



HAL
open science

A dual task learning approach to fine-tune a multilingual semantic speech encoder for Spoken Language Understanding

Gaëlle Laperrière, Sahar Ghannay, Bassam Jabaian, Yannick Estève

► To cite this version:

Gaëlle Laperrière, Sahar Ghannay, Bassam Jabaian, Yannick Estève. A dual task learning approach to fine-tune a multilingual semantic speech encoder for Spoken Language Understanding. Interspeech 2024, Sep 2024, Kos, Greece. hal-04615074

HAL Id: hal-04615074

<https://hal.science/hal-04615074v1>

Submitted on 18 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A dual task learning approach to fine-tune a multilingual semantic speech encoder for Spoken Language Understanding

Gaëlle Laperrière^{1,2}, Sahar Ghannay², Bassam Jabaian¹, Yannick Estève¹

¹LIA - Avignon Université, France

²LISN - CNRS/Université Paris-Saclay, France

¹firstname.lastname@univ-avignon.fr, ²firstname.lastname@lisn.upsaclay.fr

Abstract

Self-Supervised Learning is vastly used to efficiently represent speech for Spoken Language Understanding, gradually replacing conventional approaches. Meanwhile, textual SSL models are proposed to encode language-agnostic semantics. SAMU-XLSR framework employed this semantic information to enrich multilingual speech representations. A recent study investigated SAMU-XLSR in-domain semantic enrichment by specializing it on downstream transcriptions, leading to state-of-the-art results on a challenging SLU task. This study’s interest lies in the loss of multilingual performances and lack of specific-semantics training induced by such specialization in close languages without any SLU implication. We also consider SAMU-XLSR’s loss of initial cross-lingual abilities due to a separate SLU fine-tuning. Therefore, this paper proposes a *dual* task learning approach to improve SAMU-XLSR semantic enrichment while considering distant languages for multilingual and language portability experiments.

Index Terms: Spoken language understanding, deep learning, self-supervised model, semantic speech representations, language portability, cross-lingual

1. Introduction

Spoken Language Understanding (SLU) implies extracting semantic information from speech signal [1]. This Natural Language Processing task can be modeled as named entity recognition or slot filling in a Human-Machine dialogue context.

This study focuses on a complex speech-to-concepts task by using end-to-end neural approaches as introduced by [2, 3, 4]. End-to-end approaches distinguish themselves from cascade approaches [5, 6, 7] by combining in a single model both Automatic Speech Recognition (ASR) and Natural Language Processing modules, resulting in less error propagation.

The main challenge of such an approach resides in the lack of paired speech recordings and semantic manual annotations. Many solutions have emerged, like transfer learning [8, 9, 10], data augmentation with speech synthesis [11, 12] and the use of pre-trained self-supervised models. Self-Supervised learning (SSL) enables the use of large amounts of unlabelled data. This method, previously used for ASR [13, 14], was proven very useful for SLU cascade approaches [15] by enhancing both speech representations [5, 6, 16] and annotations [14, 17].

However, combining text and speech SSL methods in a single end-to-end architecture is found very difficult due to the complexity of unifying the speech and textual representation spaces while optimizing a huge number of parameters. To exploit textual SSL models in end-to-end architectures, [18, 19] proposed to project ASR representations to a BERT model, while [10, 20] carried out an intent classification by using

sentence-level acoustic representations tied to a BERT SLU model. [21] proposed a similar multilingual approach.

Driven by this will to make use of speech SSL models in end-to-end architectures, [22] proposed SAMU-XLSR (Semantically-Aligned Multimodal Utterance-level Cross-Lingual Speech Representation). This new method produces semantically enriched multilingual and multimodal speech representations by combining the multilingual speech encoder XLS-R [23] to the Language Agnostic BERT Sentence Embedding generator LaBSE [17]. Lately, [24] has demonstrated the pertinence of the model’s semantically aware frame-level speech representations for a specific semantic extraction task on the challenging SLU MEDIA benchmark [25]. More recently, we pursued this analysis [26] by demonstrating the pertinence of another semantic specialization of SAMU-XLSR representations on a small amount of transcribed data linked to the SLU task with language portability.

In this research work, our main objective is to leverage the availability of transcribed speech data from closely related domains in multiple languages to enhance the performance of an SLU model in languages with limited SLU annotations. The use of the SAMU-XLSR model has proven to be particularly relevant in this case, but it seems perfectible. We assume that during the fine-tuning of the SAMU-XLSR model for an SLU task, it tends to forget its capacity to generate certain semantic abstractions at the utterance level prematurely. To limit this early forgetting, we propose a *dual* task learning approach when fine-tuning the SAMU-XLSR to an SLU downstream task, further on referred to as “*dual* fine-tuning”.

This approach is considered for the original French and Italian MEDIA benchmark used in [26] and the Tunisian TARIC-SLU dataset, freshly introduced by [27]. By the use of two very low resource datasets, we investigate the multilinguality and cross-linguality gains of such *dual* architecture compared to a more classic sequential approach. These experiments lead to state-of-the-art results with an end-to-end approach with all three datasets.

2. SLU tasks in different languages

In this study, we considered three datasets of different languages to conduct experiments on language portability and multilinguality: French (MEDIA), Italian (PortMEDIA) and the more distant Tunisian language (TARIC-SLU). All datasets were annotated for a complex semantic extraction task from speech, with close-domain labels. Our contributions focus on PortMEDIA and TARIC-SLU low-resource datasets, while MEDIA is mostly used to enhance their performances.

Table 1 presents the audio duration and words distribution in each corpus for each of the three datasets described below.

Table 1: Hours and Words distribution in MEDIA, PortMEDIA and TARIC-SLU.

		lang	Train	Valid	Test
Dur.	MEDIA	fr	10h52m	1h13m	3h01m
	PortMEDIA	it	7h18m	2h32m	4h51m
	TARIC-SLU	tu	7h30m	0h29m	0h54m
Wrd	MEDIA	fr	94.5k	10.8k	26.6k
	PortMEDIA	it	21.7k	7.7k	14.7k
	TARIC-SLU	tu	58.5k	3.5k	7.0k

2.1. The French MEDIA dataset

The French MEDIA dataset [28], part of ELRA’s MEDIA Evaluation Package^{1, 2}, is freely accessible for academic research, all recorded contributors having permitted its distribution.

The MEDIA benchmark is composed of recorded phone calls for hotel booking, transcribed and annotated with semantic concepts. The recorded speech consists of Human-Machine dialogues collected with the Wizard-of-Oz method. Only the user’s turns are fully annotated and considered for this study.

This dataset contains 1258 dialogues, from approximately 250 different French speakers. Our following experiments were made with the *full* MEDIA version, containing much richer semantic annotations, for a total of 152 different concepts in comparison to the 76 concepts of the *relax* version.

2.2. The Italian PortMEDIA dataset

The Italian PortMEDIA dataset [29] is part of the same ELRA package³ as MEDIA. Its annotations and transcriptions were collected in the same way as MEDIA’s, for the exact same task.

The dataset is composed of 604 dialogues from more than 150 different Italian speakers. It uses a total of 139 semantic concepts [30] and is only available as a *full* version. Note that this is the smallest corpus used in this study, considering the number of words in the training set.

2.3. The Tunisian TARIC-SLU dataset

The Tunisian TARIC-SLU dataset was sourced from the ASR TARIC one [31], dedicated to Tunisian Dialect Automatic Speech Recognition in the context of Human-Human dialogues. It has then been semantically annotated for a semantic extraction task from speech, on top of intents of the speech segments which we do not study in this paper.

The dataset’s domain of annotation varies slightly from MEDIA’s and PortMEDIA’s. Its speech recordings are conversations for train booking, of real conditions or acted dialogues. 809 out of 2’549 concepts’ values have code-switching, with an expected 80% of French words and 20% of English words.

A first version of TARIC-SLU, now updated and soon freely distributed, was proposed recently by [27]. TARIC-SLU is composed of more than 2,000 dialogues from 108 speakers, like its ASR version, and uses 60 different semantic concepts.

2.4. Evaluation Metrics

The semantic extraction task of this study aims to output the transcription of the speech signal and its annotation with semantic concepts as follows: *I <reservation> would like to book <room-number> one <room-type> double room >*.

¹<http://catalog.elra.info/en-us/repository/browse/ELRA-E0024/>

²ISLRN: 699-856-029-354-6

³<http://www.elra.info/en/projects/archived-projects/port-media/>

We consider two conventional metrics: the Concept Error Rate (CER) and the Concept Value Error Rate (CVER). Both metrics are computed like a Word Error Rate with the CER only accounting the hypothesis’ semantic concepts while the CVER considers the correctness of each value (“double room”) and concept (“<room-type>”) pair as a single occurrence.

A relevant CER improvement should differ by 0.4% considering MEDIA’s test set, 0.7% considering PortMEDIA’s test set, and 1.0% considering TARIC-SLU’s test set. Please note a possible 0.3% variation of Error Rates observed with 5 trainings of both SLU and *dual* models presented in Section 3.

3. SAMU-XLSR

In the domain of speech representation, self-supervised approaches have recently gained popularity over more conventional approaches such as filter-banks and MFCCs. Their aptitude to leverage large masses of unlabelled speech has been proved very useful for many tasks. During the last few years, many models like wav2vec 2.0 [13], HuBERT [16], and WavLM [32] have proven their efficiency in ASR [13], speaker verification [33, 34] and emotion recognition [35, 36]. Following this surge for SSL, [22] proposed the SAMU-XLSR approach for a translation task, aiming to capture the semantics in the speech signal by fine-tuning a pre-trained multilingual SSL model.

While higher-level semantics were already proven useful in speech-to-text mining and speech-to-text translation contexts, SAMU-XLSR’s approach has been proven pertinent in other tasks such as semantic extraction from speech [26], where we specialized SAMU-XLSR and extracted its frame-level speech representations for an SLU task.




Figure 1: Training and specialization process of SAMU-XLSR.

Figure 1 illustrates how SAMU-XLSR processes audio and text paired data. The pre-trained multilingual XLS-R [23] used in this approach was designed to generate speech representations for short 20 milliseconds speech frames. To make use of this model, SAMU-XLSR performs pooling and projection to create a single sentence-level representation. In parallel, LaBSE [17] sentence-level textual representations are simply extracted. Both representations being on the same semantic space, SAMU-XLSR’s is then being pulled towards LaBSE’s with the help of a cosine similarity loss function. This means the parameters

of all SAMU-XLSR’s components are optimized to predict the textual representations generated by the frozen LaBSE model.

In consideration to its semantic extraction performances, SAMU-XLSR has however been proven less apt to transcribe speech dialogues when specialized on multiple languages [26].

3.1. SLU fine-tuning

Fine-tuning SAMU-XLSR speech encoder on a specific SLU task initially required an SLU model, such as the one used in [26] and presented in Figure 2. This end-to-end model was used after SAMU-XLSR specialization on the downstream task, while this study presents an architecture to do both specialization and fine-tuning steps in one single training. We compared our *dual* approach to the previous architecture and tried out a second fine-tuning on TARIC-SLU in section 4.2.2. This SLU model was also necessary to alleviate the lack of baseline on the TARIC-SLU dataset.




Figure 2: Fine-tuning process of SAMU-XLSR for an SLU task.

The architecture consists of a fine-tuned speech encoder (specialized or dually fine-tuned SAMU-XLSR) generating frame-level speech representations of the input signal. These are given to 3 bi-LSTM layers of 1024 neurons, contextualizing the audio segments. A DNN layer of the same dimension, activated by LeakyReLU, feeds its outputs to a softmax function. We optimize a greedy loss with an Adam optimizer of $learning_rate = 0.0001$ for both speech encoder and Bi-LSTMs, and Adadelta with $learning_rate = 1.0$ for the DNN.

After trying out other upper modules with different dimensions to decode the speech representations, we decided to keep the bi-LSTMs, proven more efficient than DNNs for this task.

3.2. Dual task learning

SAMU-XLSR approach allows better capture of semantics directly from the speech signal for multiple languages [24, 22].

This study aims to improve the specialization and fine-tuning processes on multilinguality and specific semantics with the help of distant languages and SLU modality. By merging the SLU fine-tuning and specialization, we expect to prevent suppressing SAMU-XLSR capacities on capturing internal utterance-level semantic abstractions during the fine-tuning of the SLU downstream task. To address this phenomenon of forgetting, we propose a *dual* task learning approach, wherein the SAMU-XLSR model is fine-tuned to:

- Generate utterance-level embeddings, similar to its initial

pre-training, but only for utterances relevant to the SLU downstream task, potentially across multiple languages.

- Perform the SLU downstream task in the targeted language.

Figure 3 schematizes the *dual* architecture with the speech encoder connection to both SAMU-XLSR and SLU upper modules. For comparison purposes, both architectures presented before have been reused in this single model.

The only hyper parameter search done for this architecture concerns the loss distribution between both modules. We considered this distribution as follows: $loss = loss(SAMU-XLSR) + \lambda loss(SLU)$, with multiple λ tested in the interval $[0; 20]$, especially between 0 and 1.

4. Experimental results

This section presents experimental results of the *dual* approach (SAMU-XLSR_{dual L} trained on languages L) by comparing them to a simple specialization (SAMU-XLSR_L) followed by a necessary SLU training (SLU_L). Datasets were not shuffled during training, having previously experimented on multiple shuffling methods without evident benefits from them. Note that only SAMU-XLSR was considered for this study, being proven the current most pertinent speech encoder for these SLU tasks, with presently the best results on MEDIA and PortMEDIA [26].

4.1. Task-oriented semantic enrichment

As we saw a frank improvement of CER in giving more importance to the SLU loss than SAMU-XLSR’s loss for MEDIA experiments, we focused on optimizing the speech encoder mostly on the SLU task while specializing it more slightly on the SAMU-XLSR task. However, contrary to experiments with a λ in a $[0; 1]$ interval, no linear improvements were observed when this coefficient was defined between 1 and 20. Therefore, all results were obtained by running experiments with multiple λ in the $[1; 20]$ interval and keeping the best system considering the development set. Results of monolingual *dual* fine-tuning on each dataset are presented for the test set in Table 2.

Table 2: Experimental results on dual approach for a monolingual fine-tuning, compared to a classic specialization followed by an SLU fine-tuning.

		CER	CVER
fr	SAMU-XLSR + SLU _{FR}	18.7	29.4
	SAMU-XLSR _{FR} + SLU _{FR} [26]	18.5	29.5
	SAMU-XLSR _{FR dual}	18.3	27.8
it	SAMU-XLSR + SLU _{IT} [26]	26.6	39.2
	SAMU-XLSR _{IT} + SLU _{IT} [26]	26.8	39.5
	SAMU-XLSR _{IT dual}	26.8	39.4
tu	SAMU-XLSR + SLU _{TU}	30.7	47.4
	SAMU-XLSR _{TU} + SLU _{TU}	30.3	45.2
	SAMU-XLSR _{TU dual}	32.4	48.3
	SAMU-XLSR _{TU dual} + SLU _{TU}	29.9	46.8

The *dual* task learning leads to relatively baseline equivalent CERs for monolingual experiments. The main advantage of this method resides in the number of parameters learned on a single v100-32G GPU during 100 epochs for each training. SAMU-XLSR specialization optimizes 316.2M parameters, in addition to 387.8M for SLU fine-tuning, while the *dual* fine-tuning only optimizes 385.6M parameters. Note that the Tunisian task requires a supplementary SLU fine-tuning to yield pertinent results. This can be explained by the lack of Tunisian data in SAMU-XLSR’s original training, leading to its seemingly less noteworthy fine-tuning. Indeed, where train-




Figure 3: Training process of the SLU and SAMU-XLSR modules combined in a dual architecture.

ing on French and Italian means fine-tuning SAMU-XLSR on these languages, training on Tunisian means a first-time SAMU-XLSR training on this dialect, done at the same time as an SLU training. Multilingual *dual* fine-tuning will prove itself efficient on never-seen data such as Tunisian in the following language portability experiments.

4.2. Language portability

State-of-the-art results for all three datasets were obtained with language portability. We can distinguish experiments on close languages like French and Italian, and ones on distant languages with the use of these two to improve performances on the Tunisian TARIC-SLU task.

4.2.1. Close languages

Close-language specialization was conducted in [26] to make use of in-domain data. This paper studies the use of the *dual* architecture in a multilingual context. Table 3 shows experimental results with French and Italian mixed fine-tunings.

Table 3: Experimental results on dual approach with close-language portability.

		CER	CVER
fr	SAMU-XLSR $FR \oplus IT$ + SLU FR	18.6	29.1
	SAMU-XLSR $FR \oplus IT$ <i>dual</i>	17.9	28.2
it	SAMU-XLSR $FR \oplus IT$ + SLU $FR \rightarrow IT$ [26]	25.1	38.1
	SAMU-XLSR $FR \oplus IT$ <i>dual</i>	24.1	39.0

These experiments lead to a pertinent improvement of state-of-the-art CER results on MEDIA full version task with a score of **17.9%** and on PortMEDIA with a score of **24.1%**.

4.2.2. Distant languages

The main limitation experimented in [26] was the lack of distant-language data to use during specializations. This paper studies the Tunisian language on which SAMU-XLSR has never been trained. Table 4 presents experimental results for this low-resource dataset with the use of French or Italian data.

During all *dual* fine-tunings on Tunisian data, models hardly succeeded to individually focus on Tunisian while learning to represent semantics in French and Italian. However, both languages were highly beneficial to TARIC-SLU process-

Table 4: Experimental results on dual approach with distant-language portability for the Tunisian TARIC-SLU dataset.

	CER	CVER
SAMU-XLSR $FR \oplus IT \oplus TU$ + SLU TU	31.9	47.0
SAMU-XLSR FR <i>dual</i> + SLU TU	30.3	46.4
SAMU-XLSR IT <i>dual</i> + SLU TU	30.8	46.4
SAMU-XLSR $FR \oplus IT$ <i>dual</i> + SLU TU	30.4	48.4
SAMU-XLSR $FR \oplus IT \oplus TU$ <i>dual</i> + SLU TU	29.1	46.2

ing with another SLU fine-tuning on Tunisian, as shown by the state-of-the-art **29.1%** CER obtained with a *dual* fine-tuning on all three datasets, in comparison to the 29.9% CER with the Tunisian-only *dual* fine-tuning of Table 2.

5. Conclusion

This paper investigates a new approach to improve the semantic enhancement of SAMU-XLSR speech representations for a complex SLU task. The *dual* architecture proposed combines both SAMU-XLSR’s original training process and an SLU module for semantic extraction, implying almost two times less parameters to optimize. This new architecture aims to prevent an expected loss of multilinguality and cross-linguality of SAMU-XLSR first pre-training, while filling in the lack of specific semantics focus during a classic specialization. Contrarily to the method proposed in [26], this study experiments on distant language’s portability such as French and Italian for a targeted low-resource Tunisian task. This induced another challenge through the lack of Tunisian data in SAMU-XLSR original pre-training. However, the *dual* fine-tuning, followed by a necessary SLU fine-tuning on this language leads to state-of-the-art 29.1% CER. At the same time, with a single *dual* training, MEDIA and PortMEDIA datasets outperformed state-of-the-art results with, respectively, 17.9% and 24.1% CER scores. One limitation of this paper could be the omission of out-of-domain or less specialized data during the *dual* fine-tuning, which was not experimented due to our focus on improving the already proven efficient SAMU-XLSR specialization and SLU fine-tuning on MEDIA and PortMEDIA. However, the *dual* task learning approach experimented on Tunisian SLU opens promising perspectives for other never-seen low-resource languages during the pre-training of SSL models, with a way to expand or specialize their multilinguality disposition.

6. Acknowledgements

This work was performed using HPC resources from GENCI/IDRIS (grant 2022 AD011012565) and received funding from the EU H2020 SELMA (grant No 957017), ESPERANTO (grant No 101007666) and PSPC AIDA: 2019-PSPC-09 (funded by BPI-France) projects.

7. References

- [1] G. Tur and R. De Mori, *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons, 2011.
- [2] S. Ghannay, A. Caubriere, Y. Estève, N. Camelin, E. Simonnet, A. Laurent, and E. Morin, “End-to-end named entity and semantic concept extraction from speech,” in *IEEE SLT*, 2018.
- [3] P. Haghani, A. Narayanan, M. Bacchiani, G. Chuang, N. Gaur, P. Moreno, R. Prabhavalkar, Z. Qu, and A. Waters, “From audio to semantics: Approaches to end-to-end spoken language understanding,” in *IEEE SLT*, 2018.
- [4] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, “Towards end-to-end spoken language understanding,” in *IEEE ICASSP*, 2018.
- [5] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, “Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders,” in *IEEE ICASSP*, 2020.
- [6] A. T. Liu, S.-W. Li, and H.-y. Lee, “Tera: Self-supervised learning of transformer encoder representation for speech,” *IEEE/ACM TASLP*, 2021.
- [7] S. Arora, S. Dalmia, F. Metze, A. Black, and S. Watanabe, “Token-level sequence labeling for spoken language understanding using compositional end-to-end models,” in *EMNLP*, 2022.
- [8] S. Bhosale, I. Sheikh, S. Dumpala, and S. Kopparapu, “End-to-End Spoken Language Understanding: Bootstrapping in Low Resource Scenarios,” in *Interspeech*, 2019.
- [9] A. Caubriere, N. Tomashenko, A. Laurent, E. Morin, N. Camelin, and Y. Estève, “Curriculum-Based Transfer Learning for an Effective End-to-End Spoken Language Understanding and Domain Portability,” in *Interspeech*, 2019.
- [10] Y. Huang, H. Kuo, S. Thomas, Z. Kons, K. Audhkhasi, B. Kingsbury, R. Hoory, and M. Picheny, “Leveraging unpaired text data for training end-to-end speech-to-intent systems,” in *IEEE ICASSP*, 2020.
- [11] T. Desot, F. Portet, and M. Vacher, “Corpus generation for voice command in smart home and the effect of speech synthesis on End-to-End SLU,” in *LREC*, 2020.
- [12] L. Lugosch, B. Meyer, D. Nowrouzezahrai, and M. Ravanelli, “Using speech synthesis to train end-to-end spoken language understanding models,” in *IEEE ICASSP*, 2020.
- [13] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *NeurIPS*, 2020.
- [14] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *NAACL-HLT*, 2019.
- [15] G. Laperriere, B. Jabaian, and Y. Esteve, “Where Are We in Semantic Concept Extraction for Spoken Language Understanding?” in *SPECOM*, 2021.
- [16] W. Hsu, B. Bolte, Y. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM TASLP*, 2021.
- [17] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, and W. Wang, “Language-agnostic BERT Sentence Embedding,” in *ACL*, 2022.
- [18] P. Wang, L. Wei, Y. Cao, J. Xie, and Z. Nie, “Large-scale unsupervised pre-training for end-to-end spoken language understanding,” in *IEEE ICASSP*, 2020.
- [19] Y. Chung, C. Zhu, and M. Zeng, “SPLAT: Speech-Language Joint Pre-Training for Spoken Language Understanding,” in *NAACL-HLT*, 2021.
- [20] B. Agrawal, M. Müller, S. Choudhary, M. Radfar, A. Mouchtaris, R. McGowan, N. Susanj, and S. Kunzmann, “Tie your embeddings down: Cross-modal latent spaces for end-to-end spoken language understanding,” in *IEEE ICASSP*, 2022.
- [21] M. Müller, S. Choudhary, C. Chung, A. Mouchtaris, and S. Kunzmann, “In Pursuit of Babel-Multilingual End-to-End Spoken Language Understanding,” in *IEEE ASRU*, 2021.
- [22] S. Khurana, A. Laurent, and J. Glass, “SAMU-XLSR: Semantically-Aligned Multimodal Utterance-level Cross-Lingual Speech Representation,” *IEEE JSTSP*, 2022.
- [23] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, “XLS-R: self-supervised cross-lingual speech representation learning at scale,” in *Interspeech*, 2022.
- [24] G. Laperrière, V. Pelloin, M. Rouvier, T. Stafylakis, and Y. Estève, “On the Use of Semantically-Aligned Speech Representations for Spoken Language Understanding,” in *IEEE SLT*, 2023.
- [25] F. Béchet and C. Raymond, “Benchmarking Benchmarks: Introducing New Automatic Indicators for Benchmarking Spoken Language Understanding Corpora,” *Interspeech*, 2019.
- [26] G. Laperrière, H. Nguyen, S. Ghannay, B. Jabaian, and Y. Estève, “Semantic enrichment towards efficient speech representations,” in *Interspeech*, 2023.
- [27] S. Mdhaffar, F. Bougares, R. De Mori, S. Zaiem, M. Ravanelli, and Y. Estève, “TARIC-SLU: A Tunisian Benchmark Dataset For Spoken Language Understanding,” in *LREC-COLING*, 2024.
- [28] H. Bonneau-Maynard, S. Rosset, C. Ayache, A. Kuhn, and D. Mostefa, “Semantic annotation of the French media dialog corpus,” in *Interspeech*, 2005.
- [29] F. Lefèvre, D. Mostefa, L. Besacier, Y. Estève, M. Quignard, N. Camelin, B. Favre, B. Jabaian, and L. Rojas-Barahona, “Leveraging study of robustness and portability of spoken language understanding systems across languages and domains: the PORT-MEDIA corpora,” in *LREC*, 2012.
- [30] H. Bonneau-Maynard, C. Ayache, F. Bechet, A. Denis, A. Kuhn, F. Lefevre, D. Mostefa, M. Quignard, S. Rosset, C. Servan, and J. Villaneau, “Results of the French evalda-media evaluation campaign for literal understanding,” in *LREC*, 2006.
- [31] A. Masmoudi, M. Khmekhem, Y. Estève, L. Belguith, and N. Habash, “A Corpus and Phonetic Dictionary for Tunisian Arabic Speech Recognition,” in *LREC*, 2014.
- [32] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE JSTSP*, 2022.
- [33] Z. Chen, S. Chen, Y. Wu, Y. Qian, C. Wang, S. Liu, Y. Qian, and M. Zeng, “Large-scale self-supervised speech representation learning for automatic speaker verification,” in *IEEE ICASSP*, 2022.
- [34] S. Chen, Y. Wu, C. Wang, S. Liu, Z. Chen, P. Wang, G. Liu, J. Li, J. Wu, X. Yu, and F. Wei, “Why does Self-Supervised Learning for Speech Recognition Benefit Speaker Recognition?” in *Interspeech*, 2022.
- [35] M. Macary, M. Tahon, Y. Estève, and A. Rousseau, “On the use of self-supervised pre-trained acoustic and linguistic features for continuous speech emotion recognition,” in *IEEE SLT*, 2021.
- [36] L. Pepino, P. Riera, and L. Ferrer, “Emotion Recognition from Speech Using wav2vec 2.0 Embeddings,” in *Interspeech*, 2021.