https://github.com/moses-smt/mosesdecoder>

¹⁷<http://opus.nlpl.eu>

¹⁸<https://github.com/modernmt/modernmt>

the maximum value (5×10^{-4}), after that it follows an inverse square root decay; dropout is set to 0.3. Minibatches consist of 3072 tokens and update frequency is set to 4; the total number of iterations is 200k; the last 10 saved checkpoints (one out of 1k iterations) are averaged. The model uses label smoothing with a uniform prior distribution (0.1) over the vocabulary; source and target languages share a BPE vocabulary of 32k sub-words.

	#segments	#en words	#trg words
en-de	58.2M	776.4M	723.3M
en-es	70.1M	972.5M	1024.9M
en-it	67.9M	792.6M	770.2M

Table 11: Statistics of the parallel training sets collected from the OPUS repository for the three language pairs.

The training data, whose statistics are reported in Table 11, are collected from the OPUS repository. For English-Italian, they resulted in almost 70M segment pairs and about 800M English words; after deduplication and the internal ModernMT cleaning, the actual training data is reduced to 45M pairs and 550M English words. For English- $\{\text{German, Spanish}\}$ pairs, the OPUS data were filtered through well-known data selection methods (Axelrod et al., 2011) using a general-domain seed; the resulting training data consist of, respectively, 17M and 19M segment pairs, for 270M and 330M English words. Trainings were performed on RTX 2080 Ti GPUs; for English-Italian, it was run on 7 GPUs and lasted 3 days, while for each of the other two directions, on a single GPU, it took 6 days.

The three models are then fine-tuned on MuST-C training data ($\sim 250\text{K}$ pairs, 4-5M English words) by continuing the training for 4k iterations on the adaptation data, with a learning rate reduced by a factor of 5. To mitigate error propagation and make the MT system more robust to ASR errors, similarly to (Di Gangi et al., 2019a) fine tuning is run on the concatenation of human and automatic transcripts of MuST-C, both paired with manual translations.

A.2 Direct approach

Our direct model (Gaido et al., 2020) uses the same architecture of the English ASR model described in §A.1, but it has 11 Transformer encoder layers (instead of 8) and 4 Transformer decoder layers (instead of 6) for a total of 64M parameters. The ST model’s encoder is initialized with the encoder of

the ASR model (Bansal et al., 2019), with the missing layers initialized randomly. The ST decoder is also initialized randomly.

The training settings and the data augmentation methods employed for the direct ST model are the same described in Section A.1 for the ASR component of the cascade system. In addition, we performed synthetic data generation, by automatically translating the English transcripts of the ASR training corpora (Jia et al., 2019). Furthermore, we transfer knowledge from MT through knowledge distillation (Hinton et al., 2015). Knowledge distillation is performed from a *teacher* MT model by optimizing the KL divergence between the distributions produced by the teacher and the *student* ST model being trained (Liu et al., 2019). The teacher MT model is trained on the OPUS datasets (Tiedemann, 2016) and is a plain transformer with 16 attention heads and 1024 features in encoder/decoder embeddings, resulting into 212M parameters.

The direct ST model is trained in two consecutive steps. First, it is optimized using KD. Then, the resulting model is fine-tuned on label-smoothed cross entropy (Szegedy et al., 2016). The training set is composed of the same corpora used for the ASR model, more precisely: *i*) the ST corpora and *ii*) the synthetic datasets derived from the ASR corpora.

The ST model is fed with the input utterance and a token representing the *type* of the target data, which can be: *i*) human reference translations (for the ST corpora), or *ii*) translations generated by the MT model fed with true case transcripts with punctuation, and *iii*) translations generated by the MT model fed with lower-cased transcripts without punctuation (for the ASR corpora). At inference time, the token “human reference” is always used to generate the translations. The token is added to the features extracted from the audio before they are passed to the encoder (Di Gangi et al., 2019c).

All trainings were performed on 8 K80 GPUs. The training of each direct model lasted 10 days, while the ASR and MT pre-trainings 6 days each.

The source code¹⁹ implemented to build these models is based on Fairseq (Ott et al., 2019).

¹⁹<https://github.com/mgaido91/FBK-fairseq-ST>

B Post-Editing Guidelines

In this task you are presented with (i) 550 audio segments that are recordings of portions of different English TED Talks, (ii) their transcripts, and (iii) corresponding automatic translations.

Starting from the original audio recording and its corresponding transcript (done by TED volunteer translators), you are asked to post-edit each given automatic translation by applying the minimal edits required to transform the system output into a fluent sentence with the same meaning as the audio/transcript.

While post-editing, remember the following guidelines:

- We noticed that some audio player software applications cut the beginning or the end of the audio segments. If you notice some audio-transcript out-of-sync, please try another audio player or inform us about the problem.
- The audio should be your first source of information, while transcripts are given for your convenience. It could happen that the transcript is not faithful to the spoken original: in these cases you should not consider the transcript and refer to the audio only.
- Some transcripts contain the name or initials of the speaker (typically followed by colons). Please don't add this information into the sentence you are post-editing. In general, don't include in your post-edit any text that is not present in the audio (e.g. explanation of acronyms, disambiguation of pronouns), even though this information could ease the understanding of the sentence.
- The post-edited sentence is intended as a translation of spoken language. Also, depending on the style of the source language talk, you can use the corresponding style in the target language (e.g. if the talk uses a friendly/colloquial style you can use informal words too).
- The focus is the correctness of the single sentence within the given context, not the consistency of a group of sentences. Hence, surrounding segments should be used to understand the context but not to enforce consistency on the use of terms. In particular, different but correct translations of terms across segments should not be corrected.

C Tool for Automatic Error Classification

The tool used for the automatic analysis of linguistic errors (Section 6.1) is downloadable at wit3.fbk.eu/2016-02. It is a modified version of the *tercom* script,²⁰ which requires the lemmatized versions of both systems' outputs and post-edits. To lemmatize the data we used the *TreeTagger*.²¹

²⁰www.cs.umd.edu/~snover/tercom

²¹www.cis.uni-muenchen.de/~schmid/tools/TreeTagger