

# Robust Neural Machine Translation for Clean and Noisy Speech Transcripts

Mattia Di Gangi<sup>\*†</sup> Robert Enyedi Alessandra Brusadin Marcello Federico

Amazon AI, East Palo Alto - USA

<sup>\*</sup>Fondazione Bruno Kessler, Trento - Italy

<sup>†</sup> University of Trento, Italy

## Abstract

Neural machine translation models have shown to achieve high quality when trained and fed with well structured and punctuated input texts. Unfortunately, the latter condition is not met in spoken language translation, where the input is generated by an automatic speech recognition (ASR) system. In this paper, we study how to adapt a strong NMT system to make it robust to typical ASR errors. As in our application scenarios transcripts might be post-edited by human experts, we propose adaptation strategies to train a single system that can translate either clean or noisy input with no supervision on the input type. Our experimental results on a public speech translation data set show that adapting a model on a significant amount of parallel data including ASR transcripts is beneficial with test data of the same type, but produces a small degradation when translating clean text. Adapting on both clean and noisy variants of the same data leads to the best results on both input types.

## 1. Introduction

The recent quality improvements [1, 2] of neural machine translation (NMT) [3, 4] opened the way to commercial applications that can provide high-quality translations. The assumption is that the sentences to translate will be similar to the training data, usually characterized by properly-formed text in the two translation languages. Poor-quality sentences are usually considered “noise” and removed from the training set to improve the final quality [5]. This practice is so common that a shared task has been devoted to it [6], which got attention from major industrial players in MT. Thus, the major weakness of NMT lies in coping with noisy input, which is an important feature of a real-world application such as speech translation.

The degradation of translation quality with noisy input has been widely reported in literature. Belinkov and Bisk [7] showed that the translation quality rapidly drops with both natural and synthetic noise. Karpukhin et al. [8] have observed a correlation between recognition and translation quality in the context of speech translation. In both cases the degradation is mainly due to word errors, and following works have shown that inserting synthetic noise in the train-

Most of the time, travellers worry about their luggage.  
Most of the time travellers worry about their luggage.

Table 1: Example of sentence in which the meaning is changed by a punctuation mark.

ing data increases the robustness to the same kind of noise [9, 10].

In practice, ASR transcripts are not only noisy in the choice of words, but also come without punctuation.<sup>1</sup> Thus, in order to feed MT systems with ASR output there are two main options: *i*) use a separate model that inserts punctuation (*pre-processing*) [11]; or *ii*) train the MT system on non-punctuated source data to resemble the test condition (*implicit learning*). The first option is exposed to error propagation, as punctuation inserted in the wrong position can alter completely the meaning of a sentence (see example in Table 1. The second option has shown to increase systems robustness [10] when the input is provided without punctuation. The two approaches have shown to be equivalent in handling text lacking punctuation [12], probably because they both rely only on plain monolingual text to recover the missing punctuation.

We consider here application scenarios (e.g. subtitling) where the same NMT system has to operate under both clean and noisy input conditions, namely post-edited or raw ASR outputs. We start from the hypothesis that word and punctuation errors compound and should be addressed separately for an increased robustness to errors. To verify our hypothesis we use a strong NMT model and fine-tune it on a recently released speech corpus<sup>2</sup> [13] of TED Talks.

Our findings are the following: *implicit learning* of punctuation helps to recover part of the errors in the source side, but training on ASR output is more beneficial. Training on clean and noisy data together leads to a system that can translate both clean and noisy data without supervision on the type of sentence and without degradation.

---

The first author carried out the work during an internship at Amazon.

<sup>1</sup>ASR services generally include punctuation as an option.

<sup>2</sup>Available at <https://ict.fbk.eu/must-c/>

	Train	Validation	Test
Words	5.7M	31.5K	54.5K
Segments	250K	1.3K	2.5K
Audio	457h	2.5h	4.2h

Table 2: The used English-Italian corpus of TED Talks.

	Clean	Noisy	Noisy-np
Gen	32.3 (30.7)	24.5 (24.0)	20.6 (22.4)
Ada	34.9 (32.9)	25.9 (25.6)	21.8 (24.6)

Table 3: BLEU scores of large-data generic and adapted NMT systems with clean and noisy input. Scores in parentheses do not consider punctuation in both hypothesis and reference.

## 2. Robust NMT

In this paper we are interested in building a single system to translate speech transcripts that can be either raw ASR outputs, or human post-editing of them. We define our problem in terms of domain adaptation, and use the method known as *fine-tuning* or *continued training* [14, 15, 16]. A parent model is trained on large data from several domains and used to initialize [17] models for spoken language translation (TED talks) on two input conditions: *clean*, and *noisy*. The *clean* domain is characterized by correct transcriptions of talks with proper punctuation, while the *noisy* domain can contain machine-generated errors in words and punctuation. In-domain data can be given with or without punctuation (allowing *implicit learning*). In a multi-domain setting, *i.e.* translating both *clean* and *noisy* data with a single system, models can suffer from catastrophic forgetting [18] by degrading performance on less-recently observed domains. In this work we avoid catastrophic forgetting by fine-tuning the model on both domains simultaneously.

## 3. Experimental Setting

We use as a parent model for fine-tuning a Transformer Big model [19] trained on public and proprietary data using label smoothed cross entropy [20], for about 16 million English-Italian sentence pairs. The model has layer size of 1024, hidden size of 4096 on feed forward layers, 16 heads in the multi-head attention, and 6 layers in both encoder and decoder. This model is then fine-tuned on the En→It portion of TED Talks in MuST-C. We keep the same training/validation/test set split as provided with the corpus (see Table 2). In all the experiments, we use Adam [21] with a fixed learning rate of  $2 \times 10^{-4}$ , dropout of 0.3, label smoothing with a smoothing factor of 0.1. Training is performed on 8 Nvidia V100 GPUs, with batches of 2000 tokens per GPU. Gradients are accumulated for 16 batches in each GPU [22]. All texts are tokenized and true-cased with scripts from the Moses toolkit [23], and then words are segmented with BPE [24] with 32K joint merge rules.

	Clean	Noisy
Clean	34.9	25.9
Clean-np	34.2 <sup>†</sup>	26.6 <sup>†</sup>
Clean + Clean-np	34.9	26.9 <sup>†</sup>
Noisy	34.0*	28.3*
Noisy-np	34.2*	28.4*
Noisy + Clean	35.1 <sup>†</sup>	28.1*
Noisy + Noisy-np	34.0*	28.2*
Noisy-np + Clean	35.0 <sup>†</sup>	28.2*
Noisy-np + Clean-np	34.5	27.9*
Noisy[-np] + Clean[-np]	34.9 <sup>†</sup>	27.7*

Table 4: Results of fine-tuning on different training conditions with clean and noisy input (large data). \* means statistical significant ( $p < 0.01$ ) wrt to Clean + Clean-np, <sup>†</sup> means statistical significant ( $p < 0.01$ ) wrt the first system of the block. (With randomization tests with 15K repetitions [27])

While our main goal is to verify our hypotheses on a large data condition, thus the need to include proprietary data, for the sake of reproducibility we also provide results with systems only trained on TED Talks (small data condition).

We transcribed automatically the entire TED Talks corpus with a general purpose ASR service.<sup>3</sup> The resulting word error rates on the test set is 11.0%. The used ASR service provides transcribed text with predicted punctuation. In the experiments which assume noisy input without punctuation, we simply remove the predicted punctuation.

All the results are evaluated in terms of BLEU score [25] using the multeval tool [26].

## 4. Experiments and Results

At first, we evaluate the degradation due to ASR noise for systems trained on clean data. In Table 3 we show the BLEU scores of our baseline system, respectively, trained on large out-of-domain data (Base) and fine-tuned on clean TED Talks data (In-domain), with three types of input: manual transcripts (Clean), ASR transcripts with predicted punctuation (Noisy), and ASR transcripts with no punctuation (Noisy-np). As these models will hardly generate punctuation when it does not appear in the source text, we also report (in parentheses) BLEU scores computed w/o punctuation in both hypothesis and reference.

Our results show that in-domain fine-tuning is beneficial in all scenarios, but more with clean input (+2.6 points) than with noisy input (+1.4 points). Translating noisy input results in a 26% relative drop in BLEU, which is apparently not due to punctuation errors (drop when evaluating w/o punctuation is 22%). Providing noisy input with no punctuation works even worse, probably due to the larger mismatch with the training/tuning conditions.

Next, we evaluate systems fine-tuned on data similar to the Noisy test condition. Table 4 lists the results of all fine-

<sup>3</sup>Amazon Transcribe: <https://aws.amazon.com/transcribe>

	Clean	np	Noisy	np
Clean	30.3	-	22.3	-
Clean-np	-	28.2	-	22.9
Clean + Clean-np	29.7	27.9	22.9	22.9
Noisy	25.8 <sup>†</sup>	-	23.9 <sup>†</sup>	-
Noisy-np	-	26.4	-	24.1 <sup>†</sup>
Noisy-np + Clean	30.1	27.9	24.0 <sup>†</sup>	24.2 <sup>†</sup>

Table 5: Results of fine-tuning on different training conditions with clean and noisy input (small data).<sup>†</sup>: statistical significant difference with Clean.

tuning experiments, when test input is either clean with punctuation or noisy with punctuation as in training. The first part of Table 4 shows that fine-tuning the Gen model with clean data and no punctuation (Clean-np) improves over the Ada model (Clean) when testing on Noisy-np (+0.7) but degrades when testing on Clean (-0.7). Fine-tuning the same model on Clean data with both punctuation condition (Clean + Clean-np) improvement by 1 point on Noisy input without any loss on Clean input. This result shows that is possible to make a model robust to noisy text, while preserving high quality on proper text.

The second part of Table 4 lists results of fine-tuning on noisy data, with and/or without punctuation. In all cases, BLEU scores on the noisy input improve from 1 to 2 point over the best systems tuned on Clean data only, reaching values above 28. However, scores on clean input degrade by 0.7 and 0.9 points (i.e. 34.2 and 34.0). If we adapt on both clean and noisy data, the score on the two input conditions reach a better balance. In particular, training on both Noisy-np and Clean data scores on the two input conditions 35.0 and 28.2, which results in the best overall working point on both conditions (together with Noisy+Clean). It is worth pointing out that this configuration obtains 33.2 points (not in the table) with the best possible noisy input, i.e. no errors and no punctuation, which is still 1.7 points below the score on Clean input.

Finally, if we expose the system to all types of data (Noisy-np +Noisy+ Clean + Clean-np) we do not see any improvement over our top results, which means that Clean-np data do not provide additional information to Noisy-np.

For the sake of replicability, we also trained our systems from scratch on the TED Talks data only. The results, listed in Table 5, show the same trend as the results discussed so far. We did not evaluate \*-np systems on input with punctuation as all the punctuation would represent out of vocabulary words. The main difference resides in the result on Clean input with the Noisy system (25.8 points), which is much worse than the result with the Noisy-np+Clean system (30.1 points), i.e. more than 4 points. This result suggests how training on noisy data can affect the model negatively if it is not balanced with clean data.

## 5. Manual Evaluation

We carried out a manual evaluation<sup>4</sup> to assess the quality of Noisy-np + Clean against Clean, the reference baseline, under the Noisy input condition. We ran a head-to-head evaluation on the first 10 sentences of each test talk, for a total of 260 sentences, by asking annotators to blindly rank the two system outputs (ties were also permitted). We collected three judgments for each output, from 11 annotators, for a total of 780 scores. Inter-annotator agreement measured with Fleiss’ kappa was 0.39.

Results reported in Table 7 confirm with high confidence the differences observed with BLEU: output of system Noisy-np+Clean is preferred 10% time more often than output of system Clean, while almost 68% of the time the two outputs are considered comparable.

From some manual inspection, we found that translations by the robust system that are unanimously ranked best show that error recovery most likely occurs on punctuation and non-content words like articles, prepositions and conjunctions (see Tables 6). In general, errors on content words that affect the meaning are not recovered.

## 6. Related works

A recent study [28] proposed to tackle ASR errors as a domain adaptation problem in the context of dialog systems. Domain adaptation for NMT has been widely studied in recent years [29]. In [30], fine-tuning was used to adapt NMT to multiple domains simultaneously, while in [31] adversarial adaptation is proposed for avoiding degradation in the original domain. Training on multiple domains simultaneously to prevent catastrophic forgetting is inspired by [32]. They proposed an incremental learning scheme that trains the network with the data from previous tasks when a new task is learned, which we adapt to our multi-domain scenario.

Punctuation insertion based on monolingual text has attracted research works for a long time [33, 34, 35, 36] and obtained recent improvements with deep networks [37, 38, 39, 40, 41]. However, this approach, which is meant to make the output of ASR more readable for humans, cannot solve ambiguity due to missing punctuation.

A more recent research line aims at using pauses and audio features to better predict punctuation [42, 43, 44, 45, 46] although it has been shown that the use of pauses is highly dependent on the speaker [47]. In [11], it is shown that *implicit learning* of punctuation in MT systems is at least as good as inserting punctuation either in the input or output text, but they only studied the effect on correct input that has been deprived of punctuation, not on noisy input. On the other side, [10] studied how to improve the robustness to misrecognized words, but did not study the effect of MT systems that are only robust to punctuation errors. We close the gap by studying the combined effect of misrecognized errors and missing punctuation, besides studying the robustness to

<sup>4</sup>We used crowd-sourcing via figure-eight.com.

Clean	when I ’m not fighting poverty , I ’m fighting fires as the assistant captain of a volunteer fire company .
Noisy	when I ’m not fighting poverty _ I ’m fighting fires _ <u>is</u> the assistant captain <u>with</u> volunteer fire company .
Base NMT	Quando non combatto la povertà _ Combatto gli incendi _ <u>È</u> l’assistente capitano di una compagnia di pompieri volontaria .
Robust NMT	Quando non combatto la povertà , combatto gli incendi come assistente capitano di una compagnia di pompieri volontaria .
Clean	that means we all share a common ancestor , an evolutionary grandmother , who lived around six million years ago.
Noisy	that means we all share a common ancestor _ <u>on</u> evolutionary grandmother _ who lived around six million years ago.
Base NMT	Ciò significa che tutti condividiamo un antenato comune _ <u>sulla</u> nonna evolutiva che ha vissuto circa sei milioni di anni fa.
Robust NMT	Ciò significa che tutti condividiamo un antenato comune <u>e</u> una nonna evolutiva che è vissuta circa sei milioni di anni fa.
Clean	and we in the West couldn’t understand how anybody would do this , how much this would restrict freedom of speech.
Noisy	<u>_</u> we in the West _ <u>I</u> couldn’t understand how anybody would do <u>_</u> how much this would restrict freedom of speech.
Base NMT	<u>_</u> Noi occidentali _ Non riesco a capire come chiunque avrebbe fatto <u>_</u> quanto questo avrebbe limitato la libertà di parola.
Robust NMT	<u>_</u> In Occidente non riesco a capire come chiunque avrebbe fatto <u>_</u> , quanto questo avrebbe limitato la libertà di parola.

Table 6: Examples of punctuation and substitution errors (“as → is”, “of a → with”, “an → on”) that are successfully recovered by the Robust NMT system. Notice that not all errors (underlined) are recovered. In the second example, Robust NMT introduces a spurious conjunction “e” (“and”) in place of the missing comma, while in the third example, Robust NMT is not able to recover the deleted words “and” and “this” at the begin and in the middle of the sentence, respectively.

	ASR w ties	ASR w/o ties
Clean	10.5	32.8
ASR-np + Clean	21.7 <sup>†</sup>	67.2 <sup>†</sup>

Table 7: Manual evaluation in the ASR input condition (large data). Percentage of wins with and without ties. <sup>†</sup> stands for statistically significant ( $p < 0.01$ ).

noisy data affects the translation quality on clean input.

## 7. Conclusion

We have studied the robustness to input errors of NMT systems for speech translation with a fine-tuning approach. We have observed that a system trained to learn implicitly the target punctuation can recover part of the quality degradation due to ASR errors up to 1 BLEU point. Fine-tuning on noisy input can instead improve by more than 2 BLEU points. A system tuned on ASR errors does not obtain a further improvement by more data for implicit punctuation learning. Finally, when fine-tuning on clean and noisy data, the system becomes robust to noisy input and keeps high performance on clean input.

## 8. References

- [1] O. Bojar, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, S. Huang, M. Huck, P. Koehn, Q. Liu, V. Logacheva *et al.*, “Findings of the 2017 conference on machine translation (wmt17),” in *Proceedings of the Second Conference on Machine Translation*, 2017, pp. 169–214.
- [2] O. Bojar, C. Federmann, M. Fishel, Y. Graham, B. Haddow, M. Huck, P. Koehn, and C. Monz, “Findings of the 2018 Conference on Machine Translation (WMT18),” in *Proceedings of WMT 2018*, 2018, pp. 272–307.
- [3] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to Sequence Learning with Neural Networks,” in *Proceedings of NIPS 2014*, 2014.
- [4] D. Bahdanau, K. Cho, and Y. Bengio, “Neural Machine Translation by Jointly Learning to Align and Translate,” in *Proceedings of ICLR 2015*, 2015.
- [5] M. Junczys-Dowmunt, “Microsoft’s submission to the wmt2018 news translation task: How i learned to stop worrying and love the data,” in *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, 2018, pp. 425–430.
- [6] P. Koehn, H. Khayrallah, K. Heafield, and M. L. Forcada, “Findings of the wmt 2018 shared task on parallel corpus filtering,” in *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, 2018, pp. 726–739.
- [7] Y. Belinkov and Y. Bisk, “Synthetic and natural noise both break neural machine translation,” *Proceedings of ICLR*, 2018.
- [8] N. Ruiz, M. A. Di Gangi, N. Bertoldi, and M. Federico, “Assessing the tolerance of neural machine translation

- systems against speech recognition errors,” *Proc. Inter-speech 2017*, pp. 2635–2639, 2017.
- [9] V. Karpukhin, O. Levy, J. Eisenstein, and M. Ghazvininejad, “Training on synthetic noise improves robustness to natural noise in machine translation,” *arXiv preprint arXiv:1902.01509*, 2019.
- [10] M. Sperber, J. Niehues, and A. Waibel, “Toward robust neural machine translation for noisy input sequences,” in *International Workshop on Spoken Language Translation (IWSLT)*, 2017.
- [11] S. Peitz, M. Freitag, A. Mauser, and H. Ney, “Modeling punctuation prediction as machine translation,” in *International Workshop on Spoken Language Translation (IWSLT) 2011*, 2011.
- [12] V. Vandeghinste, C.-K. Leuven, J. Pelemans, L. Verwimp, and P. Wambacq, “A comparison of different punctuation prediction approaches in a translation context,” in *21st Annual Conference of the European Association for Machine Translation*, 2018, p. 269.
- [13] M. A. Di Gangi, R. Cattoni, L. Bentivogli, M. Negri, and M. Turchi, “MuST-C: a Multilingual Speech Translation Corpus,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, Minneapolis, MN, USA, June 2019.
- [14] M.-T. Luong and C. D. Manning, “Stanford neural machine translation systems for spoken language domains,” in *Proceedings of the International Workshop on Spoken Language Translation*, 2015.
- [15] M. A. Farajian, R. Chatterjee, C. Conforti, S. Jalalvand, V. Balaraman, M. A. Di Gangi, D. Ataman, M. Turchi, M. Negri, and M. Federico, “FBK’s neural machine translation systems for IWSLT 2016,” in *Proceedings of the ninth International Workshop on Spoken Language Translation, USA*, 2016.
- [16] C. Chu, R. Dabre, and S. Kurohashi, “An empirical comparison of domain adaptation methods for neural machine translation,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2017, pp. 385–391.
- [17] B. Thompson, H. Khayrallah, A. Anastasopoulos, A. D. McCarthy, K. Duh, R. Marvin, P. McNamee, J. Gwinup, T. Anderson, and P. Koehn, “Freezing subnetworks to analyze domain adaptation in neural machine translation,” in *Proceedings of the Third Conference on Machine Translation: Research Papers*, 2018, pp. 124–132.
- [18] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, “Overcoming catastrophic forgetting in neural networks,” *Proceedings of the national academy of sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is All You Need,” in *Proceedings of NIPS 2017*, 2017.
- [20] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proc. of CVPR*, 2016.
- [21] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. of ICLR*, 2015.
- [22] M. Ott, S. Edunov, D. Grangier, and M. Auli, “Scaling neural machine translation,” in *Proceedings of the Third Conference on Machine Translation: Research Papers*, 2018, pp. 1–9.
- [23] P. Koehn, H. Hoang, A. Birch, C. Callison-Burch *et al.*, “Moses: Open source toolkit for statistical machine translation,” in *Proc. of ACL*, 2007.
- [24] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” *arXiv preprint arXiv:1508.07909*, 2015.
- [25] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: a Method for Automatic Evaluation of Machine Translation,” in *Proceedings of ACL 2002*, 2002.
- [26] J. H. Clark, C. Dyer, A. Lavie, and N. A. Smith, “Better hypothesis testing for statistical machine translation: Controlling for optimizer instability,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*. Association for Computational Linguistics, 2011, pp. 176–181.
- [27] S. Riezler and J. T. Maxwell, “On some pitfalls in automatic evaluation and significance testing for MT,” in *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*. Ann Arbor, Michigan: Association for Computational Linguistics, Jun. 2005, pp. 57–64. [Online]. Available: <https://www.aclweb.org/anthology/W05-0908>
- [28] P.-J. Chen, I.-H. Hsu, Y.-Y. Huang, and H.-Y. Lee, “Mitigating the impact of speech recognition errors on chatbot using sequence-to-sequence model,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 497–503.

- [29] C. Chu and R. Wang, “A survey of domain adaptation for neural machine translation,” in *Proceedings of the 27th International Conference on Computational Linguistics*, 2018, pp. 1304–1319.
- [30] H. Khayrallah, B. Thompson, K. Duh, and P. Koehn, “Regularized training objective for continued training for domain adaptation in neural machine translation,” in *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 36–44.
- [31] D. Britz, Q. Le, and R. Pryzant, “Effective domain mixing for neural machine translation,” in *Proceedings of the Second Conference on Machine Translation*, 2017, pp. 118–126.
- [32] S. Stojanov, S. Mishra, A. Thai, N. Dhanda, A. Humayun, L. Smith, C. Yu, and J. M. Rehg, “Incremental object learning from contiguous views,” in *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR) 2019*, 2019.
- [33] J. Huang and G. Zweig, “Maximum entropy model for punctuation annotation from speech,” in *Seventh International Conference on Spoken Language Processing*, 2002.
- [34] E. Matusov, A. Mauser, and H. Ney, “Automatic sentence segmentation and punctuation prediction for spoken language translation,” in *International Workshop on Spoken Language Translation (IWSLT) 2006*, 2006.
- [35] W. Lu and H. T. Ng, “Better punctuation prediction with dynamic conditional random fields,” in *Proceedings of the 2010 conference on empirical methods in natural language processing*, 2010, pp. 177–186.
- [36] N. Ueffing, M. Bisani, and P. Vozila, “Improved models for automatic punctuation prediction for spoken and written text,” in *Interspeech*, 2013, pp. 3097–3101.
- [37] E. Cho, J. Niehues, and A. Waibel, “Segmentation and punctuation prediction in speech language translation using a monolingual translation system,” in *International Workshop on Spoken Language Translation (IWSLT) 2012*, 2012.
- [38] E. Cho, J. Niehues, K. Kilgour, and A. Waibel, “Punctuation insertion for real-time spoken language translation,” in *Proceedings of the Eleventh International Workshop on Spoken Language Translation*, 2015.
- [39] O. Tilk and T. Alumäe, “LSTM for punctuation restoration in speech transcripts,” in *Sixteenth annual conference of the international speech communication association*, 2015.
- [40] —, “Bidirectional recurrent neural network with attention mechanism for punctuation restoration,” in *Interspeech*, 2016, pp. 3047–3051.
- [41] W. Salloum, G. Finley, E. Edwards, M. Miller, and D. Suendermann-Oeft, “Deep learning for punctuation restoration in medical reports,” in *BioNLP 2017*, 2017, pp. 159–164.
- [42] H. Christensen, Y. Gotoh, and S. Renals, “Punctuation annotation using statistical prosody models,” in *ISCA tutorial and research workshop (ITRW) on prosody in speech recognition and understanding*, 2001.
- [43] O. Klejch, P. Bell, and S. Renals, “Sequence-to-sequence models for punctuated transcription combining lexical and acoustic features,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5700–5704.
- [44] P. Żelasko, P. Szymański, J. Mizgajski, A. Szymczak, Y. Carmiel, and N. Dehak, “Punctuation prediction model for conversational speech,” *Proc. Interspeech 2018*, pp. 2633–2637, 2018.
- [45] A. Nanchen and P. N. Garner, “Empirical Evaluation and Combination of Punctuation Prediction Models Applied to Broadcast News,” Tech. Rep.
- [46] J. Yi and J. Tao, “Self-attention based model for punctuation prediction using word and speech embeddings,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7270–7274.
- [47] M. Igras-Cybulska, B. Ziółko, P. Żelasko, and M. Witkowski, “Structure of pauses in speech in the context of speaker verification and classification of speech type,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2016, no. 1, p. 18, 2016.