

KIT’s Multilingual Speech Translation System for IWSLT 2023

Danni Liu, Thai Binh Nguyen, Sai Koneru, Enes Yavuz Ugan, Ngoc-Quan Pham, Tuan-Nam Nguyen, Tu Anh Dinh, Carlos Mullov, Alexander Waibel, Jan Niehues

Karlsruhe Institute of Technology
firstname.lastname@kit.edu

Abstract

Many existing speech translation benchmarks focus on native-English speech in high-quality recording conditions, which often do not match the conditions in real-life use-cases. In this paper, we describe our speech translation system for the multilingual track of IWSLT 2023, which focuses on the translation of scientific conference talks. The test condition features accented input speech and terminology-dense contents. The tasks requires translation into 10 languages of varying amounts of resources. In absence of training data from the target domain, we use a retrieval-based approach (k NN-MT) for effective adaptation (+0.8 BLEU for speech translation). We also use adapters to easily integrate incremental training data from data augmentation, and show that it matches the performance of re-training. We observe that cascaded systems are more easily adaptable towards specific target domains, due to their separate modules. Our cascaded speech system outperforms its end-to-end counterpart on scientific talk translation, although their performance remains similar on TED talks.

1 Introduction

This paper summarizes Karlsruhe Institute of Technology’s speech translation system for the multilingual track of IWSLT 2023 (Agarwal et al., 2023). In this track, the task is to translate scientific talks in English into 10 languages: Arabic (ar), Chinese (zh), Dutch (nl), French (fr), German (de), Japanese (ja), Persian/Farsi (fa), Portuguese (pt), Russian (ru), Turkish (tr). The talks are from presentations in the 60th Annual Meeting of the Association for Computational Linguistics (ACL 2022).

Translating scientific talks presents several challenges. On the *source* side, most speakers are non-native, and the recording conditions often vary. This requires acoustic robustness to accents and noise. On the *target* side, domain-specific terminologies are frequently used, calling for accurate

translation of these words that rarely occur in the training data. The styles of the talks, e.g. formality, also differ from other domains. As no training data from the same domain is provided, effective few-shot or zero-shot adaptation is crucial.

As the task focuses on *one-to-many* translation, it is also an interesting testbed for whether multilinguality improves speech translation quality. For text-to-text translation, the gain from multilinguality is mostly concentrated in *many-to-one* translation (Aharoni et al., 2019; Fan et al., 2021), i.e., multilinguality on the source side. In contrast, for *X-to-many* translation, it remains unclear whether incorporating *more target languages* improves translation quality.

In this system description paper, we present cascaded and end-to-end systems for the English-to-many speech translation task. We leverage pretrained models, including WavLM (Chen et al., 2022), mBART50 (Tang et al., 2020), and DeltaLM (Ma et al., 2021). The systems do not use additional data beyond the allowed corpora, and therefore fall under the *constrained* data condition.

For the cascaded system, to handle the unique style of scientific talks, we use k NN-MT (Khandelwal et al., 2021) to bias the output generation towards the target domain. Moreover, as no target monolingual data is provided, we use data diversification (Nguyen et al., 2020) to enrich the existing parallel data. We also use adapters (Rebuffi et al., 2017; Bapna and Firat, 2019) as a lightweight approach for incremental learning and language adaptation. For the ASR model, we improve over last year’s performance by using a more recent audio encoder (Chen et al., 2022) and adding a dedicated decoder. To adapt the ASR system to the target domain, we use n -gram re-weighting and synthesized data for the target domain. For the end-to-end system, we use our machine translation model for knowledge distillation. We also ensemble models trained with and without synthesized speech data.

Our main findings are as follow:

- For cascaded ST systems, we can effectively adapt the model towards a target domain/style using k NN-MT (Khandelwal et al., 2021). A datastore as small as a few hundred sentence pairs was sufficient for achieving consistent gains (avg. +0.8 BLEU over 10 languages).
- Besides the common use-case of adding language-specific capacity, adapters (Bapna and Firat, 2019) is also an effective method when subsequently adding training data. Empirically, we show it matches the performance of re-training on all new data.
- For ASR, lexical constraints for domain adaptation are more easily integrated in CTC models. For encoder-decoder model, the control could be achieved by TTS-synthesized source speech, but it requires more careful tuning.

2 Data and Preprocessing

After describing the evaluation data (§2.1), we outline the training data and preprocessing steps for our automatic speech recognition (ASR; §2.2), machine translation (MT; §2.3), casing/punctuation restoration (§2.4), and speech translation (ST; §2.5) models.

2.1 Development and Test Data

In the multilingual track, the testing condition is scientific conference talks. Therefore, we primarily rely on the ACL development (dev) set for validation. It consists of English transcripts of the talks and translations into the 10 target languages. The systems are then evaluated on a blind test set. The dev and test sets consist of 5 talks each. The paper abstracts for all talks are available in English. The talks are pre-segmented. In all experiments, we use the given segmentation.

We also report performance on tst-COMMON of MuST-C (Di Gangi et al., 2019), tst2019 and tst2020 from previous years’ evaluations (Anastasopoulos et al., 2021, 2022).

An overview of the development and test data is in Table 1.

2.2 Speech Recognition Data

For the ASR training, we use Common Voice (Ardila et al., 2020), LibriSpeech (Panayotov et al., 2015), MuST-C v2 (Di Gangi et al., 2019), TED-LIUM v3 (Hernandez et al., 2018), and

Dev/Test set	Hours	# Utterances	Domain
ACL dev	1.0	468	ACL conference talks
tst-COMMON	4.9	2823	TED talks
tst2019	4.8	2279	TED talks
tst2020	4.1	1804	TED talks

Table 1: Overview of development and test data.

Corpus / Data Source	Hours	# Utterances
Common Voice	1667	1225k
LibriSpeech	963	281k
MuST-C v2	482	251k
TED-LIUM v3	452	268k
VoxPopuli	501	177k
TTS	7284	4.7M

Table 2: ASR data overview.

VoxPopuli (Wang et al., 2021). The data overview is in Table 2.

Synthesized Speech Data To adapt the ASR model to the ACL talks, we add synthesized speech created by a text-to-speech (TTS) model. Specifically, from the MT bitext English side (Table 3), we select sentences similar to the ACL domain based on similarity with the provided ACL dev bitext and abstracts. Inspired by data selection strategies for MT (Eck et al., 2005; Koneru et al., 2022), we use n -gram overlap as similarity metric. 4.7M sentences are selected and then synthesized to speech by a VITS (Kim et al., 2021) model trained on MuST-C. The synthesized data amount is shown in the last row of Table 2.

2.3 Machine Translation Data

The MT training data include the following text-to-text translation corpora: Europarl v7 and v10 (Koehn, 2005), NewsCommentary v16, OpenSubtitles v2018 (Lison and Tiedemann, 2016), Tatoeba (Tiedemann, 2012), and ELRCORDIS_News, JParaCrawl (Morishita et al., 2022) for Japanese, and TED2020 (Reimers and Gurevych, 2020) for German¹. We also include the text translation part of the following ST corpora: MuST-C (Di Gangi et al., 2019), CoVoST v2 (Wang et al., 2020), and Europarl-ST (Iranzo-Sánchez et al., 2020). The aggregated data amount per language is summarized in the “Original” column of Table 3.

¹This dataset has duplication with past evaluation sets: tst2019 tst2020 and tst-COMMON. The duplications were removed prior to training.

Lang.	Original	After Diversification	
	# sent. (M)	# sent. (M)	# tokens (M)
ar	26.0	65.2	865.0
zh	11.2	21.5	254.3
nl	33.1	82.1	1162.7
fr	38.9	91.6	1427.8
de	23.0	54.4	860.0
ja*	2.6	27.2	832.7
fa	5.8	11.3	162.1
pt	29.0	72.3	1024.3
ru	22.1	51.5	685.3
tr	36.7	89.7	1021.2
Total	228.4	566.8	8295.4

Table 3: MT data overview. *: For ja, the original data of 2.6M sentences did not include JParaCrawl, which was announced later as allowed data.

As preprocessing, we perform truecasing, deduplication, length ratio filtering, and histogram filtering using the statistics by Fan et al. (2021). Then we perform subword segmentation using Sentencepiece (Kudo and Richardson, 2018) based on the vocabulary of mBART50 (Tang et al., 2020).

Data Diversification Different from last years’ shared tasks (Anastasopoulos et al., 2021, 2022), no monolingual (non-English) data is provided. This means conventional data augmentation techniques like backward translation are not directly applicable. On the other hand, forward translation from existing English monolingual data may introduce undesirable errors in the translation targets, especially on lower-resource languages. In this light, we use data diversification (Nguyen et al., 2020), a data augmentation method that enriches existing parallel data by forward and backward translating the training bitext. As the model has seen the parallel data in training, the synthetic translations are expected to have relatively high quality. Moreover, either the source or target side of the synthetic data is from the original bitext. The diversified data amount after deduplication is shown in Table 3. Here we perform one round of forward and backward translation, as Nguyen et al. (2020) have empirically shown further rounds do not lead to substantial gains.

2.4 Casing/Punctuation Restoration Data

The ASR outputs are lower-cased and unpunctuated, while the MT model expects cased and punctuated inputs. We randomly sample 1.5 million English sentences from the MT training data (Table 3), and remove the casing and punctuation marks as

training source data. We then train a model to restore the casing and punctuation marks.

2.5 Speech Translation Data

The speech translation data are shown in Table 4. We additionally use our trained MT model to create forward translations based on the following transcript-only datasets: Common Voice, TEDLIUM, and VoxPopuli. The TTS data described in §2.2 is also used.

Lang.	Corpus / Data Source	Hours	# Utterances
ar	CoVoST	429	289k
	MuST-C	463	212k
zh	TTS	283	203k
	CoVoST	429	289k
	MuST-C	596	358k
nl	TTS	204	183k
	MuST-C	434	248k
	europarl-ST	75	32k
fr	TTS	1138	713k
	MuST-C	485	275k
	europarl-ST	76	32k
de	TTS	1768	998k
	CoVoST	429	289k
	MuST-C	440	269k
	europarl-ST	77	33k
ja	TTS	1891	779k
	CoVoST	429	289k
	MuST-C	541	329k
fa	TTS	73	56k
	CoVoST	429	289k
	MuST-C	347	182k
pt	TTS	89	88k
	MuST-C	377	206k
	europarl-ST	75	32k
ru	TTS	1678	639k
	MuST-C	482	265k
	TTS	331	331k
tr	CoVoST	429	289k
	MuST-C	446	236k
	TTS	428	511k
all	Common Voice	1488	948k
	TEDLIUM	453	268k
	VoxPopuli	502	177k

Table 4: ST data overview. The last section “all” indicates forward translated synthetic targets from transcript-only corpora, which are available for all 10 languages.

3 Cascaded System

For the cascaded system, we introduce our ASR (§3.1) and MT (§3.2) models.

3.1 Automatic Speech Recognition Module

Baseline Models The first baseline is our ASR model for last year’s offline track (Pham et al., 2022). It is a Wav2vec 2.0 (Baevski et al., 2020) with LARGE configuration pretrained on 960 hours

of Librispeech data. This year, after seeing initial favourable results compared to Wav2vec, we opt for WavLM (Chen et al., 2022) as audio encoder. We use the LARGE configuration with 24 layers. We use the mBART50 (Tang et al., 2020) decoder along with the WavLM encoder. As the ASR model only needs to transcribe English², we trim the mBART50 vocabulary from 256k down to 62k tokens by removing all non-alphabetic tokens.

In-Domain TTS Data We also use the synthesized TTS data. Compared to the same model without TTS data, the word error rate (WER) improves from 11.6% to 10.7% on ACL dev, but degrades from 8.4% to 9.0% on the TEDLIUM test set. There are two potential explanations: First, the noisy TTS speech may be helpful for handling the non-native utterances prominent in the ACL dev set. Second, the target side of the TTS data is more relevant to the ACL domain, as we selected them based on n -gram overlap with ACL data. This in turn improves ASR performance on the ACL dev set.

As shown in Table 5, compared to last year’s submission, this year’s ASR model achieves consistent gains across domains on ACL dev, tst-COMMON, and tst2020.

Model	ACL dev	tstCom.	tst2020
ASR 2022 (Pham et al., 2022)	12.5	5.4	5.6
WavLM + mBART50	10.7	3.9	4.8

Table 5: ASR results in WER(↓) in comparison to our submission last year (Pham et al., 2022) which used Wav2vec trained with CTC and a 5-gram LM. By using WavLM audio encoder and the mBART decoder, we achieve consistent gains across domains (ACL and TED, i.e., tst*).

Language Model (LM) Adaptation Aside from using TTS data, we also investigate other methods to adapt towards the ACL domain using the provided paper abstracts. On preliminary experiments with Connectionist Temporal Classification (CTC) + n -gram LM models, we integrate ACL abstract 5-grams statistics into the language models. As shown in the upper section of Table 6, this improves on ACL dev (WER 13.8% → 13.0%) while preserving the performance on TED talks (tst-COMMON WER stays at 7.6%).

²BART, the English-only predecessor of mBART, is not among the allowed pretrained models.

As our final system is an encoder-decoder model (WavLM + mBART50), adapting the LM alone is less straightforward. We create pseudo ASR training data with ACL data on the transcript side. Specifically, we use our TTS model to synthesize speech from the ACL dev and test abstracts. As the amount of ACL abstract data is very limited (less than 100 sentences in total), we heavily upsampled them, so that they consist of 60% of the training data. As shown in the lower section of Table 6, this leads to a minor improvement of WER for ACL dev. However, the gain does not carry over to ST performance when later cascading with our MT model. Therefore, our final ASR system did not use the abstracts. The lack of improvement could be related to the low amount of ACL abstract data, which requires heavy upsampling of the TTS data, and as a result hinders the ability of transcribing real speech.

The contrast between the two sets of experiments may be related to diminishing gains as WER improves, i.e., for the Wav2vec + CTC + LM model, gaining over a WER of 13.8% is easier than starting from a 10.7% WER. Another interpretation of the difference could be that adding specific constraints to “end-to-end” ASR models is more challenging than the counterparts with separate LMs.

Model	ACL dev	tst-COMMON
Wav2vec + CTC + 5-gram + ACL abstract 5-gram	13.8 13.0	7.6 7.6
WavLM + mBART50 + ACL abstract TTS (upsampled)	10.7 10.5	3.9 4.3

Table 6: ASR adaptation results in WER(↓). On preliminary experiments with Wav2vec + CTC + LM models, we improve ASR performance on ACL dev by integrating n -gram statistics from the ACL abstracts. For the WavLM + mBART 50 model, adding synthesized audio-transcript data based ACL dev abstracts does not give consistent gain.

Casing/Punctuation Restoration We take a sequence-to-sequence approach to the casing and punctuation restoration problem. Specifically, we train a punctuation model initializing from DeltaLM-base (Ma et al., 2021) to restore the casing and punctuation information, using the training data described in §2.4.

3.2 Machine Translation Module

Baseline Model We start with the pretrained DeltaLM (Ma et al., 2021) with LARGE configura-

ID	ACL dev (en→X)										TED (en→de)		
	de	ja	zh	ar	nl	fr	fa	pt	ru	tr	Avg.	tst2019	tst2020
From ground-truth transcripts (MT alone)													
(1) base	39.8	44.2	47.4	30.4	45.7	48.9	23.6	51.1	19.5	22.9	37.4	29.5	32.9
(2) data divers. all	41.6	44.5	49.8	33.6	50.7	51.1	25.4	52.5	21.5	24.6	39.5	30.0	33.7
(3) (1) + data divers.; adapter	41.4	45.8	48.8	33.3	49.8	51.5	25.2	54.1	21.9	24.1	39.6	29.5	33.2
(4) ensemble (2) + (3)	41.7	46.1	49.6	33.7	50.8	52.1	25.9	54.3	23.1	24.8	40.2	30.4	33.7
(5) (4) + k NN-MT	43.7	47.3	49.8	35.4	52.3	52.8	27.2	55.3	23.9	27.1	41.5	30.4	33.4
From ASR outputs (cascaded ST)													
(1) base	34.3	38.2	41.6	25.3	36.6	39.9	19.1	40.7	16.7	18.9	31.1	26.5	28.0
(2) data divers. all	35.4	38.6	44.3	26.8	39.2	41.5	20.5	42.6	18.7	19.5	32.7	27.0	29.3
(3) (1) + data divers.; adapter	35.5	39.0	43.6	26.4	38.9	41.9	20.2	43.0	19.3	19.6	32.7	26.7	28.3
(4) ensemble (2) + (3)	36.1	39.8	44.4	26.9	39.8	42.3	20.7	43.5	19.2	19.7	33.2	26.9	28.7
(5) (4) + k NN-MT	36.8	40.2	44.6	28.2	40.8	42.0	21.8	44.5	19.7	21.1	34.0	26.9	28.5
End-to-end ST													
(6) WavLM + mBART50 decoder	31.7	29.2	40.7	25.0	36.7	40.5	19.5	43.0	16.9	18.5	30.2	27.0	29.3
(7) (6) + TTS	33.2	29.2	40.5	25.5	37.9	41.0	20.1	43.9	16.5	18.9	30.7	27.0	29.1
(8) ensemble (6) + (7)	34.0	29.9	41.7	25.5	38.2	42.0	20.2	44.4	18.3	20.2	31.4	27.3	29.6

Table 7: MT and ST results in BLEU(\uparrow).

tion. The pretrained model has 24 and 12 encoder and decoder Transformer layers respectively. It uses postnorm layer normalization. It is a fully multilingual model where all parameters are shared across languages. The target language tokens are prepended to the source target sentences. We use temperature-based sampling (Arivazhagan et al., 2019) with $\tau = 5.0$ to counteract the data imbalance between languages. When training, we use a relatively large effective batch size of 128k as preliminary experiments with smaller batch sizes showed more instabilities in training. This might be a side effect of the postnorm layer normalization (Nguyen and Salazar, 2019). The results of the baseline are shown in Row (1) of Table 7, with an average score of 37.4 BLEU³ on ACL dev.

Data Diversification As motivated in §2.3, we use data diversification as an alternative data augmentation method in absence of monolingual target data for backtranslation. As data diversification needs forward and backward translations on the training data, we additionally train a 10-to-English model to create the backward translations. Row (2) of Table 7 shows the results after data diversification on all languages pairs. On average, this data augmentation approach improves MT quality by 2.1 BLEU and (37.4 \rightarrow 39.5), and ST quality by 1.6 BLEU (31.1 \rightarrow 32.7).

³By default using tok.13a from sacreBLEU (Post, 2018), except for zh and ja where we use tok.zh and tok.ja-mecab-0.996-IPA.

Adapters for Incremental Data Retraining on the new training data after diversification (Row (2) of Table 7) is time-consuming and costly. To adapt the initial model (Row (1) of Table 7) rapidly towards to the augmented data, we use adapters (Bapna and Firat, 2019; Philip et al., 2020). In this case, the adapters are target-language-specific. The adapters are inserted after each encoder and decoder layer. We initialize from the trained baseline (Row (1) in Table 7), freeze trained parameters and update the adapters only. We use the efficient implementation from Baziotis et al. (2022). As shown in Row (3) of Table 7, only training the adapters on the new diversified training data performs on par with the re-training setup in Row (2) (39.6 on MT and 32.7 on ST on average for ACL dev). These results demonstrate that adapters are suitable for fast and effective incremental learning when additional training data emerges later.

To our surprise, adding adapters to the model trained with full data diversification (Row (2) from Table 7) does not bring further gain. A similar observation was reported by Pires et al. (2023), who opted for training the full network from scratch along with adapters instead. In our case, it therefore would be interesting to see the impact of training on data diversification with adapters from scratch.

Multilingual vs Bilingual To investigate the impact of interference from multiple target languages, in preliminary experiments, we also compare the multilingual and bilingual translation performance for selected language pairs. As shown in Table 8,

compared to bilingual models, the multilingual model lags behind especially on higher-resource languages. Adding the adapters partly closes this gap. Note the score difference to main result table (Table 7) is because the preliminary experiments did not fully use diversified data for all languages.

Model	ACL dev			tst-COMMON		
	en-de	en-ru	en-fa	en-de	en-ru	en-fa
bilingual	41.0	20.0	24.2	34.3	22.7	16.0
multilingual	39.8	19.5	23.6	34.1	21.9	15.9
+ adapters	40.9	20.2	23.7	34.7	22.2	16.3

Table 8: Comparison of bilingual vs multilingual translation performance in BLEU (\uparrow) on German (de), Russian (ru), Farsi (fa), which are high-, mid-, low-resource in the training data (Table 3). Multilingual system falls behind bilingual system, while adapters partly closes the gap. Note the score difference to main result table (Table 7) is because the experiments here did not fully use diversification.

Ensemble Although the models in Row (2) and (3) in Table 7 are trained on the same data and share the same base architecture, we expect their representations to be sufficiently different, as (3) additionally uses adapters. We therefore ensemble these two models. The results are in Row (4) of Table 7. On MT and ST, for ACL, ensembling shows an improvement of 0.6 and 0.5 BLEU respectively over the single models in Row (2) and (3). On TED, however, ensembling does not seem to impact the scores compared to the single models. One explanation is that the adapter model from Row (3) performs worse than its non-adapter counterpart (Row (2)) on TED, which limits the overall effectiveness of ensembling.

k NN-MT We also adapt the MT model to the target domain of scientific talks. A challenge is that we do not have sufficient training data to fully fine-tune the MT model towards the desired domain or style. In this case, we use k NN-MT (Khandelwal et al., 2021) to adapt the model at inference time. In k NN-MT, bitexts are passed through a trained MT model. For each target token, its decoder hidden state is stored in a datastore. At inference time, based on the current decoder hidden state, k candidate target tokens are retrieved from the datastore using a nearest neighbor lookup. The retrieved token distribution is then interpolated with the MT target distribution, which in turn generates the output tokens. Hyperparameters for k NN-MT include

Source (ASR output): ... in a **zero shot evaluation setup**, meaning that **pre trained** word embedding models are applied out of the box without any additional fine tuning

w/o k NN-MT (Table 7 row (4)): ... in einer **Null-Shot-Bewertungs-Setup** (zero-shot evaluation setup), was bedeutet, dass **vorgebildete** (pre-educated) Wort-Einbettungsmodelle ohne zusätzliche Feinabstimmung direkt angewendet werden.

w/ k NN-MT (Table 7 row (5)): ... in einer **Null-Shot-Bewertung** (zero-shot evaluation), was bedeutet, dass **vortrainierte** (pretrained) Wort-Einbettungsmodelle ohne zusätzliche Feinabstimmung direkt angewendet werden.

Source (ASR output): **Hello**. My name is Ramachandra, and I will **present** our paper.

w/o k NN-MT (Table 7 row (4)): **你好** (Hello; addressing a single person), 我叫拉玛钱德拉 我要**发表** (publish) 我们的论文

w/ k NN-MT (Table 7 row (5)): **大家好** (Hi all; addressing a group of audience), 我叫拉玛钱德拉, 我要**介绍** (introduce) 我们的论文。

Table 9: Examples of k NN-MT improving translation quality for en \rightarrow de (upper) and en \rightarrow zh (lower). k NN-MT creates more accurate terminology translations (“pre trained” for en \rightarrow de) and create more context-appropriate translation (“Hello” for en \rightarrow zh).

the number of retrieved neighbors k , the temperature for smoothing the k NN distribution T , and the interpolation weight w .

In our experiments, we use systems (2) and (3) from Table 7 for creating the datastores. As different models’ hidden states (which serve as keys in the datastore) also differ substantially, the datastore is MT-model-dependent. To use k NN-MT when ensembling systems (2) and (3), we therefore need two datastores for systems (2) and (3) respectively. The k NN-MT candidate tokens are interpolated with the output vocabulary distribution before the ensembling operation.

We use hyperparameters $k = 8$, $T = 50$, $w = 0.3$, after an initial search with $T \in [10, 50, 100]$, $w \in [0.1, 0.3, 0.5]$. Our implementation mostly follows Zheng et al. (2021), which uses the FAISS toolkit (Johnson et al., 2019) for efficient k NN operations. Comparing the inference speed of system (4) and (5), with the same batch size of 64 sentences⁴, using k NN-MT takes roughly 50% more time on a Nvidia Titan RTX GPU with 24GB memory.

Naively using all ACL dev bitext as datastore would lead the model to copying the oracle targets. To simulate the scenario on the blind test set, when

⁴System (5) requires more GPU memory than system (4). The latter would be able to use a larger batch size of 128 sentences.

translating the i -th talk, we use the other $j_{j \neq i} \in [n]$ talks’ bitext as datastore, where n is the total number of talks.

As shown in Row (5) of Table 7, k NN-MT brings an additional gain of 1.3 BLEU on MT and 0.8 BLEU on ST. These results shows a datastore as small as hundreds of sentence pairs can be effectively used for inference-time domain adaptation.

Table 9 shows two examples of k NN-MT improving translation quality, apart from generic improvements in fluency and accuracy, in these examples k NN-MT also helps generate correct terminologies and context-appropriate greetings.

4 End-to-End System

For the end-to-end system, similar to our ASR model, after seeing initial favourable results of WavLM over Wav2vec, we choose WavLM as the audio encoder. Following last year’s submission (Pham et al., 2022), we use the mBART50 decoder. The results are shown in Row (6) of Table 7. Contrasting Row (6) and (7) reveals that adding the TTS data does not substantially change ST performance. However, ensembling the two models trained with and without TTS data (Row (8)) improves over the single models (on average +0.7 for ACL, +0.4 for TED), despite them having the identical architecture.

Compared to the strongest cascaded system (Row (5)), the end-to-end system falls behind 2.6 BLEU on ACL dev. On TED, however, it appears to slightly outperform the cascaded system. One explanation is that the MT model of the cascaded system has not been separately adapted to TED texts (although parts of the full training data do cover TED data), which was shown essential in improving performance on TED test sets (Zhang et al., 2022; Pham et al., 2022). The end-to-end system, on the other hand, has seen a larger proportion of TED data in training (Table 4).

Similar to the previous year (Polák et al., 2022), we also adapt our end-to-end offline model for simultaneous track (Polák et al., 2023).

5 Conclusion

In this paper, we described our systems for the multilingual speech translation track of IWSLT 2023, which translates English speech into 10 target languages. To tackle the task of translating scientific conference talks, which feature non-native input speech and terminology-dense contents, our

systems have several novelties. Lacking suitable training data for the target domain, we used k NN-MT for inference-time adaptation and showed an improvement of +0.8 BLEU for cascaded speech translation system. We also used adapters to integrate incremental data from augmentation, and achieved performance on-par with re-training on all data. In our experiments, we observed that cascaded systems are more easily adaptable towards desired target domains due to their separate modules. Our cascaded speech system outperforms its end-to-end counterpart on scientific talk translation, although their performance remains similar on TED talks. For future work, we are interested in the feasibility of applying the adaptation approaches shown effective on MT to end-to-end ST.

Acknowledgement We thank the anonymous reviewers for detailed and insightful feedback. Part of this work was performed on the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the Federal Ministry of Education and Research of Germany. Part of this work was supported by the Federal Ministry of Education and Research of Germany under grant agreement 01EF1803B (RELATER).

References

- Milind Agarwal, Sweta Agrawal, Antonios Anastopoulos, Ondřej Bojar, Claudia Borg, Marine Carpuat, Roldano Cattoni, Mauro Cettolo, Mingda Chen, William Chen, Khalid Choukri, Alexandra Chronopoulou, Anna Currey, Thierry Declerck, Qianqian Dong, Yannick Estève, Kevin Duh, Marcello Federico, Souhir Gahbiche, Barry Haddow, Benjamin Hsu, Phu Mon Htut, Hirofumi Inaguma, Dávid Javorský, John Judge, Yasumasa Kano, Tom Ko, Rishu Kumar, Pengwei Li, Xutai Ma, Prashant Mathur, Evgeny Matusov, Paul McNamee, John P. McCrae, Kenton Murray, Maria Nadejde, Satoshi Nakamura, Matteo Negri, Ha Nguyen, Jan Niehues, Xing Niu, Atul Ojha Kr., John E. Ortega, Proyag Pal, Juan Pino, Lonneke van der Plas, Peter Polák, Elijah Rippeth, Elizabeth Salesky, Jiatong Shi, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Yun Tang, Brian Thompson, Kevin Tran, Marco Turchi, Alex Waibel, Mingxuan Wang, Shinji Watanabe, and Rodolfo Zvallos. 2023. Findings of the IWSLT 2023 Evaluation Campaign. In *Proceedings of the 20th International Conference on Spoken Language Translation (IWSLT 2023)*. Association for Computational Linguistics.
- Roe Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Antonios Anastasopoulos, Loïc Barrault, Luisa Bentivogli, Marceley Zanon Boito, Ondřej Bojar, Roldano Cattoni, Anna Currey, Georgiana Dinu, Kevin Duh, Maha Elbayad, Clara Emmanuel, Yannick Estève, Marcello Federico, Christian Federmann, Souhir Gahbiche, Hongyu Gong, Roman Grundkiewicz, Barry Haddow, Benjamin Hsu, Dávid Javorský, Věra Kloudová, Surafel Lakew, Xutai Ma, Prashant Mathur, Paul McNamee, Kenton Murray, Maria Nădejde, Satoshi Nakamura, Matteo Negri, Jan Niehues, Xing Niu, John Ortega, Juan Pino, Elizabeth Salesky, Jiatong Shi, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Yogesh Virkar, Alexander Waibel, Changhan Wang, and Shinji Watanabe. 2022. [Findings of the IWSLT 2022 evaluation campaign](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 98–157, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Antonios Anastasopoulos, Ondřej Bojar, Jacob Bremerman, Roldano Cattoni, Maha Elbayad, Marcello Federico, Xutai Ma, Satoshi Nakamura, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Alexander Waibel, Changhan Wang, and Matthew Wiesner. 2021. [FINDINGS OF THE IWSLT 2021 EVALUATION CAMPAIGN](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 1–29, Bangkok, Thailand (online). Association for Computational Linguistics.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber. 2020. [Common voice: A massively-multilingual speech corpus](#). In *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, pages 4218–4222. European Language Resources Association.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *CoRR*, abs/1907.05019.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework for self-supervised learning of speech representations](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Christos Baziotis, Mikel Artetxe, James Cross, and Shruti Bhosale. 2022. [Multilingual machine translation with hyper-adapters](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1170–1185, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei. 2022. [Wavlm: Large-scale self-supervised pre-training for full stack speech processing](#). *IEEE J. Sel. Top. Signal Process.*, 16(6):1505–1518.
- Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [MuST-C: a Multilingual Speech Translation Corpus](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2012–2017, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthias Eck, Stephan Vogel, and Alex Waibel. 2005. [Low cost portability for statistical machine translation based on n-gram frequency and TF-IDF](#). In *Proceedings of the Second International Workshop on Spoken Language Translation*, Pittsburgh, Pennsylvania, USA.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *The Journal of Machine Learning Research*, 22:107:1–107:48.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia A. Tomashenko, and Yannick Estève. 2018. [TED-LIUM 3: Twice as much data and corpus repartition for experiments on speaker adaptation](#). In *Speech and Computer - 20th International Conference, SPECOM 2018, Leipzig, Germany, September 18-22, 2018, Proceedings*, volume 11096 of *Lecture Notes in Computer Science*, pages 198–208. Springer.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Adrià Giménez, Albert Sanchís, Jorge Civera, and Alfons Juan. 2020.

- Europarl-st: A multilingual corpus for speech translation of parliamentary debates. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 8229–8233. IEEE.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2021. **Nearest neighbor machine translation**. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021. **Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech**. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5530–5540. PMLR.
- Philipp Koehn. 2005. **Europarl: A parallel corpus for statistical machine translation**. In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.
- Sai Koneru, Danni Liu, and Jan Niehues. 2022. **Cost-effective training in low-resource neural machine translation**. *CoRR*, abs/2201.05700.
- Taku Kudo and John Richardson. 2018. **SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. **OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, Alexandre Muzio, Saksham Singhal, Hany Hassan Awadalla, Xia Song, and Furu Wei. 2021. **Deltalm: Encoder-decoder pre-training for language generation and translation by augmenting pretrained multilingual encoders**. *CoRR*, abs/2106.13736.
- Makoto Morishita, Katsuki Chousa, Jun Suzuki, and Masaaki Nagata. 2022. **JParaCrawl v3.0: A large-scale English-Japanese parallel corpus**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6704–6710, Marseille, France. European Language Resources Association.
- Toan Q. Nguyen and Julian Salazar. 2019. **Transformers without tears: Improving the normalization of self-attention**. In *Proceedings of the 16th International Conference on Spoken Language Translation, IWSLT 2019, Hong Kong, November 2-3, 2019*. Association for Computational Linguistics.
- Xuan-Phi Nguyen, Shafiq R. Joty, Kui Wu, and Ai Ti Aw. 2020. **Data diversification: A simple strategy for neural machine translation**. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. **Librispeech: An ASR corpus based on public domain audio books**. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, pages 5206–5210. IEEE.
- Ngoc-Quan Pham, Tuan Nam Nguyen, Thai-Binh Nguyen, Danni Liu, Carlos Mullov, Jan Niehues, and Alexander Waibel. 2022. **Effective combination of pretrained models - KIT@IWSLT2022**. In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 190–197, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Jerin Philip, Alexandre Berard, Matthias Gallé, and Laurent Besacier. 2020. **Monolingual adapters for zero-shot neural machine translation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4465–4470, Online. Association for Computational Linguistics.
- Telmo Pessoa Pires, Robin M. Schmidt, Yi-Hsiu Liao, and Stephan Peitz. 2023. **Learning language-specific layers for multilingual machine translation**. *CoRR*, abs/2305.02665.
- Peter Polák, Danni Liu, Ngoc-Quan Pham, Jan Niehues, Alexander Waibel, and Ondřej Bojar. 2023. **Towards efficient simultaneous speech translation: CUNI-KIT system for simultaneous track at IWSLT 2023**. In *Proceedings of the 20th International Conference on Spoken Language Translation (IWSLT 2023)*, Toronto, Canada (in-person and online). Association for Computational Linguistics.
- Peter Polák, Ngoc-Quan Pham, Tuan Nam Nguyen, Danni Liu, Carlos Mullov, Jan Niehues, Ondřej Bojar, and Alexander Waibel. 2022. **CUNI-KIT system for simultaneous speech translation task at IWSLT 2022**. In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 277–285, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. [Learning multiple visual domains with residual adapters](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 506–516.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.
- 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](#). *CoRR*, abs/2008.00401.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Changhan Wang, Morgane Riviere, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux. 2021. [VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 993–1003, Online. Association for Computational Linguistics.
- Changhan Wang, Anne Wu, and Juan Miguel Pino. 2020. [Covost 2: A massively multilingual speech-to-text translation corpus](#). *CoRR*, abs/2007.10310.
- Weitai Zhang, Zhongyi Ye, Haitao Tang, Xiaoxi Li, Xinyuan Zhou, Jing Yang, Jianwei Cui, Pan Deng, Mohan Shi, Yifan Song, Dan Liu, Junhua Liu, and Lirong Dai. 2022. [The USTC-NELSLIP offline speech translation systems for IWSLT 2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 198–207, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Xin Zheng, Zhirui Zhang, Junliang Guo, Shujian Huang, Boxing Chen, Weihua Luo, and Jiajun Chen. 2021. [Adaptive nearest neighbor machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 368–374, Online. Association for Computational Linguistics.