



The Faetar Speech Recognition Benchmark

Michael Ong¹, Sean Robertson^{2,3}, Leo Peckham^{1,2}, Alba Jorquera Jimenez de Aberasturi¹, Paula Arkhangorodsky¹, Robin Huo^{1,2}, Aman Sakhardande^{1,4}, Mark Hallap⁴, Naomi Nagy¹, Ewan Dunbar^{1,2,3}

¹Department of Linguistics, University of Toronto, Canada

²Department of Computer Science, University of Toronto, Canada

³Department of French, University of Toronto, Canada

⁴Department of Philosophy, University of Toronto, Canada

michael.ong@mail.utoronto.ca, sdrobert@cs.toronto.edu, leo.peckham@mail.utoronto.ca,
alba.jorquera@mail.utoronto.ca, p.arkhangorodsky@mail.utoronto.ca,
robin.huo@mail.utoronto.ca, aman.sakhardande@mail.utoronto.ca,
mark.hallap@mail.utoronto.ca, naomi.nagy@utoronto.ca, ewan.dunbar@utoronto.ca

Abstract

We introduce the Faetar Automatic Speech Recognition Benchmark, a benchmark corpus designed to push the limits of current approaches to low-resource speech recognition. Faetar, a Franco-Provençal variety spoken primarily in Italy, has no standard orthography, has virtually no existing textual or speech resources other than what is included in the benchmark, and is quite different from other forms of Franco-Provençal. The corpus comes from field recordings, most of which are noisy, for which only 5 hours have matching transcriptions, and for which transcriptions are inconsistent. The corpus contains an additional 20 hours of unlabelled speech. We report baseline results from multilingual speech foundation models with a best phone error rate of 30.5%, using a pipeline that continues pre-training on the foundation model using the unlabelled set.

Index Terms: Speech Recognition, Low Resource, Faetar, Self-Supervised, Unsupervised

1. Introduction

Speech technology for low-resource languages has been accelerated by foundation models. However, for the most part, demonstrations of the benefits of pre-training or cross-language transfer approaches for low-resource automatic speech recognition (ASR) have been broad-coverage evaluations aggregated over many languages and corpora at once [1]. This masks vast performance differences. The value of human linguistic diversity makes it important for the next era of speech recognition to deal with limited data, and speech in difficult conditions.

We introduce a new benchmark corpus of Faetar, a variety of the Franco-Provençal language which developed in isolation in Italy, far from other speakers of Franco-Provençal, in close contact with Italian and its dialects [2, 3]. Faetar has less than 1000 speakers, in Italy and in the diaspora [4]. The benchmark represents the majority of all archived speech recordings of Faetar in existence, and is not available from any other source.

The language has no standard written form. However, as it is endangered, preservation, learning, and documentation are a priority for many community members. As such, automatic speech recognition would provide the valuable ability to transcribe and index Faetar recordings. This situation is typical of many languages, as is the quantity and quality of the data set, around five hours of labelled speech (transcribed quasi-phonetically), and around 20 hours unlabelled, from field recordings, generally noisy, of semi-spontaneous speech.

We believe exploration and advances can be made quickly by drawing the community’s attention to a common problem. The benchmark seeks to encourage researchers interested in bringing speech recognition to more languages to closely examine issues that arise in extreme conditions, rather than focusing on average-case performance, which hides a long tail with poor performance. Building a useful ASR system for Faetar is made difficult by a confluence of factors: the recordings are noisy, there are only a few hours transcribed, and, perhaps most significantly, the transcription conventions are internally inconsistent due to the lack of a standardized orthography. Most of the world’s languages do not have extensive high-quality annotated speech corpora, and many do not have standard writing systems, which means that these issues represent a major concern moving forward, yet they are likely under-represented in available corpora, which are generally biased towards making use of the best-quality data available. As such, the Faetar ASR Benchmark is an important case study.

The benchmark, including baseline and evaluation code, is freely available subject to an online data-sharing agreement.¹ At the time of writing, the test decodings are under embargo, in order to discourage researchers from optimizing on the test set. The benchmark is thus currently available as a challenge task.

2. Related Work

Low-resource speech benchmarks began with IARPA Babel,² which aimed to facilitate the rapid development of robust ASR and keyword-spotting systems on languages with limited transcribed data [5]. In 2020 and 2021, Babel data were repurposed for the NIST OpenASR challenges, with similar goals. In 2018, a low-resource speech recognition challenge for Indian languages was held [6], using about 40 hours of labelled data per language.

While all of the languages in the above challenges are severely more resource-limited than English, substantial additional speech and textual data can be found for all of the languages used in the above challenges [7, 8, 9, 10, 11, 12]. Thus, in this sense, the Faetar corpus is very different from corpora used in previous challenges, as there is almost no additional publicly available data, either speech or text, that challenge par-

¹<https://perceptimatic.github.io/faetarspeech>

²<https://www.iarpa.gov/index.php/research-programs/babel>, last accessed April 8, 2024.

ticipants could make use of. Furthermore, the amount of labelled data is substantially smaller than most challenges. Previous challenges also generally use higher-quality data. While other endangered languages such as Yoloxóchitl Mixtec and Sierra Totonac are similarly under-resourced [13, 14, 15, 16, 17], these languages have not been explicitly presented to the community as benchmarks.

Recent benchmark evaluations have been developed for multilingual and low-resource speech recognition, focusing on evaluating on as many languages as possible [18, 12, 19, 20, 1]. This breadth represents a great advance over the field’s traditional attention to a few languages. However, aggregate scores mask radical variation across languages and corpora [1, 20, 21, 22]. Importantly, attacking the difficulties posed by the Faetar benchmark would help both with very low-resource languages without consistent orthography, and in unlocking the use of lower-quality data in higher-resource languages. This benchmark complements breadth-first approaches with the depth of a test case posing problems of significant interest.

3. Materials

3.1. Language

Faetar (pronounced [fajdar]) is spoken in Faeto, in the Apulia region of southern Italy, and is very similar to the language spoken in the neighbouring town of Celle St. Vito. The population of both towns is less than 1000, and Faetar is threatened by population loss and language shift to Italian. The language has absorbed many Italian loan words, and, because of bilingualism, Faetar speakers often code-switch with Italian. The language is very different from Italian both in its structure and its base vocabulary. It is a variety of Franco-Provençal, a member of the southern group of Gallo-Romance languages. Franco-Provençal was brought to Apulia by a group of speakers between the 12th and 15th centuries (according to tradition, soldiers fighting for Charles of Anjou), and the community has since been isolated from the rest of the Franco-Provençal speaking community [3, 4]. There is also a Faetar-speaking diaspora, including in the Greater Toronto Area in Canada [23].

Faetar has no standard orthography. Our recordings were transcribed for linguistic analysis using phonetic symbols. The transcription is a mix of phonetic and phonemic transcription. For example, the word for *table* is sometimes transcribed as [tawulina] and sometimes as [tawolin:], reflecting variability in the pronunciation; but the word for *above* is always transcribed as [indjok] even though it is sometimes pronounced [ijok]. This is unlike typical spelling systems, which tend to be phonemic, but similar to transcriptions of oral collections in other languages, which often include limited dialectal variation. Given the lack of an agreed-upon orthography, it would be a large task to standardize the transcription (see Section 6).

The phoneme inventory of Faetar is not agreed upon, but a set of core phones is given in Table 1, based on [3]. In addition to the 29 phonemes in Table 1, geminate versions of all consonants are phonemically contrastive in most cases. They are transcribed in our corpus using the length marker [:]. Some consonants have allophonic variants (/g/ can be pronounced as fricative [ɣ]; fricatives and affricates can all be voiced), and /i e u o/ have allophonic variants [ɪ ɛ ʊ ɔ] respectively, in closed syllables. There are no phonemic length distinctions for vowels, but it is possible to have hiatus (same vowel across adjacent syllables), in which case we also used the length marker in the transcription [:], and treated the resulting doubled vowel

as its own phone. There are a number of diphthongs; in our transcription, we treat these as vowel–glide sequences. There are three additional phones transcribed that are not mentioned in [3], namely [h] (which mostly occurs at the end of a word, and is rarely very audible), and both single and double [ʌ], a stressed variant of /ə/. In total, with geminates, doubled vowels, and variant phones, we have a 68-phone phone set.

3.2. Data

Data were extracted from the Faetar collection of the Heritage Language Variation and Change in Toronto (HLVC) corpus [24]. The corpus contains 184 recordings of native Faetar speakers collected in Faeto between 1992 and 1994 (the *Homeland* subset) and 37 recordings of first- and second-generation heritage Faetar speakers collected in Toronto between 2009 and 2010 (the *Heritage* subset) [24]. Some of the recordings were interviews (*Interview*), designed to elicit spontaneous speech, and, in others, participants were asked to describe scenes and objects from a picture book (*Words*). The recordings were saved at a 44.1 kHz sampling rate (in the case of the Homeland recordings, this is the sampling rate of the digitization from analog cassette tapes). Recordings can be noisy, with considerable background noise, back-channels, and interruptions.

The source data set consisted of long-form recordings, of which 68 of the Homeland recordings, and 26 of the Heritage recordings, had been at least partially transcribed at the utterance level. Some of the transcriptions (mostly those in the Heritage subset) were extracted from ELAN annotation files, which had utterance-level alignments [25], while the remainder were extracted from Microsoft Word files, by manually mapping numerical values from each of the two fonts used (PalPhon or IPA-Phon) to modern UTF-8 encodings, and removing anything that did not look like a transcription (glosses, identifiers, *etc.*).

For the portion which was transcribed but not aligned, we adapted the JHU Arabic MGB-3 recipe for segmentation and alignment [26], built on Kaldi [27]. This splits recordings and recording-level transcriptions by i-vector speaker diarization, splits each diarized segment into smaller, overlapping segments, aligns each overlapping segment to the text (allowing for minor changes to the text) with a seed ASR system, and stitches the overlapping segments back together using string matching. We trained the seed system on the aligned part of the data. To generate our automatic alignments, we used a monophone, speaker-independent Gaussian mixture model-Hidden Markov model (GMM-HMM) system. We found this worked better than triphone or speaker-dependent models, presumably because of the limited aligned data available for training. We then used PyAnnotate 3.0 [28, 29] to adjust utterance boundaries according to voice-activity detection and label them with speaker diarization. Because the pipeline is error-prone, we manually threw out utterances whose boundaries were clearly misaligned or whose transcriptions were clearly wrong. On the test set (see below), we made a more rigorous manual pass to correct alignments.

After obtaining time-alignments for the whole data set, further filtering was performed to remove the interviewer’s speech (generally clearly marked), to remove utterances with duration less than 500 ms, and to remove utterances in Italian or English. While code-switching is common in Faetar speech (and thus, a recognizer should deal with it) removing utterances with substantial parts of Italian or English did not result in a major loss of data. As such, we take it that the (simpler) problem we propose, which consists of transcribing Faetar with minimal code-switching, is a reasonable approximation to the real

	Bilabial		Labiodental		Alveolar		Postalveolar	Palatal		Velar	
Plosive	p	b			t	d				k	g
Nasal		m				n			ɲ		ŋ
Trill/flap						r					
Fricative			f	v	s		ʃ				
Affricate					ts		tʃ	dʒ			
Approximant									j		w
Lateral						l			ʎ		

Table 1: The core phonetic inventory of Faetar. See text for discussion of additional phones used in transcriptions.




Table 2: Distribution of data in (hh:)mm:ss.

		Homeland		Heritage		Total
		M	F	M	F	
dev	Int.	8:43	0:31	0:00	2:35	11:49
test	Int.	13:10	13:33	10:54	9:16	46:54
train	Words	42:31	58:29	51:44	9:07	4:30:17
	Int.	15:58	27:56	32:49	14:41	
1h	Int.	11:08	13:44	32:50	0:51	58:34
10min	Int.	0:00	4:25	5:00	0:24	9:49
unlab	Words	7:29:11		0:05:24		19:55:21
	Int.	8:23:18		0:51:31		
	Mixed	2:12:47		0:53:09		

problem. For the purposes of ASR research, it will be easier to interpret the results in a more homogeneous corpus. Nevertheless, as Faetar has been in contact with Italian for a long time, and substantially longer than it has been in contact with English (English only appears in the Heritage portion of the corpus), we took a more lenient approach to identifying Italian utterances. The presence of a single English word, drawn from a closed list that were clearly not words in Faetar, was sufficient to mark an utterance as English. On the other hand, for Italian, we required that there be three words in a row, again from a closed list.

We split the aligned data into *train*, *dev*, and *test*, ensuring a reasonable balance between male and female speakers, and between Homeland and Heritage subsets, in both *train* and *test*. We do not include the *Words* subset in the *dev* or *test* splits, as we take it that this is too easy an evaluation: many of the utterances consist of individual words, often the same word. Anticipating compatibility with the ML-SUPERB benchmark [1], we also construct *1h* and *10min* subsets of *train*. We distribute the remainder of the data, for which we did not have transcriptions, or for which the (long-form) file could not be time-aligned, as an *unlabelled* set, after obtaining VAD and speaker diarization using PyAnnote 3.0. Because we determined that many of the shorter utterances extracted in this way were not speech, we discarded all unlabelled utterances less than 1.5 seconds. Table 2 shows the distribution of data in the corpus.

4. Evaluation

The measure of performance is phone error rate (PER) on the *test* set, calculated by aligning reference and hypothesis tran-

scripts by the unit-cost Levenshtein algorithm. Given all utterances u in a partition \mathcal{U} , the number of phones in a reference transcript n_u , and the number of errors in a model’s hypothesis transcript e_u , the PER on \mathcal{U} is defined as:

$$\text{PER}_{\mathcal{U}} = \frac{\sum_{u \in \mathcal{U}} e_u}{\sum_{u \in \mathcal{U}} n_u} \times 100\%. \quad (1)$$

The phonetic transcriptions contain spaces between words, making both word error rate (WER) and character error rate (CER) possible alternative metrics. WER is too severe as the transcriptions include substantial phonetic variation that would not ordinarily be present in orthographic transcripts. PER and CER (which differ by the exclusion or inclusion of spaces, and by the treatment of affricates and long phones) do not differ by more than a few percentage points in our observations. We chose PER (no spaces, affricates and long phones single units) to avoid making the task unnecessarily difficult. We return to PER in the context of inconsistent transcriptions in Section 6.

5. Baseline Experiments

We consider *unconstrained* and *constrained* baseline systems, where by “constrained” we mean that the model is not trained on data sets other than *train* (thus precluding pre-trained models trained on other languages, as well as *unlab*). Considering constrained approaches separately allows us to focus on questions of model architecture, rather than on pre-training. For constrained models, as discussed in Section 3.2, force-alignment of transcriptions to audio was performed with Kaldi [27], initially using a monophone, speaker-independent HMM-GMM model, then refined at the end using a triphone, speaker-dependent model; both were composed with a 5-gram, character-level, modified Kneser-Ney language model [30] trained on the (post-alignment) transcriptions from the *train* partition. We report performance in Table 3, \pm the half width of a 95% bootstrap confidence interval calculated using $K = 10000$ samples.

In the *unconstrained* condition, we fine-tune existing multilingual models MMS [31] and mHuBERT-147 [32] on the *train* set. Both are masked prediction models; MMS is a wav2vec 2.0 model that takes a language-specific adapter, head, and fine-tuning strategy. Rather than initializing the adapter layers from scratch, we initialize the adapters and head with a model fine-tuned on Italian. mHuBERT-147 takes an iterative refinement approach, training first on classification according to classes derived from k-means clustering on spectral features, then, after the first iteration, on targets derived from clustering the internal representations of the previous one. As [32] obtained best results with three iterations, we fine-tune the third iteration. For

Table 3: Phone error rates (PERs, %) for each model (rows) and partition (columns), $\pm \frac{1}{2}$ -width of 95% confidence interval.

	Training	Type	Train	Dev	Test
Mono +5gr	train	Constr	47.1 ± 0.6	65.9 ± 1.8	62.6 ± 0.8
Tri +5gr	train	Constr	9.6 ± 0.4	58.2 ± 2.0	56.7 ± 0.9
MMS	train		34.4	39.9	33.0
Fine-tune (FT)			± 0.5	± 2.8	± 0.8
mHuBERT-147 FT	train		26.6 ± 0.5	41.1 ± 2.6	33.6 ± 0.8
ESPnet-MMS	train		36.0 ± 0.5	43.0 ± 2.8	35.9 ± 0.8
MMS	train	+ unlab	31.4	38.7	31.5
Pre-train (PT)			± 0.5	± 2.7	± 0.8
MMS	train	+ unlab	31.4	36.6	31.0
Self-train (ST)			± 0.5	± 2.7	± 0.8
MMS PT + ST	train	+ unlab	30.0 ± 0.5	36.4 ± 2.7	30.5 ± 0.8
ESPnet-MMS	1hr		39.3 ± 0.5	44.4 ± 2.8	37.4 ± 0.8
ESPnet-MMS	10min		48.2 ± 0.5	50.2 ± 2.4	45.1 ± 0.8

both models, we use a linear schedule with peak learning rate of $1e^{-5}$ for 200 epochs of which ten are warmup, dropout of 10%, and effective batch size of 8. For reference, we also train using the ML-SUPERB recipe [1] from ESPNet [33], which also fine-tunes the 1-billion parameter MMS model, keeping all hyperparameters as-is except to set the effective batch size to 4. For comparison with ML-SUPERB, we also include results for the restricted *1hr* and *10min* training conditions.

The data includes an *unlabelled* subset. We show two possible approaches to using it. First, *continued pre-training* of a multilingual model using Faetar audio for self-supervised training prior to fine-tuning. We use the entire *unlab* set and continue pre-training of MMS (effective batch size 32; 300000 steps ≈ 258 epochs, of which the first 30% are warmup in a linear scheduler; peak learning rate 3×10^{-4}), then fine-tune as before. As can be seen from Table 3, this leads to a large performance boost. Alternatively, we attempt *self-training*, decoding *unlab* using the fine-tuned (base) MMS model, then adding the resulting decoded speech to a second round of fine-tuning. Finally, we combine both approaches. Table 3 demonstrates that both of these approaches improve performance, and that the results are cumulative, leading to a best PER of 30.5% when combining both approaches (although the pairwise differences are within the overlap of the confidence intervals again). In sum, the unlabelled data set shows potential for improving ASR.

6. Limitations

The major limitation of this benchmark is that the transcriptions themselves are inconsistent. As discussed above, they vary unpredictably between tracking fine-grained phonetic distinctions and lexical-level (phonemic) distinctions. As is typical of field data, it was collected over many years, and different files were transcribed for different purposes. Furthermore, as is often the case when fine-grained phonetic details are transcribed but there is no clear consensus about what the phonemic form is,

the transcriptions are inconsistent in another way: specific phonetic details may not always be transcribed in exactly the same way. Genuine errors were corrected in the test set, but these two kinds of inconsistency remain, even in the test set. These two sources of noise in the transcription mean that there is an unknown lower bound (“noise threshold”) on PER.

The fact that we are limited to transcriptions of this kind is partly a consequence of the fact that Faetar has no standard orthography and that the linguistic (phonemic) analysis is not fully understood. The fact that there is no standard orthography does not detract from the utility of ASR. A major application of ASR is in indexing audio recordings and in generating textual resources from them (that may, for example, later be used to build educational materials), in addition to linguistic research. Such collections can be built without a standard orthography.

However, for most purposes—even for linguistic phonetic research—having fine-grained phonetic transcriptions (even perfectly consistent ones) is insufficient. For most practical purposes, we need a lexical transcription indicating the intended word, without all pronunciation details—for example, grouping [tawulinə] and [tawolin:] together. The serious challenge posed here is that, even armed with an extensive dictionary of Faetar, we would have no way of automatically mapping our sub-lexical transcriptions to the headwords in the dictionary, and, more importantly, of knowing whether we have accurately done so. Furthermore, existing lexica are highly incomplete. Previous approaches to inconsistent transcriptions in ASR generally rely on extensive manual labelling or elicitation of alternate transcriptions of the same utterance [34, 35, 36]. Both would be prohibitively time- and expert-intensive for us.

We take the future challenge of guessing at an ASR lexicon or a phonemic-level transcription convention based in a fully unsupervised manner, with no lexical-level grouping or annotation used in training, to be potentially ground-breaking for low-resource languages. We do not explore this except to note that informal experiments with language modelling based on lexica derived by clustering did not yield improvements to PER.

7. Summary of Contributions

We have presented the Faetar Automatic Speech Recognition Benchmark, a benchmark corpus aimed at under-resourced languages with noisy audio and transcriptions. We demonstrate that out-of-the-box application of state-of-the-art multilingual foundation models yields an improvement over spectral features, and that using unlabelled data yields further improvements. The fact that we obtain improvements on the constrained ASR task by varying models suggests that there may be further improvements to be made to the model architectures as well. Nevertheless, in all cases, the baseline PERs would generally be unacceptable without manual intervention. We leave additional approaches for the community to explore, and release the benchmark with the aim of encouraging careful attention to corpora and languages like Faetar, deepening the field’s commitment to the promise of universal speech technology.

8. Acknowledgments

This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) RGPIN-2022-04431, as well as the Data Sciences Institute and Jackman Humanities Institute at the University of Toronto. It was also enabled in part by resources provided by Compute Ontario, Calcul Québec, and the Digital Research Alliance of Canada.

9. References

- [1] J. Shi, W. Chen, D. Berrebbi, H.-H. Wang, W.-P. Huang, E.-P. Hu, H.-L. Chuang, X. Chang, Y. Tang, S.-W. Li, A. Mohamed, H.-Y. Lee, and S. Watanabe, "Findings of the 2023 ML-SUPERB challenge: Pre-training and evaluation over more languages and beyond," in *ASRU*, 2023, pp. 1–8.
- [2] N. Nagy, "Language contact and language change in the Faetar speech community," Ph.D. dissertation, University of Pennsylvania, 1996.
- [3] —, *Faetar*. Munich: Lincom Europa, 2000.
- [4] A. Zulato, J. Kasstan, and N. Nagy, "An overview of Francoprovençal vitality in Europe and North America," *International Journal of the Sociology of Language*, vol. 2018, no. 249, pp. 11–29, 2017.
- [5] K. Peterson, A. Tong, and Y. Yu, "OpenASR21: The second open challenge for automatic speech recognition of low-resource languages," in *Proc. Interspeech 2022*, 2022, pp. 4895–4899.
- [6] B. M. L. Srivastava, S. Sitaram, R. Kumar Mehta, K. Doss Mohan, P. Matani, S. Satpal, K. Bali, R. Srikanth, and N. Nayak, "Inter-speech 2018 low resource automatic speech recognition challenge for Indian languages," in *SLTU*, 2018, pp. 11–14.
- [7] S. T. Abate, W. Menzel, and B. Tafila, "An Amharic speech corpus for large vocabulary continuous speech recognition," in *Proc. Interspeech 2005*, 2005, pp. 1601–1604.
- [8] O. Kjartansson, S. Sarin, K. Pipatsrisawat, M. Jansche, and L. Ha, "Crowd-sourced speech corpora for Javanese, Sundanese, Sinhala, Nepali, and Bangladeshi Bengali," in *SLTU*, Gurugram, India, 2018, pp. 52–55.
- [9] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common Voice: A massively-multilingual speech corpus," in *LREC*. Marseille, France: European Language Resources Association, May 2020, pp. 4218–4222.
- [10] Y. Khassanov, S. Mussakhoyeva, A. Mirzakhmetov, A. Adiyev, M. Nurpeiissov, and H. A. Varol, "A crowdsourced open-source Kazakh speech corpus and initial speech recognition baseline," in *EACL*, P. Merlo, J. Tiedemann, and R. Tsarfaty, Eds. Online: Association for Computational Linguistics, 2021, pp. 697–706.
- [11] A. W. Black, "CMU Wilderness multilingual speech dataset," in *ICASSP*, 2019, pp. 5971–5975.
- [12] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "MLS: A large-scale multilingual dataset for speech research," in *Proc. Interspeech 2020*, 2020, pp. 2757–2761.
- [13] J. D. Amith and R. C. García, "Audio corpus of YoloXóchitl Mixtec with accompanying time-coded transcripts in ELAN."
- [14] J. D. Amith and O. López-Francisco, "Audio corpus of Totonac recordings from Northern Puebla and adjacent areas of Veracruz." [Online]. Available: <https://www.openslr.org/107/>
- [15] D. Berrebbi, J. Shi, B. Yan, O. López-Francisco, J. Amith, and S. Watanabe, "Combining spectral and self-supervised features for low resource speech recognition and translation," in *Proc. Interspeech 2022*, 2022, pp. 3533–3537.
- [16] I. Chizzoni and A. Vietti, "Towards an ASR System for Documenting Endangered Languages: A Preliminary Study on Sardinian," in *CLiC-it 2024: Tenth Italian Conference on Computational Linguistics*, F. Dell'Orletta, A. Lenci, S. Montemagni, and R. Sprugnoli, Eds., 2024.
- [17] C. Tsoukala, K. Kritsis, I. Douros, A. Katsamanis, N. Kokkas, V. Arampatzakis, V. Sevetlidis, S. Markantonatou, and G. Pavlidis, "ASR pipeline for low-resourced languages: A case study on Pomak," in *Proceedings of the Second Workshop on NLP Applications to Field Linguistics*, 2023, pp. 40–45.
- [18] T. Javed, K. Bhogale, A. Raman, P. Kumar, A. Kunchukuttan, and M. M. Khapra, "IndicSUPERB: A speech processing universal performance benchmark for Indian languages," *AAAI*, vol. 37, no. 11, pp. 12 942–12 950, 2023.
- [19] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *ACL-IJCNLP*, C. Zong, F. Xia, W. Li, and R. Navigli, Eds. Online: Association for Computational Linguistics, 2021, pp. 993–1003.
- [20] A. Conneau, M. Ma, S. Khanuja, Y. Zhang, V. Axelrod, S. Dalmia, J. Riesa, C. Rivera, and A. Bapna, "FLEURS: Few-shot learning evaluation of universal representations of speech," in *SLT*, 2023, pp. 798–805.
- [21] Y. Zhang, W. Han, J. Qin, Y. Wang, A. Bapna, Z. Chen, N. Chen, B. Li, V. Axelrod, G. Wang, Z. Meng, K. Hu, A. Rosenberg, R. Prabhavalkar, D. S. Park, P. Haghani, J. Riesa, G. Perng, H. Soltau, T. Strohman, B. Ramabhadran, T. Sainath, P. Moreno, C.-C. Chiu, J. Schalkwyk, F. Beaufays, and Y. Wu, "Google USM: Scaling automatic speech recognition beyond 100 languages," 2023.
- [22] E. Morris, R. Jimerson, and E. Prud'hommeaux, "One size does not fit all in resource-constrained ASR," in *Proc. Interspeech 2021*, 2021, pp. 4354–4358.
- [23] N. Nagy, "Lexical change and language contact: Faetar in Italy and Canada," *Journal of Sociolinguistics*, vol. 15, no. 3, pp. 366–382, 2011.
- [24] —, "A multilingual corpus to explore variation in language contact situations," *Rassegna Italiana di linguistica applicata*, pp. 65–84, 2011.
- [25] H. Brugman and A. Russel, "Annotating multi-media/multi-modal resources with ELAN," in *LREC*, M. T. Lino, M. F. Xavier, F. Ferreira, R. Costa, and R. Silva, Eds. Lisbon, Portugal: European Language Resources Association (ELRA), 2004.
- [26] V. Manohar, D. Povey, and S. Khudanpur, "JHU Kaldi system for Arabic MGB-3 ASR challenge using diarization, audio-transcript alignment and transfer learning," in *ASRU*, 2017, pp. 346–352.
- [27] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The Kaldi speech recognition toolkit," in *ASRU*. Hilton Waikoloa Village, Big Island, Hawaii, US: IEEE Signal Processing Society, 2011.
- [28] A. Plaquet and H. Bredin, "Powerset multi-class cross entropy loss for neural speaker diarization," in *Proc. INTERSPEECH 2023*, 2023.
- [29] H. Bredin, "pyannote.audio 2.1 speaker diarization pipeline: principle, benchmark, and recipe," in *Proc. INTERSPEECH 2023*, 2023.
- [30] S. F. Chen, D. Beeferman, and R. Rosenfeld, "Evaluation metrics for language models," Jan. 2008.
- [31] V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi, A. Baevski, Y. Adi, X. Zhang, W.-N. Hsu, A. Conneau, and M. Auli, "Scaling speech technology to 1,000+ languages," 2023.
- [32] M. Z. Boito, V. Iyer, N. Lagos, L. Besacier, and I. Calapode-scu, "mhubert-147: A compact multilingual hubert model," *arXiv preprint arXiv:2406.06371*, 2024.
- [33] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proc. Interspeech 2018*, 2018, pp. 2207–2211.
- [34] A. Ali, W. Magdy, P. Bell, and S. Renais, "Multi-reference WER for evaluating ASR for languages with no orthographic rules," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 576–580.
- [35] I. Nigmatulina, T. Kew, and T. Samarzic, "Asr for non-standardised languages with dialectal variation: the case of swiss german," in *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, 2020, pp. 15–24.
- [36] U. Barends, "Automatic orthography standardisation for under-resourced languages," Master's thesis, Stellenbosch University, 2024.