



HAL
open science

End-to-end model for named entity recognition from speech without paired training data

Salima Mdhaffar, Jarod Duret, Titouan Parcollet, Yannick Estève

► To cite this version:

Salima Mdhaffar, Jarod Duret, Titouan Parcollet, Yannick Estève. End-to-end model for named entity recognition from speech without paired training data. Interspeech 2022, Sep 2022, Incheon, South Korea. hal-03701145

HAL Id: hal-03701145

<https://hal.science/hal-03701145>

Submitted on 21 Jun 2022

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.

End-to-end model for named entity recognition from speech without paired training data

Salima Mdhaffar, Jarod Duret, Titouan Parcollet, Yannick Estève

LIA - Avignon Université, France
{firstname.lastname}@univ-avignon.fr

Abstract

Recent works showed that end-to-end neural approaches tend to become very popular for spoken language understanding (SLU). Through the term end-to-end, one considers the use of a single model optimized to extract semantic information directly from the speech signal. A major issue for such models is the lack of paired audio and textual data with semantic annotation. In this paper, we propose an approach to build an end-to-end neural model to extract semantic information in a scenario in which zero paired audio data is available. Our approach is based on the use of an external model trained to generate a sequence of vectorial representations from text. These representations mimic the hidden representations that could be generated inside an end-to-end automatic speech recognition (ASR) model by processing a speech signal. A SLU neural module is then trained using these representations as input and the annotated text as output. Last, the SLU module replaces the top layers of the ASR model to achieve the construction of the end-to-end model. Our experiments on named entity recognition, carried out on the QUAERO corpus, show that this approach is very promising, getting better results than a comparable cascade approach or than the use of synthetic voices.

Index Terms: Low resource spoken language understanding, End to end neural model, Named Entity Recognition

1. Introduction

Spoken language understanding (SLU) is a very popular topic related to artificial intelligence (AI) widely used in various real-life applications e.g., personal assistants or smart devices. Since 2018, end-to-end neural approaches have been introduced to approach these tasks [1, 2, 3, 4, 5]. While classical SLU systems are built on a cascade approach composed of at least two modules e.g., an automatic speech recognition system and another one for natural language understanding, end-to-end architectures, on the other hand, consider the use of a single neural model optimized to extract semantic information directly from the speech signal.

The two main purposes of the introduction of end-to-end approaches were (i) the joint optimization of the automatic speech recognition (ASR) and the natural language understanding (NLU) parts to the final task, and (ii) the limitation of the error propagation due to the bottleneck that constitutes the ASR output. Nevertheless, such end-to-end SLU models need to be trained on large amounts of data, while a major issue is the lack of paired audio and textual data with semantic annotation: producing a large amount of annotated speech data remains difficult and expensive.

However, for some tasks including named entity recognition (NER), large amounts of annotated data may be available in the textual format. Hence, with this paper, we propose to leverage the latter textual-only data to conceive an end-to-end

neural model extracting semantic information from speech in a named entity recognition scenario in which no paired textual-audio data is accessible.

2. Related work

The lack of resource is recurrent for SLU tasks. For end-to-end approaches this gap is even more critical than for cascaded ones due to the bi-modal nature of the necessarily paired training data. Several solutions have been recently proposed in the literature to overcome this issue. In [6], the authors propose to apply a curriculum-based transfer learning approach to take benefit of existing semantic annotations related to more generic SLU tasks than the targeted one. In [7, 8] the authors suggest to generate synthetic speech in order to extend a small training dataset of paired data of audio and semantic annotation. This approach has also been proposed for data augmentation in ASR [9] or with speech translation [10]

Other approaches aim to merge acoustic and text embeddings. In [8], acoustic embeddings for intent classification are tied to fine-tuned BERT text embeddings, while in [11] a multi-task architecture is designed to map the acoustic feature sequence into a semantic space shared for both speech recognition and SLU.

More recently, [12] presented a study that explores several ways intended to inject external speech and/or text unlabeled data for the task, to train cascaded or end-to-end named entity recognition systems. Their experiments showed that, with access to unlabeled speech or transcribed speech, end-to-end models outperform pipeline models.

To the best of our knowledge, our work is the first successful attempt at conceiving an end-to-end model for named entity recognition from speech without the use of paired training data¹.

3. Proposed approach

The approach presented below is dedicated to SLU for which no paired data (i.e., audio with transcriptions annotated with semantic information) is available. In this study, we focus on NER from speech, but we assume that this approach could be extended to other similar SLU tasks. More precisely, NER from speech consists of recognizing words from an input speech source, detecting word sequences that support a named entity, and categorizing this entity.

In our low resource SLU scenario, an end-to-end model for ASR and a corpus of textual documents with named entity annotations but without the corresponding audios are available.

Our approach is based on the use of an external model trained to generate a sequence of vectorial representations from text. These representations mimic the hidden representations that could be generated inside an end-to-end automatic speech

¹This paper has been submitted to INTERSPEECH 2022

recognition model by processing a speech signal. A SLU neural module is then trained to use these representations as input and the annotated text as output. Last, the SLU module replaces the top layers of the ASR model to achieve the construction of the end-to-end model.

To generate the simulated ASR hidden representations (or ASR embeddings), we train a sequence-to-sequence neural model, called *Text-to-ASR-Embeddings* model. Such an approach can be compared to some recent propositions mentioned in section 2 that use synthetic voices to feed an ASR end-to-end model. We motivated our proposition for different reasons. First, the use of synthetic speech introduces some artifacts in the input of the ASR model. If the ASR model is fine-tuned on such synthetic voices, these artifacts will degrade the capability of the model to process natural voices. A solution to avoid this consists of freezing the weights of the bottom layers and only update the weights of the higher layers, in which the semantic is better encoded. Since the bottom layers were optimized to process natural speech, the quality of the embeddings computed from synthetic speech is not guaranteed, and can introduce a gap between embedding computed from natural and computed from synthetic speech. With our approach, we aim to reduce this gap. In addition, our approach needs less computation at training time than the ones based on synthetic speech, since we avoid the use of a consequent number of lower layers.

To train this neural network, we must produce a training dataset composed of pairs of transcriptions, used as input, and sequences of ASR embeddings, used as output. To produce this training dataset, the end-to-end ASR model is used to transcribe its training dataset. For each transcribed utterance, we extract a sequence of ASR embeddings from a hidden layer, and associate this ASR embedding sequence to the automatic transcription. When the entire ASR training data has been processed, the ASR embedding sequences and their associated automatic transcriptions are used to train the *Text-to-ASR-Embeddings* model, as illustrated in Figure 1

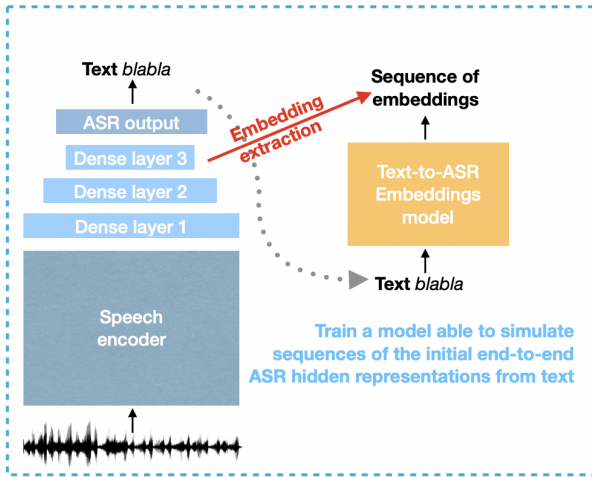


Figure 1: ASR embeddings extraction to train the *Text-to-ASR-Embeddings* model

At this stage, we obtain a module able to simulate ASR embeddings from text. Our objective is then to train a neural SLU sub-module able to convert such a sequence of ASR embedding into an automatic transcription with SLU annotation, like annotation of named entities. For this purpose, we

exploit the textual dataset with semantic annotation. For each sentence in this dataset, we first remove the semantic annotation to keep only the sequence of words. Thanks to the *Text-to-ASR-Embeddings* model, we transform this sequence of words to a sequence of ASR embeddings. We iterate this process for all the annotated sentences in the semantic textual dataset. We get a set of pairs composed of a sequence of ASR embeddings and the corresponding text sequence of words semantically annotated. Once the entire textual dataset has been processed, we use this data to train an SLU sub-module able to generate a sequence of words semantically annotated from a sequence of ASR embeddings. This step is illustrated in Figure 2.

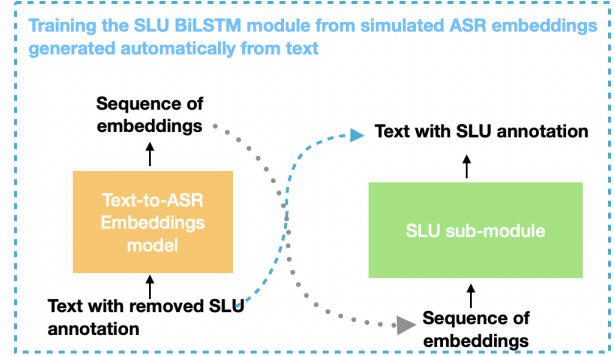


Figure 2: Training the SLU sub-module from simulated ASR embeddings

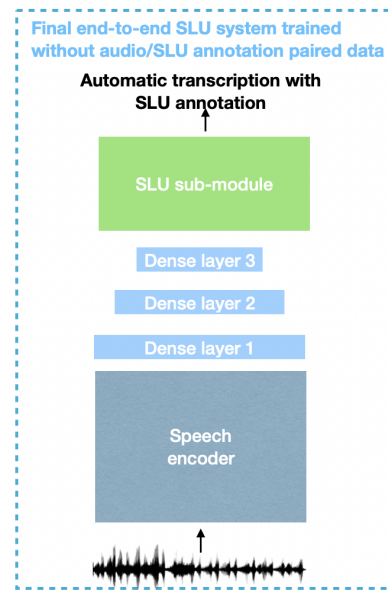


Figure 3: Final end-to-end SLU system trained without audio/SLU annotation paired data

Finally, we plug the end-to-end ASR and the SLU sub-module, as illustrated in Figure 3. In order to merge the ASR model with the SLU sub-module, we keep all the ASR hidden layers needed to generate the ASR embeddings that can be mimicked by the *Text-to-ASR-Embeddings* model. The mimicked hidden layer is then connected to the SLU sub-module. In the example of Figure 3, only the ASR output layer is removed.

The final model is an end-to-end model able to transcribe and extract semantic information from speech, while no real paired training data exists. In the next section, we show how we implemented this approach to recognize named entities from speech.

4. Experiments

Our experiments were carried out on the French QUAERO data, officially used in the framework of the ETAPE evaluation campaign [13] and in different research studies [14, 15] dedicated to named entity recognition from speech.

4.1. Data

Several publicly available corpora have been used for these experiments.

4.1.1. ASR training data

The ASR training data is composed of different corpora in French language, all collected from TV or radio stations:

- REPERE² [16]: The REPERE dataset was created by the organizers of the REPERE evaluation campaign. It is composed of articles from videos (news, political debate and talk-shows) from 7 French TV programs (2 programs from BFMTV channel and 5 from LCP channel) broadcasted between 2011 and 2012.
- ESTER2³: The ESTER2 dataset was created for the ESTER 2 evaluation campaign [17]. It consists on radio broadcast news, from French, Moroccan and African radio stations.
- EPAC⁴ [18]: this data is composed of audio recordings of French radio and TV shows, mainly debates or interviews.

The total duration of all these datasets is 197 hours of speech (35.5 hours of REPERE, 91 hours of ESTER2, 70.5 hours of EPAC).

4.1.2. Text-to-ASR-Embeddings data

The ASR embeddings used to train the *Text-to-ASR-Embeddings* model have been extracted from the ASR model by processing its own training data, described in the previous subsection. The *Text-to-ASR-Embeddings* model is trained from these ASR embeddings and their corresponding automatic transcriptions computed by the ASR model during the extraction. The word error rate of the automatic transcriptions got on the 197 hours of the ASR training data is 7.4%. No filtering was applied on these transcriptions before training the *Text-to-ASR-Embeddings* model, but 10% of the data was removed from the training data to be used as validation data to control the training process of this model.

²The REPERE corpus can be obtained from ELRA under the reference ELRA-E0044

³The ESTER2 dataset can be obtained from ELRA under the reference ELRA-S0241

⁴The EPAC corpus can be obtained from ELDA under the reference ELDA-S0305

4.1.3. Dataset for named entity recognition

To train and evaluate our named entity recognition systems, we used the QUAERO⁵ dataset [19].

This dataset contains broadcast news speech with text transcription and named entity annotation. Manual annotation of named entities is available according to eight types: *person*, *location*, *organization*, *product*, *time*, *amount*, *function* and *event*. In our scenario, we assume that we do not have a speech dataset. We use only the textual transcriptions and the named entity annotation in the text, except for a contrastive experiment. We divided the QUAERO corpus into three parts (the official distribution contains only two parts): a training corpus, a validation corpus, and a test corpus. The test corpus is composed of the ten most recent hours, in order to avoid overlaps with the training data. To ensure the reproducibility of our experimental results, we share the details of our file distribution⁶. Some statistics about this data are described in Table 1. Figure 4 shows an example of QUAERO dataset.

	#words	#vocab	#entities	#sentence
Train	993,746	33,583	89,609	90,378
Dev	81,531	8,713	5,538	10,982
Test	119,361	13,447	10,883	4,023

Table 1: QUAERO dataset

- (a) demain rfi présente le huitième festival de jazz de saint louis au sénégal
- (b) <time demain > <organisation rfi > présente le huitième festival de jazz de <location saint louis > au <location sénégal >

Figure 4: An example of a QUAERO dataset sample in French of the English sentence ‘tomorrow rfi presents the eighth saint louis jazz festival in senegal’. (a) corresponds to the transcribed sentence. (b) the same sample with its additional named entity tags. Here, ‘<time’ is an opening tag starting the support word sequence ‘demain’ (tomorrow in English) and expressing that this word sequence is associated with the time entity. The character ‘>’ represents the closing tag and it is used to close all entity tags.

4.2. Model implementation

In this section, the different models used in our experiments are described. All the models were implemented thanks to the SpeechBrain toolkit [20]. For reproducibility purpose, we will also make the recipe available⁷.

4.2.1. End-to-end ASR architecture

The ASR used in this paper is composed of a large pre-trained French wav2vec 2.0 model, with three additional hidden layers of 1024, 512 and 80 units, associated to Leaky ReLU activation functions, and last a softmax output layer. wav2vec 2.0 is a model pre-trained through self-supervision [21]. The loss function used at each fine-tuning step is the Connectionist Temporal Classification (CTC) loss function [22]. This ASR model

⁵The QUAERO dataset can be obtained from ELRA under the reference ELRA-S0349

⁶<https://github.com/mdhaffar/Named-Entity-Recognition>

⁷note for reviewers: the recipe will be released before the Inter-speech conference

is a state-of-the-art ASR model for French: the pretrained wav2vec 2.0 model is the LeBenchmark-7K model presented in [23] and publicly accessible ⁸.

4.2.2. Text-to-ASR Embeddings model architecture

In this study, the model used to mimic the ASR hidden representations is based on the well-known Tacotron2 neural architecture [24] initially designed for speech synthesis, also called text-to-speech (TTS) task. Tacotron2 is composed of a first module that maps character embeddings to mel-scale spectrogram, followed by a second module, a vocoder that transforms the mel-scale spectrogram into audio waveforms. To generate the simulated ASR hidden representations (or ASR embeddings), we only need this first module: instead of generating a mel-scale spectrogram, the recurrent sequence-to-sequence network that composes this module is trained to generate sequences of ASR embeddings. The network consists of a stack of 3 convolutional layers, followed by a batch normalization and ReLU activations. The goal of these convolutional layers is to model longer-term context in the input character sequence. Then the output is fed to a bi-directional LSTM layer to generate the encoded features.

We use text input with basic rule-based text normalization for French language. For a contrastive experiment presented below, we also trained a Tacotron2 text-to-speech model on the SynPaFlex corpus data [25]. This corpus is a French audiobook dataset containing eighty-seven hours of speech, recorded by a single speaker reading audiobooks of different literary genres.

4.2.3. SLU sub-module

The SLU sub-module is implemented through a BiLSTM architecture. This neural architecture learns to transcribe and recognize named entities from the mimicked ASR embeddings generated automatically from text. The Bidirectional LSTM model is composed of a stack of 5 biLSTM layers, with 512 dimensions each.

4.3. Experimental results

This section describes the different results we obtained. The evaluation of the NER task uses the Named Entity Error Rate (NEER), which is estimated like the classical Word Error Rate (WER) but applied to semantic concepts instead of words. Insertions, substitutions, and deletions are all counted as errors. All our experiments are evaluated on the test partition defined in Section 4.1.3.

The first row in Table 2 presents the results of our proposed approach to train an end-to-end SLU model (47.5% NEER for dev and 39.1% for test) when the *Text-to-ASR-Embeddings* model is used to mimic the 80-dimensional embeddings of the last hidden layer presented in section 4.2.1.

We also evaluate the use of synthetic speech to train the end-to-end NER model. In a first experiment, all the weights of the end-to-end ASR model are updated during the training process. On the test set, this system gets a NEER of 62.7%, very higher than the NEER got by our system. In a second experiment, we freeze all the layers of the ASR speech encoder (i.e. the wav2vec 2.0 part) before starting the NER training. The objective is to limit the possible negative effects of the synthetic speech artifacts to the speech encoder weights: the results are worse, with a NEER of 92.5%.

Then, we consider the availability of paired audio data. The

SLU sub-module is trained by using the embeddings generated from the ASR model by processing the real audio data. These results can be seen as the Oracle NEER we could reach if the *Text-to-ASR-Embeddings* model was able to produce ideal embeddings.

Training data	Dev	Test
ASR embeddings simulation (ours)	47.6	39.1
Synthetic speech (all weights are updated)	65.2	62.7
Synthetic speech (frozen speech encoder)	86.4	92.5
Oracle (real audio)	45.9	34.1

Table 2: Evaluation in NEER (%) of our approach to train an end-to-end NER model without paired training data compared to other approaches using speech synthesis, and compared to the ideal scenario when paired data is available

We have shown that our approach makes possible the training of an end-to-end model for named entity recognition from speech without paired training data (using zero audio data).

We would like now to compare its performance to the one got by a cascaded approach trained in the same conditions, since such an approach does not need paired (audio/text) training data. In this experiment, we use the same neural architecture as the SLU sub-module used on the top of the end-to-end proposed system to train a text-to-text NER model. We apply this NER model to the automatic transcription of our end-to-end ASR model. This ASR model obtains a WER of 12.85% on the development set and 8.6% on the test set, showing that the test set seems easier to process than the development set. We also apply this NER model to the manual transcriptions. Table 3 presents the results of these experiments. As expected, ASR errors impact the NER model performance, and we can notice that the cascaded approach applied to speech gets worse results than the end-to-end approach trained through the proposed approach. We are aware that a state-of-the-art cascaded system could strongly take advantage of self-supervised learning-based models like BERT [26], but this is out of the scope of this study.

	Dev	Test
Text-to-text NER on manual transcripts	39.6	30.0
Text-to-text NER on automatic transcripts	48.0	40.2

Table 3: Results in term of NEER(%) got by NER module applied to text.

5. Conclusion

In this paper, we propose an approach to train an end-to-end model for named entity recognition from speech without paired audio and textual data with semantic annotation. Our approach, based on artificial ASR embeddings generated from text, exhibits highly promising results outperforming alternative approaches based on the use of synthetic speech.

We consider that this approach can be extended to similar SLU tasks such as slot filling, and opens new perspectives in different use cases where enriching or adapting the linguistic knowledge captured by an end-to-end ASR model is needed.

6. Acknowledgements

This work was partially funded by the European Commission through the SELMA European project under grant number 957017. This work was granted access to the HPC resources

⁸<https://huggingface.co/LeBenchmark>

of IDRIS under the allocation 2021-AD011012551 made by GENCI.

7. References

- [1] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, "Towards end-to-end spoken language understanding," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5754–5758.
- [2] S. Ghannay, A. Caubrière, Y. Estève, N. Camelin, E. Simonnet, A. Laurent, and E. Morin, "End-to-end named entity and semantic concept extraction from speech," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 692–699.
- [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 *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 720–726.
- [4] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, "Speech model pre-training for end-to-end spoken language understanding," *Proc. Interspeech 2019*, pp. 814–818, 2019.
- [5] T. Desot, F. Portet, and M. Vacher, "Towards end-to-end spoken intent recognition in smart home," in *2019 International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*. IEEE, 2019, pp. 1–8.
- [6] A. Caubrière, 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 *20th Annual Conference of the International Speech Communication Association (InterSpeech)*, 2019, pp. 1198–1202.
- [7] L. Lugosch, B. H. Meyer, D. Nowrouzezahrai, and M. Ravanelli, "Using speech synthesis to train end-to-end spoken language understanding models," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 8499–8503.
- [8] Y. Huang, H.-K. 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 *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7984–7988.
- [9] A. Laptev, R. Korostik, A. Svischev, A. Andrusenko, I. Medennikov, and S. Rybin, "You do not need more data: Improving end-to-end speech recognition by text-to-speech data augmentation," in *2020 13th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*. IEEE, 2020, pp. 439–444.
- [10] T. Kano, S. Sakti, and S. Nakamura, "End-to-end speech translation with transcoding by multi-task learning for distant language pairs," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1342–1355, 2020.
- [11] X. Jia, J. Wang, Z. Zhang, N. Cheng, and J. Xiao, "Large-scale transfer learning for low-resource spoken language understanding," 2020.
- [12] A. Pasad, F. Wu, S. Shon, K. Livescu, and K. J. Han, "On the use of external data for spoken named entity recognition," *arXiv preprint arXiv:2112.07648*, 2021.
- [13] G. Gravier, G. Adda, N. Paulson, M. Carré, A. Giraudel, and O. Galibert, "The ETAPE corpus for the evaluation of speech-based tv content processing in the french language," in *LREC-Eighth international conference on Language Resources and Evaluation*, 2012, p. na.
- [14] A. Caubrière, S. Rosset, Y. Estève, A. Laurent, and E. Morin, "Where are we in named entity recognition from speech?" in *Proceedings of the 12th Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 4514–4520. [Online]. Available: <https://aclanthology.org/2020.lrec-1.556>
- [15] M. A. B. Jannet, O. Galibert, M. Adda-Decker, and S. Rosset, "How to evaluate asr output for named entity recognition?" in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [16] A. Giraudel, M. Carré, V. Mapelli, J. Kahn, O. Galibert, and L. Quintard, "The REPERE corpus: a multimodal corpus for person recognition," in *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, 2012, pp. 1102–1107.
- [17] S. Galliano, G. Gravier, and L. Chaubard, "The ester 2 evaluation campaign for the rich transcription of french radio broadcasts," in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [18] Y. Estève, T. Bazillon, J.-Y. Antoine, F. Béchet, and J. Farinas, "The EPAC corpus: manual and automatic annotations of conversational speech in french broadcast news," in *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, 2010.
- [19] C. Grouin, S. Rosset, P. Zweigenbaum, K. Fort, O. Galibert, and L. Quintard, "Proposal for an extension of traditional named entities: From guidelines to evaluation, an overview," in *Proceedings of the 5th linguistic annotation workshop*, 2011, pp. 92–100.
- [20] T. Parcollet, M. Ravanelli, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong *et al.*, "Speechbrain: A general-purpose speech toolkit," 2022.
- [21] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [22] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [23] S. Evain, H. Nguyen, H. Le, M. Z. Boito, S. Mdahaffar, S. Alisamir, Z. Tong, N. Tomashenko, M. Dinarelli, T. Parcollet, A. Allauzen, Y. Estève, B. Lecouteux, F. Portet, S. Rossato, F. Ringeval, D. Schwab, and L. Besacier, "LeBenchmark: A Reproducible Framework for Assessing Self-Supervised Representation Learning from Speech," in *Proc. Interspeech 2021*, 2021, pp. 1439–1443.
- [24] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [25] A. Sini, D. Lolive, G. Vidal, M. Tahon, and E. Delais-Roussarie, "Synpaflex-corpus: An expressive french audiobooks corpus dedicated to expressive speech synthesis," in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [26] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.