

LEARNED IN SPEECH RECOGNITION: CONTEXTUAL ACOUSTIC WORD EMBEDDINGS

Shruti Palaskar, Vikas Raunak* and Florian Metze*

Carnegie Mellon University, Pittsburgh, PA, U.S.A.
{spalaska | vraunak | fmetze }@cs.cmu.edu

ABSTRACT

End-to-end acoustic-to-word speech recognition models have recently gained popularity because they are easy to train, scale well to large amounts of training data, and do not require a lexicon. In addition, word models may also be easier to integrate with downstream tasks such as spoken language understanding, because inference (search) is much simplified compared to phoneme, character or any other sort of sub-word units. In this paper, we describe methods to construct contextual acoustic word embeddings directly from a supervised sequence-to-sequence acoustic-to-word speech recognition model using the learned attention distribution. On a suite of 16 standard sentence evaluation tasks, our embeddings show competitive performance against a word2vec model trained on the speech transcriptions. In addition, we evaluate these embeddings on a spoken language understanding task, and observe that our embeddings match the performance of text-based embeddings in a pipeline of first performing speech recognition and then constructing word embeddings from transcriptions.

Index Terms— acoustic word embeddings, contextual embeddings, attention, acoustic-to-word speech recognition

1. INTRODUCTION

The task of learning fixed-size representations for variable length data like words or sentences, either text or speech-based, is an interesting problem and a focus of much current research. In the natural language processing community, methods like word2vec [1], GLoVe [2], CoVe [3] and ELMo [4] have become increasingly popular, due to their utility in several natural language processing tasks. Similar research has progressed in the speech recognition community, where however the input is a sequence of short-term audio features, rather than words or characters. Therefore, the variability in speakers, acoustics or microphones for different occurrences of the same word or sentence adds to the challenge.

Prior work towards the problem of learning word representations from variable length acoustic frames involved either providing word boundaries to align speech and text [5], or chunking (“chopping” or “padding”) input speech into fixed-length segments that usually span only one word [6, 7, 8, 9]. Since these techniques learn acoustic word embeddings from audio fragment and word pairs obtained via a given segmentation of the audio data, they ignore the specific audio context associated with a particular word. So the resulting word embeddings do not capture the contextual dependencies in speech. In contrast, our work constructs individual acoustic word embeddings grounded in utterance-level acoustics.

In this paper, we present different methods of obtaining acoustic word embeddings from an attention-based sequence-to-sequence

model [10, 11, 12] trained for direct Acoustic-to-Word (A2W) speech recognition [13]. Using this model, we jointly learn to automatically segment and classify input speech into individual words, hence getting rid of the problem of chunking or requiring pre-defined word boundaries. As our A2W model is trained at the utterance level, we show that we can not only learn acoustic word embeddings, but also learn them in the proper context of their containing sentence. We also evaluate our contextual acoustic word embeddings on a spoken language understanding task, demonstrating that they can be useful in non-transcription downstream tasks.

Our main contributions in this paper are the following:

1. We demonstrate the usability of attention not only for aligning words to acoustic frames without any forced alignment but also for constructing Contextual Acoustic Word Embeddings (CAWE).
2. We demonstrate that our methods to construct word representations (CAWE) directly from a speech recognition model are highly competitive with the text-based word2vec embeddings [1], as evaluated on 16 standard sentence evaluation benchmarks.
3. We demonstrate the utility of CAWE on a speech-based downstream task of Spoken Language Understanding showing that pre-trained speech models could be used for transfer learning similar to VGG in vision [14] or CoVe in natural language understanding [3].

2. RELATED WORK

A2W modeling has been largely pursued using Connectionist Temporal Classification (CTC) models [15] and Sequence-to-Sequence (S2S) models [10]. Prior work shows the need for large amounts of training data for these models (thousands of hours of speech) with large word vocabularies of frequently occurring words [16, 17, 18, 19, 20]. Progress in the field showed the possibility of training these models with smaller amount of data (300 hours Switchboard corpus [21]) but restricting the vocabulary to words occurring at least 5 or 10 times [22, 23]. The solutions to generate out-of-vocabulary words have revolved around backing off to smaller units like characters or sub-words [22, 17, 18, 23, 20]. While this solves the problem of rare word generation, the models are no longer pure-word models.

[24] present one of the first S2S models for pure-word large vocabulary A2W recognition with the 300 hour Switchboard corpus with a vocabulary of about 30,000 words. [25, 13] build upon their work and improve the training of these models for the large vocabulary task. [13] is one of our previous works where we show that the direct A2W model is also able to automatically learn word boundaries without any supervision and is the current best pure-word S2S model. We use the same model in this work and expand it towards learning acoustic embeddings.

[5, 6, 8, 7, 26, 9] all explore ways to learn acoustic word embeddings. All above methods except [7] use unsupervised learning based methods to obtain these embeddings where they do not use the transcripts or do not perform speech recognition. [7] use a su-

*Equal contribution

pervised Convolutional Neural Network based speech recognition model but with short speech frames as input that usually correspond to a single word. This is the common practice in most prior work that simplifies training but prevents the models to scale to learn contextual word embeddings grounded in utterance level acoustics. [5] propose an unsupervised method to learn speech embeddings using a fixed context of words in the past and future. The drawbacks of their method are the fixed context and need for forced alignment between speech and words for training.

Learning text-based word embeddings is also a rich area of research with well established techniques such as [1, 2]. Research has further progressed into learning contextualized word embeddings [3, 4] that are useful in many text-based downstream tasks [27]. [3] learns contextual word embeddings from a fully trained machine translation model and depict re-use of their encoder in other downstream tasks. Our work ties A2W speech recognition model with learning contextual word embeddings from speech.

3. ACOUSTIC-TO-WORD RECOGNITION

Our S2S model is similar in structure to the Listen, Attend and Spell model [11] which consists of 3 components: the encoder network, a decoder network and an attention model. The encoder maps the input acoustic features vectors $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_T)$ where $\mathbf{a}_i \in \mathcal{R}^d$, into a sequence of higher-level features $\mathbf{h} = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{T'})$. The encoder is a pyramidal (sub-sampling) multi-layer bi-directional Long Short Term Memory (BLSTM) network. The decoder network is also an LSTM network that learns to model the output distribution over the next target conditioned on sequence of previous predictions i.e. $P(\mathbf{y}_t | \mathbf{y}_{t-1}^*, \mathbf{y}_{t-2}^*, \dots, \mathbf{y}_0^*, \mathbf{x})$ where $\mathbf{y}^* = (\mathbf{y}_0^*, \mathbf{y}_1^*, \dots, \mathbf{y}_{L+1}^*)$ is the ground-truth label sequence. In this work, $\mathbf{y}_i^* \in \mathcal{U}$ is from a word vocabulary. This decoder generates targets \mathbf{y} from \mathbf{h} using an attention mechanism.

We use the location-aware attention mechanism [12] that enforces monotonicity in the alignments by applying a convolution across time to the attention of previous time step. This convolved attention feature is used for calculating the attention for the current time step which leads to a peaky distribution [12, 13]. Our model follows the same experimental setup and model hyper-parameters as the word-based models described in our previous work [13] with the difference of learning 300 dimensional acoustic feature vectors instead of 320 dimensional.

4. CONTEXTUAL ACOUSTIC WORD EMBEDDINGS

We now describe our method to obtain the acoustic word embeddings from the end-to-end trained speech recognition system described in Section 3. The model is as shown in Figure 1 where the embeddings are constructed using the hidden representations obtained from the encoder and the attention weights from the decoder. Our method of constructing “contextual” acoustic word embeddings is similar to a method proposed for text embeddings, CoVe [3]. The main challenge that separates our method from CoVe [3] in learning embeddings from a supervised task, is the problem of alignment between input speech and output words. We use the location-aware attention mechanism that has the property to assign higher probability to certain frames leading to a peaky attention distribution. We exploit this property of location-aware attention in an A2W model to automatically segment continuous speech into words as shown in our previous work [13], and then use this segmentation to obtain word embeddings. In the next two subsections, we formalize this process




Fig. 1. A2W model with the CAWE representations obtained by combining the encoders representations and attention weights.

of constructing contextual acoustic word embeddings. Intuitively, attention weights on the acoustic frames hidden representations reflect their importance in classifying a particular word. They thereby provide a correspondence between the frame and the word within a given acoustic context. We can thus construct word representations by weighing the hidden representations of these acoustic frames in terms of their importance to the word i.e. the attention weight. We show this in the Figure 1 wherein the hidden representations and their attention weights are colored according to their correspondence with a particular word.

4.1. Automatic Segmentation of Speech into Words

Given that a_j represents the acoustic frame j , let $encoder(a_j)$ represent the higher-level features obtained for the frame a_j (i.e. $encoder(a_j) = \mathbf{h} = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{T'})$, as explained in Section 3). Then, for the i^{th} word w_i our model first obtains the mappings of w_i to acoustic frames a_K where K is the set such that $\forall k \in K$

$$k = \arg \max_j (attention(a_j))$$

over all utterances U containing the word w_i in the training set.

4.2. Investigating the Utility of Attention Distribution for Constructing Embeddings: CAWE

Below we describe three different ways of using attention to obtain acoustic word embeddings for a word w_i (here, $n(K)$ represents the cardinality of the set K):

$$w_i = \frac{\sum_{k \in K} encoder(a_k)}{n(K)} \quad (1)$$

$$w_i = \frac{\sum_{k \in K} attention(a_k) \cdot encoder(a_k)}{n(K)} \quad (2)$$

$$w_i = encoder(a_k) \text{ where } k = \arg \max_{k \in K} attention(a_k) \quad (3)$$

Therefore, unweighted Average (U-AVG, Equation 1) is just the unweighted combination of all the hidden representations of acoustic frames mapped to a particular word. Attention weighted Average (CAWE-W, Equation 2) is the weighted average of the hidden representations of all acoustic frames using the attention weights for a given word. Finally, maximum attention (CAWE-M, Equation 3) is the hidden representation of the acoustic frame with the highest attention score for a given word across all utterances in the training

Table 1. Comparing three methods to obtain acoustic word embeddings from an A2W model: unweighted average (U-AVG), weighted average (CAWE-W) and maximum attention (CAWE-M).

Dataset	Switchboard			How2		
	U-AVG	CAWE-W	CAWE-M	U-AVG	CAWE-W	CAWE-M
STS 2012	0.3230	0.3281	0.3561	0.3255	0.3271	0.3648
STS 2013	0.1252	0.1344	0.1969	0.2070	0.2071	0.2716
STS 2014	0.3358	0.3389	0.3888	0.3375	0.3426	0.3940
STS 2015	0.3854	0.3881	0.4275	0.3852	0.3843	0.4173
STS 2016	0.2998	0.2974	0.3833	0.3248	0.3271	0.3159
STS B	0.3667	0.3510	0.4010	0.3343	0.3440	0.4000
SICK-R	0.5640	0.5800	0.6006	0.5800	0.6060	0.6440
MR	63.86	63.75	64.69	63.46	63.19	63.64
MRPC	70.67	69.45	69.80	68.29	67.83	70.61
CR	71.42	72.13	72.93	74.12	73.99	73.03
SUBJ	82.45	82.22	81.19	81.48	81.88	81.01
MPQA	73.76	73.28	73.75	74.21	74.18	73.53
SST	66.45	66.61	65.02	63.43	63.43	65.13
SST-FG	32.81	32.04	33.53	31.95	32.35	32.03
TREC	63.80	62.40	67.60	66.60	66.00	60.60
SICK-E	74.20	73.41	74.06	75.14	75.34	75.97

data. We call the attention-weighted average and the maximum attention based techniques as Contextual Acoustic Word Embeddings (CAWE) since they are contextual owing to the use of attention scores (over all acoustic frames for a given word).

5. EXPERIMENTS AND RESULTS

We use a commonly used speech recognition setup, the 300 hour Switchboard corpus (LDC97S62) [21] which consists of 2,430 two-sided telephonic conversations between 500 different speakers and contains 3 million words of text. Our second dataset is a 300 hour subset of the How2 [28] dataset of instructional videos, which contains planned, but free speech, often outdoor and recorded with distant microphones, as opposed to the indoor, telephony, conversational speech of Switchboard. There are 13,662 videos with a total of 3.5 million words in this corpus. The A2W obtains a word error rate of 22.2% on Switchboard and 36.6% on CallHome set from the Switchboard Eval2000 test set and 24.3% on dev5 test set of How2.

5.1. Comparing Methods for Constructing Embeddings

Datasets for Downstream Tasks: We evaluate our embeddings by using them as features for 16 benchmark sentence evaluation tasks that cover Semantic Textual Similarity (STS 2012-2016 and STS B), classification: Movie Review (MR), product review (CJ), sentiment analysis (SST, SST-FG), question type (TREC), Subjectivity/Objectivity (SUBJ), and opinion polarity (MPQA), entailment and semantic relatedness using the SICK dataset for SICK-E (entailment) and SICK-R (relatedness) and paraphrase detection (MRPC). The STS and SICK-R tasks measure Spearman’s coefficient of correlation between embedding based similarity and human scores, hence the scores range from $[-1, 1]$ where higher number denotes high correlation. All the remaining tasks are measured on test classification accuracies. We use the SentEval toolkit [27] to evaluate.

Training Details: In all downstream evaluations involving classification tasks, we have used a simple logistic regression for classification since a better representation should lead to better scores

without using complicated models (hence abstracting away model complexities from our evaluations). This also means that we can use the concatenation of CAWE and CBOW as features to the logistic regression model without adding tunable embedding parameters.

Discussion: From the results in Table 1 we see that CAWE-M outperforms U-AVG by 34% and 13% and CAWE-W by 33.9% and 12% on Switchboard and How2 datasets respectively in terms of average performance on STS tasks and leads to better or slightly worse performance on the classification tasks. We observe that CAWE-W usually performs worse than CAWE-M which could be attributed to a noisy estimation of the word embeddings on the account of taking even the less confident attention scores while constructing the embedding. In contrast, CAWE-M is constructed using the most confident attention score obtained over all the occurrences of the acoustic frames corresponding to a particular word. We also observe that U-AVG performs worse than CAWE-W on STS and SICK-R tasks since it is constructed using an even noisier process in which all encoder hidden representations are weighted equally irrespective of their attention scores.

5.2. Comparing with Text-based Embeddings

Datasets for Downstream Tasks: The datasets are the same as described in Section 5.1.

Training Details: In all the following comparisons, we compare embeddings obtained only from the training set of the speech recognition model, while the text-based word embeddings are obtained by training Continuous Bag-of-Words (CBOW) word2vec model on all the transcripts (train, validation and test). This was done to ensure a fair comparison between our supervised technique and the unsupervised word2vec method. This naturally leads to a smaller vocabulary for CAWE. Further, one of the drawbacks of A2W speech recognition model is that it fails to capture entire vocabulary, recognizing only 3044 words out of 29874 (out of which 18800 words occur less than 5 times) and 4287 out of 14242 total vocabulary for Switchboard and How2 respectively. Despite this fact, the performance of CAWE is very competitive with word2vec CBOW which does not

Table 2. Sentence Evaluations on 16 benchmark datasets for Switchboard and How2 corpus. We compare the CAWE-M method with the word2vec embeddings trained with CBOW method and with CAWE-M + CBOW concatenated (Concat) embeddings.

Dataset	Switchboard			How2		
	CAWE-M	CBOW	Concat	CAWE-M	CBOW	Concat
STS 2012	0.3561	0.3639	0.3470	0.3648	0.3688	0.3790
STS 2013	0.1969	0.1960	0.2010	0.2716	0.2524	0.2675
STS 2014	0.3888	0.3745	0.3795	0.3940	0.3973	0.3971
STS 2015	0.4275	0.4459	0.4481	0.4173	0.4781	0.4710
STS 2016	0.3833	0.3471	0.3651	0.3159	0.4023	0.3388
STS B	0.401	0.4100	0.3995	0.4000	0.4720	0.4487
SICK-R	0.6006	0.6170	0.6228	0.6440	0.6550	0.6945
MR	64.69	66.24	66.89	63.64	66.03	66.89
MRPC	69.80	68.99	68.00	70.61	69.68	68.52
CR	72.93	74.49	75.39	73.03	74.89	74.84
SUBJ	81.19	84.62	84.59	81.01	84.75	85.04
MPQA	73.75	76.44	75.36	73.53	75.56	75.60
SST	65.02	68.37	68.97	65.13	67.66	68.20
SST-FG	33.53	34.71	35.79	32.08	33.62	33.67
TREC	67.60	69.80	71.40	60.60	68.40	67.40
SICK-E	74.06	75.02	76.19	75.97	76.29	78.14

suffer from reduced vocabulary problem.

Discussion: In Table 2, we see that our embeddings perform as well as the text-embeddings. Evaluations using CAWE-M extracted from Switchboard based training show that the acoustic embeddings when concatenated with the text embeddings outperform the word2vec embeddings on 10 out of 16 tasks. This concatenated embedding shows that we add more information with CAWE-M that improves the CBOW embedding as well. The gains are more prominent in Switchboard as compared to the How2 dataset since How2 is planned instructional speech whereas Switchboard is spontaneous conversational speech (thereby making the How2 characteristics closer to text leading to a stronger CBOW model).

5.3. Evaluation on Spoken Language Understanding

Dataset: In addition to generic sentence-level evaluations, we also evaluate CAWE on the widely used ATIS dataset [29] for Spoken Language Understanding (SLU). ATIS dataset is comprised of spoken language queries for airline reservations that have intent and named entities. Hence, it is similar in domain to Switchboard, making it a useful test bed for evaluating CAWE on a speech-based downstream evaluation task.

Training Details: For this task, our model is similar to the simple Recurrent Neural Network (RNN) based model architecture as investigated in [30]. Our architecture is comprised of an embedding layer, a single layer RNN-variant (Simple RNN, Gated Recurrent Unit (GRU)) along with a dense layer and softmax. In each instance, we train our model for 10 epochs with RMSProp (learning rate 0.001). We train each model 3 times with different seed values and report average performance.

Discussion: [30] concluded that text-based word embeddings trained on large text corpora consistently lead to better performance on the ATIS dataset. We demonstrate that direct speech-based word embeddings could lead to matching performance when compared to text-based word embeddings in this speech-based downstream task, thus highlighting the utility of our speech based embeddings. Specifically, we compare the test scores obtained by initializing the model

Table 3. Speech-based contextual word embeddings (CAWE-M and CAWE-W) match the performance of the text-based embeddings (CBOW) on the ATIS dataset with an RNN and GRU model

	CAWE-M	CAWE-W	CBOW
RNN	91.49	91.67	91.82
GRU	93.25	93.56	93.11

with CAWE-M, CAWE-W and CBOW embeddings and fine-tuning them based on the task.

6. CONCLUSION

We present a method to learn contextual acoustic word embeddings from a sequence-to-sequence acoustic-to-word speech recognition model that learns to jointly segment and classify speech. We analyze the role of attention in constructing contextual acoustic word embeddings, and find our acoustic embeddings to be highly competitive with word2vec (CBOW) text embeddings. We discuss two variants of such contextual acoustic word embeddings which outperform the simple unweighted average method by upto 34% on semantic textual similarity tasks. The embeddings also matched the performance of text-based embeddings in spoken language understanding, showing the use of this model as a pre-trained model for other speech-based downstream tasks. We surmise that contextual audio embeddings will generalize and improve downstream tasks in a way that is similar to their text counterparts, despite the additional complexity presented by noisy audio input. In the future, we will explore ways to scale our model to larger corpora, larger vocabularies and compare with non-contextual acoustic word embedding methods.

7. ACKNOWLEDGEMENTS

This work was supported by the Center for Machine Learning and Health (CMLH) at Carnegie Mellon University and by Facebook.

8. REFERENCES

- [1] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [2] Jeffrey Pennington, Richard Socher, and Christopher 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.
- [3] Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher, "Learned in translation: Contextualized word vectors," in *Advances in Neural Information Processing Systems*, 2017, pp. 6294–6305.
- [4] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer, "Deep contextualized word representations," in *Proc. of NAACL*, 2018.
- [5] Yu-An Chung and James Glass, "Speech2vec: A sequence-to-sequence framework for learning word embeddings from speech," *arXiv preprint arXiv:1803.08976*, 2018.
- [6] Herman Kamper, Weiran Wang, and Karen Livescu, "Deep convolutional acoustic word embeddings using word-pair side information," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 4950–4954.
- [7] Samy Bengio and Georg Heigold, "Word embeddings for speech recognition," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [8] David Harwath and James Glass, "Deep multimodal semantic embeddings for speech and images," *arXiv preprint arXiv:1511.03690*, 2015.
- [9] Wanxia He, Weiran Wang, and Karen Livescu, "Multi-view recurrent neural acoustic word embeddings," *arXiv preprint arXiv:1611.04496*, 2016.
- [10] Ilya Sutskever, Oriol Vinyals, and Quoc V Le, "Sequence to sequence learning with neural networks," in *Proc. NIPS*, 2014.
- [11] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 4960–4964.
- [12] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, "Attention-based models for speech recognition," in *Advances in neural information processing systems*, 2015, pp. 577–585.
- [13] Shruti Palaskar and Florian Metze, "Acoustic-to-word recognition with sequence-to-sequence models," *arXiv preprint arXiv:1807.09597*, 2018.
- [14] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [15] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, pp. 369–376.
- [16] Hagen Soltau, Hank Liao, and Hasim Sak, "Neural speech recognizer: Acoustic-to-word lstm model for large vocabulary speech recognition," *arXiv preprint arXiv:1610.09975*, 2016.
- [17] Jinyu Li, Guoli Ye, Amit Das, Rui Zhao, and Yifan Gong, "Advancing acoustic-to-word ctc model," *arXiv preprint arXiv:1803.05566*, 2018.
- [18] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjuli Kannan, Ron J Weiss, Kanishka Rao, Katya Gonina, et al., "State-of-the-art speech recognition with sequence-to-sequence models," *arXiv preprint arXiv:1712.01769*, 2017.
- [19] Kanishka Rao, Haşim Sak, and Rohit Prabhavalkar, "Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer," in *Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE*. IEEE, 2017, pp. 193–199.
- [20] Albert Zeyer, Kazuki Irie, Ralf Schlüter, and Hermann Ney, "Improved training of end-to-end attention models for speech recognition," *arXiv preprint arXiv:1805.03294*, 2018.
- [21] John J Godfrey, Edward C Holliman, and Jane McDaniel, "Switchboard: Telephone speech corpus for research and development," in *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*. IEEE, 1992, vol. 1, pp. 517–520.
- [22] Kartik Audhkhasi, Brian Kingsbury, Bhuvana Ramabhadran, George Saon, and Michael Picheny, "Building competitive direct acoustics-to-word models for english conversational speech recognition," *arXiv preprint arXiv:1712.03133*, 2017.
- [23] Sei Ueno, Hirofumi Inaguma, Masato Mimura, and Tatsuya Kawahara, "Acoustic-to-word attention-based model complemented with character-level ctc-based model," .
- [24] Liang Lu, Xingxing Zhang, and Steve Renals, "On training the recurrent neural network encoder-decoder for large vocabulary end-to-end speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 5060–5064.
- [25] Zhehuai Chen, Qi Liu, Hao Li, and Kai Yu, "On modular training of neural acoustics-to-word model for lvcsr," *arXiv preprint arXiv:1803.01090*, 2018.
- [26] Yougen Yuan, Cheung-Chi Leung, Lei Xie, Hongjie Chen, Bin Ma, and Haizhou Li, "Learning acoustic word embeddings with temporal context for query-by-example speech search," *arXiv preprint arXiv:1806.03621*, 2018.
- [27] Alexis Conneau and Douwe Kiela, "Senteval: An evaluation toolkit for universal sentence representations," *arXiv preprint arXiv:1803.05449*, 2018.
- [28] Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze, "How2: A large-scale dataset for multimodal language understanding," in *Submitted to Visually Grounded Interaction and Language (ViGIL) NIPS 2018 Workshop*, 2018.
- [29] Patti J Price, "Evaluation of spoken language systems: The atis domain," in *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*, 1990.
- [30] Grégoire Mesnil, Xiaodong He, Li Deng, and Yoshua Bengio, "Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding," 2013.