



Spoken Language Intent Detection using Confusion2Vec

Prashanth Gurunath Shivakumar*, Mu Yang*, Panayiotis Georgiou

Signal Processing for Communication Understanding and Behavior Analysis Laboratory
University of Southern California, Los Angeles, USA

pgurunat@usc.edu, yangmu@usc.edu, georgiou@sipi.usc.edu

Abstract

Decoding speaker's intent is a crucial part of spoken language understanding (SLU). The presence of noise or errors in the text transcriptions, in real life scenarios make the task more challenging. In this paper, we address the spoken language intent detection under noisy conditions imposed by automatic speech recognition (ASR) systems. We propose to employ confusion2vec word feature representation to compensate for the errors made by ASR and to increase the robustness of the SLU system. The confusion2vec, motivated from human speech production and perception, models acoustic relationships between words in addition to the semantic and syntactic relations of words in human language. We hypothesize that ASR often makes errors relating to acoustically similar words, and the confusion2vec with inherent model of acoustic relationships between words is able to compensate for the errors. We demonstrate through experiments on the ATIS benchmark dataset, the robustness of the proposed model to achieve state-of-the-art results under noisy ASR conditions. Our system reduces classification error rate (CER) by 20.84% and improves robustness by 37.48% (lower CER degradation) relative to the previous state-of-the-art going from clean to noisy transcripts. Improvements are also demonstrated when training the intent detection models on noisy transcripts.

Index Terms: spoken language understanding, intent detection, confusion2vec

1. Introduction

Spoken Language Understanding (SLU) systems aim at extracting semantic information from human spoken utterances. Such systems play a significant role in practical applications like personal AI voice assistants (e.g. Alexa, Siri, etc.), phone-call routing, booking system and so on. A SLU system is typically modeled as two separate components: an ASR front-end, which translates acoustic signal into text, followed by a Natural Language Understanding (NLU) module that performs inference for downstream tasks. Typical tasks include Domain classification, Intent Detection and Slot filling. In this work, we focus on the SLU system that performs Intent Detection, a task identifying speaker's intent from speech. Such task is usually treated as an utterance classification problem [1].

1.1. Prior Work

In the light of success of Deep Learning techniques, applying Deep Neural Networks on intent detection has been shown to be effective, often outperforming conventional classifiers, such as Support Vector Machines [2]. In recent years, the NLU community have applied various techniques to improve intent detection performance on manual transcripts. [3–5] jointly model

intent detection with slot filling, simplifying the NLU task by a unified model. [6, 7] extend the joint modeling with domain knowledge, which enables information from multiple tasks to benefit the individual tasks and allow the NLU model to be applied to multiple-domain tasks. Going one step further, [8, 9] involve adapting domain-specific language model (LM) while performing intent detection and slot filling, improving the performance on both LM and language understanding task. [10] explores strategies in joint modeling intent classification and slot filling using explicit alignment information provided by slot filling using attention-based encoder-decoder structure. On the basis of attention-based model, [11] connects context information from intent detection with slot filling using a gate mechanism. [12] employs a similar intent-augmented gating mechanism to guide the learning of the slot filling task. It further incorporates character-level embedding along with word-level embedding achieving state-of-the-art results in intent detection.

However, suffering from ASR front-end errors, such as misrecognized words, insertions and deletions, the performance of such systems degrades significantly, as shown in [13–15] and is still the bottleneck in SLU systems. On one hand, in order to make system more ASR-robust, ASR hypotheses can be incorporated into the model's training corpus. [16–18] exploit Word Confusion Networks to efficiently connect NLU models with ASR hypotheses. [19] simulated ASR errors by randomly substituting words with their linguistically and acoustically similar candidates. On the other hand, there have been works that aim to jointly perform NLU tasks and ASR error adaptation. [20, 21] employ Recurrent Neural Network (RNN) based Encoder-Decoder structure to reconstruct correct utterances from ASR hypotheses while performing intent detection and slot filling. [22] makes richer feature representations by adding acoustic pitch accent flags into word embedding.

In this paper, we specifically target the task of spoken language intent detection on noisy ASR transcripts. In contrast to the majority of the works which mostly deal with the innovation of classification models [8, 10–12], in our study, we concentrate on robust word feature representations. We propose to employ the confusion2vec [23] word vector representation to compensate for the errors made by an ASR and to provide enhanced and robust performance for the task of spoken language intent detection. Confusion2vec captures acoustic similarity information of words in addition to the semantic-syntactic relations and is trained in a completely unsupervised manner on ASR lattices decoded on an out-of-domain corpora. Moreover, unlike the studies which adapt the ASR to the target datasets and tasks [8, 20], we treat the ASR as a generic independent module, but contribute towards bridging the gap between the ASR and the NLU model. We demonstrate with our experiments, on the benchmark ATIS dataset [24], the vital role of the confusion2vec to the robustness of the intent classification.

The rest of the paper is structured as follows. In section 2,

*The authors contributed equally to this work

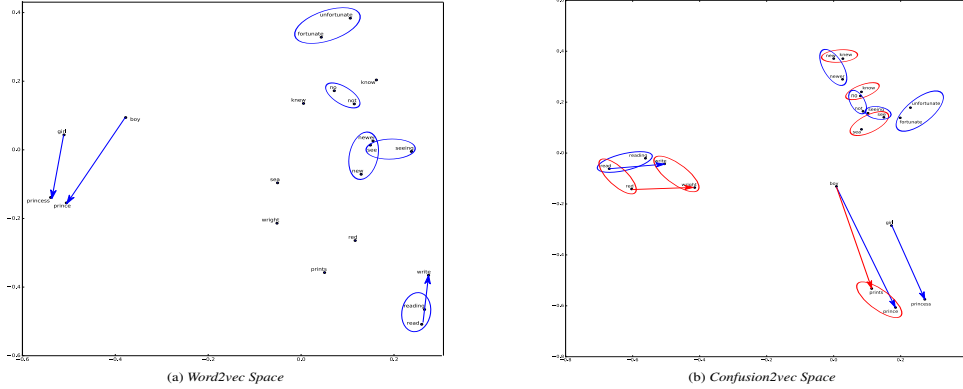


Figure 1: 2D Vector space illustration after PCA dimension reduction for Word2vec and Confusion2vec: The blue ellipses indicate syntactic word relations. The red ellipses indicate acoustic similarity relations. The blue arrows illustrate the semantic relationships. The red arrows illustrate the interaction of acoustic similarity with semantic relationships. The word2vec space is rich in semantic and syntactic word relations, however no trivial acoustic similarity is evident. The confusion2vec space preserves the semantic and syntactic word relations, moreover captures additional acoustic similarity information.

we present the proposed methodology, provide brief description of the confusion2vec word embedding and the intent classification model. Section 3 describes the databases employed, our experimental setup and the baseline systems. In section 4, we present and discuss the experimental results. Finally, section 5 concludes the study and discusses the future work.

2. Proposed Technique

In this section, we first describe the confusion2vec word vector representation for the task of spoken language intent detection and then introduce the RNN intent classification model.

2.1. Confusion2vec Word Embedding

The role of word vector representations is crucial for NLP [25]. Efficient and information rich word embeddings like word2vec [26], glove [27] are shown to capture semantics and syntactics of the language. Using such efficient word representations have proven to be beneficial in the NLU tasks like named entity detection [28], intent detection [29]. The SLU tasks like intent detection [30–33], slot-filling [34], spoken dialogue systems [35] have also benefited from using information rich word embeddings. However, they are less than optimal in the cases of erroneous transcriptions, for example ASR transcriptions [8, 20], since the errors corrupt the semantic-syntactic space over local context of occurrence and thereby introduce noise in the model.

In this work, we propose to employ recently proposed confusion2vec word vector representation [23] for the task of intent detection to counter for errors present in the spoken transcriptions. Motivated from human speech production and perception, the confusion2vec models the acoustic relations of words in addition to the semantic and syntactic relationships of words [23]. The confusion2vec uses unsupervised training techniques similar to skip-gram of word2vec, but operates on lattice-like structures or confusion networks output by the ASR. Since the confusion networks of a typical ASR exhibits confusions between words on two principle axes (i) contextual, and (ii) acoustic similarity, the confusion2vec is devised to operate on both the axes, thereby modeling local context information (like word2vec) as well as acoustic similarity information. Figure 1 illustrates the 2-dimensional word vector space for word2vec and confusion2vec after dimension reduction using principal component analysis (PCA). From the figure (and from extensive

analysis done in [23]), it is evident that confusion2vec space captures acoustic similarity between words without compromising the information captured by the word2vec. Complex meaningful, useful interactions between the acoustic subspace and the semantic-syntactic subspaces are also observed. For more information we would like to point the interested readers towards [23], which in detail describes the confusion2vec.

In application to the spoken language intent detection task, the nature of ASR errors are often acoustically related. Confusion2vec incorporates real, unsupervised, ASR output as its training corpus, thus the feature representation incorporates confusions (errors) nearby in its embedding space. In other words, we hypothesize that the embedded acoustic similarity information in confusion2vec limits the impact of errors made by the ASR, and thus allows subsequent NLP tasks to be minimally affected. We expect the following from our model with respect to the intent detection task: (i) to be less affected from ASR errors and thus achieve better performance in the case of noisy ASR transcriptions, and (ii) to be at-least on par with word2vec under clean conditions.

2.2. Intent Classification Model

Since the contribution of this work is towards word feature representations, we employ a fairly simple recurrent neural network model for the classification task. However, we believe the contributions on feature representations are orthogonal to the classification model and thus expect even better performance for more complex models like in [8, 10–12]. In this work, we use Bi-directional Long Short-Term Memory (LSTM) units, as shown in Figure 2. Given an input utterance w_0, w_1, \dots, w_T , each word in the input sequence is mapped to its word vector representation x_0, x_1, \dots, x_T by embedding look up. We formulate the model outputs as

$$\vec{h}_t = \overrightarrow{LSTM}(\vec{h}_{t-1}, x_t; \vec{\Theta}) \quad \overleftarrow{h}_t = \overleftarrow{LSTM}(\overleftarrow{h}_{t+1}, x_t; \overleftarrow{\Theta}) \quad (1)$$

$$\hat{P}_{\text{intent}} = \text{Softmax}(W [\vec{h}_T, \overleftarrow{h}_0] + b) \quad (2)$$

where h_t is the LSTM output of each direction at each time step t , Θ is the parameter of the LSTM. We feed the concatenation of two directional LSTM outputs at the last time step into the linear output layer (with weights W and bias b) which projects it into the intent label space. Finally, the intent label is predicted from the softmax-normalized probability distribution over all intents.

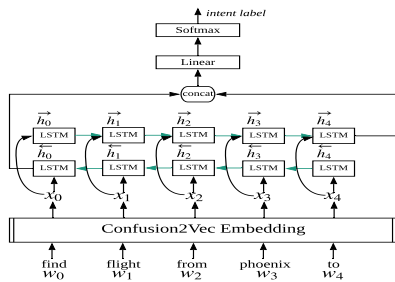


Figure 2: *Intent Classification RNN Model*

3. Database & Experimental Setup

3.1. Database

The ASR is trained on the Fisher English Training (LDC2004S13 and LDC2005S13) Speech corpora [36]. The confusion2vec is trained on the output of the ASR, i.e., the confusion networks generated via Fisher English Corpora. The database setup for the ASR and confusion2vec is identical and explained in detail in [23].

We trained the intent detection model on ATIS (Airline Travel Information Systems) dataset [24], which comes with audio recordings and corresponding manual transcripts about humans asking for flight information. Following [6, 11], we apply the same train, development, test split setup. The setup contains 4478, 500 and 893 intent-labeled reference utterances in train, development and test set respectively. The database split and the corresponding audio mappings are made available to public¹. In order to evaluate our model’s robustness to ASR outputs, we also construct our ASR output set by decoding the corresponding audio recording for each of the data splits using the ASR. In cases where an utterance is labeled with multiple intent labels, the top intent was selected as the true label, yielding 18 intents in total.

3.2. Experimental Setup

The training setup for the ASR and the confusion2vec is identical to our previous work [23]. For decoding the ATIS dataset, through our ASR, the audio samples were down-sampled from 16kHz to 8kHz. The ASR achieves a WER rate of 18.54% on the ATIS test set. We choose the confusion2vec model yielding the best performance in [23], i.e., independently trained C2V-1 and C2V-c models are concatenated and jointly optimized with intra-Confusion2vec scheme (556 dimensions).

For intent detection, we train models on the 4478 utterances in training set, and tune hyper-parameters based on the classification accuracy on the 500 reference utterances in development set. The model with the best performance on the development set is chosen and evaluated on both reference test set and ASR test set. The hyper-parameter space we experimented with is as follows: Batch size is set to 1, i.e. each sentence is viewed as an independent sample. Hidden dimension of LSTM unit is tuned over {256, 128, 64, 32}, and dropout is tuned over {0.1, 0.2, 0.25}. We select Adam optimizer, with learning rate set to be among {0.001, 0.0005}. The maximum number of epochs is set to 50 with early-stopping strategy.

3.3. Baseline Systems

The first set of baselines compare different conventional word embeddings. They include: (i) random initialization (556 di-

mensions) sampled from a uniform distribution, (ii) vanilla GloVe² (300 dimensions) as in [27], and (iii) skip-gram Word2Vec³ (556 dimensions) fine-tuned on Fisher English corpus reference transcripts (for fair comparison with confusion2vec). We also tried the vanilla Google word2vec³. However, the performance was found to be consistently lower than the fine-tuned version, thus, we don’t include it in the comparisons. Note, only the randomly initialized word embedding is trainable, while all other embeddings are fixed throughout the training. All the above baselines use identical RNN architecture for intent classification as described in section 2.2.

The second set of baselines compare our proposed model with the recent state-of-the-art models, including: (i) a joint intent detection, slot filling & LM model [8], (ii) an attention-based joint model that incorporates alignment information provided by slot filling task [10], and (iii) an intent-augmented gating mechanism based model which further incorporates character-level embedding along with word-level embedding [12]. The baselines are trained under our experiment setup using the same hyper-parameters reported in their original papers. We also reproduce the result of each model under their original experiment settings⁴, and report the obtained scores in parentheses for reference. For the adapted model trained on ASR, we consider a joint intent slot filling and intent detection model which performs sentence reconstruction from ASR hypotheses [20] as the baseline, and report the scores on ASR outputs and corresponding ASR WER from original paper.

4. Results and Discussion

4.1. Training on Reference Clean Transcripts

Table 1 and Figure 3 illustrate the results obtained training on clean transcripts. First, we compare the results between different word feature embedding (refer to upper half of Table 1). On clean reference transcripts, GloVe embeddings provides the best performance (as found in [32]). Both word2vec and random initialization provide identical results. The proposed confusion2vec gives considerably lower CER compared to the Word2Vec and random initialization. Although GloVe outperforms confusion2vec, we believe the comparison of confusion2vec is more fair with word2vec, since both use skip-gram modeling. With the proposed Confusion2vec system, we don’t expect improvements on clean reference transcripts, since the acoustic similarity/confusion is less relevant. As expected, we observe no degradation in performance with confusion2vec and is on par with the popular, leading word vector representations for the task of intent detection on clean transcripts.

On noisy ASR transcripts, we see an increase in CER with all models. Although, random initialization performed identical to Word2Vec on clean transcripts, we see Word2Vec performs relatively better on ASR transcripts. This observation confirms that better word feature representations exhibit higher robustness to errors. Similar trend is apparent with GloVe embeddings in comparison with random initialization, although we observe slightly higher CER and degradation (between clean and noisy transcripts) compared to word2vec. The proposed confusion2vec gives the least CER among all the models (a relative improvement of 20.84% over word2vec, 21.9% over GloVe and 40.65% over random initialization). Moreover, C2V displays higher robustness going from clean to noisy ASR transcripts (degradation, Δ_{diff} , is minimal). A relative improve-

¹https://github.com/pgurunath/slu_confusion2vec

²<https://nlp.stanford.edu/projects/glove/>

³<https://code.google.com/archive/p/word2vec>

Table 1: Results with Training on Reference: Classification Error Rates (CER) for Reference and ASR Transcripts: Δ_{diff} is the absolute degradation of model from clean to ASR. The numbers inside parenthesis indicate CER obtained reproducing the results under their original experiment settings⁴.

Model	Reference	ASR	Δ_{diff}
Random	2.69	10.75	8.06
GloVe [27]	1.90	8.17	6.27
Word2Vec [26]	2.69	8.06	5.37
C2V (proposed)	2.46	6.38	3.92
Liu and Lane [8]	1.90 (1.57)	9.41 (8.29) ⁵	7.51 (6.72)
Liu and Lane [10]	1.79 (1.90)	8.06 (8.40)	6.27 (6.50)
Li et al. [12]	2.02 (1.34)	9.18 (9.07)	7.16 (7.73)

ment in robustness of 37.48%, 27% and 51.36% compared to GloVe, Word2vec and random initialization respectively is observed with C2V (in terms of Δ_{diff}). The confusion2vec word feature representation is able to use the embedded acoustic similarity information to recover from errors resulting from acoustically confusable words in the ASR output transcriptions.

Further, we compare our proposed system with the recent state-of-the-art works on SLU (see bottom part of Table 1). Note, the recent works employ much more complex modeling techniques compared to ours. Thus, as expected the recent works outperform our simple RNN architecture testing on clean transcriptions. However, on noisy ASR transcriptions, even with a much simpler model, the proposed confusion2vec achieves significantly lower CER (a relative improvement of at least 20.84%) compared to state-of-the-art models. Moreover, again, the degradation with confusion2vec is the least among all the models, a relative 37.48% lesser degradation compared to the recent works. The results highlight the potent robustness of the confusion2vec word feature representation. In addition, we believe that the gains from the complex classification modeling are orthogonal to gains from confusion2vec word feature representations and thus should result in additional gains incorporating more complex models with confusion2vec.

4.2. Training on ASR

Further, we also perform additional experiments by training the intent classification models on noisy ASR transcripts. A more robust feature representation should theoretically help in reducing the noise in the model translating to better performance. From Table 2, it is evident that all the models improve with the matched noisy train and test conditions. The proposed confusion2vec based model gives the least CER among all models. The confusion2vec feature representation is better able to explain the (acoustic) errors and doing so reduces confusion and noise in the intent classification model, thereby resulting in a better and robust performance. Moreover, comparing it with




Figure 3: Comparison of CER for different systems

⁴Original settings of [8, 10, 12], make use of train + dev data for training. They pre-process data by substituting the digits with a token.

Table 2: Results with Training and Testing on ASR transcripts.

Model	WER %	CER %
Random	18.54	5.15
GloVe [27]	18.54	6.94
Word2Vec [26]	18.54	5.49
C2V (proposed)	18.54	4.70
Schumann and Angkitittrakul [20]	10.55	5.04 ⁵

the recent study by Schumann and Angkitittrakul [20], although the results are not directly comparable due to differences in the WER of the ASRs, our proposed method achieves a lower CER in spite of much worse WER conditions⁵. This suggests that explicitly modeling in-domain ASR errors as in [20] is of lesser effect compared to modeling the general acoustic signatures between words in a language as in the case with confusion2vec.

Finally, comparing the results from Table 1 and Table 2, it is encouraging to see that the proposed confusion2vec model trained on clean transcripts is able to inherit enough robustness to achieve lower (than GloVe) and comparable CER to the models (word2vec) trained on ASR, possibly reducing the need for adaptation on ASR and realizing more generalizable systems.

5. Conclusion and Future Work

In this paper, we proposed an intent detection model based on confusion2vec word vector representation targeting noisy ASR transcriptions. The proposed word embeddings significantly outperform the popular leading word vector representations like word2vec and GloVe in the cases of noisy ASR output. Comparisons are made with various recent state-of-the-art studies, and we find the proposed method improves over them by a considerable margin despite using relatively simple RNN architectures for classification. The robustness of confusion2vec also extends to models trained on noisy ASR, achieving the least CER among the conventional word embedding as well as the recent studies. Encouraging results suggest confusion2vec robustness to errors eliminates the need for adapting the intent classification models on noisy ASR outputs.

In future, we plan to apply and evaluate the proposed confusion2vec on additional SLU tasks like slot-filling, domain classification and named entity recognition. We believe the proposed model should provide similar advantages, especially under noisy conditions. Addressing multiple SLU tasks also allows us to use more complex joint-modelling systems with confusion2vec. The better, more complex, models should provide improvements orthogonal to confusion2vec feature representations, and we thus expect to see further improvements. We also plan to conduct more in-depth analysis on how the signal conditions and ASR performance affect each model; we expect confusion2vec to provide more gains as the ASR performance deteriorates. Representation of multiple path outputs from the ASR with confusion2vec instead of only the best path is also a possible future direction.

6. Acknowledgements

This research was supported in part by NSF, NIH, and DOD.

⁵We don't domain-constrain, optimize or rescore our ASR, as in [8, 20]. We treat ASR as an independent module for fair comparison with other models and for domain-generalization and portability of our system and conclusions.

7. References

- [1] S. Yaman, L. Deng, D. Yu, Y.-Y. Wang, and A. Acero, "An integrative and discriminative technique for spoken utterance classification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 6, pp. 1207–1214, 2008.
- [2] P. Haffner, G. Tur, and J. H. Wright, "Optimizing svms for complex call classification," in *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP'03).*, vol. 1. IEEE, 2003, pp. I–I.
- [3] P. Xu and R. Sarikaya, "Convolutional neural network based triangular crf for joint intent detection and slot filling," in *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*. IEEE, 2013, pp. 78–83.
- [4] D. Guo, G. Tur, W.-t. Yih, and G. Zweig, "Joint semantic utterance classification and slot filling with recursive neural networks," in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 554–559.
- [5] X. Zhang and H. Wang, "A joint model of intent determination and slot filling for spoken language understanding," in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, ser. IJCAI'16. AAAI Press, 2016, pp. 2993–2999.
- [6] 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.
- [7] Y.-B. Kim, S. Lee, and K. Stratos, "Onenet: Joint domain, intent, slot prediction for spoken language understanding," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 547–553.
- [8] B. Liu and I. Lane, "Joint online spoken language understanding and language modeling with recurrent neural networks," *arXiv preprint arXiv:1609.01462*, 2016.
- [9] H. Zhang, S. Zhu, S. Fan, and K. Yu, "Joint spoken language understanding and domain adaptive language modeling," in *International Conference on Intelligent Science and Big Data Engineering*. Springer, 2018, pp. 311–324.
- [10] B. Liu and I. Lane, "Attention-based recurrent neural network models for joint intent detection and slot filling," in *Interspeech 2016*, 2016, pp. 685–689.
- [11] C.-W. Goo, G. Gao, Y.-K. Hsu, C.-L. Huo, T.-C. Chen, K.-W. Hsu, and Y.-N. Chen, "Slot-gated modeling for joint slot filling and intent prediction," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, vol. 2, 2018, pp. 753–757.
- [12] C. Li, L. Li, and J. Qi, "A self-attentive model with gate mechanism for spoken language understanding," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 3824–3833.
- [13] Y. He and S. Young, "A data-driven spoken language understanding system," in *2003 IEEE Workshop on Automatic Speech Recognition and Understanding (IEEE Cat. No. 03EX721)*. IEEE, 2003, pp. 583–588.
- [14] A. Deoras and R. Sarikaya, "Deep belief network based semantic taggers for spoken language understanding," in *Interspeech*, 2013, pp. 2713–2717.
- [15] G. Mesnil, Y. Dauphin, K. Yao, Y. Bengio, L. Deng, D. Hakkani-Tür, 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*, vol. 23, no. 3, pp. 530–539, 2015.
- [16] G. Kurata, N. Itoh, M. Nishimura, A. Sethy, and B. Ramabhadran, "Leveraging word confusion networks for named entity modeling and detection from conversational telephone speech," *Speech Communication*, vol. 54, no. 3, pp. 491–502, 2012.
- [17] G. Tür, A. Deoras, and D. Hakkani-Tür, "Semantic parsing using word confusion networks with conditional random fields," in *INTERSPEECH*. Citeseer, 2013, pp. 2579–2583.
- [18] R. Masumura, Y. Ijima, T. Asami, H. Masataki, and R. Higashinaka, "Neural confnet classification: Fully neural network based spoken utterance classification using word confusion networks," *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6039–6043, 2018.
- [19] E. Simonnet, S. Ghannay, N. Camelin, and Y. Estève, "Simulating asr errors for training slu systems," in *Proceedings of the 11th Language Resources and Evaluation Conference*. Miyazaki, Japan: European Language Resource Association, May 2018.
- [20] R. Schumann and P. Angkitittrakul, "Incorporating asr errors with attention-based, jointly trained rnn for intent detection and slot filling," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6059–6063.
- [21] S. Zhu, O. Lan, and K. Yu, "Robust spoken language understanding with unsupervised asr-error adaptation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6179–6183.
- [22] S. Stehwien and N. T. Vu, "First step towards enhancing word embeddings with pitch accent features for dnn-based slot filling on recognized text," in *Konferenz Elektronische Sprachsignalverarbeitung 2017, Saarbrücken*, 2017.
- [23] P. G. Shivakumar and P. G. Georgiou, "Confusion2vec: Towards enriching vector space word representations with representational ambiguities," *CoRR*, vol. abs/1811.03199, 2018.
- [24] C. T. Hemphill, J. J. Godfrey, and G. R. Doddington, "The atis spoken language systems pilot corpus," in *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*, 1990.
- [25] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuska, "Natural language processing (almost) from scratch," *Journal of machine learning research*, vol. 12, no. Aug, pp. 2493–2537, 2011.
- [26] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [27] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [28] S. K. Sienčnik, "Adapting word2vec to named entity recognition," in *Proceedings of the 20th nordic conference of computational linguistics, nodalida 2015, may 11-13, vilnius, lithuania*, no. 109. Linköping University Electronic Press, 2015, pp. 239–243.
- [29] H. B. Hashemi, A. Asiaee, and R. Kraft, "Query intent detection using convolutional neural networks," in *International Conference on Web Search and Data Mining, Workshop on Query Understanding*, 2016.
- [30] J.-K. Kim, G. Tur, A. Celikyilmaz, B. Cao, and Y.-Y. Wang, "Intent detection using semantically enriched word embeddings," in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 414–419.
- [31] Y. Luan, S. Watanabe, and B. Harsham, "Efficient learning for spoken language understanding tasks with word embedding based pre-training," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [32] M. Firdaus, S. Bhatnagar, A. Ekbal, and P. Bhattacharyya, "Intent detection for spoken language understanding using a deep ensemble model," in *Pacific Rim International Conference on Artificial Intelligence*. Springer, 2018, pp. 629–642.
- [33] K. Balodis and D. Dekšne, "Intent detection system based on word embeddings," in *International Conference on Artificial Intelligence: Methodology, Systems, and Applications*. Springer, 2018, pp. 25–35.
- [34] M. Firdaus, S. Bhatnagar, A. Ekbal, and P. Bhattacharyya, "A deep learning based multi-task ensemble model for intent detection and slot filling in spoken language understanding," in *International Conference on Neural Information Processing*. Springer, 2018, pp. 647–658.
- [35] E. Ferreira, B. Jabaian, and F. Lefèvre, "Zero-shot semantic parser for spoken language understanding," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [36] C. Cieri, D. Miller, and K. Walker, "The fisher corpus: a resource for the next generations of speech-to-text," in *LREC*, vol. 4, 2004, pp. 69–71.