



End-to-end named entity and semantic concept extraction from speech

Sahar Ghannay, Antoine Caubrière, Yannick Estève, Nathalie Camelin, Edwin Simonnet, Antoine Laurent, Emmanuel Morin

► To cite this version:

Sahar Ghannay, Antoine Caubrière, Yannick Estève, Nathalie Camelin, Edwin Simonnet, et al.. End-to-end named entity and semantic concept extraction from speech. IEEE Spoken Language Technology Workshop, Dec 2018, Athens, Greece. hal-01987740v2

HAL Id: hal-01987740

<https://hal.science/hal-01987740v2>

Submitted on 25 Jun 2020

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 NAMED ENTITY AND SEMANTIC CONCEPT EXTRACTION FROM SPEECH

S. Ghannay¹, A. Caubrière¹, Y. Estève¹, N. Camelin¹, E. Simonnet¹, A. Laurent¹, E. Morin²

¹LIUM - University of Le Mans, France

²LS2N - University of Nantes, France

ABSTRACT

Named entity recognition (NER) is among SLU tasks that usually extract semantic information from textual documents. Until now, NER from speech is made through a pipeline process that consists in processing first an automatic speech recognition (ASR) on the audio and then processing a NER on the ASR outputs. Such approach has some disadvantages (error propagation, metric to tune ASR systems sub-optimal in regards to the final task, reduced space search at the ASR output level,...) and it is known that more integrated approaches outperform sequential ones, when they can be applied. In this paper, we explore an end-to-end approach that directly extracts named entities from speech, through a unique neural architecture. On a such way, a joint optimization is possible for both ASR and NER. Experiments are carried on French data easily accessible, composed of data distributed in several evaluation campaigns. The results are promising since this end-to-end approach provides similar results (F-measure=0.66 on test data) than a classical pipeline approach to detect named entity categories (F-measure=0.64). Last, we also explore this approach applied to semantic concept extraction, through a slot filling task known as a spoken language understanding problem, and also observe an improvement in comparison to a pipeline approach.

Index Terms— End-to-end approach, Named entity recognition, Spoken language understanding, Automatic speech recognition, Deep learning.

1. INTRODUCTION

Named entities are sequences of words that bring basic pre-defined semantic information that usually refers to locations, persons, organization... that can be denoted by proper nouns or that are unique in the real world, and they usually include numeric and temporal values. Named entities often constitute the first semantic bricks to extract in order to construct a structured semantic representation of a document content.

This work was supported by the French ANR Agency through the CHIST-ERA M2CR project, under the contract number ANR-15-CHR2-0006-01, and by the RFI Atlantic2020 RAPACE project. In addition, authors would like to sincerely thank Sean Naren to make his implementation of Deep Speech 2 available, as well as the contributors to the NeuroNLP2 project.

Named entity recognition (NER) is among SLU tasks that usually extract semantic information from textual documents. Until now, NER from speech is made through a pipeline process that consists in processing first an automatic speech recognition (ASR) on the audio and then processing a NER on the ASR outputs. Such approach has some disadvantages.

For instance, ASR errors have a negative impact on the NER performances, introducing noise within the text to be processed [1]. Rule-based NER systems are usually built to process written language and are not robust to ASR errors. Machine learning based systems do not have good performance when they are trained on perfect transcriptions and deployed to process ASR ones, even if that can be partially compensated by simulating ASR errors in textual training data [2]. Additionally, ASR systems are generally tuned in order to get the lowest word error rate on a validation corpus, but this metric is not optimal to the NER task. For instance, this metric does not distinguish between errors on verbs or proper nouns while such errors do not have the same impact for NER. To compensate this problem, some dedicated metrics to tune ASR systems for better NER performances have been proposed, such as in [3]. Another inconvenience is that usually no information about named entities are used in the ASR process, while such information could help to better choose the partial recognition hypotheses that are dropped away during the decoding process. As a consequence, even when confusion networks or word lattices are used to go beyond the 1-best ASR hypothesis for a better robustness to ASR errors [4], such search space have been pruned without taking into account knowledge on named entity.

In the past, an integrated approach built on a high coupling of ASR and NER modules has been proposed [5], based on the finite-state machine (FSM) paradigm (*i.e.* transducer composition), showing that such integration can offer significant improvements in terms of NER quality. The main limit of this approach concerns the FSM paradigm itself, that is not able to natively model long distant constraint without combinatory explosion and that, by nature, can only express dependencies through a regular grammar. Another proposition to inject information about named entities in the ASR consists in directly adding some expressions of named entities into the ASR vocabulary [6], and to estimate a language model for speech recognition that take into account these named entity

expressions. The main default of a such approach is that it cannot allow to detect named entity that were not injected in the ASR vocabulary.

All of these issues motivate our research work on neural end-to-end approach to extract named entities from speech. On a such way, a joint optimization is able for both ASR and NER in a NER task perspective, the architecture is more compact than the ones used in usual pipeline, and we expect to take benefit of the deep neural architecture capacities to capture long distant constraint at the sentence level. Very recently, a similar approach has been proposed by Facebook on a paper posted on the arXiv.org website [7]. This end-to-end approach is dedicated to domain and intent classification tasks, and experiments were carried on internal data close to the spirit of the ATIS corpus, as expressed by the authors. As an extension of our study on named entity recognition, we also applied the proposed approach to semantic concept extraction.

In this paper, we present a first study of an end-to-end approach to extract named entities. Our neural architecture is very similar to the Deep Speech 2 neural ASR system proposed by Baidu in [8]. To use it for named entity recognition, we apply a multi-task training and modify the sequence of characters to be recognized from speech. Experiments were carried on French data easily accessible, and so reproducible, that were distributed in the framework on evaluation campaigns and are still available. They constitute all the available data of speech with manual annotations on named entity for French.

In the last part of this paper, we explore the use of the proposed approach to process a slot filling task, that is very close to the named entity recognition task. This SLU task consists in detecting and extracting semantic concepts in speech, in the framework of a human/machine spoken dialog dedicated to hotel booking.

This paper is structured as follows. Section 2 describes the neural ASR architecture we used. Section 3 explains how we propose to exploit a such neural architecture for named entity extraction from speech. Section 4 presents some propositions to optimize the system and also compensate the lack of manually annotated audio data. Section 5 presents our experimental results on NER, and last section 6 present some preliminary results on the slot filling task before the conclusion.

2. MODEL ARCHITECTURE

The RNN architecture used in this study is similar to the Deep Speech 2 neural ASR system proposed by Baidu in [8]. This architecture is composed of nc convolution layers (CNN), followed by nr uni or bidirectional recurrent layers, a lookahead convolution layer [9], and one fully connected layer just before the softmax layer, as shown in Figure 1.

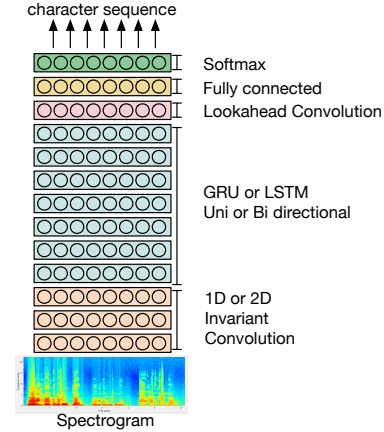


Fig. 1. Deep RNN architecture used to extract named entities from French speech.

The system is trained end-to-end using the CTC loss function [10], in order to predict a sequence of characters from the input audio. In our experiments we used two CNN layers and six bidirectional recurrent layers with batch normalization as mentioned in [8].

Given an utterance x^i and label y^i sampled from a training set $X = (x^1, y^1), (x^2, y^2), \dots$, the RNN architecture has to train to convert an input sequence x^i into a final transcription y^i s. For notational convenience, we drop the superscripts and use x to denote a chosen utterance and y the corresponding label. The RNN takes as input an utterance x represented by a sequence of log-spectrograms of power normalized audio clips, calculated on 20ms windows. As output, all the characters l of a language alphabet may be emitted, in addition to the space character used to segment character sequences into word sequences (space denotes word boundaries). The RNN makes a prediction $p(l_t|x)$ at each output time step t .

At test time, the CTC model is coupled with a language model trained on a big textual corpus. A specialized beam search CTC decoder [11] is used to find the transcription y that maximizes:

$$Q(y) = \log(p(l_t|x)) + \alpha \log(pLM(y)) + \beta wc(y) \quad (1)$$

where $wc(y)$ is the number of words in the transcription y . The weight α controls the relative contributions of the language model and the CTC network. The weight β controls the number of words in the transcription.

3. NAMED ENTITY EXTRACTION PROCESS

In the literature, many studies focus on named entity recognition from text. State-of-the-art systems are based on neural networks architectures. Some of them rely heavily on hand-crafted features and domain-specific knowledge [12, 13].

Recent approaches [14, 15] takes benefits from both word and/or character-level embeddings learned automatically, by using combination of bidirectional LSTM, CNN and CRF. However, named entities recognition from automatic transcriptions is less studied. This task is made through a pipeline process that consists in processing first an automatic speech recognition (ASR) on the audio and then processing a NER on the ASR outputs [16]. Usually, the named entity recognition task is to assign a named entity tag to every word in a sentence. A single named entity could concern several words within a sentence. For this reason, the word-level labels begin-inside-outside (BIO) encoding [17] is very often adopted.

In this preliminary study, we focus on named entity extraction from speech using the network described above, without changing the neural architecture. We would like to evaluate if this neural architecture is able to capture high level semantic information that allow it to recognize named entities. For that, we propose to modify the character sequence that the neural network has to produce: information about named entities are added in the initial character sequence. Instead of applying a BIO approach, we propose to add some tag characters in this sequence to delimit named entities boundaries, but also their category. We are interested to eight NE categories that are: *person*, *function*, *organization*, *location*, *production*, *amount*, *time* and *event*.

In our experiments, the system will attribute a starting tag or an ending one only before and after the named entities, the other words are not concerned. To distinguish the named entity category, we consider a starting tag for each NE category. Only one ending tag is used for all the NE categories, considering that since there is no overlap between named entities in a such representation, this information is sufficient to delimit the end of a named entity.

According to the eight named entity categories targeted by the task, nine NE tags has to be added to the character list emitted by the neural network: “< *pers*”, “< *func*”, “< *org*”, “< *loc*”, “< *prod*”, “< *amount*”, “< *time*”, “< *event*”, and “>”. A such neural model predicts a character at each time step: in our experiments, each NE tag is considered as one single special character. With this way, the NE tags are included in the prediction process, and are taken into account by the CTC loss function during the training process.

4. MULTI-TASK TRAINING, DATA AUGMENTATION, AND STARRED MODE

Audio recordings with both manual transcriptions and manual annotations of named entities are relatively rare, while neural end-to-end approaches are known to need large amount of data to become competitive.

To compensate this lack of data, we first propose to apply a multi-task learning approach to train the neural network. This consists in starting to train it only for the ASR task,

without emitting character used to represent named entities, on all the audio recordings available with their manual transcriptions. At the end, the softmax layer is reinitialized to take into consideration the named entity tag markers, and a new training process is realized, on the named entity recognition task, with only training data with manual annotations of named entities.

A second proposition consists in artificially increasing the training data for the named entity recognition task. For this purpose, we propose to apply a named entity recognition system dedicated to text data in order to tag the manual transcriptions used to train the ASR neural network. Then, these manual transcriptions automatically annotated with named entities can be injected in the training data used to train the neural network to extract named entities from speech.

In addition, since we want the system to focus on named entities, and since the CTC loss gives the same importance to each character, we propose to modify the character sequence that the neural network must emit to give more importance to named entities. This proposition is interesting to better understand how the CTC loss behaves on this case, and consists in replacing by a star “*” all character subsequences that do not contain a named entity. For instance, the character sequence presented in Figure 2 becomes: * < *pers* césar > * < *time* hier > * < *loc* paris > * < *amount* soixante dix sept ans >. We call this approach the *starred mode*, and we expect that it can make the neural model more sensitive to named entities.

5. EXPERIMENTS

5.1. Experimental setups

Experiments have been carried out on four different French corpora, including ESTER 1&2, ETAPE and Quaero. These corpora are composed of data recorded from francophone radio and TV stations, and are annotated with named entities. They are all the publicly available data of speech with manual annotations on named entities for French language. The ESTER corpora were divided into three parts: training, development and evaluation. ESTER 1 [18] training (73 hours) and development (17 hours) corpora are composed of data recorded from four radio stations in French. ESTER 1 test corpus is composed of 10 hours coming from the same four radio stations plus two other stations, all of which recorded 15 month after the development data. ESTER 2 [19] training corpus was not annotated with named entities and was not used in this study. The development (17 hours) and test set (10 hours) is composed of manual transcriptions of speech recorded from six radio stations (two of those radio stations were already used in ESTER 1). The ETAPE [20] data consists of manual transcriptions and annotations of TV and radio shows. It contains 36 hours of speech, recorded between 2010 and 2011, divided into three parts: training (22 hours), development (7 hours) and test (7h). QUAERO (ELRA-S0349)

word sequence	le sculpteur césar est mort hier à paris à l' âge de soixante dix sept ans
word sequence with EN tags	le sculpteur <pers césar > est mort <time hier > à <loc paris > à l' âge de <amount soixante dix sept ans >
starred mode: only EN tags, EN values, and stars	* <pers césar > * <time hier > * <loc paris > * <amount soixante dix sept ans >

Fig. 2. Example of mapping the real NE tags to character sequence. This sentence means, in English and case sensitive: "the sculptor Caesar died yesterday in Paris at the age of seventy-seven years"

data is composed of 12 hours of manual transcriptions of TV and radio shows coming from 6 different sources recorded in 2010. Our experimental corpus is the combination of those four corpora. The training corpus is composed of the training sets of ESTER 1, ETAPE and QUAERO, while the development and test sets are composed respectively of the development and test sets of ESTER 1&2, and ETAPE. It contains almost 160 hours of speech (training 107 hours, test 24 hours, development 30 hours).

The performance of our approach is evaluated in terms of precision (P), recall(R) and F-measure for named entity detection, the named entity/value detection and the accuracy of the value detection when the named entities tags are correctly detected. These evaluations are made with the help of the *sclite*¹ tool.

5.2. Multi-task training

For multi-task training, we first train the E2E architecture only for ASR task, without emitting character used to represent named entities. The system is trained on all the audio recordings available with their manual transcriptions around 297.7 hours of training set, including the data described above. It composed of two convolution layers and six BLSTM layers with batch normalization, the number of epochs was set to 35. This system achieves 20.70% word error rate (WER) and 8.01% character error rate (CER) on the development corpus (30.2 hours) and 19.95% of WER and 7.68% of CER on the test set (40.8 hours). These results were obtained by applying a CTC beam search decoding coupled with a trigram language model. Once this system is trained, the softmax layer is reinitialized to take into consideration the named entity tag markers, and a new training process is realized, on the named entity recognition task, with only training data with manual annotations of named entities. In addition, for the training of both E2E and ASR systems, each training audio samples is randomly perturbed in gain and tempo for each iteration.

5.3. Experimental results

We present in this section some experimental results. Table 1 shows the performances of the end-to-end model (E2E) to detect EN categories (among the eight ones). That means that in this evaluation we do not take care of values associated to the detected EN. The starred mode is also experimented and is called (E2E*) in the table: this mode provides better results in this task than the normal mode.

Table 1. Named entity category detection results for E2E and E2E* (starred mode) systems

System	Corpus	Precision	Recall	F-measure
E2E	dev	0.85	0.57	0.68
E2E	test	0.83	0.52	0.64
E2E*	dev	0.75	0.65	0.71
E2E*	test	0.82	0.57	0.67

Table 2 evaluates the quality of the category/value pairs that have been recognized. While precision and recall do not have the same behavior between normal and starred mode, both modes gets the same F-measure value.

Table 2. Named entity category+value pair detection results for E2E and E2E* systems

System	Corpus	Precision	Recall	F-measure
E2E	dev	0.64	0.45	0.53
E2E	test	0.55	0.36	0.44
E2E*	dev	0.57	0.47	0.52
E2E*	test	0.47	0.38	0.42

Last, we would like to compare these results to the ones obtained by a pipeline process, that consists in applying a text named entity recognition on the automatic transcripts produced by the end-to-end ASR system trained on the first step of the multi-task learning presented above.

The text named entity recognition system used for this experiment is based on the combination of bi-directional LSTM (BLSTM), CNN and CRF modules [15], and takes benefits from both word and character-level embeddings learned auto-

¹<http://www.icsi.berkeley.edu/Speech/docs/setk-1.2/sclite.htm>

matically during the training process. For this experiment, we used the NeuroNLP2 implementation². Convolutional neural network encodes character-level information of a word into its character-level embedding. Then the character- and word-level embeddings are fed into the BLSTM to model context information of each word. On top of BLSTM, the sequential CRF is used to jointly decode labels for the whole sentence. In addition, this system can be enriched with syntactic information like part of speech tagging (POS). In our experiment, NeuroNLP2 is used as a NER system and Deep Speech 2 as the ASR system. Both are trained on the experimental corpus described in section 5.1. Automatic transcriptions of development and test data have been annotated with NER system. To measure the impact of POS, we used the MACAON system [21] to tag the experimental corpus on manual and automatic transcriptions. To feed NeuroNLP2, one-hot vectors represent POS information. Word embeddings, character representations and one-hot concatenations feed the BLSTM layer. As we can see in Tables 3, the pipeline process is a little bit less competitive than the end-to-end model to recognize EN category, while it is more efficient to extract EN values. Results also confirm that linguistic information like POS is really important for the NER task. Such observation will help for future work on the continuity of this study.

Table 3. NER results for the pipeline approach (Pip) on the **test** data. When POS are used to tag ASR outputs before NER processing, the system is called Pip+POS

System	Detection	Precision	Recall	F-measure
Pip	category	0.75	0.56	0.64
Pip+POS	category	0.74	0.58	0.65
Pip	cat+value	0.58	0.43	0.49
Pip+POS	cat+value	0.57	0.45	0.50

As described in section 4, we applied NeuroNLP2 (the version using POS tagging) on the manual transcriptions of the ASR training data in order to augment the amount of "NER from speech" training data. In this experiment, the normal and starred modes were used. Table 4 shows the improvement got by the end-to-end system when training on these imperfect augmented data using the normal (E2E+) and the starred (E2E+*) modes. As we can see, the use of the augmented data was helpful for the starred mode.

6. EXTENSION TO SLOT FILLING TASKS

In order to deeper explore this end-to-end approach, we also carried on some experiments on a SLU task. In this context, semantic representation consists classically in the use of *frames* describing general concepts and their specific instances. A frame is composed by a data structure which represents a predefined concept by associating to the concept

Table 4. NER results on the **test** data for the E2E system trained with imperfect augmented data (E2E+) in comparison to the E2E system trained with imperfect augmented data and the starred mode (E2E+*)

System	Detection	Precision	Recall	F-measure
E2E+	category	0.82	0.57	0.67
E2E+*	category	0.76	0.63	0.69
E2E+	cat+value	0.55	0.40	0.46
E2E+*	cat+value	0.49	0.41	0.47

name a set of roles: these roles are represented by *slots*. In this framework, SLU corresponds to a slot filling task and, classically, this task can be defined as a concept tagging process, which is the extraction of a sequence of concepts out of a given word sequence [22]. We consider this task very close to the named entity recognition task explored above.

By applying the same end-to-end approaches presented in the previous sections, we would like to check if the starred mode is still interesting, and also explore the impact of multi-task learning in a scenario on which we own training data dedicated to two different applicative tasks.

6.1. Experimental corpora

For these experiments, we focus on two French corpora dedicated to semantic extraction from speech in a context of human/machine dialogues. The main target is the MEDIA corpus, that is dedicated to hotel booking [23]. The other one is the PORT-MEDIA corpus, dedicated to reservation of theater tickets [24]. Both are composed of telephone conversations, and were collected through a wizard of Oz approach, in which a human plays the role of the machine.

The MEDIA corpus is composed of 1257 dialogues. It is split into three subparts: a training corpus containing 17.7k sentences, a development corpus containing 1.3k sentences, and a test corpus of 3.5k sentences. 76 semantic tags have to be recognized (e.g. "number of room", "hotel name", "localization", "month (date)", "room equipment",....). The PORT-MEDIA corpus contains 10.4k sentences, and uses 35 semantic tags, very close to tags existing in MEDIA. For both MEDIA and PORTMEDIA corpora we removed short speech segments inferior to 1 second, excluding short answers like 'yes' or 'no' usually easy to detect.

Since training data for speech recognition on telephone conditions are more rare than for broadcast news, all the data used to pre-training an end-to-end ASR system were converted into telephone format. We also used the MEDIA training data and the PORT-MEDIA training data for the ASR pre-training. Last, we injected data from the DECODA corpus [25], a corpus of human/human conversations from the Parisian public transportation call-center: this corresponds to around 74h of audio signal with manual transcriptions.

In order to try to take benefits from the PORTMEDIA data, we used them in a SLU pre-training step, that uses words

²<https://github.com/XuezheMax/NeuroNLP2>

and semantic tags as outputs, just after the ASR pre-training that uses only words as outputs. Figure 3 illustrates this training process. In a contrastive experiment, we also skipped the SLU pre-training step (dashed line). All the results given in the next section were reached without beam search applied on the neural outputs: only the greedy mode was used.

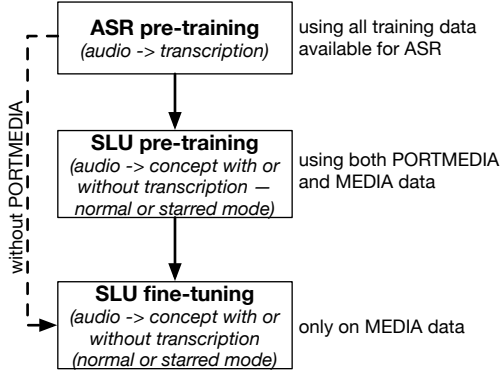


Fig. 3. Training steps to build an end-to-end SLU system dedicated to the MEDIA data.

6.2. Results

Concept error rate is a metric very commonly used in previous paper on SLU, especially when using the MEDIA corpus [5, 22, 26]: for this reason, we used this metric in our experiments.

Table 5 shows that the use of data from a different applicative domain (buying theater tickets instead of booking hotel rooms) seems really useful: on the test MEDIA corpus, the CER decreases of 4% (29.3 to 28.1) when the PORTMEDIA data are used in a SLU pre-training step.

Table 5. Comparison of semantic concept detection results for E2E systems on the test MEDIA corpus according to the use of not of the PORT-MEDIA corpus before fine-tuning on the MEDIA training data. Results are evaluated with the concept error rate (CER) metric.

System	CER (%)
E2E without PORTMEDIA	29.3
E2E pretrained (SLU) with PORTMEDIA	28.1

Table 6 shows the impact of the starred mode that is again more efficient than the normal mode, since the CER continues to decrease from 28.1% to 27.0% on the MEDIA test set.

Last, table 7 allows us to compare the end-to-end approach to a pipeline approach, as we did for named entity recognition. The end-to-end ASR system made during the first step described in figure 3 reaches a word error rate of 20.4% on the MEDIA test data. The *Pip-SLU* system consists in applying on these automatic transcriptions a neural encoder-decoder architecture with attention mechanism (NEDAM) trained to recognize semantic concepts. This SLU

Table 6. Semantic concept detection results for E2E and E2E* systems on the test MEDIA corpus. In both condition, the PORT-MEDIA corpus was used during the SLU pre-training

System	CER (%)
E2E	28.1
E2E*	27.0

architecture has been presented in [27] and more details are given in [28]. We observe that the end-to-end SLU approach, in normal mode and trained without using additional semantic training data from PORTMEDIA, reached better results than the *Pip-SLU* system. And, as seen before, using both SLU pre-training with additional data and starred mode still decreases the CER, for a global reduction of 15% (32.0 to 27.0).

Table 7. Comparison of semantic concept detection results for E2E systems and a pipeline of neural system (ASR + SLU) on the test MEDIA corpus

System	CER (%)
Pip-SLU	32.0
E2E without PORTMEDIA	29.3
E2E* pretrained (SLU) with PORTMEDIA	27.0

7. CONCLUSION

This paper presents a first study about end-to-end named entity and semantic concept extraction from speech. By integrating in the character sequence emitted by a CTC end-to-end speech recognition system some special characters to delimit and categorize named entities or semantic concepts, we showed that such extraction is feasible. To compensate the lack of training data, we propose a multi-task learning approach (ASR + NER or ASR + SLU), in addition to an artificial data augmentation of the training corpus with automatic annotation for named entities, or to a transfer learning applied on semantic data coming from another domain for semantic concept extraction. A starred mode is also proposed to make the neural network more focused on named entities or semantic concepts. Experimental results show that this end-to-end approach in starred mode with training augmentation provides better results than a pipeline approach for both named entity and semantic concept extractions. On the other side, we observed that performances of this end-to-end approach to extract named entity values are worse than the ones got by the pipeline process. As a conclusion, this study presents promising results in a first attempt to experiment an end-to-end approach to extract named entities and semantic concepts. It constitutes an interesting start point for future work that could help to improve its performances, or that could study intermediate representations within the neural architecture in order to better understand how the neural network handles the semantic level of speech.

8. REFERENCES

- [1] Mohamed Hatmi, Christine Jacquin, Emmanuel Morin, and Sylvain Meignier, “Named entity recognition in speech transcripts following an extended taxonomy,” in *First Workshop on Speech, Language and Audio in Multimedia*, 2013.
- [2] Edwin Simonnet, Sahar Ghannay, Nathalie Camelin, and Yannick Estève, “Simulating ASR errors for training SLU systems,” in *LREC 2018*, Miyazaki, Japan, May 2018.
- [3] Mohamed Ameer Ben Jannet, Olivier Galibert, Martine Adda-Decker, and Sophie Rosset, “Investigating the effect of ASR tuning on named entity recognition,” *Proc. Interspeech 2017*, pp. 2486–2490, 2017.
- [4] Dilek Hakkani-Tür, Frédéric Béchet, Giuseppe Ricciardi, and Gokhan Tur, “Beyond ASR 1-best: Using word confusion networks in spoken language understanding,” *Computer Speech & Language*, vol. 20, no. 4, pp. 495–514, 2006.
- [5] Christophe Servan, Christian Raymond, Frédéric Béchet, and Pascal Nocéra, “Conceptual decoding from word lattices: application to the spoken dialogue corpus media,” in *The Ninth International Conference on Spoken Language Processing (Interspeech 2006-ICSLP)*, 2006.
- [6] Mohamed Hatmi, Christine Jacquin, Emmanuel Morin, and Sylvain Meignier, “Incorporating named entity recognition into the speech transcription process,” in *Proceedings of the 14th Annual Conference of the International Speech Communication Association (Interspeech’13)*, 2013, pp. 3732–3736.
- [7] Dmitriy Serdyuk, Yongqiang Wang, Christian Fuegen, Anuj Kumar, Baiyang Liu, and Yoshua Bengio, “Towards end-to-end spoken language understanding,” *arXiv preprint arXiv:1802.08395*, 2018.
- [8] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al., “Deep speech 2: End-to-end speech recognition in English and Mandarin,” in *International Conference on Machine Learning*, 2016, pp. 173–182.
- [9] Chong Wang, Dani Yogatama, Adam Coates, Tony Han, Awni Hannun, and Bo Xiao, “Lookahead convolution layer for unidirectional recurrent neural networks,” 2016.
- [10] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 369–376.
- [11] Awni Y Hannun, Andrew L Maas, Daniel Jurafsky, and Andrew Y Ng, “First-pass large vocabulary continuous speech recognition using bi-directional recurrent DNNs,” *arXiv preprint arXiv:1408.2873*, 2014.
- [12] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa, “Natural language processing (almost) from scratch,” *Journal of Machine Learning Research*, vol. 12, no. Aug, pp. 2493–2537, 2011.
- [13] Jason PC Chiu and Eric Nichols, “Named entity recognition with bidirectional LSTM-CNNs,” *arXiv preprint arXiv:1511.08308*, 2015.
- [14] Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer, “Neural architectures for named entity recognition,” *arXiv preprint arXiv:1603.01360*, 2016.
- [15] Xuezhe Ma and Eduard Hovy, “End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF,” *arXiv preprint arXiv:1603.01354*, 2016.
- [16] Christian Raymond, “Robust tree-structured named entities recognition from speech,” in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 8475–8479.
- [17] Lance A. Ramshaw and Mitchell P. Marcus, “Text chunking using transformation-based learning,” *CoRR*, vol. cmp-lg/9505040, 1995.
- [18] Sylvain Galliano, Edouard Geoffrois, Djamel Mostefa, Khalid Choukri, Jean-François Bonastre, and Guillaume Gravier, “The ESTER phase II evaluation campaign for the rich transcription of French broadcast news,” in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [19] Sylvain Galliano, Guillaume Gravier, and Laura 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.
- [20] Guillaume Gravier, Gilles Adda, Niklas Paulson, Matthieu Carré, Aude Giraudel, and Olivier 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.

- [21] Alexis Nasr, Frédéric Béchet, Jean-François Rey, Benoît Favre, and Joseph Le Roux, “Macaon: An NLP tool suite for processing word lattices,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Systems Demonstrations*. Association for Computational Linguistics, 2011, pp. 86–91.
- [22] Stefan Hahn, Marco Dinarelli, Christian Raymond, Fabrice Lefevre, Patrick Lehn, Renato De Mori, Alessandro Moschitti, Hermann Ney, and Giuseppe Riccardi, “Comparing stochastic approaches to spoken language understanding in multiple languages,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 19, no. 6, pp. 1569–1583, 2011.
- [23] Hélène Bonneau-Maynard, Sophie Rosset, Christelle Ayache, Anne Kuhn, and Djamel Mostefa, “Semantic annotation of the French MEDIA dialog corpus,” in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [24] Fabrice Lefevre, Djamel Mostefa, Laurent Besacier, Yannick Esteve, Matthieu Quignard, Nathalie Camelin, Benoit Favre, Bassam Jabaian, and Lina Maria Rojas Barahona, “Leveraging study of robustness and portability of spoken language understanding systems across languages and domains: the PORTMEDIA corpora,” in *The International Conference on Language Resources and Evaluation*, 2012.
- [25] Frederic Bechet, Benjamin Maza, Nicolas Bigouroux, Thierry Bazillon, Marc El-Beze, Renato De Mori, and Eric Arbillot, “DECODA: a call-centre human-human spoken conversation corpus,” in *LREC*, 2012, pp. 1343–1347.
- [26] Vedran Vukotic, Christian Raymond, and Guillaume Gravier, “Is it time to switch to word embedding and recurrent neural networks for spoken language understanding?,” in *InterSpeech*, 2015.
- [27] Edwin Simonnet, Nathalie Camelin, Paul Deléglise, and Yannick Estève, “Exploring the use of attention-based recurrent neural networks for spoken language understanding,” in *Machine Learning for Spoken Language Understanding and Interaction NIPS 2015 Workshop (SLUNIPS 2015)*, 2015.
- [28] Edwin Simonnet, Sahar Ghannay, Nathalie Camelin, Yannick Estève, and Renato De Mori, “ASR error management for improving spoken language understanding,” in *Interspeech 2017*, 2017.