

DUB: Discrete Unit Back-translation for Speech Translation

Dong Zhang^{1,2†}, Rong Ye², Tom Ko², Mingxuan Wang^{2*}, Yaqian Zhou^{1*}

¹ School of Computer Science, Fudan University

² ByteDance

dongzhang22@m.fudan.edu.cn

{yerong, tom.ko, wangmingxuan.89}@bytedance.com

zhouyaqian@fudan.edu.cn

Abstract

How can speech-to-text translation (ST) perform as well as machine translation (MT)? The key point is to bridge the modality gap between speech and text so that useful MT techniques can be applied to ST. Recently, the approach of representing speech with unsupervised discrete units yields a new way to ease the modality problem. This motivates us to propose **Discrete Unit Back-translation (DUB)** to answer two questions: (1) Is it better to represent speech with discrete units than with continuous features in direct ST? (2) How much benefit can useful MT techniques bring to ST? With DUB, the back-translation technique can successfully be applied on direct ST and obtains an average boost of 5.5 BLEU on MuST-C En-De/Fr/Es. In the low-resource language scenario, our method achieves comparable performance to existing methods that rely on large-scale external data. Code and models are available at <https://github.com/0nutation/DUB>.

1 Introduction

Speech-to-text translation (ST) converts the spoken source language into the written target language, which is a closely related task to machine translation (MT). In recent years, direct ST that does not rely on intermediate transcription has received considerable attention due to its potential applications in unwritten language scenarios and various domains (Bérard et al., 2016; Sung et al., 2019; Han et al., 2021; Papi et al., 2021; Fang et al., 2022; Ye et al., 2022; Cheng et al., 2023). One of the major challenges faced by ST is data scarcity, which is similar to the low-resource scenarios encountered in MT. Intuitively, techniques developed for low-resource MT (Imamura et al., 2018; Xia et al., 2019; Chen et al., 2020; Liu et al., 2020; Tang

et al., 2020) should be utilized to improve ST performance. However, these techniques are hard to be transferred to ST due to the modality gap between speech and text, where ST takes continuous speech as input and MT takes discrete tokens as input. Generally speaking, if there is a way to efficiently remove the modality gap, a large number of useful NLP techniques can be applied and facilitate the improvement of ST.

Recently, representing speech with unsupervised discrete units has become popular and successful in the field of speech processing (Baevski et al., 2019, 2020; Hsu et al., 2021; Lakhota et al., 2021). Instead of losing relevant information, discretizing continuous speech has been found to have the advantage of filtering out extraneous signals (Sicherman and Adi, 2023; Lakhota et al., 2021), leading to significant improvements in the speech tasks, such as automatic speech recognition (Meng et al., 2022), text-to-speech (Dunbar et al., 2019), and speech-to-speech translation (Zhang et al., 2021; Lee et al., 2022; Inaguma et al., 2022). Based on this observation, we are motivated to explore the answers to the following two questions: (1) Is it better to represent speech with discrete units and use them as model input than with continuous features for direct ST? (2) By narrowing the modality gap with discrete speech units, how much benefit can useful MT techniques bring to direct ST?

In this paper, we propose **Discrete Unit Back-translation (DUB)**, which migrates the useful back-translation technique from MT to ST by discretizing the speech signals into unit sequences. In our proposed method, we first convert speech into discrete units using the clustering indices on HuBERT (Hsu et al., 2021) representations. To complete the translation task, we feed the discrete units into the **Unit-to-Text Translation (U2TT)** model. For the back-translation training strategy, DUB employs a text-to-unit translation model that learns to predict the source discrete units from the target

[†]Work was done while Dong Zhang was a research intern at ByteDance.

*Corresponding authors.

text. By leveraging the additional easily accessible text in the target language, we utilize the synthetic parallel data generated by the text-to-unit translation model in conjunction with the original parallel data to update the final unit-to-text model.

Our contributions include the following.

- We design a novel unit-text translation (U2TT) framework for direct ST by discretizing the speech feature in an unsupervised manner. Our analysis shows that in such a framework, the unit retains the semantic information for translation and can be used as model input.
- Based on the U2TT framework, we propose DUB, which successfully applies the back-translation technique to direct ST. Experimental results show that DUB can further yield an average 5.5 BLEU gain over the U2TT model on the MuST-C English-to-German, French, and Spanish translation directions.
- Our approach is particularly beneficial for low-resource or unwritten languages in the world because unit extraction does not require any textual supervision and only speech-translation pairs are used for training.

2 Related Work

Speech translation Without using textual transcriptions during inference or training, translating audio directly into the target language is very meaningful for languages that do not have a written form. [Bérard et al. \(2016\)](#) first proposed an end-to-end encoder-decoder architecture for such direct speech-to-text translation. Later, novel models ([Di Gangi et al., 2019b](#); [Dong et al., 2021a,b](#); [Zheng et al., 2021](#)) and training techniques, such as multi-task learning ([Indurthi et al., 2021](#); [Tang et al., 2021](#); [Ye et al., 2021](#)), knowledge distillation ([Liu et al., 2019](#); [Dong et al., 2021b](#)), and pre-training methods ([Zheng et al., 2021](#); [Zhang et al., 2022c](#)), were developed to improve end-to-end performance. However, these training methods often rely on the use of source text or knowledge from the pre-trained models. Without using transcripts or pretraining, [Zhang et al. \(2022a\)](#) proposed the parameterized distance penalty to better model speech locality in the self-attention structure and provided results on ST benchmarks covering 23 languages.

Back-translation Back-translation (BT) is a widely used method for improving machine translation systems by training a target-to-source model and creating synthetic parallel data from monolin-

gual target text. This approach has been shown to be effective in both statistical ([Bertoldi and Federico, 2009](#); [Bojar and Tamchyna, 2011](#)) and neural machine translation models ([Sennrich et al., 2016](#); [Edunov et al., 2018](#); [Hoang et al., 2018](#)), and is frequently used to improve translation performance in WMT competitions ([Farhad et al., 2021](#); [Wenzek et al., 2021](#); [Adelani et al., 2022](#)). A similar data augmentation idea through synthesizing speech from utterances can be applied to automatic speech recognition (ASR) ([Tjandra et al., 2017](#); [Hayashi et al., 2018](#); [Ueno et al., 2021](#)). However, applying BT in ST is not trivial. [Zhang et al. \(2022b\)](#) augmented the triplet data by TTS generation from transcription, but the experiment showed that such scaling yields minimal improvement to the final ST model.

Discrete speech representation Discrete speech representations are often studied in the work on self-supervised speech representation learning ([Van Den Oord et al., 2017](#); [Baevski et al., 2019, 2020](#); [Hsu et al., 2021](#); [Meng et al., 2022](#)). For example, [Van Den Oord et al. \(2017\)](#) proposed Vector Quantised-Variational AutoEncoder (VQ-VAE) to map continuous signals, like speech or image, into a discrete sequence space. [Hsu et al. \(2021\)](#) proposed HuBERT, which learns self-supervised speech representation by extracting speech features and clustering them offline, and iteratively training the clustering indexes of features at masked locations. Although the clustered discrete representations are only a by-product of HuBERT, they are used to build the generative spoken language model ([Lakhotia et al., 2021](#); [Kharitonov et al., 2022](#)), enhance speech representation ([Chung et al., 2021](#); [Meng et al., 2022](#); [Chen et al., 2022](#); [Wu et al., 2022](#); [Zhang et al., 2022c](#)), and model direct speech-to-speech translation ([Lee et al., 2022](#); [Inaguma et al., 2022](#)). In the prior literature, probably the most similar task to ours is the textless speech-to-speech translation ([Lee et al., 2022](#); [Nguyen et al., 2022](#)), but the difference is that they discretized the target-side speech and convert speech-to-speech generation into speech-to-discrete-unit generation, while we discretize the speech at the source side. [Zhang et al. \(2022c\)](#) leveraged the discrete unit as an interface to align speech and text, and proposed a unified-modal encoder-decoder pre-training model, SpeechUT. SpeechUT aims to improve speech representation via the units, while we use the units to construct a unified framework for

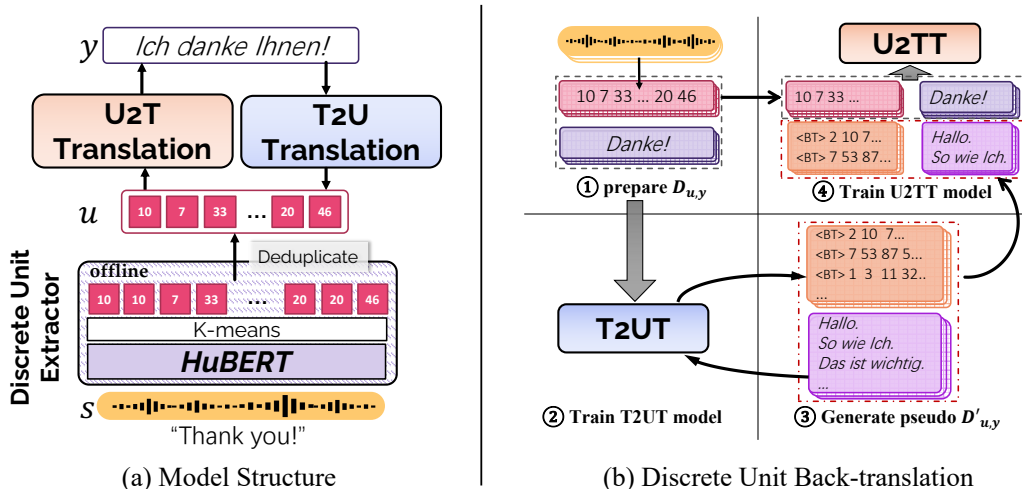


Figure 1: **Left:** The model structure of our approach. The offline discrete unit extractor converts speech into discrete units. The unit-to-text translation (U2TT) model translates the discrete units into the translation and the text-to-unit (U2TT) model does the opposite. **Right:** An illustration of the discrete unit back-translation (DUB) training procedure.

ST and MT, and to explore transferable training methods from NLP to speech.

3 Our Approach

3.1 Problem Formulation

Unlike cascade systems or existing end-to-end ST work that utilizes speech-transcription-translation triplet (s, x, y) , we aim to build and train the model that translates speech directly into text in another language without using the transcription x . The training dataset is denoted as $\mathcal{D}_{s,y} = \{(s, y)\}$. Also, we introduce the monolingual corpus of the target language $\mathcal{D}'_y = \{y'\}$, enhance the model via the discrete unit back-translation (DUB) method (described in Section 3.3).

3.2 Model Structure

As illustrated in Figure 1(a), our model consists of three main components: *discrete unit extractor*, *unit-to-text translation model*, and *text-to-unit translation model*.

Discrete Unit Extractor The discrete unit extractor converts continuous speech signals into a sequence of discrete units, which we use the Hidden-unit BERT (HuBERT) (Hsu et al., 2021). HuBERT is a self-supervised model learned by predicting discrete labels of masked audio segments from k-means clustering on the model’s intermediate representations. It consists of a stack of 1-D convolutional layers and a Transformer encoder to encode the speech into continuous intermediate representations, and a k-means model to convert the representations into a sequence of cluster in-

stances. We then remove the adjacent duplicate indices to obtain the discrete units sequence, denoted as $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_T)$, $\mathbf{u}_i \in \{0, 1, \dots, K - 1\}$, $\forall 1 \leq i \leq T$, where K is the number of clusters. Note that the discrete unit extractor used **offline** during the pre-processing stage before translation, can be considered as a **feature extractor**.

Unit-to-Text Translation (U2TT) Model The U2TT model $\theta_{u \rightarrow y}$ performs the forward translation. It consists of a discrete unit embedding layer and a Transformer. The discrete unit embedding layer converts discrete units \mathbf{u} into the embedding $\mathbf{e} = (\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_T)$. In order to retain more contextual and textual information from HuBERT, we adopt the intermediate representations of HuBERT’s k-means cluster centroids as prior knowledge to initialize the unit embedding. This initialization operation is referred to as *pre-trained embedding* in the later analysis (Section 6.1). The Transformer follows the vanilla Transformer architecture (Vaswani et al., 2017), consisting of a Transformer encoder and a Transformer decoder. The encoder takes unit embedding \mathbf{e} plus sinusoidal positional embedding as input and outputs semantic representation. The decoder generates the translation sequence $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{|y|})$ autoregressively based on the semantic representation.

Text-to-Unit Translation (T2UT) Model The T2UT model $\theta_{y \rightarrow u}$ has the same structure as the U2TT model, but with a randomly initialized text embedding layer. It is added to perform the text-to-unit translation and to incorporate the DUB training.

3.3 Discrete Unit Back-Translation (DUB)

Training Steps Given ST parallel dataset $\mathcal{D}_{s,y} = \{(s, y)\}$, extra target-language corpus $\mathcal{D}'_y = \{y'\}$, and the discrete unit extractor E . As shown in Figure 1(b), the DUB training steps are as follows.

1. Extract unit for each speech input $\mathbf{u} = E(s)$, and get unit-translation pairs $\mathcal{D}_{u,y} = \{(\mathbf{u}, y)\}$;
2. Train the T2UT model based on $\mathcal{D}_{u,y}$ with cross-entropy loss as in Eq. (2);
3. For each text $y' \in \mathcal{D}'_y$, generate corresponding synthetic units $\hat{\mathbf{u}}'$ through the BT model (generation methods will be discussed in Section 3.4). Then, add special <BT> indicator at the beginning of $\hat{\mathbf{u}}'$ (Caswell et al., 2019). The synthetic units-translation set is denoted as $\mathcal{D}'_{u,y} = \{(\hat{\mathbf{u}}', y')\}$;
4. Upsample the original training data by a rate of r and train U2TT model based on $\mathcal{D}_{u,y} \cup \mathcal{D}'_{u,y}$ and loss in Eq. (1).

Training Objective The training objectives for U2TT and T2UT models are the negative log-likelihood losses based on the unit-translation pairs:

$$\mathcal{L}_{\text{U2TT}} = -\mathbb{E}_{(u,y) \in D} \log P(y | \mathbf{u}, \theta_{u \rightarrow y}) \quad (1)$$

$$\mathcal{L}_{\text{T2UT}} = -\mathbb{E}_{(u,y) \in D_{u,y}} \log P(\mathbf{u} | y, \theta_{y \rightarrow u}) \quad (2)$$

, where D refers to $\mathcal{D}_{u,y} \cup \mathcal{D}'_{u,y}$ for DUB training, and when $D = \mathcal{D}_{u,y}$, Eq. (1) is the loss function for training U2TT from scratch.

3.4 Generation methods for back-translated units

We explore the following generation methods for producing synthetic units: beam search, sampling, and top-k sampling. We also apply a speech normalization method to remove speaker information when generating units.

Beam search tries to identify the maximum a posteriori (MAP) output and generate the sentence with the largest estimated probability given an input. **Sampling** means sampling from the distribution randomly at each step, which generates diverse outputs. **Top-k sampling** is a middle ground between beam search and sampling. At each time step, we select the k most likely tokens from the output distribution, re-normalize and then sample from this restricted set.

The discrete unit extractor produces various unit sequences for speech with the same content when delivered by multiple speakers (Lee et al., 2022). These variations pose a challenge for training the

BT model. In order to address this issue, we adopt a **Speech Normalization** module from (Lee et al., 2022), which removes speaker information from the discrete units and produces norm units. Specifically, it is an off-the-shelf HuBERT-CTC model trained on VoxPopuli (Wang et al., 2021a) that normalizes variant speech input to a single speaker to eliminate such influence (denoted as *Speech Norm*). We implement back-translation with norm units and use the resulting BT model to generate pseudo norm units.

4 Experiments

4.1 Datasets

MuST-C MuST-C¹ (Di Gangi et al., 2019a), one of the most widely-used ST benchmarks, contains translations from English to 8 languages collected from TED talks. We train and validate our approach in three ST directions: English-German (En-De), English-French (En-Fr), and English-Spanish (En-Es).

CoVoST-2 X-En CoVoST-2 (Wang et al., 2021b) is a multilingual ST corpus derived from the Common Voice project, which offers translations from English to 15 languages and from 21 languages to English. We conducted experiments on X-En, including high-resource languages (≥ 10 hours of speech) such as French (Fr) and German (De), and low-resource languages (≤ 2 hours of speech), like Arabic (Ar), Swedish (Sv), Japanese (Ja), etc. Without the need for transcription, the evaluation focuses on the capability of our method to generalize to the low-resource unwritten multi-languages.

Monolingual text corpus Monolingual target-language text corpora are introduced for back-translation. For MuST-C we include 48M German, 79M French and 64M Spanish sentences sampled from TED¹ (Duh, 2018), WMT¹ (Bojar et al., 2016) and CCMartix¹ (Schwenk et al., 2019) datasets for En-De/Fr/Es respectively. For CoVoST-2 X-En experiments, we introduce 1M extra English sentences sampled from the transcriptions of Common Voice Corpus 11.0¹ (Ardila et al., 2020).

All statistics of the datasets are in Appendix A.

4.2 Experimental setups

Pre-processing The model accepts 16-bit 16KHz mono-channel raw waveform speech and then discretizes them into units. We denote the discrete units clusters by the numbers (e.g. #1, #2), and

¹All released under CC BY NC ND 4.0 International

combining with the target-language sentences, we learn the joint vocabulary via SentencePiece (Kudo and Richardson, 2018). We set the joint vocabulary size to 8000 for both MuST-C and CoVoST-2.

Model Configuration For MuST-C experiments, we use the HuBERT-base² (pre-trained on Librispeech without fine-tuning) with a 500-cluster k-means quantizer based on the 9th layer representation as the *discrete unit extractor*. For CoVoST-2, we employ the mHuBERT³ with a 1000-cluster k-means quantizer based on the 11th layer representations pre-trained on the VoxPopuli (Wang et al., 2021a) speech in English, Spanish, and French. The U2TT and T2UT models have the same model architecture, consisting of a 12-layer Transformer encoder and a 6-layer Transformer decoder, with hidden size $d = 768$, 16 attention heads, and 4096 FFN hidden states. Additionally, we implement the BASE and LARGE versions of the model, with hidden sizes of 768 and 1024, respectively. Both versions are performed in the main experiments, while in the analysis, we primarily investigate the BASE model. More information on model size and scalability experiments can be found in Appendix D.

Evaluation We evaluate the models using case-sensitive sacreBLEU⁴ (Post, 2018) on the MuST-C $tst-COM$ sets and CoVoST-2 test sets.

See Appendix B for more details on vocabulary learning, training, and test.

4.3 Baseline models

We compare our method with the baselines as listed in Table 1 ~ 3 (Appendix C for details). In particular, we explain the following baselines that do not involve transcriptions during training.

Revisit ST (Zhang et al., 2022a) is a direct speech-to-translation model with parameterized distance penalty (PDP) and CTC regularization. Its framework and training objectives are sorely different from ours.

Unit-to-text Translation (U2TT) has the structure as described in Section 3.2 and is trained using only speech-translation supervision from the ST dataset from scratch, without applying DUB. As a

²https://dl.fbaipublicfiles.com/hubert/hubert_base_ls960.pt

³https://dl.fbaipublicfiles.com/hubert/mhubert_base_vp_en_es_fr_it3.pt

⁴<https://github.com/mjpost/sacrebleu>,
BLEU Signature: nrefs:1 | bs:1000 | seed:12345 | case:mixed | eff:no | tok:13a | smooth:exp | version:2.0.0

Method	De	Fr	Es	Avg.
<i>Methods that utilize transcriptions</i>				
Fairseq ST (Wang et al., 2020)	22.7	32.9	27.2	27.6
NeurST (Zhao et al., 2021)	22.8	33.3	27.4	27.8
Espnet ST (Inaguma et al., 2020)	22.9	32.8	28.0	27.9
E2E-ST-JT (Du et al., 2022)	23.1	32.8	27.5	27.8
Speechformer (Papi et al., 2021)	23.6	-	28.5	-
Cascaded (Inaguma et al., 2020)	23.6	33.8	28.7	28.7
MTL (Tang et al., 2021)	23.9	33.1	28.6	28.5
Self-training (Pino et al., 2020)	25.2	34.5	-	-
SpeechT5 (Ao et al., 2022)	25.2	35.3	-	-
<i>Methods that do not involve transcriptions</i>				
Revisit ST (Zhang et al., 2022a)	23.0	33.5	28.0	28.2
Transformer-ST	18.0	28.5	24.1	23.5
U2TT (BASE)	20.4	30.3	25.3	25.3
w/ DUB	25.8	34.7	30.2	30.2
U2TT (LARGE)	20.5	30.1	24.7	25.1
w/ DUB	26.2	35.3	30.4	30.6
<i>SoTA: use much more speech and various pre-training tasks</i>				
SpeechUT (Zhang et al., 2022c)*	30.1	41.4	33.6	35.0

Table 1: BLEU Scores on MuST-C En-X $tst-COM$ set. * is the state-of-the-art system, which designed various mask-predict pre-training tasks and trained using extra 1.4k hours of speech and parallel MT data from WMT. Random sampling is the decoding strategy for DUB.

baseline, comparison with this model helps to see the influence of DUB.

Transformer-ST stands for training the Speech-Transformer (Dong et al., 2018) from scratch, but without ASR pre-training as in the previous work (Wang et al., 2020; Inaguma et al., 2020; Zhao et al., 2021). The training details are in Appendix C.

4.4 Main results on Speech-to-text Translation

MuST-C As shown in Table 1, compared to the methods that do not involve the transcribed text, our method, U2TT (LARGE) with DUB, gets the best ST results by introducing extra target-language text, and DUB obtains an average boost of 5.5 BLEU compared with U2TT in the three En-X directions. Encouragingly, we find that our method achieves comparable performance to previous models that utilize transcriptions through multi-task learning or pre-training. As for the baseline, U2TT outperforms the Transformer-ST, where we believe that the discrete units still retain the semantic information of the audio feature (*e.g.* log Mel-filter bank, *abbr.* Fbank) for translation. As for the gap between our method and the SoTA system, we argue that SpeechUT (Zhang et al., 2022c) performed various mask-predict pre-training tasks using extra 1.4k hours of speech and parallel MT data from WMT, which is not included in our approach.

Methods	Aux. Data		Fr	De	Es	Ca	It	Ru	Zh	Avg.
	ASR	Text								
Transformer-ST [†]	-	-	4.3	8.4	12.0	14.4	0.2	1.2	1.4	8.8
Transformer-ST + ASR pre-train [†]	✓	-	26.3	17.1	23.0	18.8	11.3	14.8	5.8	16.7
Cascaded ST [†]	✓	-	27.6	21.0	27.4	21.3	13.5	16.8	7.0	19.2
Revisit ST (Zhang et al., 2022a)	-	-	26.9	14.1	15.7	17.2	2.4	3.6	2.0	11.7
U2TT (LARGE)	-	-	27.4	16.7	28.1	23.1	20.0	21.9	5.9	20.5
w/ DUB	-	✓1M	29.5	19.5	30.9	25.2	23.9	23.2	6.1	22.6

Table 2: Test BLEU scores on CoVoST-2 X-En language pairs with more than 10 hours of speech. Auxiliary data refers to all data at training excluding $\langle speech, translation \rangle$ pairs. [†]: Results from (Wang et al., 2021b). Random sampling is the decoding strategy for DUB.

Methods	Aux. Data		Ar 2h	Sv 2h	Lv 2h	Sl 2h	Ta 2h	Ja 1h	Id 1h	Cy 2h	Avg. -
	ASR	Text									
Transformer-ST [†]	-	-	0.3	0.2	0.1	0.3	0.3	0.3	0.4	0.3	0.3
Transformer-ST A2E [†]	-	-	0.6	0.6	0.4	1.2	0.1	0.2	0.3	2.6	0.8
Transformer-ST + ASR pre-train [†]	✓	-	4.3	2.7	2.5	3.0	0.3	1.5	2.5	2.7	2.4
<i>Larger models based on large-scale multilingual speech, text or joint pre-training, involving more data</i>											
XLS-R (0.4B)*	✓	-	8.1	5.3	3.1	5.3	0.0	2.0	3.3	3.4	3.8
Wav2seq (0.4B)*	✓	-	10.5	8.8	4.8	5.9	0.0	1.9	5.0	5.7	5.3
XLS-R + mBART-50 (0.7B) [◇]	✓	✓	3.0	10.3	6.0	6.6	0.2	0.6	1.4	2.5	3.8
LNA-E,D (0.7B) [▲]	✓	✓	3.7	5.9	4.6	4.6	0.7	1.7	2.9	2.8	3.4
U2TT (LARGE)	-	-	7.0	8.0	6.3	6.8	0.3	1.6	6.6	2.7	4.9
w/ DUB	-	✓1M	7.1	8.9	6.9	7.9	0.5	2.1	7.0	5.7	5.8

Table 3: Test BLEU scores on CoVoST-2 low-resource X-En language pairs with less than 2 hours of speech. [†]: results from (Wang et al., 2021b). *: results from (Wu et al., 2022) [◇]: results from (Babu et al., 2021). [▲]: results from (Li et al., 2021). The numbers in parentheses are their parameter sizes. Random sampling is the decoding strategy for DUB.

CoVoST-2 Our method performs similarly to MuST-C on the high-resource En-X (Table 2). Without considering auxiliary data or pre-training methods, adding only 1M additional English text, DUB improves by an average of 2.1 BLEU over 7 language pairs compared to U2TT, and by an average of 3.4 BLEU over the cascaded ST system. For the low-resource setting, our method can bring improvement on almost every language pair and achieve better performance than the large-scale multilingual speech or text pre-training models, like XLS-R+mBART-50 model (Babu et al., 2021), with much fewer parameters. The discrete unit extractor is unsupervised, so our method does not require transcriptions, which is particularly advantageous for unwritten ST. This experiment mimics such low-resource nature of unwritten languages in practice. The results also show that the U2TT model and the DUB training have the potential to translate low-resource unwritten languages.

Method	Aux. Data		BLEU
	Speech	Text	
Transformer-ST	-	-	18.0
w/ Cascaded BT	-	10M	20.3
U2TT	-	-	20.4
w/ Bimodal BART	10kh	10M	22.4
w/ DUB	-	10M	25.0
w/ Bimodal BART + DUB	10kh	10M	25.4

Table 4: MuST-C En-De t_{st-COM} BLEU scores for different methods that utilize 10M monolingual text data in ST. Transformer-ST and U2TT are described in Section 4.3.

5 Analysis on the Effect of Discrete Unit Back-translation (DUB)

5.1 Is DUB better than other methods that leverage extra raw data?

The key benefit of the DUB is to make use of a lot of monolingual text. Here, alternative techniques such as pseudo-labeling and pre-training (implemented as Cascaded BT and Bi-modal BART) are also evaluated on the MuST En-De translation, by introducing an equivalent corpus of 10 million German sentences.

- **Cascaded BT** aims to build a target-to-source MT-TTS pipeline to construct pseudo-speech translation augmented data for the training. Specifically, we use transcription-translation pairs of MuST-C to train a back-translation MT model and use the released FastSpeech2⁵ (Ren et al., 2020) and a HiFi-GAN (Kong et al., 2020) vocoder for TTS generation.
- **Bi-modal BART** has the same structure as U2TT, and is pre-trained by denoising large-scale corrupted discrete units and monolingual text, following the recipe of mBART (Liu et al., 2020). We combine the 10M additional text with 7M discrete units extracted from 10k hours of speech in GigaSpeech (Chen et al., 2021) to pre-train the model and fine-tune it based on MuST-C unit-translation pairs. See Appendix B for training details.

As shown in Table 4 introducing equivalent raw text, **DUB is superior to the above two approaches and has a greater potential to exploit monolingual raw text.** We find that the gain from *cascaded BT*-synthesized speech is limited because the synthetic speech is robotic and monotonic, making it easy to overfit the model to the synthetic pairs. Although the bi-modal BART pre-training can bring about 2 BLEU improvements, it is still inferior to DUB. We attribute this to the gap between the denoising pre-training task and the downstream generation tasks, while DUB does not have such a gap. Meanwhile, we observe that combination of bi-modal BART and DUB can bring further performance improvements, which indicates that they are complementary to each other.

5.2 The better the pseudo-unit, the more effective the DUB method?

In Section 3.3, we presented four generation methods to create synthetic pseudo-units based on the BT model, namely beam search, sampling, top-k sampling, and speech normalization. In the experiments, we set a beam size of 5 for the beam search, k=10 for top-k sampling, and use an off-the-shelf speech normalizer⁶ from Lee et al. (2022) for Speech Norm.

Does the forward model gain more from synthesized pairs when the synthesized units are of higher

⁵<https://github.com/ming024/FastSpeech2>

⁶https://dl.fbaipublicfiles.com/fairseq/speech_to_speech/speech_normalizer/en/en_10h.tar.gz

quality? We calculate the Unit Error Rate (UER) on the MuST-C validation set to assess the synthesis quality. A lower UER indicates that the generated units are closer to the directly extracted units, *i.e.* of higher quality. We systematically vary the back-translated data from 1M to 10M, and present the BLEU scores and UERs of the generation methods in Table 5 and Figure 2. The *Speech Norm* module produces the highest quality synthesized units, while the sampling-based methods have lower quality. Interestingly, the sampling method with the lowest synthesis quality has the most significant improvement over the forward model.

We conjecture that **the richness and irregularity of the synthesized data can better improve the forward ST model**, while regular pseudo-units, *e.g.* generated by MAP-based beam search, are more predictable and not conducive to performance improvement. This is consistent with previous findings of BT techniques in machine translation (Edunov et al., 2018). In addition, ***Speech Norm*, which normalizes speech to a single speaker, is not necessary for our DUB method.** Although such an operation makes the ST model easier to learn and the UER smaller, it compromises the diversity of the synthetic data, which is also not helpful for performance improvement. The model generalization ability weakens when these single-speaker synthesis units increase.

	UER(%)	Δ BLEU
Speech Norm	73.0	0.7
Beam Search	83.0	1.7
Top-10 Sampling	89.0	3.6
Sampling	92.0	4.6

Table 5: The quality of generated pseudo-units using different generation methods and their BLEU increases from 10M extra texts, evaluated by **Unit Error Rate (UER)** on MuST-C En-De_{Dev}, the smaller the better.

6 Why does DUB work? — Analysis on the Property of Discrete Unit

Why does DUB work? To answer this question, we examine the properties of the discretized speech unit. Specifically, (1) do the units make sense to replace the original speech input in the forward translation process? (2) Do the units generated by back translation also contain semantic information and can they even restore the speech?




Figure 2: MuST-C En-De t_{st-COM} BLEU scores for the increasing amounts of BT data (1M, 5M, and 10M), obtained using different generation methods. Note that the unit extractor of *Speech Norm* is mHuBERT, so without extra text, the BLEU score is 20.9, which is slightly higher than HuBERT, but the gap is quite small.

6.1 Are discrete units suitable features for ST input?

We show how much semantic information is retained for different input forms by comparing the results of the downstream ST task (shown in Table 6). Training from scratch, we find that the U2TT model translates better than Transformer-ST (19.9 vs. 18.0), indicating that **compared to the speech feature, like Fbank, the discrete unit is a better choice for input and has no information loss in terms of the semantic information required for translation.** We assume that this is strongly correlated with the HuBERT-based discrete unit extractor, since HuBERT is designed to learn a combined acoustic and language model over the continuous speech input, which preserves much textual information for the speech. But rigorously, compared to the continuous representation of HuBERT, the discretization procedure does suffer from semantic information loss. Comparing Line II and III, there is a gap of 2.9 BLEU between U2TT and HuBERT-Transformer (where frozen HuBERT Layer-9 continuous representation is taken as input to perform ST), in terms of ST metrics. Fortunately, the gap can be compensated by *a*) initializing the unit embedding as its corresponding K-means cluster centroid on continuous HuBERT representations as described in Section 3.2 (denoted as pre-trained embedding, Line IV), which can slightly close the gap by 0.5 BLEU; and *b*) simply introducing only 1M additional text and applying DUB, which can achieve 2.7 BLEU improvement (Line V vs. IV).

No.	Methods	BLEU
I	Transformer-ST	18.0
II	HuBERT + Transformer	22.8
III	U2TT	19.9
IV	\hookrightarrow w/ pre-trained embedding	20.4*
V	\hookrightarrow w/ DUB-1M	23.1

Table 6: BLEU scores on MuST-C En-De t_{st-COM} set. DUB-1M means to introduce 1M unpaired text via the DUB method. *: the improvement over U2TT is statistically significant ($p < 0.05$)

6.2 Can we recover faithful speech from pseudo-units?

Do the back-translated units capture the semantics of the target language text? Since it is difficult to directly evaluate the correctness of the pseudo-units generated by the back-translation model, we concatenate a unit-based HiFi-GAN vocoder⁷ with our back-translation model to recover speech from the generated pseudo-units, thus completing the text-to-speech (TTS) translation task. TTS generation quality is measured by ASR-BLEURT, where we transcribe the speech output using a high-quality open-source ASR model⁸ and calculate BLEURT⁹ with reference transcription. As shown in Table 7, the ASR-BLEURT of beam search and sampling is 0.6 and 0.47 respectively, indicating that **the unit sequence back-translated from a given target language text can convey its general semantic meaning, which is the guarantee for the success of DUB.** We conduct the listening test by checking 30 randomly sampled BT-recovered speeches for semantic consistency with ground-truth. 22 of 30 sentences matched ground-truth speech, while the remaining 8 had minor issues, with only 1 being of low quality and the other 7 missing or repeating 1-2 details. We also provide some generated audio samples in Appendix F to help illustrate the degree of speech restoration.

7 Conclusion

In this paper, we propose **Discrete Unit Back-translation (DUB)**, as well as the **Unit-to-Text**

⁷https://dl.fbaipublicfiles.com/fairseq/speech_to_speech/vocoder/code_hifigan/mhubert_vp_en_es_fr_it3_400k_layer11_kml000_lj/g_00500000

⁸<https://huggingface.co/facebook/wav2vec2-large-960h-1v60-self>

⁹<https://github.com/google-research/bleurt>

Generation Method	ASR-BLEURT
beam search	0.60
sampling	0.47

Table 7: MuST-C En-De $t_{st}-COM$ ASR-BLEURT for text-to-speech translation

Translation (U2TT) model for direct speech translation. Our approach successfully migrates the back-translation technique from MT to ST by discretizing the speech signals into unit sequences and making use of extra widely accessible text in the target language. Without using transcription, DUB can achieve an average increase of 5.5 BLEU on MuSTC En-De/Fr/Es over the raw U2TT framework, and achieves comparable performance to the large-scale speech-text joint pre-training models on CoVoST-2 low-resource ST. The analysis experiments also show the potential of such discrete audio units as inputs and outputs for text or speech generation tasks.

Broader Impact

Our proposed model structure with a discrete unit extractor for speech and the unit-to-text translation model, which does not need any transcriptions during training, is particularly relevant for speech translation for more than 3,000 languages and dialects in the world that cannot be transcribed. Since these unwritten languages are typically low-resource, we emphasize that boosting ST performance via text-to-unit back-translation data augmentation, *i.e.* DUB, is very promising. Meanwhile, as a by-product of DUB, TTS translation has significant implications for assisting visually impaired or dyslexic people in understanding the world as well as preserving low-resource unwritten spoken languages.

However, as exploratory work, we focus on investigating the potential of using BT to enhance ST performance, while popular large-scale pre-training methods are not employed in this paper. This makes our method slightly inferior to these methods in terms of performance, perhaps. But promisingly, in terms of structure, the model is more general across various modalities and also has more potential to integrate with the methods in NLP area (might be the topic of future research). Also, the models are still far from real industrial applications. For example, the data used for training is much smaller than the scale in reality, while the

real speech is noisier and more complex than the open-source dataset, which may require front-end processing. Moreover, the success of our method is partly attributable to the HuBERT representation, which contains certain textual information for the speech, and via experiments, we also find that the quality of discrete units influences the translation performance. Nevertheless, learning meaningful discrete units is not the primary goal of HuBERT pre-training, and how to learn discrete units or representations for speech with more *contextual* semantic information can be explored in the future.

Acknowledgements

We thank Fuliang Weng for the careful guidance and revisions to the paper and thank all the anonymous reviewers for their insightful and valuable comments.

References

- David Ifeoluwa Adelani, Md Mahfuz Ibn Alam, Antonios Anastasopoulos, Akshita Bhagia, Marta Costa-Jussá, Jesse Dodge, Fahim Faisal, Christian Federmann, Natalia Fedorova, Francisco Guzmán, et al. 2022. Findings of the wmt 2022 shared task on large-scale machine translation evaluation for african languages. In *Proc. of WMT*.
- Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang, et al. 2022. Specht5: Unified-modal encoder-decoder pre-training for spoken language processing. In *Proc. of ACL*, pages 5723–5738.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. **Common voice: A massively-multilingual speech corpus**. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4218–4222, Marseille, France. European Language Resources Association.
- Arun Babu, Changan 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, Steffen Schneider, and Michael Auli. 2019. vq-wav2vec: Self-supervised learning of discrete speech representations. In *Proc. of ICLR*.
- 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.

- Alexandre Bérard, Olivier Pietquin, Laurent Besacier, and Christophe Servan. 2016. Listen and translate: A proof of concept for end-to-end speech-to-text translation. In *NIPS Workshop on end-to-end learning for speech and audio processing*.
- Nicola Bertoldi and Marcello Federico. 2009. Domain adaptation for statistical machine translation with monolingual resources. In *Proceedings of the fourth workshop on statistical machine translation*, pages 182–189.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurelie Neveol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Ondřej Bojar and Aleš Tamchyna. 2011. Improving translation model by monolingual data. In *Proceedings of the sixth workshop on statistical machine translation*, pages 330–336.
- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. Tagged back-translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63.
- Guoguo Chen, Shuzhou Chai, Guanbo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, et al. 2021. Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio. *arXiv preprint arXiv:2106.06909*.
- Peng-Jen Chen, Ann Lee, Changhan Wang, Naman Goyal, Angela Fan, Mary Williamson, and Jiatao Gu. 2020. Facebook ai’s wmt20 news translation task submission. In *Proceedings of the Fifth Conference on Machine Translation*, pages 113–125.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Xuxin Cheng, Qianqian Dong, Fengpeng Yue, Tom Ko, Mingxuan Wang, and Yuexian Zou. 2023. M 3 st: Mix at three levels for speech translation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Yu-An Chung, Yu Zhang, Wei Han, Chung-Cheng Chiu, James Qin, Ruoming Pang, and Yonghui Wu. 2021. W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 244–250. IEEE.
- Mattia A Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019a. Must-c: a multilingual speech translation corpus. In *Proc. of the NAACL-HLT*, pages 2012–2017. Association for Computational Linguistics.
- Mattia A Di Gangi, Matteo Negri, and Marco Turchi. 2019b. Adapting transformer to end-to-end spoken language translation. In *Proc. of INTERSPEECH*, pages 1133–1137. International Speech Communication Association (ISCA).
- Linhao Dong, Shuang Xu, and Bo Xu. 2018. Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition. In *Proc. of ICASSP*, pages 5884–5888. IEEE.
- Qianqian Dong, Mingxuan Wang, Hao Zhou, Shuang Xu, Bo Xu, and Lei Li. 2021a. Consecutive decoding for speech-to-text translation. In *Proc. of AAAI*, pages 12738–12748.
- Qianqian Dong, Rong Ye, Mingxuan Wang, Hao Zhou, Shuang Xu, Bo Xu, and Lei Li. 2021b. Listen, understand and translate: Triple supervision decouples end-to-end speech-to-text translation. In *Proc. of AAAI*, pages 12749–12759.
- Yichao Du, Zhirui Zhang, Weizhi Wang, Boxing Chen, Jun Xie, and Tong Xu. 2022. Regularizing end-to-end speech translation with triangular decomposition agreement. In *Proc. of AAAI*, pages 10590–10598.
- Kevin Duh. 2018. The multitarget ted talks task. <http://www.cs.jhu.edu/~kevinduh/a/multitarget-tedtalks/>.
- Ewan Dunbar, Robin Algayres, Julien Karadayi, Mathieu Bernard, Juan Benjumea, Xuan-Nga Cao, Lucie Miskic, Charlotte Dugrain, Lucas Ondel, Alan Black, et al. 2019. The zero resource speech challenge 2019: Tts without t. In *Proc. of INTERSPEECH*.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. In *Proc. of EMNLP*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- Qingkai Fang, Rong Ye, Lei Li, Yang Feng, and Mingxuan Wang. 2022. Stem: Self-learning with speech-text manifold mixup for speech translation. In *Proc. of ACL*, pages 7050–7062.
- Akhbardeh Farhad, Arkhangorodsky Arkady, Biesialska Magdalena, Bojar Ondřej, Chatterjee Rajen, Chaudhary Vishrav, Marta R Costa-jussa, España-Bonet Cristina, Fan Angela, Federmann Christian, et al. 2021. Findings of the 2021 conference on machine translation (wmt21). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88. Association for Computational Linguistics.

- Chi Han, Mingxuan Wang, Heng Ji, and Lei Li. 2021. [Learning shared semantic space for speech-to-text translation](#). In *Proc. of ACL - Findings*, pages 2214–2225, Online. Association for Computational Linguistics.
- Tomoki Hayashi, Shinji Watanabe, Yu Zhang, Tomoki Toda, Takaaki Hori, Ramon Astudillo, and Kazuya Takeda. 2018. Back-translation-style data augmentation for end-to-end asr. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 426–433. IEEE.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative back-translation for neural machine translation. In *Proceedings of the 2nd workshop on neural machine translation and generation*, pages 18–24.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhota, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Kenji Imamura, Atsushi Fujita, and Eiichiro Sumita. 2018. Enhancement of encoder and attention using target monolingual corpora in neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 55–63.
- Hirofumi Inaguma, Shun Kiyono, Kevin Duh, Shigeki Karita, Nelson Yalta, Tomoki Hayashi, and Shinji Watanabe. 2020. ESPnet-ST: All-in-one speech translation toolkit. In *Proc. of ACL*, pages 302–311.
- Hirofumi Inaguma, Sravya Popuri, Iliia Kulikov, Peng-Jen Chen, Changhan Wang, Yu-An Chung, Yun Tang, Ann Lee, Shinji Watanabe, and Juan Pino. 2022. Unity: Two-pass direct speech-to-speech translation with discrete units. *arXiv preprint arXiv:2212.08055*.
- Sathish Indurthi, Mohd Abbas Zaidi, Nikhil Kumar Lakumarapu, Beomseok Lee, Hyejung Han, Seokchan Ahn, Sangha Kim, Chanwoo Kim, and Inchul Hwang. 2021. Task aware multi-task learning for speech to text tasks. In *Proc. of ICASSP*, pages 7723–7727. IEEE.
- Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhota, Tu Anh Nguyen, Morgane Riviere, Abdelrahman Mohamed, Emmanuel Dupoux, et al. 2022. Text-free prosody-aware generative spoken language modeling. In *Proc. of ACL*, pages 8666–8681.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proc. of EMNLP*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Kushal Lakhota, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, et al. 2021. On generative spoken language modeling from raw audio. *TACL*, 9:1336–1354.
- Ann Lee, Hongyu Gong, Paul-Ambroise Duquenne, Holger Schwenk, Peng-Jen Chen, Changhan Wang, Sravya Popuri, Juan Pino, Jiatao Gu, and Wei-Ning Hsu. 2022. Textless speech-to-speech translation on real data. In *Proc. of the NAACL-HLT*. Association for Computational Linguistics.
- Xian Li, Changhan Wang, Yun Tang, Chau Tran, Yuqing Tang, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2021. Multilingual speech translation from efficient finetuning of pretrained models. In *Proc. of ACL*, pages 827–838.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *TACL*, 8:726–742.
- Yuchen Liu, Hao Xiong, Jiajun Zhang, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2019. End-to-end speech translation with knowledge distillation. *Proc. of INTERSPEECH*, pages 1128–1132.
- Chutong Meng, Junyi Ao, Tom Ko, Mingxuan Wang, and Haizhou Li. 2022. Cobert: Self-supervised speech representation learning through code representation learning. *arXiv preprint arXiv:2210.04062*.
- Xuan-Phi Nguyen, Sravya Popuri, Changhan Wang, Yun Tang, Iliia Kulikov, and Hongyu Gong. 2022. [Improving speech-to-speech translation through unlabeled text](#).
- 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. *arXiv preprint arXiv:1904.01038*.
- Sara Papi, Marco Gaido, Matteo Negri, and Marco Turchi. 2021. Speechformer: Reducing information loss in direct speech translation. In *Proc. of EMNLP*.
- Juan Pino, Qiantong Xu, Xutai Ma, Mohammad Javad Dousti, and Yun Tang. 2020. Self-training for end-to-end speech translation. In *Proc. of INTERSPEECH*, pages 1476–1480.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191.
- Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2020. Fastspeech 2: Fast and high-quality end-to-end text to speech. *arXiv preprint arXiv:2006.04558*.

- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, and Armand Joulin. 2019. Ccmatrix: Mining billions of high-quality parallel sentences on the web. *arXiv preprint arXiv:1911.04944*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proc. of ACL*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Amitay Sicherman and Yossi Adi. 2023. Analysing discrete self supervised speech representation for spoken language modeling. *arXiv preprint arXiv:2301.00591*.
- Tzu-Wei Sung, Jun-You Liu, Hung-yi Lee, and Lin-shan Lee. 2019. Towards end-to-end speech-to-text translation with two-pass decoding. In *Proc. of ICASSP*, pages 7175–7179. IEEE.
- Yun Tang, Juan Pino, Xian Li, Changhan Wang, and Dmitriy Genzel. 2021. Improving speech translation by understanding and learning from the auxiliary text translation task. In *Proc. of ACL*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. 2017. Listening while speaking: Speech chain by deep learning. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 301–308. IEEE.
- Sei Ueno, Masato Mimura, Shinsuke Sakai, and Tatsuya Kawahara. 2021. Data augmentation for asr using tts via a discrete representation. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 68–75. IEEE.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Changhan Wang, Morgane Riviere, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux. 2021a. Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation. In *Proc. of ACL*, pages 993–1003.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020. [Fairseq S2T: Fast speech-to-text modeling with fairseq](#). In *Proc. of AACL*, pages 33–39, Suzhou, China. Association for Computational Linguistics.
- Changhan Wang, Anne Wu, Jiatao Gu, and Juan Pino. 2021b. Covost 2 and massively multilingual speech translation. In *Proc. of INTERSPEECH*, pages 2247–2251.
- Guillaume Wenzek, Vishrav Chaudhary, Angela Fan, Sahir Gomez, Naman Goyal, Somya Jain, Douwe Kiela, Tristan Thrush, and Francisco Guzmán. 2021. Findings of the wmt 2021 shared task on large-scale multilingual machine translation. In *Proceedings of the Sixth Conference on Machine Translation*, pages 89–99.
- Felix Wu, Kwangyoun Kim, Shinji Watanabe, Kyu Han, Ryan McDonald, Kilian Q Weinberger, and Yoav Artzi. 2022. Wav2seq: Pre-training speech-to-text encoder-decoder models using pseudo languages. *arXiv preprint arXiv:2205.01086*.
- Mengzhou Xia, Xiang Kong, Antonios Anastasopoulos, and Graham Neubig. 2019. Generalized data augmentation for low-resource translation. *arXiv preprint arXiv:1906.03785*.
- Rong Ye, Mingxuan Wang, and Lei Li. 2021. End-to-end speech translation via cross-modal progressive training. In *Proc. of INTERSPEECH*.
- Rong Ye, Mingxuan Wang, and Lei Li. 2022. Cross-modal contrastive learning for speech translation. In *Proc. of NAACL-HLT*.
- Biao Zhang, Barry Haddow, and Rico Sennrich. 2022a. Revisiting end-to-end speech-to-text translation from scratch. In *Proc. of ICML*.
- Chen Zhang, Xu Tan, Yi Ren, Tao Qin, Kejun Zhang, and Tie-Yan Liu. 2021. Uwspeech: Speech to speech translation for unwritten languages. In *Proc. of the AAAI*, pages 14319–14327.
- Weitai Zhang, Zhongyi Ye, Haitao Tang, Xiaoxi Li, Xinyuan Zhou, Jing Yang, Jianwei Cui, Pan Deng, Mohan Shi, Yifan Song, et al. 2022b. The ustcnslip offline speech translation systems for iwslt 2022. In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 198–207.
- Ziqiang Zhang, Long Zhou, Junyi Ao, Shujie Liu, Lirong Dai, Jinyu Li, and Furu Wei. 2022c. Speechut: Bridging speech and text with hidden-unit for encoder-decoder based speech-text pre-training. In *Proc. of EMNLP*.
- Chengqi Zhao, Mingxuan Wang, Qianqian Dong, Rong Ye, and Lei Li. 2021. NeurST: Neural speech translation toolkit. In *Proc. of ACL - System Demonstrations*.
- Renjie Zheng, Junkun Chen, Mingbo Ma, and Liang Huang. 2021. [Fused acoustic and text encoding for multimodal bilingual pretraining and speech translation](#). In *Proc. of ICML*.

A Data Statistics

En→	Hours	Samples
De	408	234K
Fr	492	280K
Es	504	270K

Table 8: Statistics of MuST-C dataset

Languages	Code	Hours	Samples
English	En	-	-
French	Fr	264	207374
German	De	184	127834
Spanish	Es	113	79015
Catalan	Ca	136	95854
Italian	It	44	31698
Russian	Ru	18	12112
Chinese	Zh	10	7085
Arabic	Ar	2	2283
Swedish	Sv	2	2160
Latvian	Lv	2	2337
Slovenian	Sl	2	1843
Tamil	Ta	2	1358
Japanese	Ja	1	1119
Indonesian	Id	1	1243
Welsh	Cy	2	1241

Table 9: Statistics of CoVoST-2 X-En involved in this paper.

Lang.	TED	WMT [†]	CCMatrix [*]	Sum
De	0.2M	4.5M	43M	48M
Fr	0.2M	39M	50M	79M
Es	-	13M	61M	64M

Table 10: Statistics of monolingual data for MuST-C experiments. [†]: De from WMT16 (Bojar et al., 2016), Fr from WMT14, Es from WMT13, ^{*}: randomly sample from CCMatrix (Schwenk et al., 2019).

B Experimental Details

Vocabulary We apply the SentencePiece¹⁰ (Kudo and Richardson, 2018) to tokenize the text and discrete units into subwords. We add all the discrete units as special symbols to the joint vocabulary.

¹⁰<https://github.com/google/sentencepiece>

The joint subword tokenizer is learned on all the translation sentences and discrete unit sequences in the ST training set. The vocabulary size is 8000 for both MuST-C and CoVoST-2 experiments. Specifically, for MuST-C experiments, since the number of K-means clusters is 500, the vocabulary is composed of 500 special unit symbols and 7500 text subwords. For CoVoST2 X-En experiments, the vocabulary consists of 1000 special unit symbols representing 1000 clusters of mHuBERT, and 7000 text subwords.

Training details We use Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.98$, and 4k warm-up updates to optimize the parameters in our model. We train the model with a batch size of 5k tokens. The learning rate is $7e-4$ and we apply an inverse square root schedule. The value of label smoothing is set to 0.1. The up-sampling rate r in DUB is set to 32, given the huge volume differences between the BT data and the original data. For MuST-C experiments, we train U2TT and T2UT models of each translation direction under bilingual settings. For CoVoST-2 X-En experiments, we train a multi-lingual X-En model covering 21 translation directions, distinguished by the language tags of the units in different languages. We implement our models based on Fairseq¹¹ (Ott et al., 2019) codebase. All models are trained on 8 Nvidia Tesla-V100 GPUs and take about 400k steps to converge. During inference, We save the checkpoint with the best BLEU on the validation set and average the last 10 checkpoints. We use beam search with a beam size of 5 for each translation direction.

Training details for Bi-modal BART The training of bi-modal BART follows the recipe of mBART (Liu et al., 2020). We implemented a mask rate of 0.3, with the replacement of the masked tokens by random tokens at a probability of 0.1. Additionally, the mask length was determined through sampling from a Poisson distribution, with a lambda parameter of 3.5.

C Baseline Models

Existing ST Systems We list the ST systems we compared with on different datasets:

- **MuST-C** In Table 1, we compare our method with the following: Fairseq ST (Wang et al., 2020), NeurST (Zhao et al., 2021), Espnet ST (Inaguma et al., 2020), E2E-ST-JT (Du et al.,

¹¹<https://github.com/facebookresearch/fairseq>

2022), Speechformer (Papi et al., 2021), Cascaded ST (Inaguma et al., 2020), MTL (Tang et al., 2021), Self-training (Pino et al., 2020), SpeechT5 (Ao et al., 2022), SpeechUT (Zhang et al., 2022c) and Revisit ST (Zhang et al., 2022a).

- **CoVoST X-En high resource** In Table 2, we compare our method with several baselines from (Wang et al., 2021b), including Transformer-ST, Transformer-ST + ASR pre-train and Cascaded ST, and Revisit ST (Zhang et al., 2022a).
- **CoVoST X-En low resource** In Table 3, we compare our method with several existing ST methods, including Transformer-ST, Transformer-ST + ASR pre-train from (Wang et al., 2021b) and large-scale multilingual speech or text pre-training methods: XLS-R (Wu et al., 2022), Wav2seq (Wu et al., 2022), XLS-R+mBART-50 (Babu et al., 2021), LNA-E,D (Li et al., 2021). Note that XLS-R is pre-trained on 436K hours of speech across 128 languages.

Transformer-ST for MuST-C For a fair comparison, we keep the parameters roughly the same size as DUB, setting two convolutional layers, a 12-layer Transformer encoder and a 6-layer Transformer decoder, with hidden size $d = 768$, 16 attention heads, and 4096 FFN hidden states, which makes the model size larger than baselines like Fairseq ST (Wang et al., 2020), NeurST (Zhao et al., 2021), and Espnet ST (Inaguma et al., 2020).

D Scalability

How does model size affect the results of our method? How much improvement does the raw text in the target language bring to our method? To this end, we take MuST-C English-German translation as an example. We set the model size to 73M, 176M and 260M parameters respectively (the specific hyperparameter settings are shown in Table 11), and introduce extra 1M, 10M, and 48M German sentences.

Figure 3 shows the BLEU scores of different sizes of models, with different amounts of monolingual back-translation data added. In general, regardless of the model size, introducing **more text brings better performance**. When we introduce a large amount of back-translated data, **the larger model gets significantly better performance**. We find that when no or less back-translated data is introduced, the performance of the large model is instead not optimal. This is because the large model

is prone to overfitting when the original training data is small, but as the monolingual data is gradually introduced, the advantage of the large model becomes obvious, without relying on the transcription, introducing 48M back-translated pairs, the model with 260M parameters can boost up to 6.1 BLEU on En-De.

	Model Params	Encoder Layers	Decoder Layers	Hidden Dim
① SMALL	73M	6	6	512
② BASE	176M	12	6	768
③ LARGE	260M	12	6	1024

Table 11: Hyper-parameter settings for the models in Figure 3.




Figure 3: MuST-C En-De $t_{\text{test-COM}}$ BLEU scores for different amounts of target-language text data are introduced, under different model sizes.

E Comparison With Cascaded System

It could be argued that our model employs a cascaded architecture, comprising a unit extractor and a unit-to-text translation model. The traditional cascade ST system (ASR+MT) can also be enhanced through applying back-translation to improve its MT model.

In Table 12, we compare the performance of DUB with the BT-enhanced cascaded ST system both utilizing 10M unpaired text. By comparison, we can find that the BLEU score of the U2TT model is inferior to that of the cascaded system when utilizing 10 million unpaired text samples. This discrepancy can likely be attributed to the higher baseline performance of the cascaded system. Additionally, DUB demonstrates a superior relative improvement in BLEU score compared to the cascaded system. Moreover, the discrete unit

Method	Extra Text	BLEU	Δ BLEU
U2TT	-	20.4	-
w/ DUB	✓	25.0	4.6
Cascaded ST	-	23.1	-
w/ MT-BT	✓	26.0	2.9

Table 12: MuST-C En-De t_{st-COM} BLEU for U2TT and cascaded system. MT-BT refers to enhancing MT model of cascaded ST through back-translation. Cascaded ST is trained by ourselves on MuST-C En-De t_{train} set. The same extra text corpus with 10 million German sentences is used for both methods.

extractor is obtained through unsupervised training on unlabeled speech, which requires no transcriptions compared with the ASR system trained on speech-transcription pairs.

F Cases of Text-to-Speech Translation

In Table 13, we show two cases of German-English text-to-speech translation on MuST-C En-DE t_{ST-COM} set. In CASE 1, our text-to-speech translation system generates speech with the same content and a similar spectrogram as reference speech. In CASE 2, the synthetic speech deviated slightly from the reference speech, but the translation is correct — “release” has the same meaning as “shoveling out” and “all the time” just means “all along”. The samples of generated audio are included in https://anonymous.4open.science/r/DUB/ttss_samples.

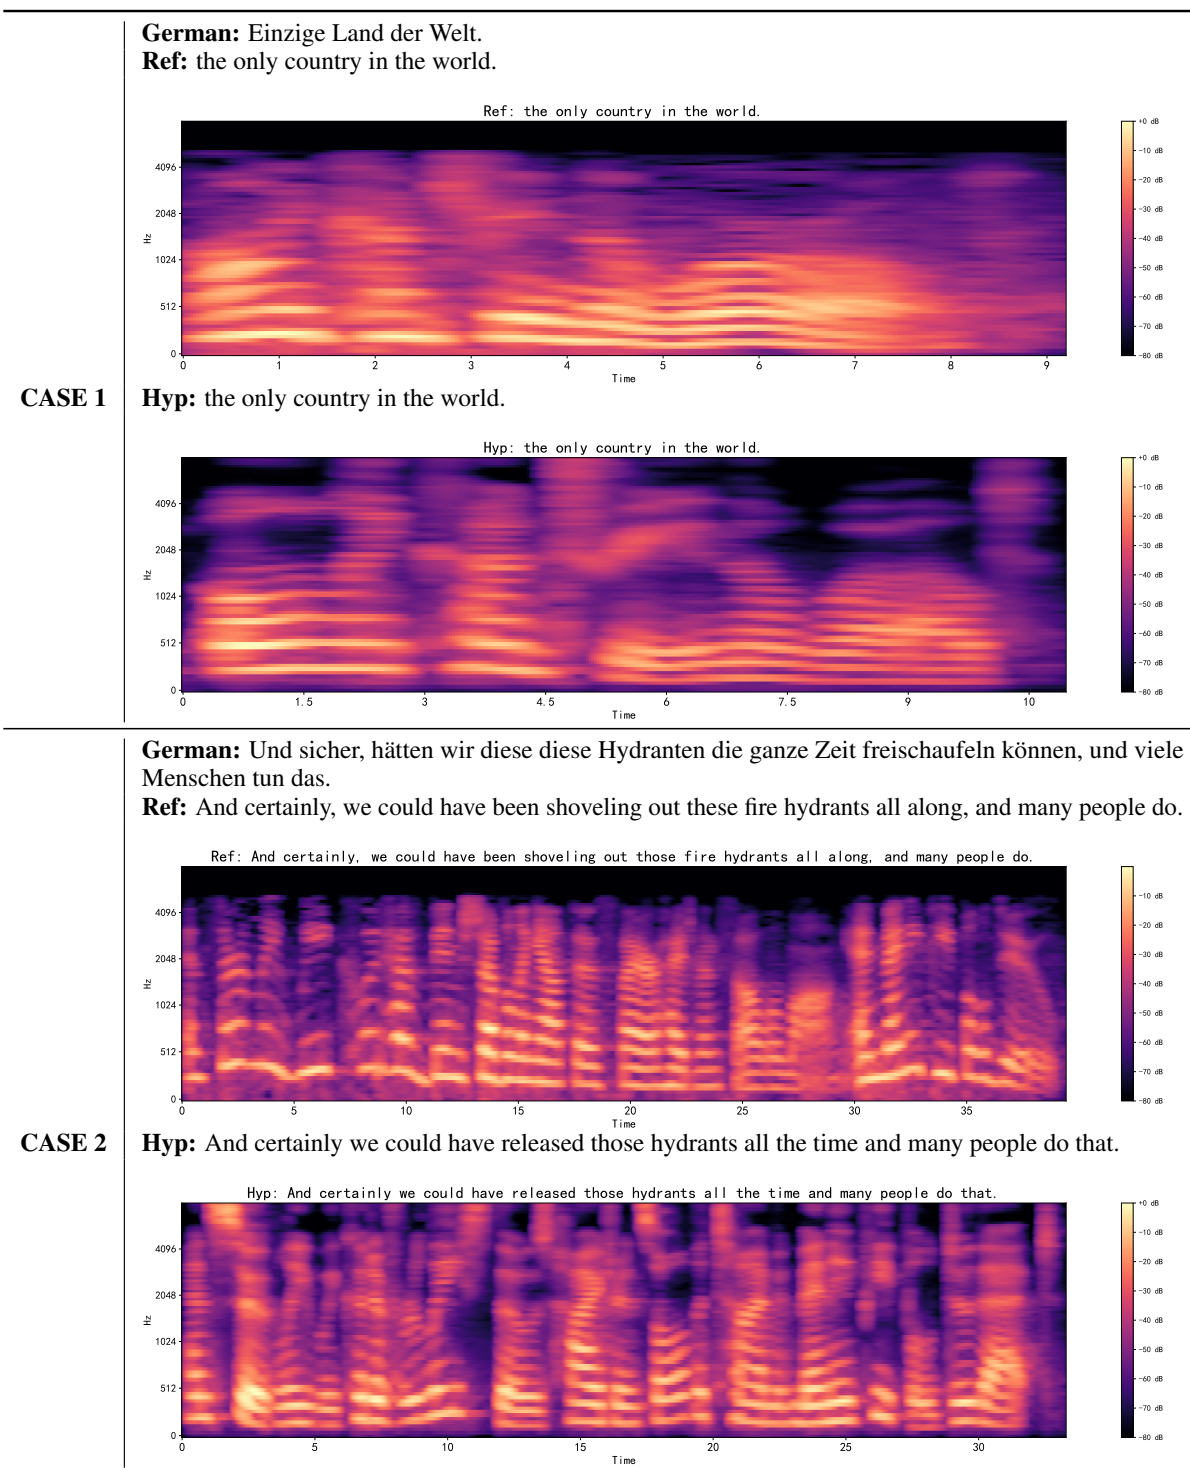


Table 13: Cases of De-En text-to-speech translation on MuST-C En-De t_{st} -COM set.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Broader Impact
- A2. Did you discuss any potential risks of your work?
Broader Impact
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

Section4 & section5 & section6 & appendix D & appendix E

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix B

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Appendix B

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 4.2

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 4.2 & appendix B

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.