Improving Cascaded Unsupervised Speech Translation with Denoising Back-translation

Yu-Kuan Fu[1][*], Liang-Hsuan Tseng[1][*], Jiatong Shi[2], Chen-An Li[1],
Tsu-Yuan Hsu[1], Shinji Watanabe[2], Hung-yi Lee[1]

[1]College of Electrical Engineering and Computer Science, National Taiwan University

[2]Language Technologies Institute, Carnegie Mellon University

[1]{r11942083,r11921067,b08902123,b08201047,hungyilee}@ntu.edu.tw

[2]{jiatongs@cs.cmu.edu,shinjiw@cmu.edu}

Abstract

Most of the speech translation models heavily rely on parallel data, which is hard to collect especially for low-resource languages. To tackle this issue, we propose to build a cascaded speech translation system without leveraging any kind of paired data. We use fully unpaired data to train our unsupervised systems and evaluate our results on CoVoST 2 and CVSS. The results show that our work is comparable with some other early supervised methods in some language pairs. While cascaded systems always suffer from severe error propagation problems, we proposed denoising back-translation (DBT), a novel approach to building robust unsupervised neural machine translation (UNMT). DBT successfully increases the BLEU score by 0.7–0.9 in all three translation directions. Moreover, we simplified the pipeline of our cascaded system to reduce inference latency and conducted a comprehensive analysis of every part of our work. We also demonstrate our unsupervised speech translation results on the established website [1].

1 Introduction

Speech translation (ST) aims to convert speech from one language to another, allowing seamless communication between individuals speaking in different languages. Conventional speech-to-text translation (S2TT) system is accomplished by concatenating automatic speech recognition (ASR) and text-to-text machine translation (MT) (Ney, 1999) modules. Meanwhile, the cascaded speech-to-speech translation (S2ST) system further appends a text-to-speech (TTS) synthesis module after the S2TT system (Lavie et al., 1997; Wahlster, 2000; Nakamura et al., 2006). Recently, direct S2TT (Bérard et al., 2016; Weiss et al., 2017) and S2ST (Jia et al., 2019, 2021) systems emerge to solve the

---

[*]Equal Contribution

[1] https://anonymous-acl2023.github.io/us2s-demo/

error propagation and inference latency problem of the cascaded systems. Some direct S2TT systems have shown comparable results or even outperform the cascaded S2TT systems (Wang et al., 2021; Bentivogli et al., 2021).

Most of the ST systems are trained on parallel data, which is extremely limited, especially for low-resource languages. This situation strongly hinders the performance of direct ST systems. Although cascaded systems could overcome this issue by collecting data for each component separately, they still face the challenge of domain mismatch caused by variations in data distribution across different corpora.

Compared to parallel data, unlabelled data is much easier to obtain regardless of modalities. The first unsupervised speech-to-text translation (US2TT) aligned spoken words with written words and then applied unsupervised word-by-word translation (Chung et al., 2018, 2019b). Moreover, with the recent progress in unsupervised automatic speech recognition (UASR) (Baevski et al., 2021), unsupervised neural machine translation (UNMT) (Lample et al., 2017; Lample and Conneau, 2019; Song et al., 2019), and unsupervised text-to-speech (UTTS) synthesis (Ni et al., 2022; Liu et al., 2022b), Wang et al. (2022) built an unsupervised speech-to-speech translation (US2ST) system. Besides building cascaded US2ST, they also generated pseudo labels for training direct US2TT systems. Their work might be considered concurrent with ours, mainly focusing on techniques to conduct simple and effective US2ST systems.

Although the idea of cascaded US2ST is simple, directly concatenating UASR, UNMT, and UTTS might suffer from severe error propagation problems. For example, UNMT is trained on clean text, and small perturbations may greatly affect the translation results (Belinkov and Bisk, 2017). To tackle the issue, research on robust NMT has been widely investigated (Di Gangi et al., 2019; Sperber

et al., 2017; Sun et al., 2020); however, improving the robustness of UNMT is rarely studied. In this paper, we proposed denoising back-translation (DBT), a novel method to build a robust UNMT system. DBT combines the idea of denoising auto-encoding and back-translation (BT), dealing with error propagation issues in a fully unsupervised fashion. Briefly speaking, the pseudo text of DBT is generated from text with some noise, and the model should learn to reconstruct the clean text from the pseudo text. According to our results, this method substantially increases the quality of the cascaded unsupervised speech translation system.

Another issue with cascaded systems is the high inference latency. To address this, we integrated two parts of our cascaded system. First, the output of the UASR is normalized (stripped of punctuation marks and cases). Then, a text detokenizer is used to reconstruct the unnormalized text and feed it into the UNMT. By fine-tuning or continually training the UNMT with normalized source language text, the model is able to translate normalized source text into unnormalized target text. While there may be some degradation in performance, this method simplifies the pipeline of the cascaded system and significantly reduces inference time.

We evaluate our cascaded US2ST system on CVSS (Jia et al., 2022), a multilingual S2ST corpus; and CoVoST 2 (Wang et al., 2020b), a multilingual ST corpus of which CVSS built on top. We demonstrate that our US2ST could yield reasonable results across multiple translation directions, some of which are even better than the previous supervised approach[2]. Moreover, the proposed DBT method can help improve the performance by mitigating error propagation and domain mismatch problems.

2 Related works

2.1 ST

Traditional S2TT system is composed of ASR and MT (Ney, 1999), and S2ST system further append a TTS model after the MT model (Lavie et al., 1997; Wahlster, 2000; Nakamura et al., 2006). However, cascaded systems might suffer from error propagation and inference latency. Recent develpments in end-to-end S2TT (Bérard et al., 2016; Weiss et al., 2017) and S2ST (Jia et al., 2019, 2021) systems have been proposed to address these issues.

The main challenge of ST systems is the lack of parallel data. The unsupervised approach for ST is limited, but some progress has been made. The first US2TT system intended to learn the cross-modal alignment to map speech into written words, and used cross-lingual embedding to align words in different languages Chung et al. (2018, 2019b). Additionally, concurrent work built a US2ST system by combining SOTA UASR, UNMT, and UTTS to build a cascaded US2ST system. Further, they generated pseudo labels from the cascaded system to train an end-to-end US2ST system (Wang et al., 2022).

2.2 UASR

UASR takes audio features or representations as input and generates phoneme sequences without supervision. To tackle the challenging problem, Liu et al. (2018) first came out with the idea of applying a Generative Adversarial Network (GAN) (Goodfellow et al., 2020). However, phoneme-level boundaries are required to segment the audio and construct embedding sequences. (Chen et al., 2019) breaks the limit by iteratively refining the audio segments with Hidden Markov Model (HMM) and GAN, achieving complete UASR.

Recently, Baevski et al. (2021) proposed wav2vec-U, building the GAN-based UASR framework on top of the representation from wav2vec 2.0 (W2V2) (Baevski et al., 2020), a self-supervised speech model. The results outperformed previous SOTA, and are even comparable with some of the best-known supervised methods. Moreover, the original paper has shown that with the cross-lingual pre-trained version of W2V2 (Conneau et al., 2020), UASR in other languages is also available. The follow-up work, wav2vec-U 2.0 (Liu et al., 2022a), enabled the model to be trained end-to-end with the simplified pipeline and the improved training objective.

2.3 UNMT

The first fully unsupervised neural machine translation model adopted a seq2seq model, and the encoder mapped monolingual corpus in two languages to a shared latent space via adversarial training. The decoder learned to reconstruct in both languages from the latent representations by denoising autoencoding loss and online back-translation loss (Lample et al., 2017).

Recently, cross-lingual language model pretraining brought large progress to UNMT. XLM (Lam-
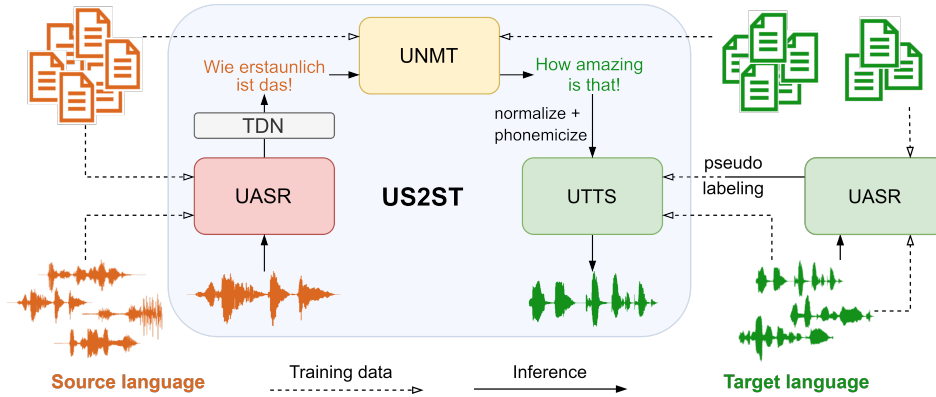
Figure 1: The framework of our cascade US2ST.

ple and Conneau, 2019) first adopted masked language modeling pretraining to initialize the encoder and decoder, and then used back-translation loss together with denoising autoencoding loss to fine-tune the whole seq2seq model. MASS (Song et al., 2019) further used masked seq2seq pretraining to pretrain encoder and decoder jointly to reduce the discrepancy between pretraining and fine-tuning. Their proposed method can align two languages with only back-translation loss. MASS outperforms XLM and all the other previous SOTA UNMT models in several language pairs.

2.4 Robust NMT

Traditional NMT and UNMT models are trained with clean input, thus small perturbations can greatly degrade the performance (Belinkov and Bisk, 2017). To improve the robustness, Di Gangi et al. (2019); Sperber et al. (2017) directly trained or fine-tuned the translation model on the target domain.

As for the unsupervised scenario, Sun et al. (2020) intended to improve the robustness of UNMT by applying some perturbation on positional embedding and word embedding to the input. The model learned to reconstruct the original input via denoising autoencoding loss and adversarial training.

2.5 TTS and UTTS

Some works intend to improve the performance of TTS through unlabeled data. For instance, pre-training the encoder/decoder (Chung et al., 2019a); utilizing the dual nature of TTS and ASR tasks (Ren et al., 2019); applying variational autoencoder to learn from speech disentanglement (Lian et al., 2022).

In spite of the improvement they brought, these methods still depend on certain levels of paired data. Directly training a UTTS without any supervision from paired data seems to be extremely hard. However, with recent success in UASR, UTTS might be accomplished in another way—training on the pseudo labels generated from UASR systems (Ni et al., 2022; Liu et al., 2022b).

3 Methods

Figure 1 shows the architecture overview of our proposed approach to unsupervised speech-to-speech translation (US2ST). We split US2ST into three stages: unsupervised speech recognition (UASR), unsupervised machine translation (UNMT), and unsupervised speech synthesis (UTTS). The modules were trained separately but all in an unsupervised manner. During inference, we form the functionality of S2ST by concatenating them. Furthermore, we proposed denoising back-translation, to mitigate the error propagation and the domain mismatch issues between UASR and UMT submodules.

3.1 Base Architecture

UASR We conducted our UASR subsystem following wav2vec-U (Baevski et al., 2021), and our code is based on their implementation in fairseq[3]. Besides its breakthrough performance on UASR in multiple languages, the robustness and stabilities across different corpora have also been well analyzed (Lin et al., 2022). Thus, We mainly follow the data preparation procedure, model architecture, and the training objective of wav2vec-U. During inference, the model takes the preprocessed audio features as input and generates phoneme sequences. To further obtain word-level sequences, we adopt

---

[3] https://github.com/facebookresearch/fairseq

different decoding strategies, such as lexicon-based kenlm decoder and the weighted finite-state transducer (WFST; Mohri et al. (2002)). Self-training techniques on HMM are also applied for better performance.

TDN The outputs from the UASR are normalized word sequences. The UNMT model, however, may rely on punctuation marks and capital letters. For better performance, we also learned a text denormalizer to transform the generated word sequences back into sequences with punctuation marks and capitalized words, namely, denormalization. The module was a transformer-based seq2seq model. We first formed the paired data by normalizing raw text sentences and then trained the model with cross-entropy loss.

UNMT UNMT aims to map sentences from source language $\mathcal{S}$ to target language $\mathcal{T}$ without leveraging any paired data. We conduct UNMT by following the architecture and pretrain process of MASS (Song et al., 2019), which is a transformer-based seq2seq language model. During the fine-tuning process, we use denoising back-translation plus denoising autoencoder objective to align two languages and increase robustness.

UTTS We conducted UTTS by following the architecture of Variational Inference with adversarial learning for end-to-end Text-to-Speech (VITS), which has shown a significant performance gain over Tacotron2 and Transformer TTS in both subjective and objective evaluation (Kim et al., 2021; Hayashi et al., 2021), and use the UASR system mentioned in section 3.1 to generate pseudo labels for training.

3.2 Mitigating Error Propagation

Denoising back-translation. In our setting, the input of our translation model comes from the output of UASR, which might contain some noise, and worsen the performance significantly. In this paper, we introduce denoising back-translation, a novel approach for robust UNMT.

Given a source sentence $x \in \mathcal{S}$, a target sentence $y \in \mathcal{T}$, and $u^*(\cdot)$, $v^*(\cdot)$ are translation functions with directions of $\mathcal{S} \to \mathcal{T}$ and $\mathcal{T} \to \mathcal{S}$ respectively. we generate the pseudo parallel data by passing x and y through a noise function $f(\cdot)$, transcribed them into $u^*(f(x)) \in \mathcal{T}$ and $v^*(f(y)) \in \mathcal{S}$ respectively. $f(\cdot)$ can be some artificial data augmentations including deletion, insertion, or other
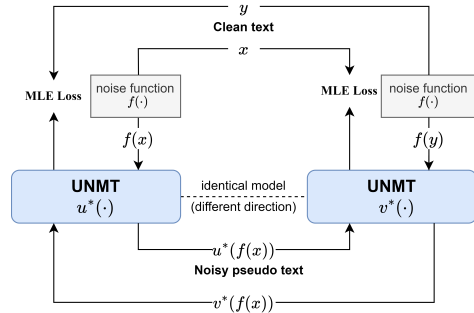

Figure 2: The illustration of our proposed method: DBT.

modules like a language model. The objective of denoising back-translation is to reconstruct x and y from $u^*(f(x))$ and $v^*(f(y))$ respectively. The denoising back-translation loss is as follows:

$$\mathcal{L}_{DBT} = \mathbb{E}_{x \in \mathcal{S}}[-\log P_{\mathcal{T} \to \mathcal{S}}(x|u^*(f(x)))] \\ + \mathbb{E}_{y \in \mathcal{T}}[-\log P_{\mathcal{S} \to \mathcal{T}}(y|v^*(f(y)))] \quad (1)$$

Unlike the original back-translation, the pseudo labels of DBT are generated from noisy sentences, so the model should learn to transcribe the noisy pseudo sentences into clean sentences, and thus become more robust. The whole process of DBT was illustrated in Figure 2.

4 Experiments

4.1 Data

To demonstrate our US2ST systems across different languages, we evaluate our S2ST results on CVSS, which is a multilingual corpus built on top of CoVoST 2 (Wang et al., 2020b) and CommonVoice ver.4 (CV4; Ardila et al., 2019). Thus, we are also available to evaluate our S2TT results on CoVoST 2 and ASR results on CV4.

However, **we do not utilize any paired data from the corpus during training**; instead, we use audio and text data from different corpora, constructing an unpaired scenario for our US2ST. For audio, we adopt Common Voice ver.4 for UASR and LJspeech (Ito and Johnson, 2017) for UTTS without using any transcriptions from them; and for text, we extract sentences from Wikipedia[4], WMT'14, CC100 (Conneau et al., 2019), and LibriSpeech LM data[5]. All of the data we used is open-sourced and public-available.

---

[4] We extract the data from wiki using WikiExtractor (Attardi, 2015)

[5] Following Liu et al. (2022b), we exclude the transcriptions of LJspeech to form fully unpaired scenario

4.2 System setups

UASR We used audio from CV4 and text data from Wikipedia to train our UASR models. More precisely, we use 100 hours of audio and about 1–3M sentences for each language. After training, we evaluate the results with the transcriptions from CV4. For the pre-trained W2V2 model, we directly use the cross-lingual version (XLSR; Conneau et al., 2020) without finetuning. XLSR had pre-trained in many different languages thus suiting our needs for training UASR in languages other than English.

During preprocessing, we adopted the same configuration in wav2vec-U, except for the silence insertion rate of French. We found that our French model converged better when <SIL> token insertion rate is 0.5 instead.

As for the GAN training configuration, we chose the coefficients of the loss function according to the original paper as follows: the gradient penalty weight $\lambda = 1.5$ or 2.0, the smoothness penalty weight $\gamma = 0.5$, and the phoneme diversity loss weight $\eta = 4$. We trained 3 seeds for each configuration, conducting 6 models for each language.

TDN For the text denormalizer (TDN), we adopted transformer encoder-decoder architecture, both of which contained 4 layers of transformer blocks. We constructed the input data by normalizing 8–10M of plain text data from CC100, and the objective is to reconstruct the unnormalized data.

UNMT We used the back-translation fine-tuned MASS model released by Microsoft[6] to initialize our German-English and French-English UNMT models. For the Spanish-English UNMT model, we followed the same pretraining and fine-tuning steps as the standard MASS model.

For denoising back-translation as discussed in Section 3.2, the artificial noise $f(\cdot)$ included random drop, substitution, and insert, whose probability were 0.05, 0.01, and 0.05 respectively. We continually fine-tune the model by denoising back-translation loss plus denoising autoencoding loss to build a robust UNMT model.

UTTS In this paper, we conducted single-speaker UTTS, so we trained a UASR model on the audio of LJSpeech (Ito and Johnson, 2017). Using the UASR-generated pseudo labels, we then proceeded with the training of VITS[7], but the transcription of

LJSpeech was replaced by the pseudo label.

4.3 Supervised cascaded S2ST

First, we constructed our upper bound model by training a supervised cascaded S2ST (ASR→MT→TTS) which shares similar model architecture with our US2ST.

For ASR, we finetuned the whole XLSR instead of treating it as a feature extractor. We adopted letter-based training and followed the configuration from fairseq (Ott et al., 2019). The amount of audio data was exactly the same as those in UASR. Furthermore, we finetuned the XLSR models individually for each language. MT is achieved by training the same initial model with UNMT, but the training data were the transcriptions of CoVoST 2 plus CC100. We supervised trained with CoVoST 2, and CC100 was used for back-translation training, which can boost the performance of supervised machine translation models. Finally, instead of training on pseudo labels from UASR, the supervised TTS model directly uses the reference phoneme sequences and their corresponding utterance audio as paired data.

By constructing cascaded supervised S2ST, we can discuss the performance individually for each component.

4.4 Evaluation

The evaluation metric of ASR was normalized word error rate (WER), which removed all punctuation marks and converted all characters to lowercase. We used sacreBLEU[8] to calculate the BLEU score of S2TT. For the final S2ST results, Whisper[9] (Radford et al., 2022), a supervised ASR model released by OpenAI, is adopted to transcribe the hypothesized audio and calculated the BLEU score.

4.5 Results

We show our results in Table 1, including the results of ASR, S2TT, and S2ST. To compare our US2ST system with others' works, we collect some results from the previous studies on CoVoST 2 and CVSS ((a)–(f), (h)–(i)). To get better comparisons, both cascaded and direct systems are included.

Next, we discuss the details of the methods in the table. *(a)* comes from Wang et al. (2020b); among all experiments in their paper, we only report the results of the cascaded S2TT system constructed

---
[6]https://github.com/microsoft/MASS
[7]https://github.com/jaywalnut310/vits

[8]https://github.com/mjpost/sacrebleu
[9]we used the cross-lingual large-v2 model released by https://github.com/openai/whisper

Table 1: The results are evaluated on CoVoST 2 for S2TT, and CVSS for S2ST. C-T stands for cascaded S2TT; D-S for direct S2ST.

Method	Type	ASR ↓			S2TT (X→En) ↑			S2ST (X→En) ↑		
		Fr	De	Es	Fr	De	Es	Fr	De	Es
SUPERVISED LEARNING										
(a) Wang et al. (2020b)	C-T	18.3	21.4	16.0	27.6	21.0	27.4	-	-	-
(b) fairseq S2T (T-Sm) (Wang et al., 2020a)	D-T	-	-	-	26.3	17.1	23.0	-	-	-
(c) fairseq S2T (Multi. T-Md)	D-T	-	-	-	26.5	17.5	27.0	-	-	-
(d) XLS-R (2B) (Babu et al., 2021)	D-T	-	-	-	37.6	33.6	39.2	-	-	-
(e) Translatotron (Jia et al., 2019)	D-S	-	-	-	-	-	-	15.5	6.9	14.1
(f) Translatotron 2 (Jia et al., 2021)	D-S	-	-	-	-	-	-	28.3	19.7	23.5
(g) Our upper bound	C-S	16.2	14.1	11.0	29.5	25.6	31.0	23.8	21.8	26.2
UNSUPERVISED LEARNING										
(h) Wang et al. (2022) cascaded US2ST		-	-	-	24.4	-	23.4	21.6	-	21.2
(i) Wang et al. (2022) end-to-end US2ST		-	-	-	24.2	-	24.0	21.2	-	20.1
(j) Our cascaded US2ST	C-S	33.2	23.8	17.4	20.0	19.5	23.8	13.4	13.8	16.7
(k) Our cascaded US2ST with DBT (artificial noise)	C-S				20.8	20.4	24.5	14.4	14.7	17.4

by monolingual ASR and bilingual MT for fair comparison. According to the table, our US2TT performances in De–En are just having small degradation from theirs (*(j)* and *(k)* vs *(a)*), indicating that our US2TT system can be comparable to some early supervised cascaded S2TT works in some language pairs.

Rows *(b)* and *(c)* are the results of the direct S2TT systems from Wang et al. (2020a). They have developed a tool-kit for S2TT and demonstrated it on CoVoST 2 with different model backbones. We compare our US2TT results (row *(j) and (k)*) with their Transformer-based models. Our results have not only outperformed their small model (*(b)*) in De-En and Es-En tasks but also outperformed their large model in De-En (row *(c)*). Our works are comparable to some early direct S2TT systems (*(b)*, *(c)*) except for the Fr-En pair.

Row *(d)* is the SOTA direct S2TT model, which is a cross-lingual speech model based on wav2vec 2.0 architecture. Taking advantage of the large pretrained self-supervised model, the performance is significantly better than all the other works in all language pairs, showing there is still a huge gap between SOTA supervised S2TT and our US2TT systems.

Rows *(h)* and *(i)* come from the concurrent US2ST system (Wang et al., 2022); *(h)* is constructed by concatenating UASR, UNMT, and UTTS, and *(i)* is an end2end S2ST systems trained on pseudo labels generated by *(h)*. Our model architecture is similar to theirs, while they fine-tune the wav2vec 2.0 with target languages in UASR to mitigate the domain mismatch between pretraining and downstream tasks. For Fr-En S2TT, their per-

formance is superior to ours. This may be due to the difficulty of French UASR, which has a higher error rate than other languages. This phenomenon was also observed in Wang et al. (2022), and we think that it might be more severe in our case since we did not fine-tune the wav2vec 2.0 models on the target languages. Suffering from domain mismatch, our systems still outperform their works in Es-En S2TT.

The performance of our S2ST systems drops even more than S2TT. However, this is because we did not utilize the data from CVSS when training our UTTS models for the fairness concern. Other studies (*(e)*, *(f)*, *(h)*, *(i)*) directly trained the whole model or UTTS with the data from CVSS. Since our models are trained with LJSpeech (row *(g)*, *(j)*, *(k)*), they could have severe domain mismatch problems during inference on out-domain data. We had also investigated the cause of the performance drop in the later section and came out with the same conclusion. In spite of this, our unsupervised S2ST results still outperform Translatotron (row *(j)*, *(k)* v.s. *(e)*) in De-En and Es-EN translation directions.

Last but not least, Our cascaded system with DBT *(k)* outperforms *(j)* by 0.7 to 0.9 in both US2TT and US2ST, indicating a model can better translate the noisy input under the guidance of DBT, and thus mitigate the problem of error propagation in cascaded systems.

5 Analysis

Stabilities of UASR cross different languages First of all, we found that the stabilities of our UASR models vary between languages. The measurement of the stability is by calculating the per-

centage of the converged rate among the models leveraging the same amount of text and speech data. We consider a UASR model is converged if its $PER < 50\%$. We summarize the discoveries in Table 2. According to our experiments, German and Spanish are easier to converge; while French usually can not converge well. However, we found that it might be more suitable for French UASR models to converge if we change the <SIL> token insertion rate from 0.25 to 0.5.

Table 2: Stabilities of UASR across different languages.

Lang.	<SIL> ins. rate	Best PER (*Viterbi*)	%-converged (PER < 50%)
De	0.25	25.3%	66%
Es	0.25	27.0%	50%
Fr	0.25	49.2%	<10%
	0.50	35.2%	17%

Decoding and self-training in UASR The original outputs of wav2vec-U are in phoneme-level, which are incompatible with the UNMT. However, with the integration with LM, we are available to obtain word-level output sequences. As shown in the part **(I)** of Table 3, we demonstrate that the two decoding methods, Kenlm and WFST can both generate word sequences. The second part **(II)** in the table illustrates the effectiveness of self-training on HMM. Among all the methods, we considered that the best strategy we found was by conducting self-training on HMM with the pseudo labels from WFST decoding. More surprisingly, even if the pseudo labels come from Viterbi decoding, using these labels on HMM can make huge improvements. After self-training, the performance gap between Viterbi and WFST decoding became relatively small. Note that for simplicity, we only show the results on the testing set of CV4-German; while the results on other languages also share similar trends.

Integration of text denormalization and UNMT In this section, we try to integrate text denormalization into UNMT; reducing pipelines of cascaded system might mitigate error propagation, and reduce inference time.

We introduced normalized fine-tuning (NFT) to direct translate normalized source text into unnormalized target text. NFT initializes the model as

Table 3: Comparison of different decoding strategies and the improvement brought by HMM self-training. We use the same 4-gram LM (phoneme-level or word-level) across different methods.

Method	LM	PER(%)	WER(%)
(I) *Without self-training*			
Viterbi	✗	25.2	-
Kenlm	✓	29.5	39.5
WFST	✓	21.3	34.4
(II) *With self-training*			
Viterbi→ HMM	✓	15.2	25.3
WFST→ HMM	✓	**14.4**	**23.8**

original DBT fine-tuning, but fine-tunes on normalized source text and unnormalized target text. While directly fine-tune the checkpoint pretrained on unnormalized text might induce mismatch between pretraining and fine-tuning. To address this problem, we further introduce normalized continual training (NCT), which continual pretrains the checkpoint on normalized source text and unnormalized target text, and follow NFT for downstream task fine-tuning.

The results are shown in Table 5. we compare NFT (I) and NCT (II) with our two baselines: (I) UDN + UNMT (our original setting), (II) UNMT (only use the translation model of (I)). The performance of (II) drops a lot, for normalized text never appear during the training process of the model, directing translating on that induces severe domain mismatch. (III) and (IV) have outperformed (II) a lot, but they still decrease the BLEU score by about 2.6 and respectively ((III), (IV) v.s. (I)), indicating mismatch between pretraining and downstream task training is severe, and NCT has only minor improvement.

Integrating text denormalization and UMT, while it did not performs better due to the mismatch between pretraining and fine-tuning, can still reduce the inference latency. To address this mismatch, a potential solution would be pretraining the model on normalized source text and unnormalized target text from scratch.

Robustness of UNMT DBT has been shown to improve the performance of cascaded S2TT and S2ST systems. In this section, we investigate the robustness of UNMT. Table 4 shows the BLEU scores of translating the ground truth of CoVoST ("Clean"), and that of translating the output of

Table 4: Robustness of UNMT across languages. "Clean" and "ASR" refer to the BLEU score of translating on the ground truth and UASR output respectively.

Direction	Clean		ASR	
	BT	DBT	BT	DBT
Fr-En	35.3	35.0	20.0	20.8
De-En	27.1	28.0	19.5	20.4
Es-En	33.4	33.1	23.8	24.5

Table 5: The BLEU score of integrating text normalization and UNMT for De-En S2TT.

Model	S2TT
(I) TDN + UNMT	20.8
(II) UNMT	13.0
(III) NFT	18.2
(IV) NCT	18.6

UASR ("ASR").

The results indicate that for Fr-En and Es-En on "Clean", the performance drops slightly compared to that of BT, but the score drop from "Clean" to "ASR" of DBT decreased by about 1 BLEU score. For De-En, DBT even performs better than BT on "Clean", and the performance drops from "Clean" to "ASR" are the same, which means that DBT can be regraded as a new data-augmentation method to boost the performance of a UNMT model. This result from avoiding the model from directly copying the input during generating pseudo-label for back-translation.

Overall, DBT increases the robustness without sacrificing its performance on clean input too much, and it even outperforms BT on "Clean" in some cases.

Performance analysis of UTTS In this section, we present more analytical results of our UTTS submodule. In part *(I)* in Table 6, we evaluate our supervised TTS and UTTS on in-domain testing set (LJspeech) and out-domain testing set (US2ST, Fr→En). After we got the synthesis speech, we send the audio to whisper and then calculate the WER. We used the base model for this experiment to further accentuate the performance differences. Our results indicated that the performance drop between supervised TTS and UTTS is much lower than the error induced by the domain mismatch problem. The results also emphasized that the do-

Table 6: Analysis of our UTTS models. In part *(I)*, we the performance drop due to the domain mismatch problem. In part *(II)*, we further investigate the effectiveness of our UTTS by using the same testing data as the supervised TTS.

Testing data	TTS	UTTS
(I) WER on in-domain / out-domain data.		
in-domain	23.5%	31.5%
out-domain	46.5%	54.2%
(II) BLEU score of using ST / UST (Fr→En)		
sup. ST	23.8	20.6
unsup. ST (BT)	15.5	13.4

main mismatch between the training data of TTS models and the testing data is one of the main reasons for our S2ST performance drop. Leveraging the data from CVSS for UTTS training might be a solution, but it may also induce fairness concerns from our point of view.

Next, in section *(II)*, we evaluate our supervised TTS and UTTS models on the outputs from supervised ST and UST. According to the table, we can infer that the gap between supervised TTS and UTTS might be overestimated. The performance drop induced by pseudo-labeling is acceptable or at least reasonable.

6   Conclusion

In this work, we build cascaded unsupervised speech-to-speech translation (US2ST) systems in several translation directions. To further improve the performance and mitigate the error propagation problems, we propose denoising back-translation (DBT), which is a novel method to improve the robustness of UNMT. DBT generally improves the performance of unsupervised speech translation (UST) across all the language pairs that we have experimented on. Without leveraging any paired data, our speech translation results are even better than some previous supervised methods. Additionally, we analyze the performance of each part in different settings individually; and we also attempt to integrate the TDN into the UNMT to reduce inference latency. In the future, we may investigate more techniques that can reduce the error propagation problems between different unsuper-

vised cascaded modules; or conduct direct UST or US2ST.

Limitations

In this work, we have handled the problem of error propagation among UASR, TDN, and UNMT. Nevertheless, we didn't resolve that between UTTS and other modules, which may lead to a lower score of the S2ST result.

Our methodology works for most languages, however, our US2ST is based on UNMT for unpaired text data. Therefore, it is limited to written languages. We believe that our denoise back-translation brings new insights to US2ST and can extend to unwritten language setups.

Ethics Statement

Our works build an effective UST cascaded system and try to mitigate the error propagation and inference latency. The communities might be interested in how to build a direct US2TT or even US2ST system or how to further improve the performance of the UST system.

References

Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. 2019. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*.

Giuseppppe Attardi. 2015. Wikiextractor. https://github.com/attardi/wikiextractor.

Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhotia, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, et al. 2021. Xls-r: Self-supervised cross-lingual speech representation learning at scale. *arXiv preprint arXiv:2111.09296*.

Alexei Baevski, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2021. Unsupervised speech recognition. *Advances in Neural Information Processing Systems*, 34:27826–27839.

Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460.

Yonatan Belinkov and Yonatan Bisk. 2017. Synthetic and natural noise both break neural machine translation. *arXiv preprint arXiv:1711.02173*.

Luisa Bentivogli, Mauro Cettolo, Marco Gaido, Alina Karakanta, Alberto Martinelli, Matteo Negri, and Marco Turchi. 2021. Cascade versus direct speech translation: Do the differences still make a difference? *arXiv preprint arXiv:2106.01045*.

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. *arXiv preprint arXiv:1612.01744*.

Kuan-Yu Chen, Che-Ping Tsai, Da-Rong Liu, Hung-Yi Lee, and Lin-shan Lee. 2019. Completely unsupervised speech recognition by a generative adversarial network harmonized with iteratively refined hidden markov models. *arXiv preprint arXiv:1904.04100*.

Yu-An Chung, Yuxuan Wang, Wei-Ning Hsu, Yu Zhang, and RJ Skerry-Ryan. 2019a. Semi-supervised training for improving data efficiency in end-to-end speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6940–6944. IEEE.

Yu-An Chung, Wei-Hung Weng, Schrasing Tong, and James Glass. 2018. Unsupervised cross-modal alignment of speech and text embedding spaces. *Advances in neural information processing systems*, 31.

Yu-An Chung, Wei-Hung Weng, Schrasing Tong, and James Glass. 2019b. Towards unsupervised speech-to-text translation. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7170–7174. IEEE.

Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. 2020. Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979*.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

Mattia Antonino Di Gangi, Robert Enyedi, Alessandra Brusadin, and Marcello Federico. 2019. Robust neural machine translation for clean and noisy speech transcripts. *arXiv preprint arXiv:1910.10238*.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144.

Tomoki Hayashi, Ryuichi Yamamoto, Takenori Yoshimura, Peter Wu, Jiatong Shi, Takaaki Saeki, Yooncheol Ju, Yusuke Yasuda, Shinnosuke Takamichi, and Shinji Watanabe. 2021. Espnet2-tts: Extending the edge of tts research. *arXiv preprint arXiv:2110.07840*.

Keith Ito and Linda Johnson. 2017. The lj speech dataset. https://keithito.com/LJ-Speech-Dataset/.

Ye Jia, Michelle Tadmor Ramanovich, Tal Remez, and Roi Pomerantz. 2021. Translatotron 2: Robust direct speech-to-speech translation. *arXiv preprint arXiv:2107.08661*.

Ye Jia, Michelle Tadmor Ramanovich, Quan Wang, and Heiga Zen. 2022. CVSS corpus and massively multilingual speech-to-speech translation. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6691–6703, Marseille, France. European Language Resources Association.

Ye Jia, Ron J Weiss, Fadi Biadsy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu. 2019. Direct speech-to-speech translation with a sequence-to-sequence model. *arXiv preprint arXiv:1904.06037*.

Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, pages 5530–5540. PMLR.

Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.

Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.

Alon Lavie, Alex Waibel, Lori Levin, Michael Finke, Donna Gates, Marsal Gavalda, Torsten Zeppenfeld, and Puming Zhan. 1997. Janus-iii: Speech-to-speech translation in multiple languages. In *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 99–102. IEEE.

Jiachen Lian, Chunlei Zhang, Gopala Krishna Anumanchipalli, and Dong Yu. 2022. Utts: Unsupervised tts with conditional disentangled sequential variational auto-encoder. *arXiv preprint arXiv:2206.02512*.

Guan-Ting Lin, Chan-Jan Hsu, Da-Rong Liu, Hung-Yi Lee, and Yu Tsao. 2022. Analyzing the robustness of unsupervised speech recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8202–8206. IEEE.

Alexander H Liu, Wei-Ning Hsu, Michael Auli, and Alexei Baevski. 2022a. Towards end-to-end unsupervised speech recognition. *arXiv preprint arXiv:2204.02492*.

Alexander H Liu, Cheng-I Jeff Lai, Wei-Ning Hsu, Michael Auli, Alexei Baevskiv, and James Glass. 2022b. Simple and effective unsupervised speech synthesis. *arXiv preprint arXiv:2204.02524*.

Da-Rong Liu, Kuan-Yu Chen, Hung-yi Lee, and Lin-shan Lee. 2018. Completely unsupervised phoneme recognition by adversarially learning mapping relationships from audio embeddings. *arXiv preprint arXiv:1804.00316*.

Mehryar Mohri, Fernando Pereira, and Michael Riley. 2002. Weighted finite-state transducers in speech recognition. *Computer Speech & Language*, 16(1):69–88.

Satoshi Nakamura, Konstantin Markov, Hiromi Nakaiwa, Gen-ichiro Kikui, Hisashi Kawai, Takatoshi Jitsuhiro, J-S Zhang, Hirofumi Yamamoto, Eiichiro Sumita, and Seiichi Yamamoto. 2006. The atr multilingual speech-to-speech translation system. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(2):365–376.

Hermann Ney. 1999. Speech translation: Coupling of recognition and translation. In *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No. 99CH36258)*, volume 1, pages 517–520. IEEE.

Junrui Ni, Liming Wang, Heting Gao, Kaizhi Qian, Yang Zhang, Shiyu Chang, and Mark Hasegawa-Johnson. 2022. Unsupervised text-to-speech synthesis by unsupervised automatic speech recognition. *arXiv preprint arXiv:2203.15796*.

Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale weak supervision. Technical report, Technical report, OpenAI, 2022. URL https://cdn.openai.com/papers/whisper. pdf.

Yi Ren, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2019. Almost unsupervised text to speech and automatic speech recognition. In *International Conference on Machine Learning*, pages 5410–5419. PMLR.

Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.

Matthias Sperber, Jan Niehues, and Alex Waibel. 2017. Toward robust neural machine translation for noisy input sequences. In *Proceedings of the 14th International Conference on Spoken Language Translation*, pages 90–96.

Haipeng Sun, Rui Wang, Kehai Chen, Xugang Lu, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2020. Robust unsupervised neural machine translation with adversarial denoising training. *arXiv preprint arXiv:2002.12549*.

Wolfgang Wahlster. 2000. Verbmobil: Foundations of speech-to-speech translation. In *Artificial Intelligence*.

Changhan Wang, Hirofumi Inaguma, Peng-Jen Chen, Ilia Kulikov, Yun Tang, Wei-Ning Hsu, Michael Auli, and Juan Pino. 2022. Simple and effective unsupervised speech translation. *arXiv preprint arXiv:2210.10191*.

Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020a. fairseq s2t: Fast speech-to-text modeling with fairseq. *arXiv preprint arXiv:2010.05171*.

Changhan Wang, Anne Wu, and Juan Pino. 2020b. Covost 2 and massively multilingual speech-to-text translation. *arXiv preprint arXiv:2007.10310*.

Changhan Wang, Anne Wu, Juan Pino, Alexei Baevski, Michael Auli, and Alexis Conneau. 2021. Large-scale self-and semi-supervised learning for speech translation. *arXiv preprint arXiv:2104.06678*.

Ron J Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. Sequence-to-sequence models can directly translate foreign speech. *arXiv preprint arXiv:1703.08581*.