

EFFECTIVENESS OF TEXT, ACOUSTIC, AND LATTICE-BASED REPRESENTATIONS IN SPOKEN LANGUAGE UNDERSTANDING TASKS

Esau Villatoro-Tello^{*1}, Srikanth Madikeri^{*1}, Juan Zuluaga-Gomez^{1,2}, Bidisha Sharma⁵,
Seyyed Saeed Sarfjoo¹, Iuliia Nigmatulina^{1,3}, Petr Motlicek^{1,4}, Alexei V. Ivanov⁵, Aravind Ganapathiraju⁵

¹ Idiap Research Institute, Martigny, Switzerland.

² Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland.

³ Institute of Computational Linguistics, University of Zürich, Switzerland.

⁴ Brno University of Technology, Brno, Czech Republic

⁵ Uniphore Software Systems Inc., Palo Alto, CA, USA.

ABSTRACT

In this paper, we perform an exhaustive evaluation of different representations to address the intent classification problem in a Spoken Language Understanding (SLU) setup. We benchmark three types of systems to perform the SLU intent detection task: 1) text-based, 2) lattice-based, and a novel 3) multimodal approach. Our work provides a comprehensive analysis of what could be the achievable performance of different state-of-the-art SLU systems under different circumstances, e.g., automatically- vs. manually-generated transcripts. We evaluate the systems on the publicly available SLURP spoken language resource corpus. Our results indicate that using richer forms of Automatic Speech Recognition (ASR) outputs, namely word-consensus-networks, allows the SLU system to improve in comparison to the 1-best setup (5.5% relative improvement). However, crossmodal approaches, i.e., learning from acoustic and text embeddings, obtains performance similar to the oracle setup, a relative improvement of 17.8% over the 1-best configuration, being a recommended alternative to overcome the limitations of working with automatically generated transcripts.

Index Terms— Speech Recognition, Human-computer Interaction, Spoken Language Understanding, Word Consensus Networks, Cross-modal Attention

1. INTRODUCTION

Spoken Language Understanding (SLU) is the underlying key component of interactive smart devices such as voice assistants, social bots, and intelligent home devices. Effectively interpreting human interactions through classification of intent and slot filling plays a crucial role in SLU. Therefore, it is not surprising that the SLU problem has received substantial attention in industry and academia.

Typically, SLU aims at parsing spoken utterances into corresponding structured semantic concepts through a pipeline approach. First, the spoken utterances are transcribed by an automatic speech recognition (ASR) system, and the transcribed audio is subsequently processed by a natural language understanding (NLU) module to identify the intent and extract slots from the utterance.¹ The main disadvantages of the pipeline approach include: (1) errors in the

ASR transcripts are directly propagated to the NLU module, normally trained only on correct transcriptions; (2) prosodical and non-phonetic aspects present in the spoken utterance are not taken into account. Even though, the classical text-based approach is mostly used in industrial applications and is still an active research area [1].

More recently, end-to-end (E2E) SLU systems have gained popularity [2–4]. E2E SLU acts as an individual single model, and it directly predicts the intent from speech without exploiting an intermediate text representation. In particular, it directly optimizes the performance metrics of SLU. Due to the complex structure of speech signals, a large SLU database along with high-end computational resources are required for training E2E models. In [4], several E2E SLU encoder-decoder solutions are investigated. For instance, instead of directly mapping speech to SLU target [2], pre-trained acoustic and language models can be used for downstream SLU tasks, showing to be an effective paradigm [5, 6]. Similarly, there have been efforts to design tighter integration of ASR and NLU systems beyond 1-best ASR results, e.g., by means of encoding several ASR hypotheses through lattice-based representations. A lattice is a compact representation encoding multiple ASR hypotheses obtained at the decoding step. Its use has shown to be key in boosting the performance of IR systems [7]. In this direction, there are several works adopting word confusion networks (WCNs) as input to NLU systems to preserve information in possible hypotheses [8–10]. The main advantage of WCN-based approaches is that they are less sensitive to the ASR errors. Finally, recent approaches based on multi-modal information have been proposed [3]. The main motivation behind this idea is founded on how humans interpret, in the real world, the meaning of an utterance and corresponding semantics from various cues, thus, assuming that the acoustic and linguistic content of a speech signal may carry complementary information for deriving robust semantic information of an utterance.

Overall, despite the promising results, there still exists a gap between the demonstrated capability of SLU systems and the requirements of an industrial application, e.g., a generalized voice assistant. For instance, E2E approaches mostly focus on databases with limited semantic complexity and structural diversity [11]. Additionally, most of the current benchmarks on SLU are widely saturated, where the obtained performances (F1-scores) are near perfect. Examples of such cases are results reported on ATIS [12], Fluent Speech Commands [6], or SNIPS [13] datasets. Hence, in order to validate the robustness of recent SLU approaches under a more realistic scenario, it becomes necessary to focus on SLU tasks that incorporate more

^{*}Corresponding authors: {*esau.villatoro, srikanth.madikeri*}@idiap.ch

This work was supported by the Idiap & Uniphore collaboration project.

¹In practice, the 1-best transcript representation is the one sent to the natural language understanding (NLU) model for intent detection.

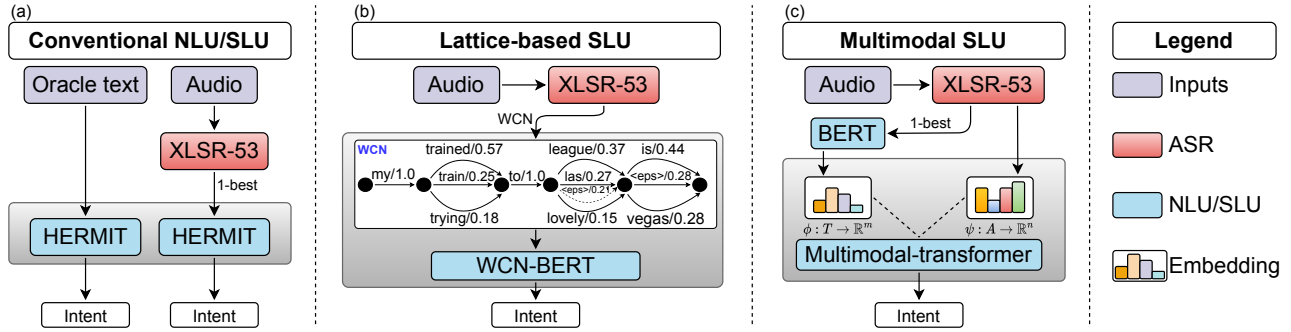


Fig. 1. Overview of the considered NLU/SLU methodologies for our performed experiments.

complex semantics and numerous intent classes and slots. To the best of our knowledge, such benchmarking, comparing the wide variety of SLU approaches, has not been performed recently.

In this paper, we present an extensive analysis of different SLU techniques, ranging from pure text-based alternatives to methods that are able to process richer forms of ASR outputs. Overall, our work has three salient features: (1) we evaluate and compare under the same circumstances four big families of SLU approaches, namely: text-based, lattice-based, multimodal, and end-to-end, (2) our performed evaluation is done considering a more realistic scenario, i.e., where no access to manual transcriptions exists, but instead, ASR transcripts are given as input to the SLU systems, and (3) we describe several inconsistencies found in the SLURP [14] dataset, one of the most challenging test beds for SLU systems.²

2. METHODOLOGY

Figure 1 depicts an overview of the considered SLU techniques in our experiments: (a) conventional or pipeline-oriented NLU/SLU approaches, (b) Lattice-based SLU architectures, and (c) multimodal (text+acoustic) architectures. Although not shown in the figure, we also report the performance of very recent E2E methods.

2.1. Conventional NLU/SLU systems

We selected the HERMIT architecture [15] as the representative approach for this category of systems. HERMIT, a **HiER**archical **Mu**lti-**T**ask Natural Language Understanding architecture, was designed for effective semantic parsing of domain-independent user utterances, extracting meaning representations in terms of high-level intents and frame-like semantic structures. According to the authors, HERMIT stands out for being a cross-domain, multi-task architecture, capable of recognizing multiple intents in human-machine interactions. The central motivation behind the design of the HERMIT architecture is the modeling of the dependence among three tasks, namely, dialogue acts identification, intents, and slots. For this, the authors addressed the NLU problem using a seq2seq model employing BiLSTM encoders and self-attention mechanisms, followed by CRF tagging layers. HERMIT was validated in two large datasets with a high number of intent labels (58 to 68 classes), reporting a performance of F1=86%.

We re-implemented HERMIT in PyTorch [16], with the following changes: we exchanged the encoder layer based on ELMO em-

beddings with a BERT [17] encoder, we replace the BiLSTM encoders by GRU modules and used the AdamW optimizer. We evaluate the performance of our implementation of HERMIT when either, manual transcriptions (1-best) ASR outputs (extracted from the XLS-R model, see Section 2.4) are given as inputs (see Figure 1.a).

2.2. Lattice-based SLU

As described earlier, one main limitation of pipeline SLU systems is their sensitivity to the errors present in the ASR transcriptions. Consequently, there have been proposed several approaches for building SLU systems robust against ASR errors based on lattices and WCNs [9, 10]. Although both, word lattices and WCNs contain more information than N-best lists, WCNs have been proven more efficient in terms of size and structure, thus representing a more plausible alternative when designing SLU systems that receive as input a graph-based structure.

We re-implemented a very recent WCN-based approach, namely WCN-BERT [9]. Originally, the WCN-BERT architecture consists of three parts: a BERT encoder for jointly encoding, an utterance representation model, and an output layer for predicting semantic tuples. The BERT encoder exploits posterior probabilities of word candidates in WCNs to inject ASR confidences. Multi-head self-attention is applied over both WCNs and system acts to learn context-aware hidden states. The utterance representation model produces an utterance-level vector by aggregating final hidden vectors. Finally, WCN-BERT adds both discriminative and generative output layers to predict semantic tuples. WCN-BERT stands out for being able to leverage the timing information and confidence scores that are part of standard WCNs. Authors evaluated the performance of WCN-BERT on DSTC2 dataset [18], a corpus of dialogs in the restaurant search domain between a human and automatic agent (i.e., human-machine conversations) reporting and overall F1=87.91%.

For our experiments, we dropped the semantic tuple classifier and the transformer-based hierarchical decoder proposed in the original WCN-BERT paper [9]. We only preserve the WCN-BERT encoder and the multi-head attention layers to generate the utterance-level representation. On top of this, we concatenate a fully-connected layer to perform intent classification.

2.3. Multimodal SLU

Multimodal SLU refers to the process of embeddings alignment for explicitly minimizing the distance between speech embeddings and the text embeddings from state-of-the-art text encoders like BERT [17]. Thus, the speech embeddings that are used for downstream

²Our code is publicly available: https://github.com/idiap/slu_representations

tasks are made to share a common embedding space with the textual embeddings, leading to better performance in SLU tasks, e.g., intent detection. However, there are a few challenges involved in the process of modeling such multimodal human language time-series, namely: 1) inherent non-aligned data due to variable sampling rates for the sequences from each modality; and 2) long-range dependencies between elements across modalities.

In order to address these problems, we implemented a solution based on a recent approach named Multimodal Transformer (MulT) [19]. MulT depicts an end-to-end model that extends the standard Transformer network [20] to learn representations directly from unaligned multimodal streams. At the heart of MulT, there is a cross-modal attention module, which attends to the crossmodal interactions at the scale of the entire utterances. It merges multimodal time-series via a feed-forward fusion process from multiple directional pairwise crossmodal transformers. Specifically, each cross-modal transformer serves to repeatedly reinforce a target modality with the low-level features from another source modality by learning the attention across the two modalities’ features.

In our experiments, we adopted the ideas proposed in MulT [19]. Hence, given two input modalities, each crossmodal transformer block (one for each modality) keeps updating its sequence. Thus, the crossmodal transformer learns to correlate meaningful elements across modalities. As a final step, outputs are concatenated and passed through a self-attention module to collect temporal information to make predictions. The last elements of the sequence are passed through fully-connected layers to make the intent prediction.

2.4. XLSR-53: Automatic Speech Recognition Module

As shown in Figure 1, we consider as main ASR component the XLSR-53 pre-trained acoustic model [21]. XLSR-53 learns cross-lingual speech embeddings by pretraining a single generic model from raw waveform of speech in multiple languages. The structure of XLSR is similar to Wav2Vec 2.0 [22], which is trained using contrastive loss over masked latent speech representations and jointly learns a quantization of the latent embeddings shared across languages. XLSR-53 model is trained using 56,000 hours of untranscribed audio data from 53 languages. We then fine-tune XLSR-53 model [21] with 390 hours of English data from AMI and Switchboard datasets using E2E-LFMMI loss function [23, 24] with biphone units [25–27]. A grapheme-based lexicon of size 1M was used, and the language model (LM) was trained with 34M utterances from publicly available English datasets including People’s speech, Fisher, Switchboard, AMI, Wikitext103, and subsets of Common Crawl and Reddit datasets. For improving the generalization of XLSR-53 model to conversational speech in English language, we fine-tuned it using 560 hours of untranscribed data crawled from YouTube. This subset was selected from conversational video calls in English language. On the YouTube data, we followed an incremental semi-supervised learning approach with four iterations [28]. For decoding, we use the WFST decoder from Kaldi [29] toolkit with a beam width of 15.

3. EXPERIMENTS AND RESULTS

3.1. SLURP Dataset

To perform our experiments we used the SLURP dataset [14], a publicly available multi-domain dataset for E2E-SLU, which is substantially bigger and more diverse than other SLU resources. SLURP is

Table 1. SLURP statistics. **SLURP_O** the original dataset, while **SLURP_F** a cleaner version of the original SLURP data.

Statistics	SLURP _O	SLURP _F
Audio Files	72,277	50,568
↔ Close range	34,603	25,799
↔ Far range	37,674	24,769
Duration [hr]	58	37.2
Av. length [s]	2.9	2.6
Nb. of intents	48	47

a collection of audio recordings of single-turn user interactions with a home assistant. Table 1 contains a few statistics about SLURP.

During a manual analysis, we found many inconsistencies in SLURP annotations. Basically, we identified cases where the manual transcription does not correspond to what is being said in its corresponding audio file. Thus, we considered as erroneous those audio files for which the manual transcription did not match with the automatic transcription, or whose transcripts were inconsistent (i.e., not the same) in the corresponding metadata files. By following this approach, we detected that nearly 30% (20 hours) of the original data contains some type of inconsistency. We refer as SLURP_F to the subset of SLURP without these inconsistent files (see Table 1).

3.2. Experiments

Experiments from **EXP1 - EXP7** correspond to the results of conventional NLU/SLU techniques. As described in Section 2.1, these experiments use our implementation of the HERMIT architecture [15]. As can be observed in Table 2, differences among these experiments are on the type of data used for training and evaluating the HERMIT model, i.e., combinations of either manual transcripts or 1-best ASR outputs. **EXP8 - EXP9** correspond to the experiments done using WCN-based representations (Section 2.2). Notice that for both set of experiments, i.e., conventional NLU and WCN-based, we use the XLS-R model, not-adapted and adapted to SLURP, to obtain the transcripts and WCNs respectively. Finally, **EXP10-EXP12** corresponds to the experiments done using the crossmodal transformer described in Section 2.3. Similarly, we evaluate the performance of this approach under circumstances where the XLS-R model is not adapted to the target domain (EXP10), and when adapted to SLURP (EXP11), and one last experiment using acoustic embeddings obtained from HuBERT (EXP12) model.³

3.3. Results

Table 2 shows the obtained experimental results for all the benchmarked architectures. Column “Exp” indicates the name of the experiment, “Input Type” describes what type of data was used for training and evaluating (dev and test) the corresponding experiment. For those experiments with the tag *manual* it means ground truth transcriptions were used, while *1-best* refers to the automatically generated transcriptions using the XLSR-53 model. Column “XLSR-53 adaptation” indicates whether or not the XLSR-53 model was fine-tuned to the SLURP dataset. In order to do the XLSR-53 adaptation, the English ASR model described in Section 2.4 was fine-tuned with the train subset of the SLURP_F data without changing the LM. ASR performances before and after fine-tuning to

³To generate the acoustic embeddings we followed the SpeechBrain [31] SLURP recipe: <https://github.com/speechbrain/speechbrain/tree/develop/recipes/SLURP>

Table 2. Accuracy (ACC) and F1-scores (F1) on intent classification for different representations. We test our approach with either manual or 1-best approaches. Manual refers to ground truth evaluation, while 1-best is obtained by using our ASR module XLSR-53. Thus, manual \rightarrow manual represents the oracle scenario (upper bound), while 1-best \rightarrow 1-best depicts a more real-world case scenario.

Exp.	Input Type	XLSR-53	SLURP _O				SLURP _F			
			Dev (\uparrow)		Test (\uparrow)		Dev (\uparrow)		Test (\uparrow)	
	<i>train</i> \rightarrow <i>dev-test</i>	adaptation	ACC	F1	ACC	F1	ACC	F1	ACC	F1
Conventional NLU/SLU										
EXP1	manual \rightarrow manual	NA	0.89	0.88	0.85	0.84	0.88	0.87	0.82	0.82
EXP2	manual \rightarrow 1-best	\times	0.70	0.65	0.69	0.65	0.74	0.69	0.71	0.67
EXP3	1-best \rightarrow manual	\times	0.85	0.85	0.86	0.85	0.86	0.86	0.85	0.83
EXP4	1-best \rightarrow 1-best	\times	0.72	0.68	0.73	0.69	0.76	0.71	0.77	0.73
EXP5	manual \rightarrow 1-best	\checkmark	0.82	0.81	0.80	0.79	0.84	0.82	0.86	0.84
EXP6	1-best \rightarrow manual	\checkmark	0.88	0.87	0.87	0.86	0.88	0.87	0.88	0.87
EXP7	1-best \rightarrow 1-best	\checkmark	0.84	0.83	0.83	0.83	0.85	0.84	0.85	0.84
Lattice-based SLU										
EXP8	WCN	\times	0.68	0.67	0.68	0.68	0.69	0.68	0.68	0.68
EXP9	WCN	\checkmark	0.78	0.77	0.79	0.79	0.80	0.80	0.78	0.77
Multimodal SLU										
EXP10	multimodal	\times	0.75	0.75	0.74	0.73	0.75	0.74	0.76	0.76
EXP11	multimodal	\checkmark	0.82	0.82	0.83	0.83	0.83	0.83	0.82	0.82
EXP12	multimodal (HuBERT [30])	\checkmark	0.87	0.88	0.84	0.84	0.88	0.88	0.86	0.86

Table 3. WER% on SLURP Test sets with the XLSR-53 English model before and after adaptation with SLURP_F train subset.

System	Dev (WER%)		Test (WER%)	
	Headset	All	Headset	all
No adaptation	23.4	34.0	23.0	34.4
Adapted to SLURP	13.3	16.1	13.0	15.5

SLURP are given in Table 3. And finally, SLURP_O and SLURP_F depict what version of the SLURP dataset was used.

Notice that the results obtained in the test partition from the cleaned version of the data, i.e., SLURP_F, are usually better than those obtained in the original version of the SLURP dataset. To some extent, this is an indicator that the identified inconsistencies in the original SLURP dataset were affecting the benchmarked models, resulting in a miss classification of some intents types.

Experiments EXP1, EXP3, and EXP6 represent artificial scenarios, as in a real-world application we do not expect to have manual transcripts for test partitions. Nevertheless, the best performance under this configuration, e.g., F1=87% in the SLURP_F for EXP6, represents an upper bound value. Interestingly, this value is even better than the performance obtained in the EXP1, i.e., considering only ground truth data. This may be due to an (unexpected) regularization effect caused by the noise contained in the 1-best transcripts from all the audios of the SLURP. Thus, it becomes really relevant that WCN-based approaches (EXP8 & EXP9) are able to obtain a competitive performance against the pipeline 1-best \rightarrow 1-best NLU experiments (EXP4). Even though the not-adapted WCN model (EXP8) does not outperform EXP4, this result validates the existence and the impact of richer ASR hypotheses in the lattice, which helps improve the performance of the SLU system, especially in noisy data (SLURP_O).

Although WCN experiment EXP9 showed a good improvement against the 1-best scenario, multimodal experiments obtain a

remarkable performance, comparable to the performance of the oracle experiment (EXP1). The main difference between EXP11 and EXP12 is that the former uses XLSR-53 adapted to SLURP, while the latter fine-tunes HuBERT toward intent and slot classification in SLURP. This is the explanation for why EXP12 performance is slightly better than EXP11 (XLSR-53 adapted). Finally, as reference results from an E2E approach, the SLURP recipe reports F1 values of F1= 0.77 and F1= 0.88 under configurations referred to as *direct* and *direct Hubert* respectively. More details can be found in the respective repository.³ Overall, using a multi-modal approach seems to be the recommended option as it guarantees the best performance. Although it should be considered that it represents a costly solution in terms of computational power. On the contrary, if access to manual transcripts is guaranteed for training an NLU/SLU system, independently of having (or not) the possibility to adapt the ASR model toward the target domain, the recommended solution would be to follow a traditional NLU pipeline.

4. CONCLUSIONS

In this work, we successfully benchmark several neural architectures to perform NLU, pipeline SLU and multi-modal SLU. Our analysis includes SOTA NLU/SLU techniques and compares them in more realistic scenarios. The presented analysis shed light on state-of-the-art architectures in the SLU domain, helping future researchers to define more clearly the application scenario of their proposed solutions. As an additional contribution, we also put together a cleaner version of the well-known SLURP dataset. During our experimentation process, we found many inconsistencies between the manual annotations and what was really spoken in the audio files. This rise concern on the SLU field, as several papers have already reported results on this dataset without being aware of it. We release the filtered version of SLURP in our git repository, in order to allow others to replicate our experiments.²

5. REFERENCES

- [1] Edwin Simonnet, Sahar Ghannay, Nathalie Camelin, Yannick Estève, and Renato De Mori, “Asr error management for improving spoken language understanding,” in *Interspeech*, 2017, pp. 3329–3333.
- [2] Dmitriy Serdyuk, Yongqiang Wang, et al., “Towards end-to-end spoken language understanding,” in *Proc. of ICASSP*. IEEE, 2018, pp. 5754–5758.
- [3] Bidisha Sharma, Maulik C. Madhavi, and Haizhou Li, “Leveraging acoustic and linguistic embeddings from pretrained speech and language models for intent classification,” in *Proc. of ICASSP, 2021*. IEEE, 2021.
- [4] Parisa Haghani, Arun Narayanan, Michiel Bacchiani, et al., “From audio to semantics: Approaches to end-to-end spoken language understanding,” in *IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 720–726.
- [5] Pavel Denisov and Ngoc Thang Vu, “Pretrained semantic speech embeddings for end-to-end spoken language understanding via cross-modal teacher-student learning,” *arXiv preprint arXiv:2007.01836*, 2020.
- [6] Loren Lugosch, Mirco Ravanelli, et al., “Speech Model Pre-Training for End-to-End Spoken Language Understanding,” in *Proc. Interspeech 2019*, 2019, pp. 814–818.
- [7] Esaú Villatoro-Tello, Srikanth Madikeri, Petr Motlicek, Aravind Ganapathiraju, and Alexei V. Ivanov, “Expanded lattice embeddings for spoken document retrieval on informal meetings,” in *Proceedings of the 45th ACM SIGIR Conference*, 2022, p. 2669–2674.
- [8] Gökhan Tür, Anoop Deoras, and Dilek Hakkani-Tür, “Semantic parsing using word confusion networks with conditional random fields,” in *Interspeech*, 2013, pp. 2579–2583.
- [9] Chen Liu, Su Zhu, et al., “Jointly Encoding Word Confusion Network and Dialogue Context with BERT for Spoken Language Understanding,” in *Proc. Interspeech 2020*, 2020, pp. 871–875.
- [10] Ye Zou, Huiping Sun, and Zhong Chen, “Associated latticebert for spoken language understanding,” in *International Conference on Neural Information Processing*, 2021, pp. 579–586.
- [11] Joseph P McKenna, Samridhi Choudhary, Michael Saxon, Grant P Strimel, and Athanasios Mouchtaris, “Semantic complexity in end-to-end spoken language understanding,” in *Interspeech*, 2020, pp. 4273–4277.
- [12] Charles T Hemphill, John J Godfrey, and George R Doddington, “The ATIS spoken language systems pilot corpus,” in *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*, 1990.
- [13] Alice Coucke, Alaa Saade, et al., “Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces,” *arXiv preprint arXiv:1805.10190*, 2018.
- [14] Emanuele Bastianelli, Andrea Vanzo, et al., “SLURP: A spoken language understanding resource package,” in *Proceedings of the 2020 EMNLP*, 2020, pp. 7252–7262.
- [15] Andrea Vanzo, Emanuele Bastianelli, and Oliver Lemon, “Hierarchical multi-task natural language understanding for cross-domain conversational AI: HERMIT NLU,” in *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, Sept. 2019, pp. 254–263.
- [16] Adam Paszke, Sam Gross, Francisco Massa, et al., “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. 2019.
- [17] Jacob Devlin, Ming-Wei Chang, et al., “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *NAACL*, 2019, pp. 4171–4186.
- [18] Matthew Henderson, Blaise Thomson, and Jason D Williams, “The second dialog state tracking challenge,” in *Proceedings SIGDIAL*, 2014, pp. 263–272.
- [19] Yao-Hung Hubert Tsai, Shaojie Bai, et al., “Multimodal transformer for unaligned multimodal language sequences,” in *Proceedings of the ACL*, 2019, vol. 2019, p. 6558.
- [20] Ashish Vaswani, Noam Shazeer, et al., “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [21] Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli, “Unsupervised cross-lingual representation learning for speech recognition,” *arXiv preprint arXiv:2006.13979*, 2020.
- [22] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12449–12460, 2020.
- [23] Yiwen Shao, Yiming Wang, Daniel Povey, and Sanjeev Khudanpur, “Pychain: A fully parallelized pytorch implementation of lf-mmi for end-to-end asr,” *arXiv preprint arXiv:2005.09824*, 2020.
- [24] Hossein Hadian, Hossein Sameti, Daniel Povey, and Sanjeev Khudanpur, “End-to-end speech recognition using lattice-free mmi,” in *Interspeech*, 2018, pp. 12–16.
- [25] Apoorv Vyas, Srikanth Madikeri, and Hervé Bourlard, “Comparing CTC and LFMMI for out-of-domain adaptation of wav2vec 2.0 acoustic model,” in *Proc. of Interspeech 2021*, 2021, pp. 2861–2865.
- [26] Yiming Wang et al., “Espresso: A fast end-to-end neural speech recognition toolkit,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019.
- [27] Srikanth Madikeri et al., “Pkwrap: a pytorch package for lf-mmi training of acoustic models,” *arXiv preprint arXiv:2010.03466*, 2020.
- [28] Banriskhem Khonglah et al., “Incremental semi-supervised learning for multi-genre speech recognition,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7419–7423.
- [29] Daniel Povey et al., “The kald speech recognition toolkit,” in *IEEE 2011 workshop on automatic speech recognition and understanding*, 2011.
- [30] Wei-Ning Hsu, Benjamin Bolte, et al., “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [31] Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, et al., “SpeechBrain: A general-purpose speech toolkit,” in *arXiv preprint arXiv:2106.04624*, 2021.