

# LEVERAGING DATA COLLECTION AND UNSUPERVISED LEARNING FOR CODE-SWITCHED TUNISIAN ARABIC AUTOMATIC SPEECH RECOGNITION

Ahmed Amine Ben Abdallah<sup>1\*</sup>, Ata Kabboudi<sup>2</sup>, Amir Kanoun<sup>3</sup>, Salah Zaiem<sup>4\*</sup>

<sup>1</sup>Tunis Business School; <sup>2</sup>University of Michigan-Dearborn;  
<sup>3</sup>Abshore; <sup>4</sup>LTCI, Télécom Paris, Institut Polytechnique de Paris

## ABSTRACT

Crafting an effective Automatic Speech Recognition (ASR) solution for dialects demands innovative approaches that not only address the data scarcity issue but also navigate the intricacies of linguistic diversity. In this paper, we address the aforementioned ASR challenge, focusing on the Tunisian dialect. First, textual and audio data is collected and in some cases annotated. Second, we explore self-supervision, semi-supervision and few-shot code-switching approaches to push the state-of-the-art on different Tunisian test sets, covering different acoustic, linguistic and prosodic conditions. Finally, and given the absence of conventional spelling, we produce a human evaluation of our transcripts avoiding the noise coming from spelling inadequacies in our training annotations. Our models, allowing to transcribe audio samples in a linguistic mix involving Tunisian Arabic, English and French, and all the data used during training and testing are released for public use and further improvements.

*Index Terms*— Speech recognition, code-switching

## 1. INTRODUCTION

Several recent works have been trying to extend the number of languages and dialects covered by high-performance speech recognition technology, with models covering hundreds and even thousands of languages [1, 2, 3]. We evaluated a few released models on Tunisian data and collected the results in Table 1. The results show that even massively multilingual models fail at reaching reasonable performance on code-switched or non code-switched Tunisian ASR test sets.

This justifies the need for local models, tackling the needs of specific idioms. In this context, Tunisian ASR has been explored in the last decade, mainly by Tunisian scholars. Linguists focused first on developing orthographic conventions for annotators [4, 5]. Then, from hybrid techniques [6] to end-to-end approaches [7], the models have been suffering from the lack of annotated resources, and thus poor generalization. This work tries to overcome these issues through, first, an effort on multi-source data collection and annotation, and the

Test Sets	TARIC	TunSwitch TO	TunSwitch CS
MMS 1B All. [2]	139.4	104.7	102.0
Wav2vec2.0 Ar. [8]	95.3	89.7	96.4
Whisper Large v2 [1]	119.5	127.3	105.8
Whisper Large v2 Ar.	81.8	74.1	85.9

**Table 1.** Failure of multilingual or Standard Arabic models on a few Tunisian Arabic testing settings. The results are showing the Word Error Rate (WER). “TunSwitch CS” contains code-switching while the other two do not.

exploitation of recent unsupervised representations and self-training techniques.

Getting closer to realistic Tunisian Speech, this work also proposes a first dataset for Tunisian Code-Switched ASR. A large part of Tunisians, use French and English words and expressions in formal or informal settings [9]. The dataset, collected from radio broadcasts and podcasts, shows the extent of this phenomenon and offers a challenging real-conditions ASR task for the code-switching community.

Code-Switching, *i.e.* speakers alternating between two or more languages in their discourse, is present in several regions and countries, and has been an active field of research for the speech-community [10]. With the exceptions of a few works involving Arabic dialects [11, 12], a major part of code-switching research has been focusing on English-Mandarin or English-Hindi situations [13]. This work presents datasets and methods handling dialectal code-switching with three languages involved, in real-world spontaneous conversations.

Thus our contributions are fourfold :

- We collect and release Tunisian audio, annotated or not, and textual data. These cover different conditions; spontaneous versus read speech, code-switched versus non code-switched, allowing the establishment of diverse open benchmarks helping foster research in the community.
- We explore self-supervision, semi-supervision and few-shot zero-switching techniques, pushing the boundaries of Tunisian ASR and reaching reasonable perfor-

\*These authors contributed equally to this work.

	Dataset	Prosody	Code-Switching	Train(H)	Dev(H)	Test(H)
<b>Public data</b>	IWSLT	Spontaneous	✗	15h 24m 47s	4h 55m 51s	4h 36m 28s
	STAC	Spontaneous	✓(slight)	2h 29m 8s	n/a	n/a
	TARIC	Spontaneous	✗	9h 25m 44s	17 m 29s	12m 5s
<b>Collected Data</b>	TunSwitch TO	Read	✗	2h 29m 29s	4m 25s	23m 39s
	TunSwitch CS	Spontaneous	✓	8h 15m 35s	15m 43s	25m 12s
<b>Unlabeled Data</b>	TunSwitch TO	Spontaneous	✓	153h 18m 22s	n/a	n/a

**Table 2.** Description of public and newly collected Tunisian Speech datasets.

mance in code-switched scenarios.

- All the models trained are released together with their code and can be used publicly without commercial restrictions.
- We conduct human evaluations to estimate the noise coming from the absence of spelling conventions.

First, Section 2 describes first the data collection and annotation process, and the public data used in our experiments. Second, Section 3 details the training approaches and choices leading to the released baseline models. Finally, Section 4 covers the results obtained, and describes the human evaluation process and its conclusions.

## 2. DATA COLLECTION AND PREPROCESSING

This section presents the collection process of the textual and audio data used in the remaining of the paper.

### 2.1. Textual Data

Given the scarcity of good quality Tunisian textual data, previous ASR works have only been using data from the training and validation sets for language model training [7]. In this work, we incorporate Tunisian text data sourced from Tunisiya [14], a vast corpus of Tunisian Arabic that is openly accessible. We also scrapped code-switched data from various online sources and public forums. To refine the dataset, we systematically eliminate diacritics, punctuation, special characters, and phrases containing numerical values. Statistics about the two resulting sets are available in Table 3.

### 2.2. Audio Data

#### 2.2.1. TunSwitch Collection tool

We developed a tool for collecting Tunisian dialect data, prompting users to record themselves reading provided phrases. We sourced sentences from Tunisiya [14]. These

sentences are consequently removed from the LM training corpus. 89 persons have participated leading to the collection of 2631 distinct phrases. This set will be called TunSwitch TO, “TO” standing for Tunisian Only, as these sentences do not have non-Tunisian words.

#### 2.2.2. TunSwitch CS

In response to the limited availability of paired Text-Speech Tunisian datasets with code-switching, we have built a corpus through meticulous manual annotation. This process was facilitated by using the Doccano annotation tool [15]. Whenever encountered, French and English words are enclosed within “<fr></fr>” or “<en></en>” tags, and left Tunisian words without any enclosing tags. While these tags have not been used in the proposed models, they allow to have language-usage statistics and may be useful for further approaches handling code-switching. The resulting set is released as TunSwitch CS, “CS” standing for Code-Switched. Figure 1 shows the proportions of the three languages within the TunSwitch CS dataset. It contains 5 times more code-switching than the STAC dataset, the only available code-switched resource.


The TunSwitch CS dataset samples come from a set of radio shows and podcasts, representing diverse topics and a large number of unique speakers. The audio are first segmented into chunks, prioritizing word integrity using the WebRTC-VAD algorithm for silence detection. Afterward, we used a Pyannote [16] overlap detection model to remove overlapping speech sections. Then, a music detection model is employed to eliminate music-containing chunks that could disrupt ASR model accuracy.

#### 2.2.3. Unlabeled Data

To explore self-training with unlabeled audio data, we have curated a vast collection of national TV shows videos spanning a total duration of 317 hours. This dataset encompasses a diverse range of topics, speakers and accents faithfully mirroring the diversity of speech encountered in real-world sce-

Text Corpus	Words	Unique Words
Non Code-Switched	4331540	186209
Code-Switched	23938	5543

**Table 3.** A few statistics on the released textual data.



**Fig. 1.** Code-switching presence in train sets : TunSwitch CS vs STAC. The released TunSwitch CS presents way more English and French parts.

narios. From these initial 317 hours, only 153 hours are kept after VAD-based segmentation and the removal of audio samples containing music or overlapping speech.

#### 2.2.4. Publicly Available Datasets

In addition to the TunSwitch dataset, we included three additional publicly available datasets: TARIC [4] a dataset of conversations in train stations, STAC [7] a radio-broadcast-based dataset with slight code-switching, and the IWSLT challenge speech-translation dataset [17] which consists in telephonic conversations.

Table 2 summarizes the different datasets used in our experiments. Now the datasets are introduced, the remaining of the paper develops a solid baseline involving various unsupervised approaches.

## 3. MODELS

This section describes the different architectures and training policies adopted for the development of the “Tunisian only” and “Code-switched” models released.

### 3.1. Base Model

Given the Tunisian-only annotated training data described in the previous section, we first train a model handling only non code-switched audios, outputting therefore only Arabic characters. Building on other works involving low-resource languages, we opt for a pretrained encoder, trained with a self-supervision objective. While wav2vec2.0 XLSR [18], trained on 53 languages, seems to be the go-to option in the literature,

the WavLM [19] model, although trained only on English data, performed better in our experiments. The downstream head, mapping the representations to the Arabic characters consists in three dense layers with LeakyReLU activations, and batch normalization [20] between layers, and is trained with Connectionist Temporal Classification (CTC) [21] loss. The WavLM encoder parameters are fine-tuned, except for the convolutional front-end that is kept frozen. During evaluation, candidate sentences are rescored using a 4-gram language model trained with the KenLM toolkit [22] and implemented with the PyCTCDecode library. Different language models based on different textual corpora are tested as we will detail in section 4.

### 3.2. Self-Training

Given a first trained Tunisian ASR model, the unlabeled collected and cleaned data samples can be used within a semi-supervised approach. Transcriptions are obtained using the aforementioned model, and added to the training set. Two options are tested, fine-tuning the previous model with the new training points or training all from scratch. The latter led to the best results. This remains a very naive approach for self-training, with the recent literature exploring better schedules for unlabeled data incorporation [23]. It is performed to show that the released data can lead to improvements. We leave more advanced techniques on the exploitation of these unlabeled resources for further works.

### 3.3. Few-Shot Code-Switching

As stated in the description of our datasets, Tunisian speech often involves a dynamic interplay between three distinct languages: Tunisian Arabic, French, and English. The code-switched data in our training set is not sufficient to offer robust large-vocabulary transcriptions in English and French. To overcome this issue, we followed the Few-Shot Code-Switching approach developed by Yan *and al.* [24]. This approach allows the combination of Tunisian, French and English ASR models, individually trained on monolingual datasets. The three models are similar, consisting in a self-supervised encoder and the same decoder outputting character-level posteriorgrams. The three posteriorgrams are first concatenated along with the encoder outputs. Then, a “mixer” model, consisting in our case of two layers of BiLSTM followed with a linear layer, generates a final posteriorgram that encapsulates aggregated character probabilities across all three languages.

During this phase of training, the three models are frozen and the “mixer” is trained only using the Code-Switched train data. The French and English models are trained in a previous phase, using their respective CommonVoice 14.0 datasets. The CommonVoice dataset consists of challenging crowd-collected read sentences. On monolingual datasets,

Reference	Decoding Results
they were they were helping me out بالرغم ما يعرفونيش معناها معناها فما واقع مرير الحقيقة لي يعيشوا فيه هوما	the war they were helping me out رغم ما يعرفونيش معناها فما واقع مرير الحقيقة لي يعيشوا فيه هوما

**Table 4.** Examples of two transcriptions with their reference sentences.

the performance, for French and English, reached, respectively, 10.24 and 18.01 WER. The choice of models for foreign languages, more adapted to the spontaneous or accented character of our test samples, is also one track for improvement left for further work.

## 4. RESULTS AND DISCUSSION

### 4.1. Non Code-Switched Results

Table 5 shows the performance obtained with our models trained and tested without code-switching. For the TARIC and IWSLT datasets, we reported the best results we found in the literature, while we introduce the first results on the collected dataset TunSwitch TO. The upper part of the table shows the performance without self-training, *i.e.* without the weakly supervised samples in the training set. Every line corresponds to a different textual corpus used for the LM training. “InDomain” indicates that the textual data only comes from the train and validation sets of the different considered audio corpus, while “OutDomain” indicates that external textual sentences have been added to the LM training corpus.

First, significant discrepancies in results between datasets are witnessed. This is natural given the different settings. The TARIC dataset, consisting of very similar discussions around buying train tickets, display a reduced vocabulary leading to low WERs. The IWSLT consists in telephonic (8khz) spontaneous conversations with multiple hesitations, and represents the hardest task in our benchmark. The TunSwitch TO dataset, although read, contains the richest vocabulary, and is openly crowd-sourced leading to very different recording and noise conditions. It is the one closest to an industrial user-oriented ASR use-cases.

In our setting, the self-training improves the performance on the three datasets, especially when no LM is used for rescoring, gaining little above 1 point of WER on TARIC and 3 on the collected data. When using language-modelling, this gain is reduced, reaching 1.1 WER progress on the TunSwitch TO set. Concerning language modelling, all three datasets see a substantial gain in performance when adding InDomain LM rescoring, respectively, 2, 3.9, and 5.3 WER absolute improvement, for TARIC, IWSLT and TunSwitch. Adding the external textual sentences to the LM training corpus, improves significantly the performance for the TunSwitch set, with 9.5 absolute WER gain in the self-training setting. This is expected, as the read testing sentences were sampled from

the same source, and, thus, may cover similar topics.

### 4.2. Code-Switching Results

Table 6 shows the performance obtained with the “Mixer” based approach on code-switched data. The table shows again the importance of properly calibrated language models for rescoring. Using more easily available “Tunisian Only” corpora for training LMs harms the ASR performance. Using our released code-switched textual data, allows for 10 points of absolute WER progress. For the final line, we enrich the textual corpus with ten thousands English and French monolingual sentences, leading to around 1 point of WER improvement. Our best model leads to 29.47 WER on a very challenging, spontaneous trilingual code-switched radio broadcasts data, establishing a solid baseline on the collected TunSwitch CS dataset.

### 4.3. Human evaluation

Table 4 shows two examples, one with code-switching and one without, for references and transcriptions. The two examples display spelling errors, one in the English part, and the other one in the Tunisian Arabic one. However, a Tunisian reader is likely to accept the second transcript. Indeed, because the Tunisian dialect does not have a clear spelling convention, annotators, especially in the case of multiple datasets, may choose to write words differently. Reading the error reports, we observed that a non-negligible part of the errors were due to the absence of spelling conventions and may not be considered false by a human evaluator.

This motivated a human evaluation of the model outputs. 25 Tunisian evaluators, reasonably fluent in English and French were recruited, and tasked to evaluate the transcriptions of 50 audios each. To make it easier for evaluators, they were only tasked to judge whether the full transcription of the sentence was correct or not, in a binary decision for each test sample. Evaluators have been handled a document showing how to use the validation website and a few examples showing good and bad transcriptions with the corresponding audios. Every audio in the test set of TunSwitch (TO and CS) is proposed to two different annotators. One sentence is considered correct if the two evaluators agreed on accepting it. The Results are reported in Table 7 showing large differences between the human and automatic sentence-level evaluation. Human Sentence Error Rate (SER) is 42% and 29.5% lower, respectively, for the sentences without and with code-

	TARIC		IWSLT		Annotation	
	CER	WER	CER	WER	CER	WER
Previous works	N/A	22.6 [6]	N/A	41.5 [25]	N/A	N/A
Without Self-Training	CER	WER	CER	WER	CER	WER
Without LM	6.44	12.84	20.28	42.74	13.34	41.45
With InDomainLM	6.23	10.81	20.27	<b>38.86</b>	12.50	36.18
With OutDomainLM	<b>6.13</b>	<b>10.55</b>	<b>20.32</b>	39.01	10.08	26.64
With Self-Training	CER	WER	CER	WER	CER	WER
Without LM	6.33	11.82	20.49	42.49	12.65	38.25
With InDomainLM	6.29	10.83	21.18	39.46	12.42	36.07
With OutDomainLM	6.22	<b>10.55</b>	21.18	39.53	<b>9.67</b>	<b>25.54</b>

**Table 5.** ASR Results on non code-switched data. Character and Word Error Rates are shown (CER and WER) for models trained with or without self-training. Proper language modelling appears to be crucial towards better performance.

	TunSpeech	
	CER	WER
Without LM	13.71	40.65
With TunisianOnly LM	17.57	47.45
With CodeSwitched LM	12.77	30.41
With EN-FR enriched LM	<b>12.44</b>	<b>29.47</b>

**Table 6.** Results on non code-switched data. Character and Word Error Rates are shown (CER and WER).

Dataset	TunSwitch TO	TunSwitch CS
Automatic SER	76.45%	96.0%
Human SER	34.5%	66.5%

**Table 7.** Sentence Error Rates (SER), automatically computed using references, or with human validation. Large differences between Human and Automatic SER are witnessed with or without code-switching.

switching. However, human evaluations should be taken with a pinch of salt, as agreement between annotators reached only 80.4% and evaluators may not be attentive enough to small errors. We think the large difference is still imputable in part to the absence of spelling conventions, questioning the way dialectal ASR should be properly evaluated.

## 5. CONCLUSION

This paper introduces new resources for code-switched Tunisian Arabic defining a very challenging ASR task in spontaneous audio involving mixing three languages. Using self-supervised representations, self-training and other monolingual ASR models, a solid baseline is proposed. We hope the code-switched speech recognition community will find this resource useful and builds upon the baseline.

## 6. REFERENCES

- [1] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, “Robust speech recognition via large-scale weak supervision,” 2022.
- [2] Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, Alexei Baevski, Yossi Adi, Xiaohui Zhang, Wei-Ning Hsu, Alexis Conneau, and Michael Auli, “Scaling speech technology to 1,000+ languages,” *arXiv*, 2023.
- [3] Yu Zhang, Wei Han, James Qin, Yongqiang Wang, Ankur Bapna, Zhehuai Chen, Nanxin Chen, Bo Li, Vera Axelrod, Gary Wang, Zhong Meng, Ke Hu, Andrew Rosenberg, Rohit Prabhavalkar, Daniel S. Park, Parisa Haghani, Jason Riesa, Ginger Perng, Hagen Soltau, Trevor Strohman, Bhuvana Ramabhadran, Tara Sainath, Pedro Moreno, Chung-Cheng Chiu, Johan Schalkwyk, Françoise Beaufays, and Yonghui Wu, “Google usm: Scaling automatic speech recognition beyond 100 languages,” 2023.
- [4] Inès Zribi, Rahma Boujelbane, Abir Masmoudi, Mariem Ellouze, Lamia Hadrich Belguith, and Nizar Habash, “A conventional orthography for tunisian arabic,” in *LREC*, 2014, pp. 2355–2361.
- [5] Hala Mulki, Hatem Haddad, Chedi Bechikh Ali, and Ismail Babaoğlu, “Tunisian dialect sentiment analysis: a natural language processing-based approach,” *Computación y Sistemas*, vol. 22, no. 4, pp. 1223–1232, 2018.
- [6] Abir Masmoudi, Fethi Bougares, Mariem Ellouze, Yannick Estève, and Lamia Belguith, “Automatic speech recognition system for tunisian dialect,” *Language Resources and Evaluation*, vol. 52, pp. 249–267, 2018.
- [7] Abir Messaoudi, Hatem Haddad, Chayma Fourati, Moez BenHaj Hmida, Aymen Ben Elhaj Mabrouk, and Mohamed Graiet, “Tunisian dialectal end-to-end speech recognition based on deepspeech,” *Procedia Computer Science*, vol. 189, pp. 183–190, 2021.
- [8] Jonas Grosman, “Fine-tuned XLSR-53 large model for speech recognition in Arabic,” <https://huggingface.co/jonatasgrosman/wav2vec2-large-xlsr-53-arabic>, 2021.
- [9] Lotfi Sayahi, “Code-switching and language change in tunisia,” vol. 2011, no. 211, pp. 113–133, 2011.
- [10] Xian Shi, Qiangze Feng, and Lei Xie, “The asru 2019 mandarin-english code-switching speech recognition challenge: Open datasets, tracks, methods and results,” *arXiv preprint arXiv:2007.05916*, 2020.
- [11] Djegdjiga Amazouz, Martine Adda-Decker, and Lori Lamel, “Addressing Code-Switching in French/Algerian Arabic Speech,” in *Interspeech 2017*, Stockholm, Sweden, Aug. 2017, pp. 62–66.
- [12] Injy Hamed, Pavel Denisov, Chia-Yu Li, Mohamed Elmahdy, Slim Abdennadher, and Ngoc Thang Vu, “Investigations on speech recognition systems for low-resource dialectal arabic–english code-switching speech,” *Computer Speech & Language*, vol. 72, pp. 101278, 2022.
- [13] Sunayana Sitaram, Khyathi Raghavi Chandu, Sai Krishna Rallabandi, and Alan W Black, “A survey of code-switched speech and language processing,” *arXiv preprint arXiv:1904.00784*, 2019.
- [14] Karen McNeil, “Tunisian arabic corpus: Creating a written corpus of an ‘unwritten’ language,” *Arabic corpus linguistics*, vol. 30, 2018.
- [15] Hiroki Nakayama, Takahiro Kubo, Junya Kamura, Yasufumi Taniguchi, and Xu Liang, “doccano: Text annotation tool for human,” 2018, Software available from <https://github.com/doccano/doccano>.
- [16] Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill, “Pyannote. audio: neural building blocks for speaker diarization,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7124–7128.
- [17] Anastasopoulos Antonios, Barrault Loc, Luisa Benvivogli, Marcelly Zanon Boito, Bojar Ondřej, Roldano Cattoni, Currey Anna, Dinu Georgiana, Duh Kevin, Elbayad Maha, et al., “Findings of the iwslt 2022 evaluation campaign,” in *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*. Association for Computational Linguistics, 2022, pp. 98–157.
- [18] Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, et al., “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [19] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al., “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.

- [20] Sergey Ioffe and Christian Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International conference on machine learning*. pmlr, 2015, pp. 448–456.
- [21] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton, “Speech recognition with deep recurrent neural networks,” in *2013 IEEE international conference on acoustics, speech and signal processing*. Ieee, 2013, pp. 6645–6649.
- [22] Kenneth Heafield, “KenLM: Faster and smaller language model queries,” in *Proceedings of the Sixth Workshop on Statistical Machine Translation*, Edinburgh, Scotland, July 2011, pp. 187–197, Association for Computational Linguistics.
- [23] Dan Berrebbi, Ronan Collobert, Samy Bengio, Navdeep Jaitly, and Tatiana Likhomanenko, “Continuous pseudo-labeling from the start,” *arXiv preprint arXiv:2210.08711*, 2022.
- [24] Brian Yan, Matthew Wiesner, Ondřej Klejch, Preethi Jyothi, and Shinji Watanabe, “Towards zero-shot code-switched speech recognition,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [25] Marcely Zanon Boito, John Ortega, Hugo Riguidel, Antoine Laurent, Loic Barrault, Fethi Bougares, Firas Chaabani, Ha Nguyen, Florentin Barbier, Souhir Gahbiche, et al., “On-trac consortium systems for the iwslt 2022 dialect and low-resource speech translation tasks,” *arXiv preprint arXiv:2205.01987*, 2022.