



A High-Quality and Large-Scale Dataset for English-Vietnamese Speech Translation

Linh The Nguyen*, Nguyen Luong Tran*, Long Doan*, Manh Luong, Dat Quoc Nguyen

VinAI Research, Hanoi, Vietnam

{v.linhnt140, v.nguyentl12, v.longdct, v.manhlt3, v.datnq9}@vinai.io

Abstract

In this paper, we introduce a high-quality and large-scale benchmark dataset for English-Vietnamese speech translation with 508 audio hours, consisting of 331K triplets of (sentence-lengthened audio, English source transcript sentence, Vietnamese target subtitle sentence). We also conduct empirical experiments using strong baselines and find that the traditional “Cascaded” approach still outperforms the modern “End-to-End” approach. To the best of our knowledge, this is the first large-scale English-Vietnamese speech translation study. We hope both our publicly available dataset and study can serve as a starting point for future research and applications on English-Vietnamese speech translation.

Index Terms: Benchmark dataset, English-Vietnamese, Speech translation, Automatic speech recognition, Machine translation, Cascaded, End-to-End.

1. Introduction

Speech translation—the task of translating speech in one language typically to text in another—has attracted interest for many years [1, 2, 3]. However, the development of speech translation systems has been largely limited to the high-resource language pairs because most public available datasets for speech translation are exclusively for the high-resource languages [4, 5, 6, 7, 8, 9, 10, 11, 12]. For example, bilingual speech translation datasets [4, 5, 6, 7, 8] are all large-scale ones with about 200+ audio hours, which offer translations either from English to a target language (e.g. French, German and Japanese) or from a source language (e.g. Chinese, Japanese and Spanish) to English. The one-to-many dataset MuST-C [9] is created to facilitate the training of speech translation systems from the source language English into 14 target languages. On the contrary, the many-to-one dataset CoVoST [10] provides speech translations from 11 languages into English. EuroParl-ST [11] is a many-to-many speech translation dataset that provides translations between 6 European languages, with a total of 30 translation directions, using the debates held in the European Parliament from 2008 to 2012. The Multilingual TEDx corpus [12] is created from TEDx talks in 8 source languages with their manual transcriptions and translations in 6 target languages.

From a societal, linguistic, machine learning, cultural and cognitive perspective [13], it is also worth investigating the speech translation task for low-resource languages, e.g. Vietnamese with about 100M speakers [14]. Vietnam is now an attractive destination for trade, investment and tourism [15]. Thus the demand for high-quality speech translation from the global language English to Vietnamese has rapidly increased. To our best knowledge, there is no existing research work that focuses

solely on speech translation to Vietnamese. The only available resource for speech translation to Vietnamese is the 441-hour English-Vietnamese speech translation data from MuST-C [9]. However, MuST-C is a TED-talk-based multilingual dataset, thus its authors could not pay attention to any specific language pairs. Thus it might contain timestamp alignment errors between English audio-transcript pairs and misalignment between English-Vietnamese source-target sentence pairs.

We use a total of 50 working hours to manually inspect both the validation set (including 1350 triplets with 2.5 audio hours in total) and the test set (including 2361 triplets with 4.1 audio hours) from the MuST-C English-Vietnamese data. Here, each triplet is checked by two out of our first three authors independently. After cross-checking, we find that: **5.63%** of the validation set and **4.10%** of the test set have an incorrect audio start or end timestamp of an English source sentence; **16.15%** of the validation set and **9.36%** of the test set have misaligned English-Vietnamese sentence pairs (i.e. completely different sentence meaning or partly preserving the meaning). Note that performing a similar manual inspection on the MuST-C English-Vietnamese training set would take $(441 - 2.5 - 4.1) / (2.5 + 4.1) \times 50 = 3290$ working hours, beyond our current human resource. As training/validation/test split is random, the substantial rates of incorrect timestamps and misalignment on the validation and test sets imply that the MuST-C English-Vietnamese training set might also not reach a high-quality level.

Our work is to help tackle that issue of low-quality data for English-Vietnamese speech translation. Our contributions are summarized as follows:

- As the first contribution, we present a high-quality and large-scale English-Vietnamese speech translation dataset containing 508 audio hours. We hope that our dataset construction process can be further adapted to create more speech translation data for other low-resource languages.
- As the second contribution, on our dataset, we empirically investigate strong baselines to compare traditional “Cascaded” and modern “End-to-End” approaches. We find that the “Cascaded” approach still does better than the “End-to-End” approach. To the best of our knowledge, this is the first large-scale empirical study for English-Vietnamese speech translation.
- As our final contribution, we publicly release our dataset for non-commercial use under the CC BY-NC-ND 4.0 license at: <https://github.com/VinAIRResearch/PhoST>. We hope both our dataset and empirical study can serve as a starting point for future English-Vietnamese speech translation research and applications.

*The first three authors contributed equally to this work.

2. Our dataset

Our dataset construction process consists of five phases, as detailed in the following subsections.

2.1. Collecting audio files and transcripts

We employ the TED2020-v1 corpus [16] of 3125 parallel English-Vietnamese transcript-subtitle pairs of TED talks. Here, the English transcript document (i.e. source document) was human-translated to the Vietnamese subtitle document (i.e. target document). For each talk, we download the corresponding video with the highest quality version and extract the audio from the downloaded video. Here, we extract the audio from each downloaded video by using the FFmpeg library.¹ All extracted audios are formatted as 16kHz WAV files. There are also 5 talks without publicly available videos. We thus obtain $3125 - 5 = 3120$ TED-talk-based triplets of (audio file, English transcript document, Vietnamese subtitle document).

2.2. Pre-processing and sentence segmentation

We manually check the obtained triplets and find that there are 23 non-English audio files of TED talks, i.e. 23 English source documents are subtitles of these TED talks, not transcripts. There are also 10 audio files using almost all of their time to display songs. We thus remove those $23 + 10 = 33$ triplets, resulting in $3120 - 33 = 3087$ remaining triplets. To extract sentences for parallel audio-sentence alignments in next phases, we perform sentence segmentation by using Stanford CoreNLP [17] for the English documents and VnCoreNLP [18, 19] for the Vietnamese documents.

Note that TED-talk transcripts and subtitles contain non-speech artifacts of audience-related information marked with parentheses (e.g. “(applause)”, “(laugh)” and the like) and speaker identity. We thus remove all of the audience-related information. For removing the speaker identity, we use an effective heuristic of removing words/phrases that start at the beginning of a sentence, consist of at most 10 characters or 3 words, and are followed by the punctuation mark colon (here, the colon is also removed together with the speaker identity).

2.3. Extracting the audio start and end timestamps for each English sentence

From a pair of TED-talk audio and its English transcript document, to extract the start and the end timestamps from the audio for each sentence in the document, following the MuST-C process [9], we employ the Gentle forced aligner [20] that is based on the Kaldi ASR toolkit [21]. The Gentle forced aligner takes the audio and transcript document pair as input and outputs a timestamp for each word token in the document. We use the timestamp of the first word in each sentence as the start timestamp of the sentence. We determine the end timestamp of the sentence as $\min(x + 0.5 \text{ second}, y - 0.01 \text{ second})$ in which x is the timestamp of the last word in the sentence while y is the timestamp of the first word in the next sentence.

Note that there are 10K English sentences where the Gentle forced aligner cannot detect a timestamp of the first or the last word in a sentence due to the corresponding background noises (in most cases), multiple-speaker overlaps or non-English sounds in the audio file. We thus manually correct the start and the end timestamps of these sentences. Here, we develop a small tool based on the PyDub library, which is to play a

¹<http://ffmpeg.org/>

TED talk audio and to display the audio timestamps, and also to play an audio span given the audio and a start and end timestamp pair.² Given a sentence that we have to determine the missing timestamps of its first and/or last words, we first locate an audio span in the corresponding TED talk audio, which covers the sentence’s part with timestamp-detected words. Through listening to the audio (via using the tool) and looking at the transcript sentence, we then expand the audio span to exactly match the whole given sentence, thus extracting the timestamps of its first and/or last words.

We divide those 10K sentences into three equal parts. Each of the first three authors manually corrects one of these parts. After that, the fourth author verifies each sentence and makes further revisions if needed. This manual correction process takes a total of 210 working hours from four authors.

2.4. Aligning parallel English-Vietnamese sentence pairs

To align parallel sentences within a parallel English-Vietnamese document pair, following [22], we first use Google Translate to translate English source sentences into Vietnamese.³ Then, to produce parallel English-Vietnamese sentence pairs, we use three alignment toolkits of Hunalign [24], Gargantua [25] and Bleualign [26] to perform an intermediate alignment between the Vietnamese Google-translated versions of the English source sentences and the Vietnamese target sentences. Note that Bleualign only performs alignment between the target sentences and the variants translated into the target language of the source sentences. And in a preliminary experiment, we find that for both Hunalign and Gargantua, performing sentence alignment between the Google-translated variants of the English source sentences and the Vietnamese target sentences produces better quality than performing direct alignment between the English sentences and the Vietnamese sentences. Only sentence pairs, aligned by at least two out of three toolkits, are selected to be included in the output of this sentence alignment process (here, 99.2% of both the English and the Vietnamese sentences are included in the output).

2.5. Post-processing

From a dataset of 3087 triplets of (TED-talk audio, English transcript document, Vietnamese subtitle document) as shown in Section 2.2, we merge the outputs of the previous two phases to produce 331284 triplets of (sentence-lengthed audio, English source sentence of transcript, Vietnamese target sentence of subtitle). We normalize punctuations in sentence pairs. We split our dataset into training/validation/test sets with a 98.8/0.6/0.6 ratio on the TED-talk level, resulting in 327370 training, 1936 validation and 1978 test triplets at the sentence level.

We manually inspect each triplet in the validation and test sets to qualify our dataset, for checking whether there is: (i) a timestamp misalignment between English audio-transcript pairs,⁴ (ii) a misalignment between English-Vietnamese source-target sentence pairs, and (iii) a low-quality translation from the English source sentence to the Vietnamese target sentence.

²<https://github.com/jiaaro/pydub>

³We use the English-to-Vietnamese translation direction because Google-translating from English to Vietnamese produces better translation than from Vietnamese to English. This is confirmed via BLEU scores in the first two rows in Tables 3 and 4 from [23] or human-evaluation results for Google Translate in Table 2 from [22].

⁴For checking the timestamp misalignment, we reuse our PyDub-based tool developed for correcting the timestamps of the first and/or last words of sentences in the third phase (Section 2.3).

Table 1: *Our dataset statistics.* “#triplets”, “#hours”, “#en/s” and “#vi/s” denote the number of triplets, the number of audio hours, the average number of word tokens per English sentence and the average number of syllable tokens per Vietnamese sentence, respectively.

Split	#triplets	#hours	#en/s	#vi/s
Training	327370	501.59	16.55	20.94
Validation	1933	3.13	17.24	22.22
Test	1976	3.77	19.23	25.65

Each triplet is inspected by two out of the first three authors independently: one inspector checks about $(1936 + 1978) \times 2/3 \approx 2609$ triplets in 18 working hours on average.

After cross-checking, we do not find on both validation and test sets any error w.r.t. the start and the end timestamps of each English source sentence. We also find that 3 triplets (**0.15%**) on the validation set and 2 triplets (**0.10%**) on the test set contain misaligned source-target sentence pairs. In addition, all the remaining source-target sentence pairs on both validation and test sets are at a high-quality translation level. Without any start or end timestamp error and with very small proportions of sentence pair misalignment on the validation and test sets, we believe that our training set reaches a high-quality standard. Therefore, our dataset obtains a substantially better quality than the TED-talk based MuST-C English-Vietnamese speech translation dataset discussed in Section 1.

We remove those $3 + 2 = 5$ triplets, resulting in 1933 validation and 1976 test triplets for final use. Table 1 presents the statistics of our dataset.

3. Speech translation approaches

On our dataset, we compare two speech translation approaches: *Cascaded* vs. *End-to-End*.

3.1. Cascaded

The “Cascaded” approach combines two main components of English automatic speech recognition (ASR) and English-to-Vietnamese machine translation (MT). For ASR, we train the Fairseq’s S2T Transformer model [27, 28] on our English audio-transcript training data.⁵ For MT, we fine-tune the pre-trained sequence-to-sequence model mBART [32] that obtains the state-of-the-art performance for English-Vietnamese machine translation [22, 33].

Note that we also perform a data augmentation to extend our MT training data. In particular, we employ the NeMo toolkit [34] to perform a normalization process of 2 steps, including: (i) the first step of inverse text normalization is to convert our trained S2T Transformer’s automatic ASR output into its written form (e.g. “in nineteen seventy” is converted into “in 1975”), and (ii) the second step is to recover capitalization and punctuation marks (here, the output of the first step is the input for the second step). Given the input of 327370 sentence-lengthed training audios to the trained S2T Transformer, each of steps (i) and (ii) then produces 327370 additional English sentences, thus resulting in an extended training set of $327370 * 3 = 982110$ parallel English-Vietnamese sentence pairs for fine-

⁵In preliminary experiments, we find that training S2T Transformer produces a slightly better ASR result (i.e. lower word error rate) than both training Conformer [29, 30] and fine-tuning Wav2Vec2 [31].

tuning mBART. We also fine-tune mBART on a combination of the 3M-sentence-pair dataset PhoMT [22] and our extended set (here, we remove PhoMT’s pairs that appear in our validation and test sets).

3.2. End-to-End

For an “End-to-End” approach that directly translates English speech into Vietnamese text, we study two baselines, including the S2T Transformer model and the UPC’s speech translation system Adaptor [35] that is the only top performance system at IWSLT 2021 [3] with publicly available implementation at the time of our empirical investigation. The UPC’s Adaptor system employs a Length Adaptor module [36] to adjust representations from the pre-trained Wav2Vec2-based speech encoder [31] to the mBART-based language-specific decoder [32]. We experiment with the best setting of “Layer-Norm Attention – Encoder-Decoder with 2-step Adaptor” of the UPC’s Adaptor.⁶

3.3. Implementation details

In the “Cascaded” approach, for validation or test, we also apply the two-step normalization process on the output of the ASR component before feeding into the MT component. We initialize the S2T Transformer model used for “End-to-End” speech translation by the S2T Transformer trained for the ASR component in the “Cascaded” approach.

Except for the UPC’s Adaptor system that requires the input of raw WAV audio file, we augment the audio data with SpecAugment [37] and extract the log Mel-filterbank with 80 dimensions as input features. We employ SentencePiece [38] to learn a vocabulary of 5K subword types for English and a vocabulary of 8K types for Vietnamese, which are used in the S2T Transformer models. We employ the S2T Transformer and mBART implementations from the fairseq library [39]. Here, we train the S2T Transformer models with a configuration of `s2t-transformer-s`.

We optimize all systems on 4 V100 GPUs (32GB each), using Adam [40] with different initial learning rates. In particular, we perform a grid search to select the initial Adam learning rate. The optimal initial learning rates for the “Cascaded” ASR component, the “Cascaded” MT component, the “End-to-End” S2T Transformer and the UPC’s Adaptor system are $2e-3$, $5e-5$, $2e-3$ and $4e-4$ respectively.

For evaluation, we use beam search with beam size of 5 as decoding algorithm. Following the IWSLT 2021 evaluation campaign [3], we use the BLEU score as our main evaluation metric in our speech translation experiments. We use SacreBLEU [41] to report the detokenized and case-sensitive BLEU score. In our ASR experiments, we use fairseq to calculate WER with lowercased, punctuation removal and using `l3a` tokenizer. For both “Cascaded” and “End-to-End”, we choose the model checkpoint that produces the highest BLEU score on the validation set to apply to the test set.

4. Experimental results

4.1. Main Results

Table 2 presents final BLEU score results obtained by each speech translation approach on our test set. For “Cascaded” results, with the same automatic ASR output as the model input, it is not surprising that fine-tuning mBART on a combination

⁶<https://github.com/mt-upc/iwslt-2021>

Table 2: BLEU scores of each approach on our test set. “Casc.” and “E2E” denote the “Cascaded” and “End-to-End” approaches, respectively. (I): mBART fine-tuned on our extended training set. (II): mBART fine-tuned on a combination of the 3M-pair dataset PhoMT and our extended training set. In the “Cascaded” approach, the word error rate (WER) computed for the ASR component is 7.06, while BLEU scores of (I) and (II) computed for the text-to-text MT component with “gold” English source transcript sentences are 36.48 and 37.41, respectively. Each score difference between two out of four approaches is statistically significant with p -value < 0.01 based on bootstrap resampling (except the difference between 33.65 and 33.30 is with p -value < 0.05).

Model		BLEU \uparrow
Casc.	(I) mBART w/ our extended dataset	33.65
	(II) mBART w/ PhoMT combination	34.31
E2E	S2T Transformer	29.98
	UPC’s Adaptor	33.30

of PhoMT and our extended training set helps produce a better BLEU score than fine-tuning mBART on only our extended training set (34.31 vs. 33.65), thus showing the effectiveness of a larger training size. For “End-to-End” results, it is also not surprising that with the ability to leverage strong pre-trained models Wav2Vec2 for ASR and mBART for machine translation, the UPC’s Adaptor system obtains a 3+ absolute higher BLEU score than the S2T Transformer model trained solely without pre-training (33.30 vs. 29.98). In a comparison between the “Cascaded” and “End-to-End” results, we find that the “Cascaded” approach does better than the “End-to-End” approach.

Note that our “Cascaded” results in Table 2 are reported w.r.t. data augmentation to extend our original training set via the 2-step normalization process. We conduct an ablation study on the validation set to investigate the usefulness of this data augmentation strategy. We find that not using data augmentation leads to a decrease in the BLEU performance score of the “Cascaded” approach. In particular, on the validation set, the “Cascaded” approach with mBART fine-tuned on our extended training set (i.e. with data augmentation) obtains a BLEU score of 30.88, while the one with mBART fine-tuned on our original training set (i.e. without data augmentation) obtains a lower BLEU score at 30.18 (i.e. 30.88 \rightarrow 30.18).

4.2. Comparison with the MuST-C English-Vietnamese

Recall that sections 1 and 2.5 detail the rates of incorrect timestamps and misalignment on the validation and test sets in the MuST-C English-Vietnamese data and our dataset, respectively; showing that our dataset obtains a substantially better quality than the MuST-C English-Vietnamese data.

To further investigate the impact of our curation effort, we also conduct experiments comparing model performances on the MuST-C English-Vietnamese (En-Vi) training set and our training set. For a fair comparison, as not all of our data are overlapping with MuST-C, we filter from the MuST-C En-Vi training set all TED talks that appear in our validation and test sets. We then sample a subset of TED talks from our training set so that our sampled subset has the same total number of speech hours as the filtered MuST-C En-Vi training set. We first train two S2T Transformer models for ASR and then two end-to-end Adaptor models for speech translation, using our sampled

Table 3: ASR WER scores and end-to-end speech translation BLEU scores on our test set. Score differences are statistically significant with p -value < 0.01 .

Training data	WER \downarrow	BLEU \uparrow
MuST-C En-Vi training set	9.09	31.66
Our sampled training subset	7.44	32.37

subset and the filtered MuST-C En-Vi training set. We tune hyper-parameters using our validation set in the same manner as presented in Section 3.3. We evaluate these models using our test set. Experimental results obtained for these experiments are presented in Table 3.

On the ASR task, the S2T Transformer trained using our sampled subset obtains a lower WER than the S2T Transformer trained using the MuST-C data (7.44 vs. 9.09), showing our better audio-transcript alignment. Similarly, on the speech translation task, the end-to-end Adaptor model trained using our sampled subset does better than the one trained using the MuST-C data (32.37 vs. 31.66), demonstrating the effectiveness of our curation effort w.r.t. the end-to-end setting. Note that when trained using our full training set, as shown in Table 2, S2T Transformer obtains a lower WER at 7.06 for ASR, while Adaptor obtains a higher BLEU score at 33.30 for speech translation, thus reconfirming the positive impact of a larger training size.

We provide a qualitative example to demonstrate the qualitative differences between the end-to-end Adaptor models trained on MuST-C and on the subset of our dataset. With input audio of the English sentence “But on a long wavelength sea, you’d be rolling along, relaxed, low energy.”, we have:

- The end-to-end Adaptor model trained on MuST-C provides an output (1) of “Nhưng trên sóng biển dài, bạn sẽ lăn dọc theo, thư giãn, năng lượng thấp.” (here, using Google Translate to translate output (1) to English obtains: “But on the long ocean waves, you’ll roll along, relaxed, low energy.”).
- The end-to-end Adaptor model trained on the subset of our dataset provides an output (2) of “Nhưng trên một vùng biển có bước sóng dài, bạn sẽ lăn dọc, thư giãn, ít tốn năng lượng hơn.” (here, using Google Translate to translate output (2) to English obtains: “But on a sea with long wavelengths, you will roll along, relax, with less energy.”).

Output (2) is more natural, smooth and better preserving the source sentence’s meaning than output (1), thus also implying a higher quality accounted for our dataset.

5. Conclusion

In this paper, we have presented a high-quality and large-scale dataset with 508 audio hours for English-Vietnamese speech translation. On our dataset, we empirically conduct experiments using strong baselines to compare the “Cascaded” and “End-to-End” approaches. Experimental results show that the “Cascaded” approach does better than the “End-to-End” approach. We hope that the public release of our dataset can serve as the starting point for further English-Vietnamese speech translation research and applications.

6. References

- [1] M. Sperber and M. Paulik, “Speech Translation and the End-to-End Promise: Taking Stock of Where We Are,” in *ACL*, 2020, pp. 7409–7421.
- [2] E. Ansari, A. Axelrod *et al.*, “FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN,” in *IWSLT*, 2020, pp. 1–34.
- [3] A. Anastasopoulos, O. Bojar *et al.*, “FINDINGS OF THE IWSLT 2021 EVALUATION CAMPAIGN,” in *IWSLT*, aug 2021, pp. 1–29.
- [4] H. Tohyama, S. Matsubara, N. Kawaguchi, and Y. Inagaki, “Construction and utilization of bilingual speech corpus for simultaneous machine interpretation research,” in *INTERSPEECH*, 2005.
- [5] M. Post, G. S. Kumar, A. Lopez, D. G. Karakos, C. Callison-Burch, and S. Khudanpur, “Improved speech-to-text translation with the Fisher and Callhome Spanish-English speech translation corpus,” in *IWSLT*, 2013.
- [6] N. Jan, R. Cattoni, S. Sebastian, M. Cettolo, M. Turchi, and M. Federico, “The IWSLT 2018 Evaluation Campaign,” in *IWSLT*, 2018, pp. 2–6.
- [7] A. C. Kocabiyikoglu, L. Besacier, and O. Kraif, “Augmenting Librispeech with French Translations: A Multimodal Corpus for Direct Speech Translation Evaluation,” in *LREC*, 2018.
- [8] R. Zhang, X. Wang, C. Zhang, Z. He, H. Wu, Z. Li, H. Wang, Y. Chen, and Q. Li, “BSTC: A Large-Scale Chinese-English Speech Translation Dataset,” in *AutoSimTrans*, 2021, pp. 28–35.
- [9] R. Cattoni, M. A. Di Gangi, L. Bentivogli, M. Negri, and M. Turchi, “MuST-C: A multilingual corpus for end-to-end speech translation,” *Computer Speech & Language*, vol. 66, p. 101155, 2021.
- [10] C. Wang, J. M. Pino, A. Wu, and J. Gu, “CoVoST: A Diverse Multilingual Speech-To-Text Translation Corpus,” in *LREC*, 2020.
- [11] J. Iranzo-Sánchez, J. A. Silvestre-Cerdà, J. Jorge, N. Roselló, A. Giménez, A. Sanchís, J. C. Saiz, and A. Juan-Ciscar, “Europarl-ST: A Multilingual Corpus for Speech Translation of Parliamentary Debates,” in *ICASSP*, 2020, pp. 8229–8233.
- [12] E. Salesky, M. Wiesner, J. Bremerman, R. Cattoni, M. Negri, M. Turchi, D. W. Oard, and M. Post, “The Multilingual TEDx Corpus for Speech Recognition and Translation,” in *INTERSPEECH*, 2021, pp. 3655–3659.
- [13] S. Ruder, “Why You Should Do NLP Beyond English,” <https://ruder.io/nlp-beyond-english/>, 2020. [Online]. Available: <https://ruder.io/nlp-beyond-english/>
- [14] D. M. Eberhard, G. F. Simons, and C. D. Fennig, *Ethnologue: Languages of the World, 22nd edition*. United States: SIL International, 2019.
- [15] A. Baum, “Vietnam’s Development Success Story and the Unfinished SDG Agenda,” *IMF Working Papers*, vol. 20, no. 31, pp. 1–31, 2020.
- [16] N. Reimers and I. Gurevych, “Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation,” in *EMNLP*, 2020, pp. 4512–4525.
- [17] C. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. Bethard, and D. McClosky, “The Stanford CoreNLP Natural Language Processing Toolkit,” in *ACL: Demonstrations*, 2014, pp. 55–60.
- [18] D. Q. Nguyen, D. Q. Nguyen, T. Vu, M. Dras, and M. Johnson, “A Fast and Accurate Vietnamese Word Segmenter,” in *LREC*, 2018, pp. 2582–2587.
- [19] T. Vu, D. Q. Nguyen, D. Q. Nguyen, M. Dras, and M. Johnson, “VnCoreNLP: A Vietnamese Natural Language Processing Toolkit,” in *NAACL: Demonstrations*, 2018, pp. 56–60.
- [20] R. M. Ochshorn and M. Hawkins, “Gentle forced aligner,” <https://github.com/lowerquality/gentle>.
- [21] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, “The Kaldi Speech Recognition Toolkit,” in *ASRU*, 2011.
- [22] L. Doan, L. T. Nguyen, N. L. Tran, T. Hoang, and D. Q. Nguyen, “PhoMT: A High-Quality and Large-Scale Benchmark Dataset for Vietnamese-English Machine Translation,” in *EMNLP*, 2021, pp. 4495–4503.
- [23] T. Kudo, “Subword Regularization: Improving Neural Network Translation Models with Multiple Subword Candidates,” in *ACL*, 2018, pp. 66–75.
- [24] D. Varga, P. Halácsy, A. Kornai, V. Nagy, L. Németh, and V. Trón, “Parallel corpora for medium density languages,” *Current Issues in Linguistic Theory*, vol. 292, pp. 247–258, 2007.
- [25] F. Braune and A. Fraser, “Improved Unsupervised Sentence Alignment for Symmetrical and Asymmetrical Parallel Corpora,” in *COLING: Posters*, 2010, pp. 81–89.
- [26] R. Sennrich and M. Volk, “Iterative, MT-based Sentence Alignment of Parallel Texts,” in *NODALIDA*, 2011, pp. 175–182.
- [27] C. Wang, Y. Tang, X. Ma, A. Wu, D. Okhonko, and J. Pino, “Fairseq S2T: Fast Speech-to-Text Modeling with Fairseq,” in *AAACL-IJCNLP: Demonstrations*, 2020, pp. 33–39.
- [28] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko, E. Grave, V. Pratap, A. Sriram, V. Liptchinsky, and R. Collobert, “End-to-end ASR: from Supervised to Semi-Supervised Learning with Modern Architectures,” in *ICML 2020 Workshop SAS*, 2020.
- [29] P. Guo, F. Boyer *et al.*, “Recent Developments on ESPnet Toolkit Boosted by Conformer,” in *ICASSP*, 2021, pp. 5874–5878.
- [30] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *INTERSPEECH*, 2020, pp. 5036–5040.
- [31] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations,” in *NeurIPS*, 2020.
- [32] Y. Liu, J. Gu, N. Goyal, X. Li, S. Edunov, M. Ghazvininejad, M. Lewis, and L. Zettlemoyer, “Multilingual Denoising Pre-training for Neural Machine Translation,” *Transactions of the ACL*, vol. 8, pp. 726–742, 2020.
- [33] T. H. Nguyen, T.-D. H. Nguyen, D. Phung, D. T.-C. Nguyen, H. M. Tran, M. Luong, T. D. Vo, H. H. Bui, D. Q. Phung, and D. Q. Nguyen, “A Vietnamese-English Neural Machine Translation System,” in *INTERSPEECH*, 2022, p. to appear.
- [34] O. Kuchaiev, J. Li, H. Nguyen, O. Hrinchuk, R. Leary, B. Ginsburg, S. Kriman, S. Beliaev, V. Lavrukhin, J. Cook *et al.*, “NeMo: A toolkit for building ai applications using neural modules,” *arXiv preprint*, vol. arXiv:1909.09577, 2019.
- [35] G. I. Gállego, I. Tsiamas, C. Escolano, J. A. R. Fonollosa, and M. R. Costa-jussà, “End-to-End Speech Translation with Pre-trained Models and Adapters: UPC at IWSLT 2021,” in *IWSLT*, Aug. 2021, pp. 110–119.
- [36] X. Li, C. Wang, Y. Tang, C. Tran, Y. Tang, J. Pino, A. Baevski, A. Conneau, and M. Auli, “Multilingual Speech Translation from Efficient Finetuning of Pretrained Models,” in *ACL-IJCNLP*, 2021, pp. 827–838.
- [37] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *INTERSPEECH*, 2019, pp. 2613–2617.
- [38] T. Kudo and J. Richardson, “SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing,” in *EMNLP: Demonstrations*, 2018, pp. 66–71.
- [39] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, “fairseq: A Fast, Extensible Toolkit for Sequence Modeling,” in *NAACL-HLT 2019: Demonstrations*, 2019.
- [40] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *arXiv preprint*, vol. arXiv:1412.6980, 2014.
- [41] M. Post, “A Call for Clarity in Reporting BLEU Scores,” in *WMT: Research Papers*, 2018, pp. 186–191.