

Is “moby dick” a Whale or a Bird? Named Entities and Terminology in Speech Translation

Marco Gaido^{1,2}, Susana Rodríguez³, Matteo Negri¹, Luisa Bentivogli¹ and Marco Turchi¹

¹Fondazione Bruno Kessler

²University of Trento

³Independent Researcher

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

Abstract

Automatic translation systems are known to struggle with rare words. Among these, named entities (NEs) and domain-specific terms are crucial, since errors in their translation can lead to severe meaning distortions. Despite their importance, previous speech translation (ST) studies have neglected them, also due to the dearth of publicly available resources tailored to their specific evaluation. To fill this gap, we *i*) present the first systematic analysis of the behavior of state-of-the-art ST systems in translating NEs and terminology, and *ii*) release NEuRoparl-ST, a novel benchmark built from European Parliament speeches annotated with NEs and terminology. Our experiments on the three language directions covered by our benchmark (en→es/fr/it) show that ST systems correctly translate 75–80% of terms and 65–70% of NEs, with very low performance (37–40%) on person names.

1 Introduction

The translation of rare words is one of the main challenges for neural machine translation (NMT) models (Senrich et al., 2016; Koehn and Knowles, 2017). Among rare words, named entities (NEs) and terminology are particularly critical: not only are they important to understand the meaning of a sentence (Li et al., 2013), but they are also difficult to handle due to the small number of valid translation options. While common words can be rendered in the target language with synonyms or paraphrases, NEs and terminology offer less expressive freedom, which is typically limited to one valid option. Under these conditions, translation errors often result in blatant (meaningless, hilarious, or even offensive) errors, which jeopardize users’ trust in the translation system. One example is “moby dick” (in lower case, as in the typical output of a speech recognition system): Google Translate¹ returns *mazikó poulí* (massive bird) for

¹Accessed on the 27th April 2021.

Greek, while the translation contains profanities for other languages like Hungarian or Italian.

Previous works in NMT tried to mitigate the problem by: *i*) integrating knowledge graphs (KG) into the models (Lu et al., 2019; Moussallem et al., 2019; Zhao et al., 2020a,b; Ahmadnia et al., 2020), *ii*) exploiting dedicated modules for NE translation (Yan et al., 2019), or *iii*) adding NE tags to the source (Ugawa et al., 2018; Zhou et al., 2020). For terminology, the so far proposed dictionary-based approaches (Hokamp and Liu, 2017; Chatterjee et al., 2017; Hasler et al., 2018; Dinu et al., 2019; Song et al., 2020; Dougal and Lonsdale, 2020) share the idea of enriching the source sentence with terms’ translations found in a dictionary (e.g. “this is a term”→“this is a #term#*término*”).

The problem is even more challenging in automatic speech recognition (ASR) and speech-to-text translation (ST), where a lookup into KGs or dictionaries is not feasible due to the different modality of the input (an audio signal rather than text). As regards NEs, the few existing studies (Ghannay et al., 2018; Caubrière et al., 2020) are all limited to ASR, for which two benchmarks are available (Galibert et al., 2014; Yadav et al., 2020), while suitable benchmarks do not even exist for ST. The situation is similar for terminology: few annotated test sets exist for MT (Dinu et al., 2019; Scansani et al., 2019; Bergmanis and Pinnis, 2021), but none for ST, which so far has remained unexplored.

In light of the above, the contribution of this work is twofold: (1) we present the first investigation on the behavior of state-of-the-art ST systems in translating NEs and terms, discussing their weaknesses and providing baseline results for future comparisons; (2) we release the annotated data that made our study possible. Our test set – NEuRoparl-ST – is derived from Europarl-ST (Iranzo-Sánchez et al., 2020) and covers three language pairs: en→es/fr/it. It relies on the Europarl-ST (audio, transcript, translation) triplets and en-

riches their textual portions with NE and terminology annotation. Besides being the first benchmark of this type for ST, it can also be used for the evaluation of NE/terminology recognition (ASR) and translation (MT). The dataset is available at: ict.fbk.eu/neuroparl-st/.

2 Speech Translation Models

Our goal is to assess the capability of state-of-the-art ST systems to properly translate NEs and terminology present in an utterance. To this aim, we compare instances of the two main approaches. One is the traditional *cascade* approach (Stentiford and Steer, 1988; Waibel et al., 1991), which consists of a pipeline where an ASR model produces a transcript of the input audio and an MT model generates its translation. The other is the so-called *direct* approach (Bérard et al., 2016; Weiss et al., 2017), which relies on a single neural network that maps the audio into target language text bypassing any intermediate symbolic representation. The two approaches have inherent strengths and weaknesses (Sperber and Paulik, 2020). Cascade solutions can exploit sizeable datasets for the ASR and MT sub-components, but rely on a complex architecture prone to error propagation. Direct models suffer from the paucity of training data, but avoid error propagation and can take advantage of unmediated access to audio information (e.g. prosody) during the translation phase. In recent years, after a long dominance of the cascade paradigm, the initially huge performance gap between the two approaches has gradually closed (Ansari et al., 2020).

Our **cascade** system integrates competitive Transformer-based (Vaswani et al., 2017) ASR and MT components built from large training corpora. Specifically, the ASR model is trained on LibriSpeech (Panayotov et al., 2015), TEDLIUM v3 (Hernandez et al., 2018) and Mozilla Common Voice,² together with (*utterance, transcript*) pairs extracted from MuST-C (Cattoni et al., 2021), Europarl-ST (Iranzo-Sánchez et al., 2020), and CoVoST 2 (Wang et al., 2020) ST corpora. ASR outputs are post-processed to add true-casing and punctuation. The MT model is trained on data collected from the OPUS repository,³ amounting to about 19M, 28M, and 45M parallel sentence pairs respectively for en-es, en-fr, and en-it.

Our **direct** model has the same Transformer-

²<http://commonvoice.mozilla.org/en/>

³<http://opus.nlpl.eu>

based architecture of the ASR component used in the cascade system. It exploits data augmentation and knowledge transfer techniques successfully applied by participants in the IWSLT-2020 evaluation campaign (Ansari et al., 2020; Potapczyk and Przybysz, 2020; Gaido et al., 2020a) and it is trained on MuST-C, Europarl-ST and synthetic data (~1.5M pairs for each language direction).

Systems’ performance is shown in Table 2 and discussed in Section 4. Complete details about their implementation and training procedures are provided in the Appendix. All the related code is available at <https://github.com/mgaido91/FBK-fairseq-ST/tree/emnlp2021>.

3 Evaluation Data: NEuRoparl-ST

To the best of our knowledge, freely available NE/term-labelled ST benchmarks suitable for our analysis do not exist. The required resource should contain *i*) the audio corresponding to an utterance, *ii*) its transcript, *iii*) its translation in multiple target languages (three in our case), and *iv*) NE/term annotation in both transcripts and target texts. Currently available MT, ST, ASR, NE and terminology datasets lack at least one of these key components. For example, most MT corpora (e.g. Europarl) lack both the audio sources and NE/terminology annotations. The very few available MT corpora annotated with NE/terminology still lack the audio portion, and extending them to ST would require generating synthetic audio, which is known to be problematic for models’ performance. For these reasons, we preferred to create a the en→es/fr/it transcripts and translations of the Europarl-ST test sets, which are mainly derived from the same original speeches. The result is a multilingual benchmark featuring very high content overlap, thus enabling cross-lingual comparisons.

NE annotation. We used the 18 tags and the annotation scheme defined by the guidelines (“OntoNotes Named Entity Guidelines - Version 14.0”) used to annotate the OntoNotes5 corpus (Weischedel et al., 2012). The annotation was carried out manually by a professional interpreter with a multi-year experience in translating from English, French and Italian into Spanish the verbatim reports of the European Parliament plenary meetings. This guarantees the high level of language knowledge and domain expertise required to ensure maximum quality and precision. To ease the task, the annotator was provided with transcripts

	en-es		en-fr		en-it	
	en	es	en	fr	en	it
NEs	1,637 (2,703)	1,638 (3,003)	1,578 (2,604)	1,562 (2,949)	1,523 (2,497)	1,466 (2,649)
TERMS	2,571 (3,174)	2,662 (3,294)	2,797 (3,502)	2,947 (3,659)	2,166 (2,669)	2,202 (2,645)
Num. of sentences	1,267		1,214		1,130	

Table 1: Total number of named entities and terms annotated in the test sets (and corresponding number of tokens).

and translations automatically pre-annotated with the BERT-based NER model⁴ available in DeepPavlov (Burtsev et al., 2018). Human annotation was then conducted in parallel on the three test sets by labelling, for each audio segment, the English transcript and the three corresponding translations. To check annotations’ reliability, all the English transcripts were also independently labelled by a second annotator with a background in linguistics and excellent English knowledge. Inter-annotator agreement was calculated in terms of *complete* agreement, i.e. the exact match of the whole NE in the two annotations. The resulting Dice coefficient⁵ (Dice, 1945) amounts to 93.87% and can be considered highly satisfactory. For the subset of NEs for which complete agreement was found (1,409 in total), we also computed the agreement on labels’ assignment with the *kappa coefficient* (in Scott’s π formulation) (Scott, 1955; Artstein and Poesio, 2008). The resulting value is 0.94, which corresponds to “almost perfect” agreement according to its standard interpretation (Landis and Koch, 1977).

Terminology annotation. Similar to (Dinu et al., 2019), terminology was automatically extracted by exploiting the IATE term base.⁶ Each entry in IATE has an identifier and a language code. Entries with the same identifier and different language codes represent the translations of a term in the corresponding languages. To annotate our parallel texts, we first removed stop-words and lemmatized the remaining words and IATE entries.⁷ Then, for each parallel sentence, we marked as terms only those words in the source and the target side that were present in IATE with the same identifier. This source/target match is essential to avoid the annotation of words that are used with a generic, common meaning but, being polysemic, can be technical terms in different contexts (e.g. the word “board”

can refer to a tool or to a committee). Checking the presence of the corresponding translation in the target language disambiguates these cases, leading to a more accurate annotation.

NE and term annotations were merged into a single test set using BIO (Ramshaw and Marcus, 1995) as span labeling format. Had a word been tagged both as term and NE, the latter was chosen favoring the more reliable manual annotation. Table 1 presents the total number of NEs and terms for the three language pairs, together with their corresponding number of tokens.⁸ These numbers differ between source and target texts and across pairs due to the peculiarities of the Europarl-ST data. Specifically, *i*) sometimes translations are not literal and NEs are omitted in the translation (e.g. when a NE is repeated in the source, one of the occurrences may be replaced by a pronoun in the target text), *ii*) the professional interpreters and translators “localize” the target translations, i.e. adapt them to the target culture (e.g. while the English source simply contains the name and surname of mentioned European Parliament members, in Italian the first name is omitted and the surname is preceded by “onoverole” - honorable), and *iii*) the number of words a NE is made of can vary across languages (e.g. “European Timeshare Owners Organisation” becomes “Organización Europea de Socios de Tiempo Compartido” in Spanish).

4 Results

We use our benchmark to measure systems’ ability to handle NEs and terminology. Besides comparing the two ST models described in Section 2, we extend our evaluation to the ASR and MT sub-components (the latter being fed with human transcripts) of the cascade system. As shown in Table 2, all models are evaluated in terms of overall output quality and accuracy in rendering the two categories of rare words subject of our study. Transcription and translation quality are respectively measured with WER and SacreBLEU⁹ (Post,

⁴<http://docs.deeppavlov.ai/en/master/features/models/ner.html>

⁵Note that Dice coefficient has the same value of the F1 measure computed considering either annotator as reference.

⁶<http://iate.europa.eu>

⁷Preprocessing made with Spacy: <http://spacy.io/>

⁸The statistics for each NE type are given in the Appendix.

⁹BLEU+c.mixed+#.1+s.exp+tok.13a+v.1.5.0

	en-es				en-fr				en-it			
	WER	BLEU	NE	Term	WER	BLEU	NE	Term	WER	BLEU	NE	Term
ASR	12.6	–	84.6	92.6	12.7	–	84.5	92.1	12.6	–	84.3	92.4
MT	–	48.8	83.5	88.8	–	36.2	78.8	85.7	–	33.8	80.2	86.9
Cascade	–	37.6	70.9	82.5	–	28.3	66.7	80.4	–	26.5	66.9	80.4
Direct	–	37.7	71.4	79.2	–	30.1	67.3	77.7	–	26.0	67.3	76.3

Table 2: WER/BLEU and NE/term case-insensitive accuracy for ASR, MT and ST (cascade and direct) models.

2018). Similarly to the Named Entity Weak Accuracy proposed in (Hermjakob et al., 2008), we compute NE/term accuracy¹⁰ as the ratio of entities that are present in the systems’ output in the correct form. We present case-insensitive accuracy scores to fairly compare the different models, as the ASR produces lowercase text.

For the sake of completeness, case-sensitive NE/term accuracy is also given in Table 3 for ST and MT models (we do not include ASR since it generates lowercase text). Comparing these results with those reported in Table 2, for all language pairs we see that the drop in NEs accuracy with respect to case-insensitive scores is higher for the cascade model – around 5 points – than for the direct one – around 2 points (e.g. for en-es, from 70.9 to 65.8 for the cascade model and from 71.4 to 69.4 for the direct model). We posit the reason is the propagation of errors in the module in charge to restore casing on the ASR output in the cascade architecture.

	en-es		en-fr		en-it	
	NE	Term	NE	Term	NE	Term
MT	81.0	88.0	75.5	85.3	77.5	86.2
Cascade	65.8	81.6	61.3	79.9	62.6	79.5
Direct	69.4	78.7	65.9	77.3	65.1	75.9

Table 3: Case sensitive accuracy scores of MT and ST (cascade and direct) models on en→es/fr/it.

4.1 ASR and MT results

The WER of the ASR is similar across the three language directions. This is not surprising because the three test sets differ only in very few debates. In terms of accuracy, it is evident that transcribing NEs is more difficult than transcribing terms (84.5 vs 92.4 on average). Besides lower frequency, the higher difficulty to transcribe NEs can be ascribed to the variety of different pronunciations by non-native speakers (in particular for person, product

¹⁰Scores have been computed with the script available at: https://github.com/mgaido91/FBK-fairseq-ST/blob/emnlp2021/scripts/eval/ne_terms_accuracy.py

and organization names). Concerning the MT performance, the BLEU differences between language directions (en-es \gg en-fr $>$ en-it) reflect the results reported in the Europarl-ST paper (Iranzo-Sánchez et al., 2020). The main reason is that the translations are less literal for some language directions. For instance, the French references are 20% longer than the human source transcripts. Analyzing NE and term translation quality, we notice that NEs are, again, harder to handle compared to terminology (average accuracy: 80.8 vs 87.1). It is worth to notice that accuracy does not strictly depend on translation quality. For instance, en-fr has a higher translation quality than en-it (+2.4 BLEU points), but NE and term accuracy scores are lower.

4.2 ST results

Unsurprisingly, when it comes to combining transcription and translation in a single task, performance decreases significantly. In particular, the results of the cascade model are a direct consequence of cumulative ASR and MT errors. As such, like for its sub-components, NEs are harder to handle than terms. Compared to MT results computed on manual transcripts, we see large drops in all languages on both translation quality (-13.2 BLEU on average) and NE/term accuracy (-12.8/-6.0).

Comparing cascade and direct models, the BLEU scores are on par for en-es and en-it (differences are not statistically significant¹¹), while the direct one is significantly better for en-fr. This is explained by the aforementioned peculiarity of the French reference translations in Europarl-ST that, unlike in common training corpora (Europarl included), are on average 20% longer than the source transcripts. The MT model of the cascade, trained on massive corpora including Europarl, tends to produce translations that are similar in length to the transcripts and shorter than Europarl-ST references, being thus penalized. Having Europarl-ST among its training corpora, the direct model produces outputs more similar in length to the references, resulting in a 2.8 BLEU gain.

¹¹Computed with bootstrap resampling (Koehn, 2004).




Figure 1: Accuracy scores on PERSON and LOCATION of MT, ASR and ST systems on en-es.

In terms of NE and term translation quality, the trend is clear and coherent on all languages: the cascade outperforms the direct on terminology (+3.5 on average), while the direct has an edge (+0.5) in handling NEs. The advantage of the cascade on terminology can be explained by the higher reliability of its MT component in selecting domain-specific target words compared to the direct models built on much smaller ST training corpora. One example is the English term “plastic explosive”, which is correctly translated into Italian by the cascade (“esplosivo plastico”), and wrongly by the direct (“esplosivo di plastica” - En: “explosive made of plastic”). Concerning NEs, instead, the unmediated access to the audio helps the direct to avoid both *i*) error propagation (e.g. the NE “Lamfalussy” is correctly translated by the direct, while the MT component of the cascade is not able to recover the wrong ASR output “blunt Hallucy”), and *ii*) the translation of NEs that are homographs of common nouns in the source language but should be copied *as is* (e.g. the English surname “Parish” is translated into Italian as “Parrocchia” by the cascade, but correctly preserved in the direct’s output).

Looking at NE types (complete results in the Appendix), the two ST systems are always close to each other, reflecting the global accuracy scores in Table 2. For both approaches, the differences across the NE types depend on their capability to *recognize* entities in the audio and properly *translate* them. Two types are paradigmatic (see Figure 1). PERSON names (the worst category, with 37–40% ST accuracy) are difficult to recognize in the audio, as shown by the poor performance of ASR and both ST systems, while their translation from manual transcripts (MT) is trivial as it only requires copying them from the source. Conversely, ST and MT results are very close on the

more frequent and normally easier to pronounce LOCATION names, for which the problem lies more in translation than in recognition.

5 Conclusions

While previous ST research has focused on improving overall systems’ performance, little has been done to evaluate the existing paradigms in relation to well known specific problems in automatic translation at large. Translating rare words is no exception, also due to the dearth of suitable labelled benchmarks. To fill this gap, we focused on named entities and terminology, which combine the problems inherent to low frequency in the training data with the difficulty of recognizing them in the audio and mapping their meaning into few valid options. We created NEuRoparl-ST, an annotated benchmark covering three language directions, and used it for the first comparison of state-of-the-art cascade and direct ST systems on NE and term translation. Our results show that NEs, especially person names, are in general more difficult to handle than terminology.

Acknowledgement

This work has been carried out as part of the project Smarter Interpreting (<https://kunveno.digital/>) financed by CDTI Neotec funds.

References

- Benyamin Ahmadnia, Bonnie J. Dorr, and Parisa Kordjamshidi. 2020. [Knowledge Graphs Effectiveness in Neural Machine Translation Improvement](#). *Computer Science*, 21(3).
- 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. Association for Computational Linguistics.
- Ron Artstein and Massimo Poesio. 2008. [Inter-coder agreement for computational linguistics](#). *Computational Linguistics*, 34(4):555–596.
- 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*, pages 355–362, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- 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 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 58–68, Minneapolis, Minnesota. Association for Computational Linguistics.
- Toms Bergmanis and Mārcis Pinnis. 2021. **Facilitating terminology translation with target lemma annotations.** In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3105–3111, Online. Association for Computational Linguistics.
- Nicola Bertoldi, Davide Caroselli, and Marcello Federico. 2018. The modernmt project. In *Proceedings of the Conference of the European Association for Machine Translation (EAMT)*.
- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, Alexey Litinsky, Varvara Logacheva, Alexey Lymar, Valentin Malykh, Maxim Petrov, Vadim Polulyakh, Leonid Pugachev, Alexey Sorokin, Maria Vikhрева, and Marat Zaynutdinov. 2018. **DeepPavlov: Open-Source Library for Dialogue Systems.** In *Proceedings of ACL 2018, System Demonstrations*, pages 122–127, Melbourne, Australia. Association for Computational Linguistics.
- Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation. In *NIPS Workshop on end-to-end learning for speech and audio processing*, Barcelona, Spain.
- Roldano Cattoni, Mattia Antonino Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. **Must-c: A multilingual corpus for end-to-end speech translation.** *Computer Speech & Language*, 66:101155.
- Antoine Caubrière, Sophie Rosset, Yannick Estève, Antoine Laurent, and Emmanuel Morin. 2020. **Where are we in Named Entity Recognition from Speech?** In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4514–4520, Marseille, France. European Language Resources Association.
- Rajen Chatterjee, Matteo Negri, Marco Turchi, Marcello Federico, Lucia Specia, and Frédéric Blain. 2017. **Guiding Neural Machine Translation Decoding with External Knowledge.** In *Proceedings of the Second Conference on Machine Translation*, pages 157–168, Copenhagen, Denmark. Association for Computational Linguistics.
- Mattia A. Di Gangi, Matteo Negri, Roldano Cattoni, Roberto Dessi, and Marco Turchi. 2019. Enhancing Transformer for End-to-end Speech-to-Text Translation. In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pages 21–31, Dublin, Ireland. European Association for Machine Translation.
- Lee Raymond Dice. 1945. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302.
- Georgiana Dinu, Prashant Mathur, Marcello Federico, and Yaser Al-Onaizan. 2019. **Training Neural Machine Translation to Apply Terminology Constraints.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3063–3068, Florence, Italy. Association for Computational Linguistics.
- Duane K. Dougal and Deryle Lonsdale. 2020. **Improving NMT Quality Using Terminology Injection.** In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4820–4827, Marseille, France. European Language Resources Association.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2020a. **End-to-end speech-translation with knowledge distillation: FBK@IWSLT2020.** In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 80–88, Online. Association for Computational Linguistics.
- Marco Gaido, Mattia Antonino Di Gangi, Matteo Negri, and Marco Turchi. 2020b. **On Knowledge Distillation for Direct Speech Translation.** In *Proceedings of the Seventh Italian Conference on Computational Linguistics, CLiC-it 2020, Bologna, Italy, March 1-3, 2021*, volume 2769 of *CEUR Workshop Proceedings*.
- Olivier Galibert, Jeremy Leixa, Gilles Adda, Khalid Choukri, and Guillaume Gravier. 2014. **The ETAPE speech processing evaluation.** In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 3995–3999, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Sahar Ghannay, Antoine Caubrière, Yannick Estève, Antoine Laurent, and Emmanuel Morin. 2018. **End-to-end named entity extraction from speech.**
- 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.

- Eva Hasler, Adrià de Gispert, Gonzalo Iglesias, and Bill Byrne. 2018. [Neural Machine Translation Decoding with Terminology Constraints](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 506–512, New Orleans, Louisiana. Association for Computational Linguistics.
- Ulf Hermjakob, Kevin Knight, and Hal Daumé III. 2008. [Name Translation in Statistical Machine Translation - Learning When to Transliterate](#). In *Proceedings of ACL-08: HLT*, pages 389–397, Columbus, Ohio. Association for Computational Linguistics.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia A. Tomashenko, and Yannick Estève. 2018. [TED-LIUM 3: Twice as much data and corpus repartition for experiments on speaker adaptation](#). In *Speech and Computer - 20th International Conference, SPECOM 2018, Leipzig, Germany, September 18-22, 2018, Proceedings*, volume 11096 of *Lecture Notes in Computer Science*, pages 198–208. Springer.
- Chris Hokamp and Qun Liu. 2017. [Lexically Constrained Decoding for Sequence Generation Using Grid Beam Search](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1535–1546, Vancouver, Canada. Association for Computational Linguistics.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Adrià Giménez, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. [Europarl-ST: A Multilingual Corpus for Speech Translation of Parliamentary Debates](#). In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233.
- 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*.
- Philipp Koehn. 2004. [Statistical Significance Tests for Machine Translation Evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. [Six Challenges for Neural Machine Translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Solomon Kullback and Richard Arthur Leibler. 1951. [On information and sufficiency](#). *Ann. Math. Statist.*, 22(1):79–86.
- J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1).
- Haibo Li, Jing Zheng, Heng Ji, Qi Li, and Wen Wang. 2013. [Name-aware Machine Translation](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 604–614, Sofia, Bulgaria. Association for Computational Linguistics.
- 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 Interspeech 2019*, pages 1128–1132, Graz, Austria.
- Yu Lu, Jiajun Zhang, and Chengqing Zong. 2019. Exploiting knowledge graph in neural machine translation. In *Machine Translation*, pages 27–38, Singapore. Springer Singapore.
- Diego Moussallem, Axel-Cyrille Ngonga Ngomo, Paul Buitelaar, and Mihael Arcan. 2019. [Utilizing knowledge graphs for neural machine translation augmentation](#). In *Proceedings of the 10th International Conference on Knowledge Capture, K-CAP '19*, page 139–146, New York, NY, USA. Association for Computing Machinery.
- 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 2020 International Conference on Acoustics, Speech, and Signal Processing - IEEE-ICASSP-2020*, Barcelona, Spain.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An asr corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- 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 Interspeech 2019*, pages 2613–2617, Graz, Austria.
- 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. Association for Computational Linguistics.
- Tomasz Potapczyk and Pawel Przybysz. 2020. [SR-POL's System for the IWSLT 2020 End-to-End Speech Translation Task](#). In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 89–94, Online. Association for Computational Linguistics.

- Lance Ramshaw and Mitch Marcus. 1995. [Text Chunking using Transformation-Based Learning](#). In *Third Workshop on Very Large Corpora*.
- Randy Scansani, Luisa Bentivogli, Silvia Bernardini, and Adriano Ferraresi. 2019. [MAGMATic: A multi-domain academic gold standard with manual annotation of terminology for machine translation evaluation](#). In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pages 78–86, Dublin, Ireland. European Association for Machine Translation.
- William A. Scott. 1955. [Reliability of Content Analysis: The Case of Nominal Scale Coding](#). *Public Opinion Quarterly*, 19(3):321–325.
- 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. Association for Computational Linguistics.
- Kai Song, Kun Wang, Heng Yu, Yue Zhang, Zhongqiang Huang, Weihua Luo, Xiangyu Duan, and Min Zhang. 2020. [Alignment-Enhanced Transformer for Constraining NMT with Pre-Specified Translations](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8886–8893.
- 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.
- Frederick 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.
- Arata Ugawa, Akihiro Tamura, Takashi Ninomiya, Hiroya Takamura, and Manabu Okumura. 2018. [Neural machine translation incorporating named entity](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3240–3250, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of Advances in Neural Information Processing Systems 30 (NIPS)*, pages 5998–6008, Long Beach, California.
- 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 International Conference on Acoustics, Speech and Signal Processing, ICASSP 1991*, pages 793–796, Toronto, Canada.
- Changhan Wang, Juan Pino, Anne Wu, and Jiatao Gu. 2020. [CoVoST: A diverse multilingual speech-to-text translation corpus](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4197–4203, Marseille, France. European Language Resources Association.
- Ralph Weischedel, Sameer Pradhan, Lance Ramshaw, Jeff Kaufman, Michelle Franchini, Mohammed El-Bachouti, Nianwen Xue, Martha Palmer, Jena D. Hwang, Claire Bonial, Aous Mansouri Jinho Choi, Maha Foster, Abdel aati Hawwary, Mitchell Marcus, Ann Taylor, Craig Greenberg, Eduard Hovy, Robert Belvin, and Ann Houston. 2012. [OntoNotes Release 5.0](#).
- 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.
- Hemant Yadav, Sreyan Ghosh, Yi Yu, and Rajiv Ratn Shah. 2020. [End-to-End Named Entity Recognition from English Speech](#). In *Proceedings of Interspeech 2020*, pages 4268–4272.
- Jinghui Yan, Jiajun Zhang, JinAn Xu, and Chengqing Zong. 2019. The Impact of Named Entity Translation for Neural Machine Translation. In *Machine Translation*, pages 63–73, Singapore. Springer.
- Yang Zhao, Lu Xiang, Junnan Zhu, Jiajun Zhang, Yu Zhou, and Chengqing Zong. 2020a. [Knowledge Graph Enhanced Neural Machine Translation via Multi-task Learning on Sub-entity Granularity](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4495–4505, Online. International Committee on Computational Linguistics.
- Yang Zhao, Jiajun Zhang, Yu Zhou, and Chengqing Zong. 2020b. [Knowledge Graphs Enhanced Neural Machine Translation](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4039–4045. International Joint Conferences on Artificial Intelligence Organization.
- Leiying Zhou, Wenjie Lu, Jie Zhou, Kui Meng, and Gongshen Liu. 2020. Incorporating Named Entity Information into Neural Machine Translation. In *Natural Language Processing and Chinese Computing*, pages 391–402, Cham. Springer International Publishing.

A Models and Trainings

A.1 Cascade ST Model

The **ASR** component of our cascade is a Transformer-based (Vaswani et al., 2017) model consisting of 11 encoder layers, 4 decoder layers, 8 attention heads, 512 features for the attention layers and 2,048 hidden units in the feed-forward layers. Its encoder has been adapted for processing speech by means of two initial 2D convolutional layers that reduce the input sequence length by a factor of 4. Also, the encoder self-attentions are biased using a logarithmic distance penalty that favors the local context (Di Gangi et al., 2019). Similar to (Gaido et al., 2020a), the model is trained with an additional Connectionist Temporal Classification (CTC) loss (Graves et al., 2006), which is added as a linear layer to the 8th encoder layer. As training data, we used LibriSpeech (Panayotov et al., 2015), TEDLIUM v3 (Hernandez et al., 2018) and Mozilla Common Voice,¹² together with (*utterance, transcript*) pairs extracted from three ST corpora: MuST-C (Cattoni et al., 2021), Europarl-ST, and CoVoST 2 (Wang et al., 2020). We augment data with SpecAugment (Park et al., 2019) and, after lowercasing and punctuation removal, text is split into sub-words with 8,000 BPE (Sennrich et al., 2016) merge rules. We set the dropout to 0.1. We optimize label smoothed cross entropy with smoothing factor 0.1 with Adam (Kingma and Ba, 2015). The learning rate is increased for 5,000 steps from 0.0003 up to 0.0005 and then decays with inverse square root policy. Our mini-batches are composed of up to 12K tokens or 8 samples and we delay parameter updates for 8 mini-batches. We train on 8 GPU K80 (11GB RAM).

Before feeding the MT with the ASR outputs, the transcripts are post-processed by an additional model to restore casing and punctuation. This model is a Transformer-based system trained on data from the OPUS repository, where the source text is lowercased and without punctuation and the target text is a normally formatted sentences.

The **MT** component is a Transformer model with 6 layers for both the encoder and the decoder, 16 attention heads, 1,024 features for the attention layers and 4,096 hidden units in the feed-forward layers. Training data were collected from the OPUS repository,¹³ and cleaned with the ModernMT framework

(Bertoldi et al., 2018). At the end of this process, the actual training data is reduced to 45M of segment pairs (550M of English words) for English-Italian. For English-Spanish, the training data is further filtered with data selection methods (Axelrod et al., 2011) using a general-domain seed resulting in 19M segment pairs (330M English words). Finally, for English-French we have 28M sentence pairs (550M of English words). Models are optimized on label-smoothed cross entropy (Szegedy et al., 2016) with Adam, with a learning rate that linearly increases for 8,000 updates up to 0.0005, after which decays with inverse square root policy. Each batch is composed of 4 mini-batches made of 3072 tokens. Dropout is set to 0.3. We train for 200,000 updates and average the last 10 checkpoints. Source and target languages share a BPE (Sennrich et al., 2016) vocabulary of 32k sub-words.

A.2 Direct ST Model

Our **direct** model has the same architecture of the ASR component described above, which is also used to initialise its encoder weights (Bansal et al., 2019). In addition, we exploit data augmentation and knowledge transfer techniques successfully applied by participants in the IWSLT-2020 evaluation campaign (Ansari et al., 2020; Potapczyk and Przybysz, 2020; Gaido et al., 2020a). For data augmentation, we use SpecAugment and time stretch (Nguyen et al., 2020), together with synthetically generated data obtained by translating with our NMT model the transcripts contained in the ASR training corpora. Besides encoder pre-training, for knowledge transfer we also apply knowledge distillation (KD): as in (Liu et al., 2019; Gaido et al., 2020b), our *student* ST model is trained by computing the KL divergence (Kullback and Leibler, 1951) with the output probability distribution of the NMT model used as *teacher*. The whole training procedure is carried out in three phases: starting from the synthetically generated data (with the KD loss function), continuing with MuST-C and Europarl-ST (still with KD), and concluding with fine-tuning on the same ST data, but switching to the label-smoothed cross entropy loss.

B Statistics for the Annotated Test Sets

Table 4 presents the number of named entities (NEs) and terms annotated in the test sets, divided by category. Since both NEs and terms can be com-

¹²<https://commonvoice.mozilla.org/en/datasets>

¹³<http://opus.nlpl.eu>

	en-es		en-fr		en-it	
	en	es	en	fr	en	it
CARDINAL	91 (105)	85 (104)	87 (101)	90 (105)	86 (100)	85 (98)
DATE	149 (314)	152 (321)	145 (303)	144 (377)	141 (300)	141 (294)
EVENT	8 (26)	9 (27)	7 (22)	7 (27)	8 (26)	9 (31)
FAC	18 (31)	19 (38)	18 (31)	21 (52)	18 (31)	16 (33)
GPE	241 (338)	240 (361)	232 (322)	221 (312)	222 (316)	209 (300)
LANGUAGE	2 (2)	2 (2)	2 (2)	2 (2)	2 (2)	2 (2)
LAW	146 (478)	141 (622)	136 (448)	143 (608)	137 (439)	128 (509)
LOC	96 (121)	91 (122)	92 (118)	86 (128)	89 (111)	83 (109)
MONEY	10 (34)	11 (45)	10 (34)	11 (49)	6 (20)	6 (25)
NORP	135 (151)	126 (147)	136 (151)	156 (182)	123 (139)	143 (194)
ORDINAL	64 (64)	65 (65)	57 (57)	40 (40)	62 (62)	52 (53)
ORG	565 (857)	582 (989)	550 (844)	533 (906)	520 (773)	485 (851)
PERCENT	4 (10)	4 (6)	3 (8)	3 (9)	4 (10)	4 (14)
PERSON	92 (134)	96 (122)	88 (129)	88 (122)	89 (130)	87 (101)
PRODUCT	1 (1)	1 (1)	1 (1)	1 (1)	1 (1)	1 (1)
QUANTITY	3 (7)	3 (7)	3 (7)	3 (7)	3 (7)	3 (7)
TIME	11 (26)	10 (20)	10 (22)	9 (18)	11 (26)	11 (24)
WORK_OF_ART	1 (4)	1 (4)	1 (4)	1 (4)	1 (4)	1 (3)
TERM	2571 (3174)	2662 (3294)	2797 (3502)	2947 (3659)	2166 (2669)	2202 (2645)

Table 4: Number of named entities and terms annotated in the test sets (and corresponding number of tokens).

	en-es				en-fr				en-it			
	ASR	MT	Casc.	Dir.	ASR	MT	Casc.	Dir.	ASR	MT	Casc.	Dir.
CARDINAL	92.31	88.24	80.00	76.47	94.25	86.67	80.00	81.11	93.02	89.41	78.82	80.00
DATE	90.60	78.95	73.68	72.37	89.66	57.64	52.08	56.25	89.36	76.60	67.38	68.09
EVENT	37.50	33.33	33.33	33.33	28.57	71.43	28.57	57.14	37.50	66.67	44.44	55.56
FAC	77.78	73.68	63.16	57.89	77.78	57.14	52.38	47.62	77.78	75.00	62.50	50.00
GPE	94.61	84.17	79.17	82.50	94.40	89.19	81.98	86.88	94.62	89.00	83.25	82.30
LANGUAGE	100.00	100.00	100.00	100.00	100.00	100.00	100.00	50.00	100.00	100.00	100.00	50.00
LAW	69.86	63.12	46.10	42.55	70.59	65.73	46.15	43.36	69.34	69.53	53.13	47.66
LOCATION	93.75	85.71	79.12	81.32	92.39	81.40	77.91	74.42	93.26	79.52	79.52	74.70
MONEY	20.00	54.55	27.27	72.73	20.00	18.18	27.27	27.27	16.67	66.67	16.67	66.67
NORP	87.41	79.37	70.63	69.84	87.50	69.43	63.06	62.18	86.99	70.63	60.84	56.64
ORDINAL	90.63	81.54	72.31	69.23	89.47	80.00	65.00	70.00	90.32	80.77	65.38	65.38
ORG	89.38	89.00	77.49	78.69	89.09	84.08	73.97	73.36	89.23	79.59	67.63	71.96
PERCENT	0.00	100.00	0.00	75.00	0.00	66.67	0.00	66.67	0.00	25.00	0.00	75.00
PERSON	40.22	93.75	40.63	38.54	39.77	93.18	38.64	38.64	39.33	98.85	42.53	41.38
PRODUCT	0.00	100.00	0.00	0.00	0.00	100.00	0.00	0.00	0.00	100.00	0.00	0.00
QUANTITY	0.00	66.67	0.00	0.00	0.00	100.00	33.33	33.33	0.00	66.67	0.00	33.33
TIME	63.64	100.00	80.00	70.00	60.00	77.78	77.78	66.67	63.64	63.64	63.64	45.45
WORK_OF_ART	0.00	100.00	0.00	0.00	0.00	100.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 5: Case insensitive accuracy scores for all the NE types on the three language pairs. We report the results for ASR, MT, Cascade (Casc.) and Direct (Dir.) systems.

posed of more than one word (e.g. for a person it is common to have both the name and surname), the total number of tokens per category is also given.

C Results for each NE Type

Table 5 shows the accuracy scores for all NE types on all language directions. First, we can notice that some types show a high variability (e.g. LANGUAGE, PRODUCT, QUANTITY), which is caused by the limited number of examples with that label. Otherwise, the performance of the two ST systems (cascade and direct) are similar on all categories, demonstrating that their different architecture does not bring to a different ability in handling specific type of entities.