



HAL
open science

End2End Acoustic to Semantic Transduction

Valentin Pelloin, Nathalie Camelin, Antoine Laurent, Renato de Mori,
Antoine Caubrière, Yannick Estève, Sylvain Meignier

► **To cite this version:**

Valentin Pelloin, Nathalie Camelin, Antoine Laurent, Renato de Mori, Antoine Caubrière, et al..
End2End Acoustic to Semantic Transduction. ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Jun 2021, Toronto, ON, Canada.
10.1109/ICASSP39728.2021.9413581 . hal-03128163

HAL Id: hal-03128163

<https://hal.science/hal-03128163v1>

Submitted on 11 Dec 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.

END2END ACOUSTIC TO SEMANTIC TRANSDUCTION

Valentin Pelloin¹ Nathalie Camelin¹ Antoine Laurent¹ Renato De Mori^{2,3}
Antoine Caubrière¹ Yannick Estève² Sylvain Meignier¹

¹ LIUM - Le Mans University - France

² LIA - Avignon University - France

³ McGill University - Montréal, Canada

ABSTRACT

In this paper, we propose a novel end-to-end sequence-to-sequence spoken language understanding model using an attention mechanism. It reliably selects contextual acoustic features in order to hypothesize semantic contents. An initial architecture capable of extracting all pronounced words and concepts from acoustic spans is designed and tested. With a shallow fusion language model, this system reaches a 13.6 concept error rate (CER) and an 18.5 concept value error rate (CVER) on the French MEDIA corpus, achieving an absolute 2.8 points reduction compared to the state-of-the-art. Then, an original model is proposed for hypothesizing concepts and their values. This transduction reaches a 15.4 CER and a 21.6 CVER without any new type of context.

Index Terms— spoken language understanding, neural networks, attention mechanisms, sequence-to-sequence, transfer learning

1. INTRODUCTION

Spoken Language Understanding (SLU) systems extract semantic contents mentioned in spoken sentences for task-oriented dialogues, question answering, and other conversational applications.

Semantic contents are fragments of a domain application ontology that can be defined with a frame language as described in [1]. Frame structures are rich representations of knowledge structures with important properties, such as type models of slot fillers and inheritance of slot properties through chains of frame-slot relations. With a frame language, a concept is represented as a relation between a *frame name* (e.g. ADDRESS), a frame property, often called *slot*, (e.g. city), and a frame *value*, often called *slot filler* (e.g. Paris). Properties can be inherited through a *chain of slots*. Using this semantic representation, speech acts such as *request* can be defined as frames with slots filled by instances of frame struc-

ture fragments. For the sake of simplicity, in the following, slots and slot fillers are denoted as *concept* and *value*.

Concepts can be annotated with word sequences called *supports*. Supports can be instantiated with hypotheses obtained by an Automatic Speech Recognition (ASR) system as described in [2, 3, 4, 5] or by direct transduction that generates instances of frame structures from acoustic features.

Supports are word spans expressed by a span of acoustic features. Such a span may have variable duration and depends on the speaker and other factors. If a value is mentioned, then it is often included in a support. Spans of acoustic feature supports may not be sufficient for unambiguously mentioning a concept instance, even if they contain relevant semantic features. This equivocation can be reduced by selecting relevant spans of acoustic feature contexts using an attention mechanism.

Human experience in understanding spoken sentences, especially in foreign languages, provides evidence that concepts can be hypothesized from spoken language tokens without thinking about, or even knowing, the exact spelling of some relevant words. This motivates the use of a model based on a transduction of acoustic into semantic content supports. In this paper, an end-to-end (E2E) SLU architecture is proposed in order to hypothesize fragments of frame instances with acoustic features. It is applied on user turns of the French MEDIA corpus, corresponding to a complex negotiation dialogue task (see Section 4.1). This task is particularly difficult since a dialogue turn may contain a large variety of concept-value instances, including repetitions and self-corrections with specific semantic relations between these contents.

A novelty of the proposed model is the introduction of an attention mechanism that selects contextual acoustic features in order to hypothesize directly concept symbols and their values. An initial architecture is first designed and tested in order to hypothesize words and concepts. An enriched architecture is then proposed in order to hypothesize concepts and delimiting their supports of underlying acoustic representations while limiting the generation of character hypotheses only for values. As a consequence, no human prior-knowledge is needed to extract values from words of the concept support.

This work has been funded by the AISSPER project supported by the French National Research Agency (ANR) under contract ANR-19-CE23-0004-01.

The paper is organized as follows. Section 2 presents related works in the SLU domain. Section 3 presents our E2E SLU architecture, while corpora, experiments and results are described in Section 4. Finally, conclusions and perspectives are given in Section 5.

2. RELATED WORK

Early speech understanding systems based on natural language semantic parsers and artificial intelligence approaches to beam search and knowledge representation are reviewed in [6]. Further developments stressing comparisons of different methods using a common annotated corpus, new parsing techniques and statistical models are reviewed in [7]. Later, advanced spoken language understanding systems with advanced statistical models and paradigms in various types of applications are described in [1]. More recently, deep neural networks (DNN) for generating semantic domain, intent and slot filler hypotheses were proposed using a pipeline of automatic speech recognition and SLU models. These architectures provided better or comparable results to those obtained with previous architectures. Examples can be found in [2, 3, 4]. Result examples and comparison using SLU in negotiation dialogues are reported in [5]. Examples of attempts to integrate ASR and SLU functions in end-to-end (E2E) compact trainable DNN architectures are proposed in [8, 9, 10, 11, 12, 13].

In this paper, a novel E2E SLU system is introduced and tested on the French MEDIA corpus. It is motivated by the evidence that the perception of spoken semantic entities and relations is associated with the perception of word spans expressing semantic properties and their values. A basic architecture is introduced with attention on suitable acoustic contexts having relevance dependent on decoder semantic representations. Then, a novel architecture is proposed in order to detect acoustic spans expressing lexical features for clue words of semantic entities and their values. This introduction is motivated by explaining, with prior knowledge, frequent errors observed in the annotated development set.

Recently, methods have been proposed to learn high-level representations from surface acoustic features for speech-to-speech translation [14, 15, 16, 17].

In this paper, we aim to use high-level representations from surface acoustic features for speech-to-semantic transduction.

In [18], relations in semantic knowledge graphs have been used for inferring answers in question answering dialogues.

In this paper, application domain model relations, explicit or inferred by inheritance, are considered for selecting useful context for the interpretation of acoustic spans of domain relevant concepts.

3. E2E SLU ARCHITECTURE

In this work, we use an encoder-decoder neural network. An attention mechanism as described by [19] was designed in order to align the input with the output. Hypothesizing spoken phonemes with an ASR system may be difficult if a training set of phonemes aligned with speech is not available. For this reason, a first attempt using experience on character recognition is used in this paper. Acoustic to concept hypothesizing through to phoneme is left for future work.

The input of the network are 40 dimensional MelFBanks extracted with a Hamming window of 25 ms and 10 ms strides. The system output is a sequence of characters predicting words, normalized values, or concept labels. Note that each concept label is represented by one special character (e.g. \hat{g} for the concept *hotel-services*).

Our attention based encoder-decoder architecture, depicted in Figure 1, is based on the *Espresso* recipe initially developed for the WSJ ASR task, as described in [20].

Let $X = (x_1, \dots, x_{T_x})$ be the vector of the input features and $Y = (y_1, \dots, y_{T_y})$ be the outputs of the model. Our model computes the outputs Y from an input sequence X . The encoder first uses 4 2-dimensional convolutional blocks (each convolution layer is followed by a batch norm). Then, 4 biLSTM layers are used to obtain the encoder hidden states as:

$$\text{Encoder}(X) = (\overleftarrow{h}_1, \dots, \overleftarrow{h}_{T_x}) \quad (1)$$

The decoder uses 4 LSTM followed by 2 fully-connected layers and a softmax. The decoder states are computed using the attention mechanism by aligning the encoder hidden states and previous decoder hidden states, as:

$$Y = \text{Decoder}(H^{\text{att}}) \quad (2)$$

The attention mechanism is applied as described in [19]:

$$H^{\text{att}} = (H_1^{\text{att}}, \dots, H_{T_y}^{\text{att}}) \quad (3)$$

$$H_i^{\text{att}} = \sum_j^{T_x} \alpha_{ij} \overleftarrow{h}_j \quad \text{with} \quad \alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k=1}^{T_x} \exp(e_{ik})} \quad (4)$$

e_{ij} is the relevance of the acoustic representation \overleftarrow{h}_j and the i^{th} semantic output. Note that the encoder vectors \overleftarrow{h} are computed using only the information from the top-encoder hidden states.

In order to improve the encoder-decoder performance, we use an additional LSTM-based Language Model (LM) that uses *look-ahead* word probabilities to obtain both word and character probabilities ([20, 21]).

Shallow-fusion [22] score combination is performed in order to take into account the output of the encoder-decoder SLU system and the LSTM-based LM. During the decoding phase with the beam search, hypotheses are scored using the weighted sum of both the encoder-decoder model and the LM scores.

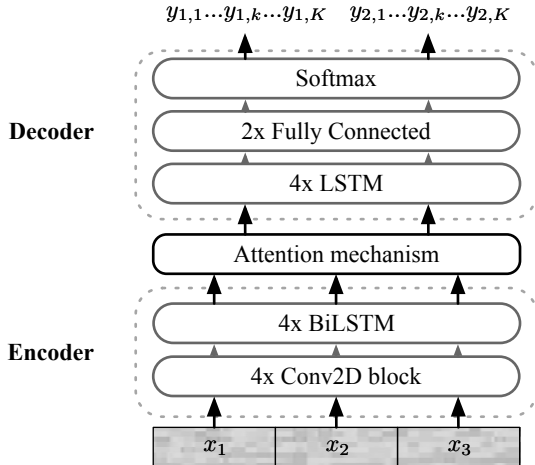


Fig. 1. E2E Encoder-Decoder SLU Model

4. EXPERIMENTS

4.1. Datasets

Our architecture is pretrained for ASR and then adapted to SLU. Training is detailed in the next section while we briefly describe the different corpora in this section.

For the transcription model, we use several French generic broadcast news corpora described in [23], in addition to the training sets of in-domain SLU telephone conversation corpora: PORTMEDIA [24], MEDIA [25] and DECODA [26]. As a result, we obtained a training corpus of 414 hours for ASR. Such a model needs a lot of audio data along with their manual transcription even if a frame acoustic alignment with words is not needed. The system will train itself to align each output character with input acoustic features thanks to its attention mechanism [27, 28, 29].

For the SLU models, only the 24 hours of MEDIA (M) and PORTMEDIA (PM) training sets are used during training. These two corpora contain both telephone conversations between two humans in the *Wizard of Oz* mode: one plays a computer, the other a user who wants to obtain information or make a reservation. MEDIA contains requests and lodging booking in France and PORTMEDIA contains conversations about theater shows during the french Avignon Festival. All conversations are manually transcribed and user turns are semantically annotated with concept labels, their supports and the corresponding values. Considering the Table 1 example, the concept label *link-coref* is supported by the word *that* annotated with the normalized value *singular*.

Turns of these datasets contain word spans expressing semantic content mentions mixed with other spans that do not express application domain contents. For this task, user turns may express concepts such as logical AND OR connectors, references to previously mentioned concepts and mentions to entities relative to other entities. Furthermore, system turns

tend to solicit the user to provide useful information for composing a *state of the world* to obtain a frame structure in an application domain database.

In the experiments section, we will present SLU results on the MEDIA corpus, split accordingly to the official training, development (*Dev*) and test sets, containing respectively 727, 79 and 208 dialogues and 13k, 1.2k and 3.6k user utterances. User transcripts are annotated using 76 semantic concepts, along with their supports and the corresponding normalized value.

4.2. Training and outputs

Inspired by the *curriculum-based transfer learning* procedure for SLU proposed in [23]¹, we first train a model for the ASR downstream task. For that purpose, we use out-of-domain and in-domain ASR data (as described in the previous section), and then fine-tune using our in-domain corpora. This first model is denoted as *ASR* model.

Thereafter, the model is retrained for the SLU task using all the in-domain SLU data (MEDIA and PORTMEDIA train sets, denoted M+PM), and then fine-tuned with the target MEDIA data. The transition from ASR to SLU is made by enriching the output with concept symbols. In contrast with [23, 30], the last fully connected layer is just extended. As a consequence, the model benefits entirely from its previous training. This second model is denoted as *AllWords-C* model.

[30] showed that hypothesizing only words belonging to concept support enhanced the SLU performance evaluation. Thus, we introduce a third model that adapts the *AllWords-C* model on the specific outputs: concepts and words from supports. A new character is then added to the last layer in order to output a global character * instead of all out-of-concept-support words. This third *SupWords-C* model is obtained with transfer learning on these outputs.

Finally, a novel solution is introduced to perform a direct transduction from acoustic features to concept-value pairs. To do so, we re-adapt a last time the *AllWords-C* model in order to output directly the normalized values and the concept. In that case, characters no longer represent pronounced words. They represent either one concept or one character of a normalized value. This model is denoted as *NormValues-C* model. In that last architecture, frame structure fragments are obtained with a fully automatic E2E sequence 2 sequence architecture.

All models and their output formats are summarized in Table 1. Note that, for the sake of clarity, concepts are written in full instead of their special character output.

4.3. Evaluation Protocol

The architecture shown in Figure 1 is used in order to generate outputs with different components, all evaluated according

¹Unlike [23], we did not choose to train our model on the Named Entity Recognition (NER) task.

<i>ASR</i>	
ASR → ASR M+PM	
is there a swimming-pool in that one	
<i>AllWords-C</i>	
ASR model → SLU M+PM → SLU M	
is there a <hotel-services> swimming-pool </hotel-services> in <linkref-coref> that </linkref-coref> <objectbd> one </objectbd>	
<i>SupWords-C</i>	
<i>AllWords-C</i> model → SLU M*	
* <hotel-services> swimming-pool </hotel-services> * <linkref-coref> that </linkref-coref> <objectbd> one </objectbd>	
<i>NormValues-C</i>	
<i>AllWords-C</i> model → SLU Norm M*	
* <hotel-services> swimming-pool </hotel-services> * <linkref-coref> singular </linkref-coref> <objectbd> hotel </objectbd>	

Table 1. Model (chain of training) and output for each proposed configuration, based on the user utterance "Is there a swimming-pool in that one?"

to the concept error rate (CER) and concept-value error rate (CVER) metrics. Insertion, deletion, and substitution errors are used to compute evaluation measures and to perform the error analysis.

Normalized values, *i.e.* the slot fillers, needed for the evaluation in CVER, are obtained with a set of manually designed regular expressions in *AllWords-C* and *SupWords-C* configurations. Expressions are applied to the outputs of each concept support, as in [12, 23, 30, 5]. Using the *NormValues-C* outputs, these handmade rules are not necessary anymore.

4.4. Results

Table 2 summarizes the results obtained with the three considered architecture solutions, without and with LM, on the MEDIA corpus on both the Dev and Test sets.

Without LM, the best CER results are obtained with *NormValues-C* model, both for Dev and Test. This supports the conjecture that concepts can be perceived without considering orthographic word transcriptions. As we could expect, when CER are similar, best CVER are achieved using human rules. Nevertheless, very encouraging results are observed, notably on Dev, when the CER is lower than others. For the sake of comparison, [23] state-of-the-art results related to the use of an E2E approach without LM, are reported with 21.6% CER and 27.7% CVER, showing an absolute gain of 6 points.

Unlike in [23] where best results are obtained with the *SupWords-C* configuration, our architecture does not benefit from this kind of representation. We suppose that this is due to the architectural difference: Attention Mechanism (AM) vs. biLSTM/CTC. In the latter, the star is used as a correspondence to the input acoustic frame when predicting outside-

	Dev		Test		
	%	CER	CVER	CER	CVER
Without a Language Model					
<i>AllWords-C</i> [23]	–	–	21.6	27.7	
<i>AllWords-C</i>	18.1	22.5	15.6	20.4	
<i>SupWords-C</i>	17.3	22.0	15.6	20.5	
<i>NormValues-C</i>	16.0	21.9	15.4	21.7	
With a Language Model					
<i>AllWords-C</i> [23]	–	–	18.1	22.1	
<i>SupWords-C</i> [23]	–	–	16.4	20.9	
<i>AllWords-C</i>	16.1	20.4	13.6	18.5	
<i>SupWords-C</i>	17.6	22.5	15.5	20.5	
<i>NormValues-C</i>	16.1	22.0	15.4	21.6	

Table 2. Results obtained on the MEDIA Dev and Test Corpora by our models compared to the State of the art.

of-concept characters. In our AM system, such a star is not necessarily needed, as the output sequence is not constrained to have the same length as the input. The AM is designed to directly select relevant acoustic spans.

Using word hypotheses and LM, the best 13.6% CER and 18.5% CVER are observed for *AllWords-C*, showing a gain of 2 points using the LM and 2.8 points CER considering the best [23] state-of-the-art result even without significant contribution of the LM on *NormValues-C* and *SupWords-C*. This can be explained by the LM weight, optimized on the Dev set, which is close to zero. This needs further investigation.

A first analysis of detailed concept errors on Dev shows that the most frequent ones are insertions and deletions of logical connectors and co-references.

5. CONCLUSIONS AND PERSPECTIVES

An E2E SLU architecture is introduced, based on an encoder-decoder model with attention mechanism focusing on acoustic representations, useful for generating hypotheses of words, concepts or values. Several combinations have been investigated using or not LM reaching better results than state-of-the-art.

The results indicate that word knowledge can be a relevant context for semantic interpretation, particularly if it is selected by an appropriate attention mechanism. The solutions evaluated in this paper are based either on word or value representations using characters. In future work we also plan to use two decoders, one that outputs each character (*allWords-C*) and the other one which focuses on concept-value pairs. This could, at least, improve our CVER without degrading the CER using the *NormValues-C* representation. Finally, new types of LM including semantic hypotheses will also be investigated in future work.

6. REFERENCES

- [1] G. Tur and R. De Mori, "Chapter 1: Spoken language understanding for human/machine interactions," in *Spoken Language Understanding: Systems for Extracting Semantic Information from Speech*. 2011.
- [2] D. Hakkani-Tür, G. Tür, A. Celikyilmaz, Y. N. Chen, J. Gao, L. Deng, and Y. Y. Wang, "Multi-domain joint semantic frame parsing using bi-directional rnn-lstm.," in *INTERSPEECH*, 2016, pp. 715–719.
- [3] X. Zhang and H. Wang, "A joint model of intent determination and slot filling for spoken language understanding," in *IJCAI*, 2016, p. 2993–2999.
- [4] B. Liu and I. Lane, "Attention-based recurrent neural network models for joint intent detection and slot filling," in *INTERSPEECH*, 2016, pp. 685–689.
- [5] E. Simonnet, S. Ghannay, N. Camelin, Y. Estève, and R. De Mori, "ASR error management for improving spoken language understanding," in *INTERSPEECH*, 2017.
- [6] D. H. Klatt, "Review of the ARPA speech understanding project," *The Journal of the Acoustical Society of America*, vol. 62, pp. 1345, 1977.
- [7] R. Kuhn and R. De Mori, "The application of semantic classification trees to natural language understanding," *TPAMI* 17, pp. 449–460, 1995.
- [8] Y. Qian, R. Ubale, V. Ramanaryanan, P. Lange, D. Suendermann-Oeft, K. Evanini, and E. Tsuprun, "Exploring ASR-free end-to-end modeling to improve spoken language understanding in a cloud-based dialog system," *ASRU*, pp. 569–576, 2017.
- [9] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, "Towards End-to-end Spoken Language Understanding," *ICASSP*, pp. 5754–5758, 2018.
- [10] R. Price, M. Mehrabani, and S. Bangalore, "Improved End-To-End Spoken Utterance Classification with a Self-Attention Acoustic Classifier," *ICASSP*, pp. 8504–8508, 2020.
- [11] 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," *SLT*, pp. 720–726, 2018.
- [12] N. Tomashenko, C. Raymond, A. Caubriere, R. De Mori, and Y. Estève, "Dialogue History Integration into End-to-End Signal-to-Concept Spoken Language Understanding Systems," *ICASSP*, pp. 8509–8513, 2020.
- [13] P. Wang, L. Wei, Y. Cao, J. Xie, and Z. Nie, "Large-scale unsupervised pre-training for end-to-end spoken language understanding," *ICASSP*, pp. 7994–7998, 2020.
- [14] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *ICLR*, 2020.
- [15] J. Chorowski, R. J. Weiss, S. Bengio, and A. van den Oord, "Unsupervised speech representation learning using wavenet autoencoders," *TASLP*, pp. 2041–2053, 2019.
- [16] Y. A. Chung and J. Glass, "Improved speech representations with multi-target autoregressive predictive coding," in *ACL*, 2020, pp. 2353–2358.
- [17] S. Khurana, A. Laurent, and J. Glass, "Cstnet: Contrastive speech translation network for self-supervised speech representation learning," 2020, arXiv:2006.02814.
- [18] Z. He, Y. He, Q. Wu, and J. Chen, "Fg2seq: Effectively encoding knowledge for end-to-end task-oriented dialog," in *ICASSP*, 2020, pp. 8029–8033.
- [19] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *ICLR*, 2015.
- [20] Y. Wang, T. Chen, H. Xu, S. Ding, H. Lv, Y. Shao, N. Peng, L. Xie, S. Watanabe, and S. Khudanpur, "Espresso: A fast end-to-end neural speech recognition toolkit," in *ASRU*, 2019, pp. 136–143.
- [21] T. Hori, J. Cho, and S. Watanabe, "End-to-end Speech Recognition with Word-Based Rnn Language Models," in *SLT*, 2018, pp. 389–396.
- [22] C. Gulcehre, O. Firat, K. Xu, K. Cho, L. Barrault, H. C. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On using monolingual corpora in neural machine translation," 2015, arXiv:1503.03535.
- [23] 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," *INTERSPEECH*, pp. 1198–1202, 2019.
- [24] F. Lefèvre, D. Mostefa, L. Besacier, Y. Estève, M. Quignard, N. Camelin, B. Favre, B. Jabaian, and Lina M. Rojas B., "Leveraging study of robustness and portability of spoken language understanding systems across languages and domains: the PORTMEDIA corpora," in *LREC*, 2012.
- [25] L. Devillers, H. Maynard, S. Rosset, P. Paroubek, K. McTait, D. Mostefa, K. Choukri, L. Charnay, C. Bousquet, N. Vigouroux, F. Béchet, L. Romary, J. Y. Antoine, J. Villaneau, M. Vergnes, and J. Gouliian, "The French MEDIA/EVALDA project: the evaluation of the understanding capability of spoken language dialogue systems," in *LREC*, 2004.
- [26] F. Bechet, B. Maza, N. Bigouroux, T. Bazillon, M. El-Beze, R. De Mori, and E. Arbillot, "Decoda: a call-centre human-human spoken conversation corpus," in *LREC*, 2012, pp. 1343–1347.
- [27] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-to-end continuous speech recognition using attention-based recurrent nn: First results," in *NIPS*, 2014.
- [28] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *ICASSP*, 2016.
- [29] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *ICASSP*, 2016, pp. 4945–4949.
- [30] 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 *SLT*, Athens, Greece, 2018.