



ASR error management for improving spoken language understanding

Edwin Simonnet¹, Sahar Ghannay¹, Nathalie Camelin¹, Yannick Estève¹, Renato De Mori^{2,3}

¹LIUM, Le Mans University, France

²LIA, University of Avignon, France

³McGill University, Montreal, Canada

firstname.lastname@univ-lemans.fr, rdemori@cs.mcgill.ca

Abstract

This paper addresses the problem of automatic speech recognition (ASR) error detection and their use for improving spoken language understanding (SLU) systems. In this study, the SLU task consists in automatically extracting, from ASR transcriptions, semantic concepts and concept/values pairs in a *e.g.* touristic information system. An approach is proposed for enriching the set of semantic labels with error specific labels and by using a recently proposed neural approach based on word embeddings to compute well calibrated ASR confidence measures. Experimental results are reported showing that it is possible to decrease significantly the Concept/Value Error Rate with a state of the art system, outperforming previously published results performance on the same experimental data. It also shown that combining an SLU approach based on conditional random fields with a neural encoder/decoder attention based architecture, it is possible to effectively identifying confidence islands and uncertain semantic output segments useful for deciding appropriate error handling actions by the dialogue manager strategy.

Index Terms: spoken language understanding, speech recognition, robustness to ASR errors

1. Introduction

In spite of impressive research efforts and recent results, systems for semantic interpretation of text and speech still make errors. Some of the problems common to text and speech are: difficulty of concept mention localization, ambiguities intrinsic in localized mentions, deficiency to identify sufficient contextual constraints for solving interpretation ambiguities. Additional problems are introduced by the interaction between a spoken language understanding (SLU) system and an error prone automatic speech recognition (ASR) system. ASR errors may affect the mention of a concept and the value of a concept instance. Automatically detected context features for disambiguating concept mentions may be insufficient and different interpretation methods may produce different concept hypotheses. Furthermore, context features that are effective on manual transcriptions may be affected by ASR errors causing SLU interpretation errors.

The focus of this paper¹ is to detect some errors of these two types. For this purpose, suitable ASR confidence measures are introduced for localizing ASR word errors that may affect SLU performance. They are used as additional SLU features to be combined with lexical and syntactic features useful for characterizing concept mentions. For this purpose, an ASR error detection sub-system has been endowed with confidence features based on syntactic dependencies and other se-

mantically relevant word features. Two different SLU architectures equipped with other sets of confidence and specific word features are introduced for generating multiple hypotheses from different representations of mentions and contexts. The architectures are based on conditional random fields (CRF) and an encoder-decoder neural network structure with a mechanism of attention (NN-EDA). Experimental results showing significant reduction for the French MEDIA corpus on concept error rate (CER) and concept-value error rate (CVER) confirm the expected benefit of introducing semantic specific ASR features. Optimal combinations of these architectures provide additional improvements with a CER relative reduction of 18.9% and a CVER relative reduction of 10.3% with respect to a baseline described in [1] not using these features and based only on CRFs.

2. Related work

Some SLU errors are caused by certain types of ASR errors. In general, ASR errors are reduced by minimizing the expected word error rate in parameter estimation [2]. The effect of word errors can be controlled by associating a single sentence hypothesis with word confidence measures. In [3] methods are proposed for constructing confidence features for improving the quality of a semantic confidence measure. Methods proposed for confidence calibration are based on a maximum entropy model with distribution constraints, artificial neural networks [4], and the deep belief network (DBN). More recently, new features and bidirectional recurrent neural networks (RNN) have been proposed for ASR error detection [5]. Most SLU systems reviewed in [6] generate hypotheses of semantic frame slot tags expressed in a spoken sentence analyzed by an ASR system. The use of deep neural networks (DNN) appeared in more recent systems [7]. Bidirectional RNNs with long-short term memory (LSTM) have been used for semantic frame slot tagging [8]. In [9], LSTMs have been proposed with a mechanism of attention for parsing text sentences to logical forms. Following [10], in [11] a convolutional neural network (CNN) is proposed for encoding the representation of knowledge expressed in a spoken sentence. This encoding is used as an attention mechanism for constraining the hypothesization of slot tags expressed in the same sentence. Most recent papers using sophisticated SLU architectures based on RNNs use as input just the best sequence of word hypotheses passed by an ASR system. In this paper, two SLU architectures are considered. The first one is based on a neural machine translation architecture [12] and integrates context information with an attention based decoder as in [13]. The second one integrates context information in the same architecture used in [1] based on conditional random fields (CRF). Both SLU systems receive word hypotheses generated by the same ASR sub-system and scored with confidence measures computed by a neural architecture embedding new types of semantically relevant confidence features.

¹This work was funded by the ANR agency through the CHIST-ERA ERA-Net JOKER contract number ANR-13-CHR2-0003-05.

3. ASR error detection and confidence measure

Two different confidence measures are used for error detection. The first one is the word *posterior* probability computed with confusion networks as described in [2]. The other one is a variant of a new approach, introduced in [14, 15]. The latter measure is computed with a Multi-Stream Multi-Layer Perceptron (MS-MLP) architecture, fed by different heterogeneous confidence features. Among them, the most relevant for SLU are word embeddings of the targeted word and its neighbors, length of the current word, language model backoff behavior, part of speech (POS) tags, syntactic dependency labels and word governors. Other features, such as prosodic features and acoustic word embeddings described in [16] and [15] could also be used but were not considered in the experiments described in this paper. A particular attention was carried on the word embeddings computation, which is the result of a combination of different well known word embeddings (CBOW, Skip-gram, GloVe) made through the use of a neural auto-encoder in order to improve the performances of this ASR error detection system [17].

The MS-MLP proposed here for ASR error detection has two output units. They compute scores for *Correct* and *Error* labels associated with an ASR generated hypothesis. This hypothesis is evaluated by the softmax value of the *Correct* label scored with the MS-MLP. Experiments have shown that this is a calibrated confidence measure more effective than word posterior probability when comparison is based on the Normalized Cross Entropy (NCE) [16], which measures the information contribution provided by confidence knowledge.

Table 1 shows the NCE values obtained by these two confidence measures on the MEDIA test data whose details can be found in section 5.

Table 1: Comparison of ASR error prediction capabilities of probability a posteriori and confidence measure derived from the MS-MLP ASR error detection system in terms of Normalized Cross Entropy on the MEDIA test set.

	PAP	MS-MLP
NCE	0.147	0.462

Figure 1 shows the predictive capability of the confidence measure based on MLP-MS compared to word posterior probability on the MEDIA test data. The curve shows the predicted percentage of correct words as a function of confidence intervals. The best measure is the one for which percentages are the closest to the diagonal line.




Figure 1: Comparison of ASR error prediction capabilities of probability a posteriori and confidence measure derived from the MS-MLP ASR error detection system on the MEDIA test set.

Thanks to these confidences measures, we expect to get relevant information in order to better handle ASR errors in a spoken language understanding framework.

4. SLU features and architectures

Two basis SLU architectures are considered to carry experiments on the MEDIA corpus (described in sub-section 5.1). The first one is an encoder/decoder recurrent neural architecture with a mechanism of attention (NN-EDA) similar to the one used for machine translation proposed in [13]. The second one is based on conditional random fields (CRF - [18]). Both architectures build their training model on the same features encoded with continuous values in the first one and discrete values in the second one

4.1. Set of Features

Word features, including those defined for facilitating the association of a word with a semantic content, are defined as follows:

- the word itself
- its pre-defined semantic categories which belongs to:
 - MEDIA specific categories: like names of the streets, cities or hotels, lists of room equipments, food type, ... *e.g.*: TOWN for Paris
 - more general categories: like figures, days, months, ... *e.g.*: FIGURE for thirty-three.
- a set of syntactic features: the MACAON tool [19] is applied to the whole turn in order to obtain for each word its following tags: the lemma, the POS tag, its word governor and its relation with the current word.
- a set of morphological features: the 1-to-4 first letter ngrams, the 1-to-4 letter last ngrams of the word and a binary feature that indicates if the first letter is an upper one.
- the two ASR confidence measures : the ASR *posterior* probability (*pap*) and the MS-MLP confidence measure as described in section 3.

These word features added in input with the words help the SLU to achieve better understanding, inspired from [1]. The two SLU architectures take all those features except the two confidence measures that can be taken partially: one or another or both according to the most powerful configuration. The SLU architectures also need to be calibrated on their respective hyper-parameters in order to give the best results. The way the best configuration is chosen is described in 5.3.

4.2. Neural EDA system

The proposed RNN² encoder-decoder architecture with an attention-based mechanism (NN-EDA) is inspired from a machine translation architecture and depicted in figure 2. The concept tagging process is considered as a translation problem from words (source language) to semantic concept tags (target language). This bidirectional RNN encoder is based on Gated Recurrent Units (GRU) and computes an annotation h_i for each word w_i from the input sequence w_1, \dots, w_T . This annotation is the concatenation of the matching forward hidden layer state and the backward hidden layer state obtained respectively by the forward RNN and the backward RNN comprising the bidirectional RNN. Each annotation contains the summaries of the dialogue turn contexts respectively preceding and the following a considered word.

The sequence of annotations h_1, \dots, h_T is used by the decoder to compute a context vector c_t (represented as a crossed circle in figure 2). A context vector is recomputed after each emission of an output label. This computation takes into account a weighted sum of all the annotations computed by the

²Implemented from nmtpty (<https://github.com/lium-lst/nmtpty>).

encoder. This weighting depends on the current output target, and is the core of the attention mechanism: a good estimation of these weights α_{tj} allows the decoder to choose parts of the input sequence to pay attention to. This context vector is used by the decoder in conjunction with the previous emitted label output y_{t-1} and the current state s_t of the hidden layer of a RNN to make a decision about the current label output y_t . A more detailed description of recurrent neural networks and attention based ones can be found in [12].




Figure 2: Architecture of the NN-EDA SLU system.

4.3. CRF system

Past experiments described in [1] have shown that the best semantic annotation performance on manual and automatic transcriptions of the MEDIA corpus were obtained with CRF systems. More recently in [20], this architecture has been compared to popular bi-directional RNN (bi-RNN). The results was that CRF systems outperform a bi-RNN architecture on the MEDIA corpus, while better results were observed by bi-RNN on the ATIS [21] corpus. This is probably explained by the fact that MEDIA contains semantic contents whose mentions are more difficult to disambiguate and CRFs make it possible to exploit complex contexts more effectively.

For the sake of comparison with the best SLU system proposed in [1], the Wapiti toolkit was used [22]. Nevertheless, the set of input features used by the system proposed in this paper is different from the one used in [1]. Among the novelties used in our system, we consider syntactic and ASR confidence features and our configuration template is different. After many experiments performed on DEV, our final feature template includes the previous and following instances for words and POS in a unigram or a bigram to associate a semantic label with the current word. Also associated with the current word are semantic categories of the two previous and two following instances. The other features are only considered at the current position.

Furthermore, the tool discretize4CRF³ is used to apply a discretization function to the ASR confidence measures in order to obtain several discrete values that can be accepted as input features by the CRF.

5. Experimental setup and results

Experiments were carried with the MEDIA corpus as in [1]. For the sake of comparison, the results of their best understanding system is reported in this paper as *baseline*. However, as the

³<https://gforge.inria.fr/projects/discretize4crf/>

WER of the ASR used in this paper is lower (23.5%) than the one used in the *baseline*, rigorous conclusions can be drawn only on comparisons between the different SLU components introduced in this paper.

5.1. The MEDIA corpus

The MEDIA corpus was collected in the French Media/Evalda project [23] and deals with negotiation of tourist services. It contains three sets of telephone human/computer dialogues, namely: a training set (TRAIN) with approximately 17.7k sentences, a development set (DEV) with 1.3k sentences and an evaluation set (TEST) containing 3.5k sentences. The corpus was manually annotated with semantic concepts characterized by a label and its value. Other types of semantic annotations (such as mode or specifiers) are not considered in this paper to be consistent with the experimental results provided in [1]. Annotations also associate a word sequence to the concepts. Annotations add a B (beginning) or I (inside) letter to each word of a sequence annotated with a concept tag while words not expressing any domain concept are annotated with the null (Outside) tag. These sequences have to be considered as estimations of concept localized mentions. Evaluations are performed with the DEV and TEST sets and report concept error rates (CER) for concept labels only and concept-value error rates (CVER) for concept-value pairs. It is worth mentioning that the number of concepts annotated in a turn has a large variability and may include more than 30 annotated concepts. Among the concepts types there are some, such as three different types of REFERENCE and CONNECTOR of application domain entities. The mentions of these concepts are often short confusable sequences of words.

5.2. LIUM ASR system dedicated to MEDIA

For these experiments, a variant of the ASR system developed by LIUM that won the last evaluation campaign on French language has been used [24]. This system is based on the Kaldi speech recognition toolkit [25]. The training set used to estimate the DNN (Deep Neural Networks) acoustic models parameters consists of 145,781 speech segments from several sources: the radiophonic broadcast ESTER [26] and ESTER2 [27] corpora, which accounts for about 100 hours of speech each; the TV broadcast ETAPE corpus [28], accounting for about 30 hours of speech; the TV broadcast REPERE train corpus, accounting for about 35 hours of speech and other LIUM radio and TV broadcast data for about 300 hours of speech. As a total, 565 hours of speech composes the training corpus. These recordings were converted to 8kHz before training the acoustic models in order to be more appropriate to the MEDIA telephone data. As inputs, DNN are fed (for training and decoding) with MFCCs (Mel-Frequency Cepstrum Coefficients) concatenated to i -vectors, in order to adapt acoustic models to speakers.

The vocabulary of the ASR system contains all the words present in the MEDIA training and development corpora, so about 2.5K words. A first bigram language model (LM) is applied during the decoding process to generate word-lattices. These lattices are then rescored by applying a 3-gram language model. In order to get an SLU training corpus close to the test corpus, SLU models are trained with ASR transcriptions. To avoid to deal with errors made by an LM over-trained on the MEDIA training corpus, a leave-one-out approach was followed: all the dialogue files in the training and the development corpora were randomly split into 4 subsets. Each subset was transcribed by using an LM trained on the manual transcrip-

tions present in the 3 other blocks and linearly interpolated to a 'generic' language model trained on a large set of French newspaper crawled on the web, containing 77 millions of words. The test data was transcribed with an LM trained on the MEDIA training corpus and the same generic language model. As shown in table 2, word error rates for the training, development, and test corpora were around 23.5%.

Table 2: Word error rates of transcriptions produced by the ASR systems in order to train, to tune and to test the SLU systems.

train.	dev.	test
23.7%	23.4%	23.6%

5.3. Results

Tests were performed for both architectures with the MEDIA DEV set. The best configuration is chosen with respect to the best results observed on the DEV set and applied for obtaining the TEST results. These results in terms of error rate, precision and recall for concepts (C) and concept value (CV) are reported for the best configuration of each architecture in Table 3.

Table 3: Comparison on error rate, precision and recall for concepts (C) and concept value (CV) pairs obtained on TEST.

	C			CV		
	%Error	P	R	%Error	P	R
baseline	23.8	-	-	27.3	-	-
NN EDA	22.3	0.88	0.84	28.8	0.81	0.77
CRF	19.9	0.90	0.85	25.1	0.85	0.80

It appears that the CRF architecture significantly outperforms NN EDA that shown minor improvements with respect to the baseline.

In order to evaluate the impact of the use of confidence measures among the input features, we made some experiments summarized in Table 4. As we can see, the confidence measure provided by the MS-MLP architecture brings relevant information to reduce the CER and the CVER.

Table 4: Impact of the use of confidence measures (probability a posteriori and MS-MLP values) on the performances of CRF on TEST

	without CM		+pap		+pap +MS-MLP	
	C	CV	C	CV	C	CV
CRF	20.9	26.0	20.5	25.7	19.9	25.1

Other versions of the two systems were considered by adding to the usual MEDIA concept labels two more output tags. During training, these tags are replacing the usual one when the hypothesized word is erroneous. If the erroneous hypothesized word is supporting a concept, it is associated to the ERROR-C (for Concept) tag, ERROR-N (for Null) otherwise. During evaluation, ERROR-C and ERROR-N hypothesized tags are replaced by *null* (tag informing that the word does not convey any MEDIA information) in order to perform the usual MEDIA evaluation protocol. Results on TEST, obtained with the best configuration observed on DEV, are reported in Table 5.

Results in Table 5 are similar to those in Table 3, but we can notice some small differences. For instance, precision is now better, even if the CER is not reduced for CRF while it is for NN-EDA. Using these four SLU systems that can be executed in parallel, it is worth trying to see if improvements can be obtained by their combination with weight estimated by optimal performance on the DEV set. The results are reported in Table 6 and compared with the ROVER [29] combination applied to the six SLU systems described in [1].

Table 5: Evaluation results on TEST obtained by adding to the usual MEDIA concept label the ASR error detection tags ERROR-C and ERROR-N.

	C			CV		
	%Error	P	R	%Error	P	R
NN EDA	22.1	0.90	0.82	27.8	0.84	0.77
CRF	20.6	0.91	0.84	25.4	0.86	0.79

Table 6: Performance of the weighted combination or the consensus of outputs from the four systems on TEST.

	C			CV		
	%Err.	P	R	%Err.	P	R
baseline comb.	23.1	-	-	27.0	-	-
CRF+NN comb.	19.3	0.91	0.85	24.5	0.86	0.80
CRF+NN cons.	-	0.96	0.72	-	0.89	0.68

The results show 0.6% and 0.6% absolute reductions for CER and CVER with respect to the best CRF architecture and 4.5% and 2.8% with respect to the baseline. Considering that the best results on manual transcriptions are above 10% on the TEST set, one may conclude that, with the solutions presented in this paper, the contribution of ASR errors to the overall SLU errors is inferior to errors observed for manual transcriptions. A detailed analysis of the errors observed in the automatic and manual transcriptions show a common large error contributions for concepts such as three different types of reference, connectors between domain relevant entities, and proper names that can be values of different attributes. These concepts are expressed by fusible words whose disambiguation requires complex context relations that cannot be automatically characterized (at least with the available amount of train data) by CRFs nor by the type of attention mechanisms used in NN EDA.

Considering the case in which all the four systems provided the same output (consensus) for each word, a 0.96 precision with 0.72 recall were observed on the TEST set. Lack of consensus in the DEV and the TEST sets appears to correspond in most cases to mentions of only few types of concepts. This is a very interesting result since it suggests that further investigation on these particular cases is an important challenge for future work.

6. Conclusions

Two variations of two SLU architectures respectively based on CRFs and NN-EDA have been considered. Using the MEDIA corpus, they were compared with the CRF SLU, considered as baseline that provided the best results among seven different approaches as reported in [1]. The main novelties of the proposed SLU architectures are the use, among others, of semantically relevant confidence and input features. The CRF architectures outperformed the NN-EDA architectures with significant improvement over the baseline. Nevertheless, NN-EDA architectures appeared to be useful when combined with the CRF ones. The results show that the interaction between the ASR and SLU components is beneficial. Furthermore, all the architectures show that most of the errors are for concepts whose mentions are made of short fusible sequences of words that remain ambiguous even if they can be localized. These concept types are difficult to detect, even on manual transcription, indicating that the interpretation of the MEDIA corpus is particularly difficult. Thus, suggested directions for future work should consider new structured mechanisms of attention capable of selecting features of distant contexts in a conversation history. The objective is to identify a sufficient set of context features for disambiguating local concept mentions.

7. References

- [1] S. Hahn, M. Dinarelli, C. Raymond, F. Lefevre, P. Lehnen, R. De Mori, A. Moschitti, H. Ney, and G. Riccardi, "Comparing stochastic approaches to spoken language understanding in multiple languages," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 6, pp. 1569–1583, 2011.
- [2] L. Mangu, E. Brill, and A. Stolcke, "Finding consensus in speech recognition: word error minimization and other applications of confusion networks," *Computer Speech & Language*, vol. 14, no. 4, pp. 373–400, 2000.
- [3] D. Yu, J. Li, and L. Deng, "Calibration of confidence measures in speech recognition," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 8, pp. 2461–2473, 2011.
- [4] R. Zhang, Z. Al Bawad, A. Chan, A. Chotimongkol, D. Huggins-Daines, and A. I. Rudnicky, "Investigations on ensemble based semi-supervised acoustic model training," *Proc. of Eurospeech, Lisbon Portugal*, 2005.
- [5] A. Ogawa and T. Hori, "Asr error detection and recognition rate estimation using deep bidirectional recurrent neural networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4370–4374.
- [6] G. Tur and R. De Mori, *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons, 2011.
- [7] G. Mesnil, Y. Dauphin, K. Yao, Y. Bengio, L. Deng, D. Hakkani-Tur, X. He, L. Heck, G. Tur, D. Yu *et al.*, "Using recurrent neural networks for slot filling in spoken language understanding," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 23, no. 3, pp. 530–539, 2015.
- [8] D. Hakkani-Tür, G. Tur, A. Celikyilmaz, Y.-N. Chen, J. Gao, L. Deng, and Y.-Y. Wang, "Multi-domain joint semantic frame parsing using bi-directional rnn-1stm," in *Proceedings of The 17th Annual Meeting of the International Speech Communication Association*, 2016.
- [9] S. Reddy, O. Täckström, M. Collins, T. Kwiatkowski, D. Das, M. Steedman, and M. Lapata, "Transforming dependency structures to logical forms for semantic parsing," *Transactions of the Association for Computational Linguistics*, vol. 4, pp. 127–140, 2016.
- [10] M. Ma, L. Huang, B. Xiang, and B. Zhou, "Dependency-based convolutional neural networks for sentence embedding," in *The 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. Association for Computational Linguistics, 2015.
- [11] Y.-N. Chen, D. Hakkani-Tür, G. Tur, A. Celikyilmaz, J. Guo, and L. Deng, "Syntax or semantics? knowledge-guided joint semantic frame parsing," in *IEEE Workshop on Spoken Language Technology (SLT 2016)*, San Diego, USA, 2016.
- [12] E. Simionnet, N. Camelin, P. Deleglise, and Y. Esteve, "Exploring the use of attention-based recurrent neural networks for spoken language understanding," in *NIPS*, 2015.
- [13] K. Cho, B. van Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *EMNLP*, 2014.
- [14] S. Ghannay, Y. Esteve, and N. Camelin, "Word embeddings combination and neural networks for robustness in asr error detection," in *Signal Processing Conference (EUSIPCO), 2015 23rd European*. Nice, France: IEEE, 2015, pp. 1671–1675.
- [15] S. Ghannay, Y. Esteve, N. Camelin *et al.*, "Acoustic word embeddings for asr error detection," *Interspeech 2016*, pp. 1330–1334, 2016.
- [16] S. Ghannay, Y. Estève, N. Camelin, C. Dutrey, F. Santiago, and M. Adda-Decker, "Combining continuous word representation and prosodic features for asr error prediction," in *International Conference on Statistical Language and Speech Processing*. Springer, 2015, pp. 84–95.
- [17] S. Ghannay, B. Favre, Y. Esteve, and N. Camelin, "Word embedding evaluation and combination," in *of the Language Resources and Evaluation Conference (LREC 2016), Portoroz (Slovenia)*, 2016, pp. 23–28.
- [18] J. Lafferty, A. McCallum, F. Pereira *et al.*, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Proceedings of the eighteenth international conference on machine learning, ICML*, vol. 1, 2001, pp. 282–289.
- [19] A. Nasr, F. Bchet, and J.-F. Rey, "Macaon : Une chane linguistique pour le traitement de graphes de mots," in *Traitement Automatique des Langues Naturelles - session de demonstrations*, Montréal, 2010.
- [20] V. Vukotic, C. Raymond, and G. Gravier, "Is it time to switch to word embedding and recurrent neural networks for spoken language understanding?" in *Interspeech*, 2015.
- [21] C. T. Hemphill, J. J. Godfrey, G. R. Doddington *et al.*, "The atis spoken language systems pilot corpus," in *Proceedings of the DARPA speech and natural language workshop*, 1990, pp. 96–101.
- [22] T. Lavergne, O. Cappé, and F. Yvon, "Practical very large scale CRFs," in *Proceedings the 48th Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics, July 2010, pp. 504–513. [Online]. Available: <http://www.aclweb.org/anthology/P10-1052>
- [23] H. Bonneau-Maynard, S. Rosset, C. Ayache, A. Kuhn, and D. Mostefa, "Semantic annotation of the french media dialog corpus," in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [24] A. Rousseau, G. Boulianne, P. Deléglise, Y. Estève, V. Gupta, and S. Meignier, "LIUM and CRIM ASR system combination for the REPERE evaluation campaign," in *International Conference on Text, Speech, and Dialogue*. Springer, 2014, pp. 441–448.
- [25] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. EPFL-CONF-192584. IEEE Signal Processing Society, 2011.
- [26] S. Galliano, E. Geoffrois, G. Gravier, J. f. Bonastre, D. Mostefa, and K. Choukri, "Corpus description of the Ester evaluation campaign for the rich transcription of French broadcast news," in *5th international Conference on Language Resources and Evaluation (LREC)*, 2006, pp. 315–320.
- [27] S. Galliano, G. Gravier, and L. Chaubard, "The Ester 2 evaluation campaign for the rich transcription of french radio broadcasts," in *Interspeech*, 2009.
- [28] G. Gravier, G. Adda, N. Paulsson, M. Carré, A. Giraudel, and O. Galibert, "The ETAPE corpus for the evaluation of speech-based TV content processing in the French language," in *Eighth International Conference on Language Resources and Evaluation (LREC)*, Istanbul, Turkey, 2012, pp. 114–118.
- [29] J. G. Fiscus, "A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (rover)," in *Automatic Speech Recognition and Understanding, 1997. Proceedings., 1997 IEEE Workshop on*. IEEE, 1997, pp. 347–354.