

Gated Embeddings in End-to-End Speech Recognition for Conversational-Context Fusion

Suyoun Kim¹, Siddharth Dalmia² and Florian Metze²

¹Electrical & Computer Engineering

²Language Technologies Institute, School of Computer Science
Carnegie Mellon University

{suyoung1, sdalmia, fmetze}@andrew.cmu.edu

Abstract

We present a novel conversational-context aware end-to-end speech recognizer based on a gated neural network that incorporates conversational-context/word/speech embeddings. Unlike conventional speech recognition models, our model learns longer conversational-context information that spans across sentences and is consequently better at recognizing long conversations. Specifically, we propose to use text-based external word and/or sentence embeddings (i.e., fastText, BERT) within an end-to-end framework, yielding significant improvement in word error rate with better conversational-context representation. We evaluated the models on the Switchboard conversational speech corpus and show that our model outperforms standard end-to-end speech recognition models.

1 Introduction

In a long conversation, there exists a tendency of semantically related words, or phrases reoccur across sentences, or there exists topical coherence. Existing speech recognition systems are built at individual, isolated utterance level in order to make building systems computationally feasible. However, this may lose important conversational context information. There have been many studies that have attempted to inject a longer context information (Mikolov et al., 2010; Mikolov and Zweig, 2012; Wang and Cho, 2016; Ji et al., 2016; Liu and Lane, 2017; Xiong et al., 2018), all of these models are developed on text data for language modeling task.

There has been recent work attempted to use the conversational-context information within a *end-to-end* speech recognition framework (Kim and Metze, 2018; Kim et al., 2018; Kim and Metze, 2019). The new end-to-end speech recognition approach (Graves et al., 2006; Graves and

Jaitly, 2014; Hannun et al., 2014; Miao et al., 2015; Bahdanau et al., 2015; Chorowski et al., 2015; Chan et al., 2016; Kim et al., 2017) integrates all available information within a single neural network model, allows to make fusing conversational-context information possible. However, these are limited to encode only one preceding utterance and learn from a few hundred hours of annotated speech corpus, leading to minimal improvements.

Meanwhile, neural language models, such as fastText (Bojanowski et al., 2017; Joulin et al., 2017, 2016), ELMo (Peters et al., 2018), OpenAI GPT (Radford et al., 2019), and Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019), that encode words and sentences in fixed-length dense vectors, embeddings, have achieved impressive results on various natural language processing tasks. Such general word/sentence embeddings learned on large text corpora (i.e., Wikipedia) has been used extensively and plugged in a variety of downstream tasks, such as question-answering and natural language inference, (Devlin et al., 2019; Peters et al., 2018; Seo et al., 2017), to drastically improve their performance in the form of transfer learning.

In this paper, we create a conversational-context aware end-to-end speech recognizer capable of incorporating a conversational-context to better process long conversations. Specifically, we propose to exploit external word and/or sentence embeddings which trained on massive amount of text resources, (i.e. fastText, BERT) so that the model can learn better conversational-context representations. So far, the use of such pre-trained embeddings have found limited success in the speech recognition task. We also add a gating mechanism to the decoder network that can integrate all the available embeddings (word, speech, conversational-context) efficiently with increase

representational power using multiplicative interactions. Additionally, we explore a way to train our speech recognition model even with text-only data in the form of pre-training and joint-training approaches. We evaluate our model on the Switchboard conversational speech corpus (Godfrey and Holliman, 1993; Godfrey et al., 1992), and show that our model outperforms the sentence-level end-to-end speech recognition model. The main contributions of our work are as follows:

- We introduce a contextual gating mechanism to incorporate multiple types of embeddings, word, speech, and conversational-context embeddings.
- We exploit the external word (fastText) and/or sentence embeddings (BERT) for learning better conversational-context representation.
- We perform an extensive analysis of ways to represent the conversational-context in terms of the number of utterance history, and sampling strategy considering to use the generated sentences or the true preceding utterance.
- We explore a way to train the model jointly even with text-only dataset in addition to annotated speech data.

2 Related work

Several recent studies have considered to incorporate a context information within a end-to-end speech recognizer (Pundak et al., 2018; Alon et al., 2019). In contrast with our method which uses a conversational-context information in a long conversation, their methods use a list of phrases (i.e. play a song) in reference transcription in specific tasks, *contact names*, *songs names*, *voice search*, *dictation*.

Several recent studies have considered to exploit a longer context information that spans multiple sentences (Mikolov and Zweig, 2012; Wang and Cho, 2016; Ji et al., 2016; Liu and Lane, 2017; Xiong et al., 2018). In contrast with our method which uses a single framework for speech recognition tasks, their methods have been developed on text data for language models, and therefore, it must be integrated with a conventional acoustic model which is built separately without a longer context information.

Several recent studies have considered to embed a longer context information within a end-to-end framework (Kim and Metze, 2018; Kim et al., 2018; Kim and Metze, 2019). In contrast with our method which can learn a better conversational-context representation with a gated network that incorporate external word/sentence embeddings from multiple preceding sentence history, their methods are limited to learn conversational-context representation from one preceding sentence in annotated speech training set.

Gating-based approaches have been used for fusing word embeddings with visual representations in genre classification task or image search task (Arevalo et al., 2017; Kiros et al., 2018) and for learning different languages in speech recognition task (Kim and Seltzer, 2018).

3 End-to-End Speech Recognition Models

3.1 Joint CTC/Attention-based encoder-decoder network

We perform end-to-end speech recognition using a joint CTC/Attention-based approach with graphemes as the output symbols (Kim et al., 2017; Watanabe et al., 2017). The key advantage of the joint CTC/Attention framework is that it can address the weaknesses of the two main end-to-end models, Connectionist Temporal Classification (CTC) (Graves et al., 2006) and attention-based encoder-decoder (Attention) (Bahdanau et al., 2016), by combining the strengths of the two. With CTC, the neural network is trained according to a maximum-likelihood training criterion computed over all possible segmentations of the utterance’s sequence of feature vectors to its sequence of labels while preserving left-right order between input and output. With attention-based encoder-decoder models, the decoder network can learn the language model jointly without relying on the conditional independent assumption.

Given a sequence of acoustic feature vectors, \mathbf{x} , and the corresponding graphemic label sequence, \mathbf{y} , the joint CTC/Attention objective is represented as follows by combining two objectives with a tunable parameter $\lambda : 0 \leq \lambda \leq 1$:

$$\mathcal{L} = \lambda \mathcal{L}_{\text{CTC}} + (1 - \lambda) \mathcal{L}_{\text{att}}. \quad (1)$$

Each loss to be minimized is defined as the negative log likelihood of the ground truth character

sequence \mathbf{y}^* , is computed from:

$$\mathcal{L}_{CTC} \triangleq -\ln \sum_{\pi \in \Phi(\mathbf{y})} p(\pi|\mathbf{x}) \quad (2)$$

$$\mathcal{L}_{att} \triangleq -\sum_u \ln p(y_u^*|\mathbf{x}, y_{1:u-1}^*) \quad (3)$$

where π is the label sequence allowing the presence of the *blank* symbol, Φ is the set of all possible π given u -length \mathbf{y} , and $y_{1:u-1}^*$ is all the previous labels.

Both CTC and the attention-based encoder-decoder networks are also used in the inference step. The final hypothesis is a sequence that maximizes a weighted conditional probability of CTC and attention-based encoder-decoder network (Hori et al., 2017):

$$\mathbf{y}^* = \operatorname{argmax} \{ \gamma \log p_{CTC}(\mathbf{y}|\mathbf{x}) + (1 - \gamma) \log p_{att}(\mathbf{y}|\mathbf{x}) \} \quad (4)$$

3.2 Acoustic-to-Words Models

In this work, we use word units as our model outputs instead of sub-word units. Direct acoustic-to-word (A2W) models train a single neural network to directly recognize words from speech without any sub-word units, pronunciation model, decision tree, decoder, which significantly simplifies the training and decoding process (Soltau et al., 2017; Audhkhasi et al., 2017, 2018; Li et al., 2018; Palaskar and Metze, 2018). In addition, building A2W can learn more semantically meaningful conversational-context representations and it allows to exploit external resources like word/sentence embeddings where the unit of representation is generally words. However, A2W models require more training data compared to conventional sub-word models because it needs sufficient acoustic training examples per word to train well and need to handle out-of-vocabulary(OOV) words. As a way to manage this OOV issue, we first restrict the vocabulary to 10k frequently occurring words. We then additionally use a single character unit and start-of-OOV (sunk), end-of-OOV (eunk) tokens to make our model generate a character by decomposing the OOV word into a character sequence. For example, the OOV word, *rainstorm*, is decomposed into (sunk) *r a i n s t o r m* (eunk) and the model tries to learn such a character sequence rather than generate the OOV token. From this method, we

obtained 1.2% - 3.7% word error rate (WER) relative improvements in evaluation set where exists 2.9% of OOVs.

4 Conversational-context Aware Models

In this section, we describe the A2W model with conversational-context fusion. In order to fuse conversational context information within the A2W, end-to-end speech recognition framework, we extend the decoder sub-network to predict the output additionally conditioning on conversational context, by learning a conversational-context embedding. We encode single or multiple preceding utterance histories into a fixed-length, single vector, then inject it to the decoder network as an additional input at every output step.

Let say we have K number of utterances in a conversation. For k -th sentence, we have acoustic features $(x_1, \dots, x_T)^k$ and output word sequence, (w_1, \dots, w_U) . At output timestamp u , our decoder generates the probability distribution over words (w_u^k) , conditioned on 1) speech embeddings, attended high-level representation (e_{speech}^k) generated from encoder, and 2) word embeddings from all the words seen previously (e_{word}^{u-1}), and 3) conversational-context embeddings ($e_{context}^k$), which represents the conversational-context information for current (k) utterance prediction:

$$e_{speech}^k = \text{Encoder}(x^k) \quad (5)$$

$$w_u^k \sim \text{Decoder}(e_{context}^k, e_{word}^k, e_{speech}^k) \quad (6)$$

We can simply represent such contextual embedding, $e_{context}^k$, by mean of one-hot word vectors or word distributions, $\text{mean}(e_{word_1}^{k-1} + \dots + e_{word_U}^{k-1})$ from the preceding utterances.

In order to learn and use the conversational-context during training and decoding, we serialize the utterances based on their onset times and their conversations rather than random shuffling of data. We shuffle data at the conversation level and create mini-batches that contain only one sentence of each conversation. We fill the "dummy" input/output example at positions where the conversation ended earlier than others within the mini-batch to not influence other conversations while passing context to the next batch.

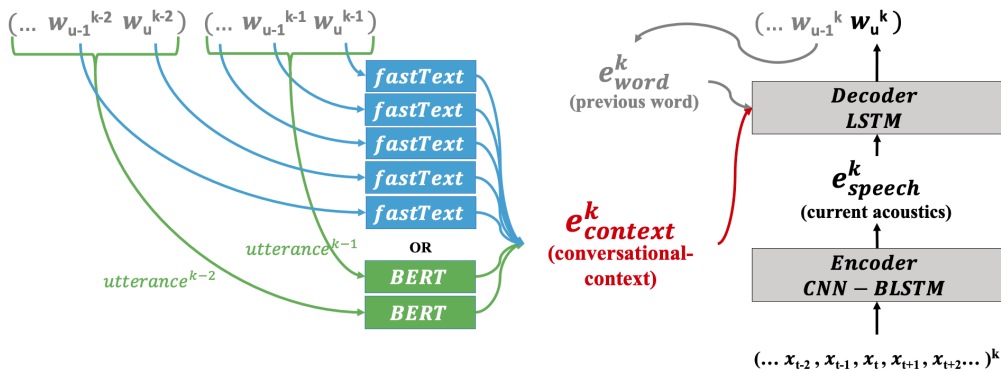


Figure 1: Conversational-context embedding representations from external word or sentence embeddings.

4.1 External word/sentence embeddings

Learning better representation of conversational-context is the key to achieve better processing of long conversations. To do so, we propose to encode the general word/sentence embeddings pre-trained on large textual corpora within our end-to-end speech recognition framework. Another advantage of using pre-trained embedding models is that we do not need to back-propagate the gradients across contexts, making it easier and faster to update the parameters for learning a conversational-context representation.

There exist many word/sentence embeddings which are publicly available. We can broadly classify them into two categories: (1) non-contextual word embeddings, and (2) contextual word embeddings. Non-contextual word embeddings, such as Word2Vec (Mikolov and Zweig, 2012), GloVe (Pennington et al., 2014), fastText (Bojanowski et al., 2017), maps each word independently on the context of the sentence where the word occur in. Although it is easy to use, it assumes that each word represents a single meaning which is not true in real-world. Contextualized word embeddings, sentence embeddings, such as deep contextualized word representations (Peters et al., 2018), BERT (Devlin et al., 2019), encode the complex characteristics and meanings of words in various context by jointly training a bidirectional language model. The BERT model proposed a masked language model training approach enabling them to also learn good “sentence” representation in order to predict the masked word.

In this work, we explore both types of embeddings to learn conversational-context embeddings as illustrated in Figure 1. The first method is to use word embeddings, fastText, to generate 300-

dimensional embeddings from 10k-dimensional one-hot vector or distribution over words of each previous word and then merge into a single context vector, $e_{context}^k$. Since we also consider multiple word/utterance history, we consider two simple ways to merge multiple embeddings (1) mean, and (2) concatenation. The second method is to use sentence embeddings, BERT. It is used to a generate single 786-dimensional sentence embedding from 10k-dimensional one-hot vector or distribution over previous words and then merge into a single context vector with two different merging methods. Since our A2W model uses a restricted vocabulary of 10k as our output units and which is different from the external embedding models, we need to handle out-of-vocabulary words. For fastText, words that are missing in the pretrained embeddings we map them to a random multivariate normal distribution with the mean as the sample mean and variance as the sample variance of the known words. For BERT, we use its provided tokenizer to generates byte pair encodings to handle OOV words.

Using this approach, we can obtain a more dense, informative, fixed-length vectors to encode conversational-context information, $e_{context}^k$ to be used in next k -th utterance prediction.

4.2 Contextual gating

We use contextual gating mechanism in our decoder network to combine the conversational-context embeddings with speech and word embeddings effectively. Our gating is contextual in the sense that multiple embeddings compute a gate value that is dependent on the context of multiple utterances that occur in a conversation. Using these contextual gates can be beneficial to decide how to weigh the different embed-

dings, conversational-context, word and speech embeddings. Rather than merely concatenating conversational-context embeddings (Kim and Metze, 2018), contextual gating can achieve more improvement because its increased representational power using multiplicative interactions.

Figure 2 illustrates our proposed contextual gating mechanism. Let $e_w = e_w(y_{u-1})$ be our previous word embedding for a word y_{u-1} , and let $e_s = e_s(x_{1:T}^k)$ be a speech embedding for the acoustic features of current k -th utterance $x_{1:T}^k$ and $e_c = e_c(s_{k-1-n:k-1})$ be our conversational-context embedding for n -number of preceding utterances $s_{k-1-n:k-1}$. Then using a gating mechanism:

$$g = \sigma(e_c, e_w, e_s) \quad (7)$$

where σ is a 1 hidden layer DNN with sigmoid activation, the gated embedding e is calculated as

$$e = g \odot (e_c, e_w, e_s) \quad (8)$$

$$h = \text{LSTM}(e) \quad (9)$$

and fed into the LSTM decoder hidden layer. The output of the decoder h is then combined with conversational-context embedding e_c again with a gating mechanism,

$$g = \sigma(e_c, h) \quad (10)$$

$$\hat{h} = g \odot (e_c, h) \quad (11)$$

Then the next hidden layer takes these gated activations, \hat{h} , and so on.

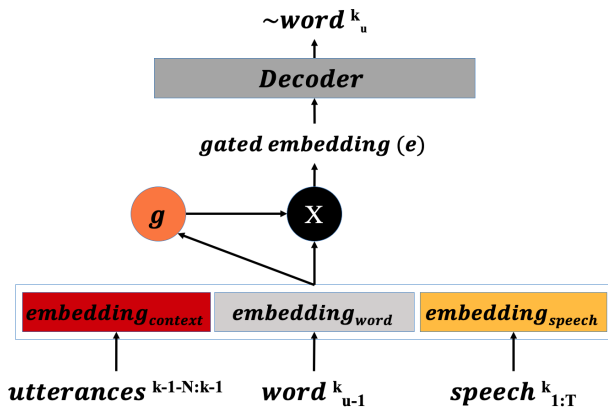


Figure 2: Our contextual gating mechanism in decoder network to integrate three different embeddings from: 1) conversational-context, 2) previous word, 3) current speech.

Dataset	# of utter.	# of conversations	avg. # of utter. /conversation
training	192,656	2402	80
validation	4,000	34	118
eval.(SWBD)	1,831	20	92
eval.(CH)	2,627	20	131

Table 1: Experimental dataset description. We used 300 hours of Switchboard conversational corpus. Note that any pronunciation lexicon or Fisher transcription was not used.

5 Experiments

5.1 Datasets

To evaluate our proposed conversational end-to-end speech recognition model, we use the Switchboard (SWBD) LDC corpus (97S62) task. We split 300 hours of the SWBD training set into two: 285 hours of data for the model training, and 5 hours of data for the hyper-parameter tuning. We evaluate the model performance on the HUB5 Eval2000 which consists of the Callhome English (CH) and Switchboard (SWBD) (LDC2002S09, LDC2002T43). In Table 1, we show the number of conversations and the average number of utterances per a single conversation.

The audio data is sampled at 16kHz, and then each frame is converted to a 83-dimensional feature vector consisting of 80-dimensional log-mel filterbank coefficients and 3-dimensional pitch features as suggested in (Miao et al., 2016). The number of our word-level output tokens is 10,038, which includes 47 single character units as described in Section 3.2. Note that no pronunciation lexicon was used in any of the experiments.

5.2 Training and decoding

For the architecture of the end-to-end speech recognition, we used joint CTC/Attention end-to-end speech recognition (Kim et al., 2017; Watanabe et al., 2017). As suggested in (Zhang et al., 2017; Hori et al., 2017), the input feature images are reduced to $(1/4 \times 1/4)$ images along with the time-frequency axis within the two max-pooling layers in CNN. Then, the 6-layer BLSTM with 320 cells is followed by the CNN layer. For the attention mechanism, we used a location-based method (Chorowski et al., 2015). For the decoder network, we used a 2-layer LSTM with 300 cells. In addition to the standard decoder network, our proposed models additionally require extra parameters for gating layers in order to fuse

Model	Output Units	Trainable Params	External LM	SWBD (WER%)	CH (WER%)
Prior Models					
LF-MMI (Povey et al., 2016)	CD phones	N/A	✓	9.6	19.3
CTC (Zweig et al., 2017)	Char	53M	✓	19.8	32.1
CTC (Sanabria and Metze, 2018)	Char, BPE-{300,1k,10k}	26M	✓	12.5	23.7
CTC (Audhkhasi et al., 2018)	Word (Phone init.)	N/A	✓	14.6	23.6
Seq2Seq (Zeyer et al., 2018)	BPE-10k	150M*	✗	13.5	27.1
Seq2Seq (Palaskar and Metze, 2018)	Word-10k	N/A	✗	23.0	37.2
Seq2Seq (Zeyer et al., 2018)	BPE-1k	150M*	✓	11.8	25.7
Our baseline					
	Word-10k	32M	✗	18.2	30.7
Our Proposed Conversational Model					
Gated Contextual Decoder	Word-10k	35M	✗	17.3	30.5
+ Decoder Pretrain	Word-10k	35M	✗	16.4	29.5
+ fastText for Word Emb.	Word-10k	35M	✗	16.0	29.5
(a) fastText for Conversational Emb.	Word-10k	34M	✗	16.0	29.5
(b) BERT for Conversational Emb.	Word-10k	34M	✗	15.7	29.2
(b) + Turn number 5	Word-10k	34M	✗	15.5	29.0

Table 2: Comparison of word error rates (WER) on Switchboard 300h with standard end-to-end speech recognition models and our proposed end-to-end speech recognition models with conversational context. (The * mark denotes our estimate for the number of parameters used in the previous work).

conversational-context embedding to the decoder network compared to baseline. We denote the total number of trainable parameters in Table 2.

For the optimization method, we use AdaDelta (Zeiler, 2012) with gradient clipping (Pascanu et al., 2013). We used $\lambda = 0.2$ for joint CTC/Attention training (in Eq. 1) and $\gamma = 0.3$ for joint CTC/Attention decoding (in Eq.4). We bootstrap the training of our proposed conversational end-to-end models from the baseline end-to-end models. To decide the best models for testing, we monitor the development accuracy where we always use the model prediction in order to simulate the testing scenario. At inference, we used a left-right beam search method (Sutskever et al., 2014) with the beam size 10 for reducing the computational cost. We adjusted the final score, $s(\mathbf{y}|\mathbf{x})$, with the length penalty 0.5. The models are implemented using the PyTorch deep learning library (Paszke et al., 2017), and ESPnet toolkit (Kim et al., 2017; Watanabe et al., 2017, 2018).

6 Results

Our results are summarized in the Table 2 where we first present the baseline results and then show the improvements by adding each of the individual components that we discussed in previous sections, namely, gated decoding, pretraining decoder network, external word embedding, external

conversational embedding and increasing receptive field of the conversational context. Our best model gets around 15% relative improvement on the SWBD subset and 5% relative improvement on the CallHome subset of the eval2000 dataset.

We start by evaluating our proposed model which leveraged conversational-context embeddings learned from training corpus and compare it with a standard end-to-end speech recognition models without conversational-context embedding. As seen in Table 2, we obtained a performance gain over the baseline by using conversational-context embeddings which is learned from training set.

6.1 Pre-training decoder network

Then, we observe that pre-training of decoder network can improve accuracy further as shown in Table 2. Using pre-training the decoder network, we achieved 5% relative improvement in WER on SWBD set. Since we add external parameters in decoder network to learn conversational-context embeddings, our model requires more efforts to learn these additional parameters. To relieve this issue, we used pre-training techniques to train decoder network with text-only data first. We simply used a mask on top of the Encoder/Attention layer so that we can control the gradients of batches contains text-only data and do not update the En-

coder/Attention sub-network parameters.

6.2 Use of words/sentence embeddings

Next, we evaluated the use of pretrained external embeddings (fastText and BERT). We initially observed that we can obtain 2.4% relative improvement over (the model with decoder pretraining) in WER by using fastText for additional word embeddings to the gated decoder network.

We also extensively evaluated various ways to use fastText/BERT for conversational-context embeddings. Both methods with fastText and with BERT shows significant improvement from the baseline as well as vanilla conversational-context aware model.

6.3 Conversational-context Receptive Field

We also investigate the effect of the number of utterance history being encoded. We tried different $N = [1, 5, 9]$ number of utterance histories to learn the conversational-context embeddings. Figure 3 shows the relative improvements in the accuracy on the Dev set (5.2) over the baseline “non-conversational” model. We show the improvements on the two different methods of merging the contextual embeddings, namely mean and concatenation. Typically increasing the receptive field of the conversational-context helps improve the model. However, as the number of utterance history increased, the number of trainable parameters of the concatenate model increased making it harder for the model to train. This led to a reduction in the accuracy.

We also found that using 5-utterance history with concatenation performed best (15%) on the SWBD set, and using 9-number of utterance history with mean method performed best (5%) on CH set. We also observed that the improvement diminished when we used 9-utterance history for SWBD set, unlike CH set. One possible explanation is that the conversational-context may not be relevant to the current utterance prediction or the model is overfitting.




Figure 3: The relative improvement in Development accuracy over sets over baseline obtained by using conversational-context embeddings with different number of utterance history and different merging techniques.

6.4 Sampling technique

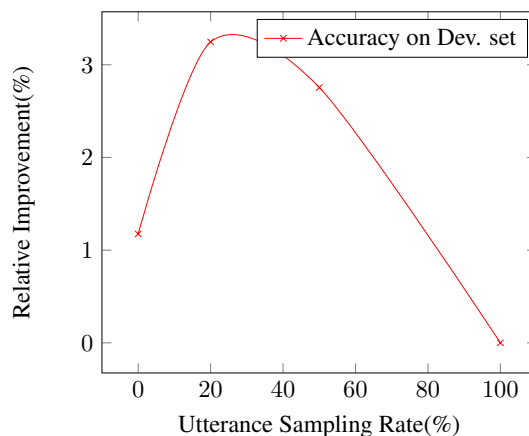


Figure 4: The relative improvement in Development accuracy over 100% sampling rate which was used in (Kim and Metze, 2018) obtained by using conversational-context embeddings with different sampling rate.

We also experiment with an utterance level sampling strategy with various sampling ratio, $[0.0, 0.2, 0.5, 1.0]$. Sampling techniques have been extensively used in sequence prediction tasks to reduce overfitting (Bengio et al., 2015) by training the model conditioning on generated tokens from the model itself, which is how the model actually do at inference, rather than the ground-truth tokens. Similar to choosing previous word tokens from the ground truth or from the model output, we apply it to choose previous utterance from the ground truth or from the model output for learning conversational-context embeddings. Fig-

ure 4 shows the relative improvement in the development accuracy (5.2) over the 1.0 sampling rate which is always choosing model’s output. We found that a sampling rate of 20% performed best.

6.5 Analysis of context embeddings

We develop a scoring function, $s(i, j)$ to check if our model conserves the conversational consistency for validating the accuracy improvement of our approach. The scoring function measures the average of the conversational distances over every consecutive hypotheses generated from a particular model. The conversational distance is calculated by the Euclidean distance, $\text{dist}(e_i, e_j)$ of the fixed-length vectors e_i, e_j which represent the model’s i, j -th hypothesis, respectively. To obtain a fixed-length vector, utterance embedding, given the model hypothesis, we use BERT sentence embedding as an oracle. Mathematically it can be written as,

$$s(i, j) = \frac{1}{N} \sum_{i, j \in \text{eval}} (\text{dist}(e_i, e_j))$$

where, i, j is a pair of consecutive hypotheses in evaluation data eval , N is the total number of i, j pairs, e_i, e_j are BERT embeddings. In our experiment, we select the pairs of consecutive utterances from the reference that show lower distance score at least baseline hypotheses.

From this process, we obtained three conversational distance scores from 1) the reference transcripts, 2) the hypotheses of our vanilla conversational model which is not using BERT, and 3) the hypotheses of our baseline model. Figure 5 shows the score comparison.




Figure 5: Comparison of the conversational distance score on the consecutive utterances of 1) reference, 2) our proposed conversational end-to-end model, and 3) our end-to-end baseline model.

We found that our proposed model was 7.4% relatively closer to the reference than the baseline.

This indicates that our conversational-context embedding leads to improved similarity across adjacent utterances, resulting in better processing a long conversation.

7 Conclusion

We have introduced a novel method for conversational-context aware end-to-end speech recognition based on a gated network that incorporates word/sentence/speech embeddings. Unlike prior work, our model is trained on conversational datasets to predict a word, conditioning on multiple preceding conversational-context representations, and consequently improves recognition accuracy of a long conversation. Moreover, our gated network can incorporate effectively with text-based external resources, word or sentence embeddings (i.e., fasttext, BERT) within an end-to-end framework and so that the whole system can be optimized towards our final objectives, speech recognition accuracy. By incorporating external embeddings with gating mechanism, our model can achieve further improvement with better conversational-context representation. We evaluated the models on the Switchboard conversational speech corpus and show that our proposed model using gated conversational-context embedding show 15%, 5% relative improvement in WER compared to a baseline model for Switchboard and CallHome subsets respectively. Our model was shown to outperform standard end-to-end speech recognition models trained on isolated sentences. This work is easy to scale and can potentially be applied to any speech related task that can benefit from longer context information, such as spoken dialog system, sentimental analysis.

Acknowledgments

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. This work also used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

References

Uri Alon, Golan Pundak, and Tara N Sainath. 2019. Contextual speech recognition with difficult negative training examples. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and*

- Signal Processing (ICASSP)*, pages 6440–6444. IEEE.
- John Arevalo, Thamar Solorio, Manuel Montes-y Gómez, and Fabio A González. 2017. Gated multimodal units for information fusion. *arXiv preprint arXiv:1702.01992*.
- Kartik Audhkhasi, Brian Kingsbury, Bhuvana Ramabhadran, George Saon, and Michael Picheny. 2018. Building competitive direct acoustics-to-word models for english conversational speech recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4759–4763. IEEE.
- Kartik Audhkhasi, Bhuvana Ramabhadran, George Saon, Michael Picheny, and David Nahamoo. 2017. Direct acoustics-to-word models for english conversational speech recognition. *CoRR*, abs/1703.07754.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *ICLR*.
- Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio. 2016. End-to-end attention-based large vocabulary speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4945–4949. IEEE.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. 2016. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964. IEEE.
- Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. 2015. Attention-based models for speech recognition. In *Advances in neural information processing systems*, pages 577–585.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL*.
- John Godfrey and Edward Holliman. 1993. Switchboard-1 release 2 ldc97s62. *Linguistic Data Consortium, Philadelphia*, LDC97S62.
- John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, volume 1, pages 517–520. IEEE.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376. ACM.
- Alex Graves and Navdeep Jaitly. 2014. Towards end-to-end speech recognition with recurrent neural networks. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1764–1772.
- Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Takaaki Hori, Shinji Watanabe, Yu Zhang, and William Chan. 2017. Advances in joint ctc-attention based end-to-end speech recognition with a deep cnn encoder and rnn-lm. *Interspeech*.
- Yangfeng Ji, Trevor Cohn, Lingpeng Kong, Chris Dyer, and Jacob Eisenstein. 2016. Document context language models. *ICLR (Workshop track)*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Herve Jégou, and Tomas Mikolov. 2016. Fasttext. zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431. Association for Computational Linguistics.
- Suyoun Kim, Siddharth Dalmia, and Florian Metze. 2018. Situation informed end-to-end asr for chime-5 challenge. *CHiME5 workshop*.
- Suyoun Kim, Takaaki Hori, and Shinji Watanabe. 2017. Joint ctc-attention based end-to-end speech recognition using multi-task learning. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 4835–4839. IEEE.
- Suyoun Kim and Florian Metze. 2018. Dialog-context aware end-to-end speech recognition. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 434–440. IEEE.

- Suyoun Kim and Florian Metze. 2019. Acoustic-to-word models with conversational context information. *NAACL*.
- Suyoun Kim and Michael L Seltzer. 2018. Towards language-universal end-to-end speech recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4914–4918. IEEE.
- Jamie Kiros, William Chan, and Geoffrey Hinton. 2018. Illustrative language understanding: Large-scale visual grounding with image search. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 922–933.
- Jinyu Li, Guoli Ye, Amit Das, Rui Zhao, and Yifan Gong. 2018. Advancing acoustic-to-word ctc model. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5794–5798. IEEE.
- Bing Liu and Ian Lane. 2017. Dialog context language modeling with recurrent neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 5715–5719. IEEE.
- Yajie Miao, Mohammad Gowayyed, and Florian Metze. 2015. EESN: End-to-end speech recognition using deep RNN models and WFST-based decoding. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 167–174. IEEE.
- Yajie Miao, Mohammad Gowayyed, Xingyu Na, Tom Ko, Florian Metze, and Alexander Waibel. 2016. An empirical exploration of ctc acoustic models. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2623–2627. IEEE.
- Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Eleventh Annual Conference of the International Speech Communication Association*.
- Tomas Mikolov and Geoffrey Zweig. 2012. Context dependent recurrent neural network language model. *SLT*, 12:234–239.
- Shruti Palaskar and Florian Metze. 2018. Acoustic-to-word recognition with sequence-to-sequence models. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 397–404. IEEE.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations.
- Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahremani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur. 2016. Purely sequence-trained neural networks for asr based on lattice-free mmi. In *Interspeech*, pages 2751–2755.
- Golan Pundak, Tara N Sainath, Rohit Prabhavalkar, Anjali Kannan, and Ding Zhao. 2018. Deep context: end-to-end contextual speech recognition. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 418–425. IEEE.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Ramon Sanabria and Florian Metze. 2018. Hierarchical multitask learning with ctc. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 485–490. IEEE.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. *CoRR*, abs/1611.01603.
- Hagen Soltau, Hank Liao, and Hasim Sak. 2017. Neural speech recognizer: Acoustic-to-word lstm model for large vocabulary speech recognition. *Interspeech*.
- Ilya Sutskever, Oriol Vinyals, and Quoc VV Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Tian Wang and Kyunghyun Cho. 2016. Larger-context language modelling. *ACL*.
- Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai. 2018. *EspNet: End-to-end speech processing toolkit*. In *Interspeech*, pages 2207–2211.
- Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R Hershey, and Tomoki Hayashi. 2017. Hybrid ctc/attention architecture for end-to-end speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11(8):1240–1253.

- Wayne Xiong, Lingfeng Wu, Jun Zhang, and Andreas Stolcke. 2018. Session-level language modeling for conversational speech. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2764–2768.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Albert Zeyer, Kazuki Irie, Ralf Schlüter, and Hermann Ney. 2018. Improved training of end-to-end attention models for speech recognition. *Interspeech*.
- Yu Zhang, William Chan, and Navdeep Jaitly. 2017. Very deep convolutional networks for end-to-end speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 4845–4849. IEEE.
- Geoffrey Zweig, Chengzhu Yu, Jasha Droppo, and Andreas Stolcke. 2017. Advances in all-neural speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 4805–4809. IEEE.