

# AUDIO-ATTENTION DISCRIMINATIVE LANGUAGE MODEL FOR ASR RESCORING

Ankur Gandhe, Ariya Rastrow

{aggandhe, arastrow}@amazon.com

## ABSTRACT

End-to-end approaches for automatic speech recognition (ASR) benefit from directly modeling the probability of the word sequence given the input audio stream in a single neural network. However, compared to conventional ASR systems, these models typically require more data to achieve comparable results. Well-known model adaptation techniques, to account for domain and style adaptation, are not easily applicable to end-to-end systems. Conventional HMM-based systems, on the other hand, have been optimized for various production environments and use cases. In this work, we propose to combine the benefits of end-to-end approaches with a conventional system using an attention-based *discriminative language model* that learns to rescore the output of a first-pass ASR system. We show that learning to rescore a list of potential ASR outputs is much simpler than learning to generate the hypothesis. The proposed model results in up to 8% improvement in word error rate even when the amount of training data is a fraction of data used for training the first-pass system.

**Index Terms**— language modeling, end-to-end, attention, minimum word error rate

## 1. INTRODUCTION

Conventional automatic speech recognition (ASR) systems model the probability of transcription given the input audio as three separate models [1] - acoustic, pronunciation and language models - each trained with a different object function. In general, the acoustic models are based on a hybrid DNN-HMM structure and language models are  $n$ -gram based models. Lately, recurrent neural network (RNN) language models have been shown to outperform  $n$ -gram based models, due to their ability to capture long-term context [2, 3]. However, due to their computational complexity, it is often not feasible to use recurrent neural networks in the first pass of a real-time ASR system. Instead, they are often used as second-pass rescorsers, as the set of candidate word sequences to rescore is usually small in size [4–6].

More recently, a lot of work has been done in training a single neural network that directly learns how to map input speech signal to word sequences (or graphemes) [7–9], without the need for training separate models. Although the results are

state-of-the-art or close to it on certain tasks, most end-to-end models are trained on large amounts of transcribed data and it is unclear if these models will perform better than conventional systems in a low data regime [10] or in (domain) mismatch scenarios where the underlying distribution for (paired speech and text) training data does not match the test distribution [11]. Various techniques have been explored to overcome these problems and bring some of the existing advantages of a conventional system - fusing external language model trained on abundantly available matched text corpora [11, 12], contextual biasing [13, 14] and customizing pronunciations [15] - to the end-to-end models.

In this paper, we combine the simplicity of rescoring conventional ASR models, using neural language models, with the benefits of an attention-based end-to-end model. More specifically, an RNNLM style model is trained using word-level contextual input, while simultaneously attending to audio. The model is trained using a minimum word error rate (MWER) criterion, which learns to rescore the  $N$ -best hypotheses list from a first pass system. The proposed model can learn, using a *discriminative criterion*, to focus and overcome the errors of the conventional first pass ASR model, having access to both word-level and acoustic-based embeddings. To avoid overfitting and data sparsity issues, we propose to use embeddings generated by a first pass acoustic model. This choice can be relaxed depending on the availability of training data and one can decide to learn the acoustic embeddings (or fine tune the first pass acoustic model embeddings) directly using the MWER criterion. However, using first pass acoustic model embeddings can practically speed up both training and inference. Additionally and since the model is trained using a MWER criterion, it does not have to be confined to a normalized generative form. This can drastically speed up inference, especially when dealing with very large vocabulary ASR systems.

## 2. RELATED WORK

MWER criterion has been previously used in [16] to improve the performance of LSTM LMs. They train the model using the MWER criterion but they do not explore using an unnormalized model or a model with attention to audio. [17] also explored training a spelling correction model to learn the

errors made by a first pass system. However, they did not explore using this model for a conventional system or condition on input audio (e.g. by attending to input audio). Besides, their spelling correction model is a generative model (decoder generates corrected characters in a left-to-right manner). Recently, [18] proposed to use a listen-attend-and-spell (LAS) model as a second pass rescoring model for output of an RNN-Transducer (RNN-T) first pass. They show significant improvements with LAS as rescoring component, even with a shared encoder. However, the amount of training data used for LAS rescoring was the same as the training data used for first pass model while we explore training the rescoring model with much less amount of data and is important for low-resource languages or domain adaptation with limited amount of in-domain data. The best performance was also achieved when the shared encoder was fine tuned in multi-task fashion (combined loss for both RNN-T and LAS). Additionally, they do not report results on using such a rescoring scheme with a conventional ASR model.

### 3. AUDIO-ATTENTION RECURRENT NEURAL NETWORK LANGUAGE MODEL

A language model estimates the probability of a given sequence of words  $\vec{w}$ . N-gram LMs make the assumption that the probability of word  $w_i$  depends only on previous  $n - 1$  words, so that the probability can be written as:

$$p(\vec{w}) = \prod_{w_i} p(w_i | w_{i-1}, w_{i-2}, \dots, w_{i-n+1}) \quad (1)$$

However, recurrent neural networks can model the probability of word given its entire history by applying a compact recurrence architecture on the hidden layer of the model.

$$p(\vec{w}) = \prod_{w_i} p(w_i | h_i) \quad (2)$$

where  $h_i$  is the hidden representation for the word  $w_i$ . In a simple recurrent network, the forward step for predicting the probability of word  $w_i$  can be described as [16]:

$$\begin{aligned} x_i &= \text{OneHot}(w_{i-1}) \\ h_i &= \sigma(W_{ih}^T x_i + W_{hh}^T h_{i-1}) \\ p(w_i | w_{<i}) &= \text{Softmax}(W_{ho}^T h_i) \end{aligned} \quad (3)$$

where  $\text{OneHot}$  returns a low dimensional representation for a word,  $\sigma()$  is an element-wise sigmoid function and the final model probabilities are obtained by applying a softmax function on the activation of the final output. The above set of equation can be easily extended to other types of recurrent models such as LSTMs [19].

An RNN-based language model is trained by minimizing the cross-entropy loss, which directly maximizes the log-

likelihood of training data sequences:

$$\mathcal{L}_{CE} = \sum_w \sum_{i=1}^L -\log p(w_i | w_{i-1}, w_{i-2}, \dots, w_0) \quad (4)$$

#### 3.1. Minimum word error rate training

Using a RNN-based language model during first pass of a real-time ASR system is often infeasible (due to lack of state merging as opposed to  $n$ -gram models) and most applications use RNNLMs to rescore the output of a first pass ASR system. The most common approach is to generate an  $n$ -best list from the first pass and replace or interpolate the language model score of the first pass with a score from the RNNLM.

As a rescoring model, an RNNLM can be directly trained to rescore the first pass  $n$ -best list, as opposed to a maximum likelihood criterion over text-only data. MWER criterion, which estimates an expectation of word errors over an output hypotheses list, has been proposed for both language model training [16] as well as attention-based models [9]. Let  $\mathcal{E}(\vec{w}_i, \vec{w}^*)$  be the edit distance between hypothesis  $\vec{w}_i$  and reference sequence  $\vec{w}^*$ , let  $\bar{\mathcal{E}}$  be the average edit distance for an  $n$ -best list, and let  $\hat{\mathcal{E}}() = \mathcal{E}() - \bar{\mathcal{E}}$  be the relative edit distance, then MWER loss is defined as:

$$\mathcal{L}_{werr} = \sum_{\vec{w}_i \in \mathbf{w}} p_{\theta}(\vec{w}_i | O) \hat{\mathcal{E}}(\vec{w}_i, \vec{w}^*) \quad (5)$$

where  $p_{\theta}(\vec{w}_i | O)$  is the normalized posterior probability (over  $n$ -best list) of the hypothesis given the input acoustics, computed considering both the acoustic model and the language model probabilities:

$$p_{\theta}(\vec{w}_i | O) = \frac{\exp(g_i)}{\sum_{j=1}^N \exp(g_j)} \quad (6)$$

where

$$g_i = \alpha \log p_{lm}(\vec{w}_i) + \log p_{am}(O | \vec{w}_i) \quad (7)$$

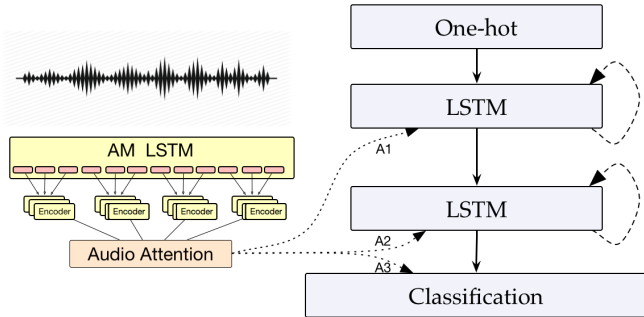
and  $p_{am}(O | \vec{w}_i)$  is the sequence probability computed by the acoustic model,  $p_{lm}(\vec{w}_i)$  is the sequence probability computed by the rescoring RNNLM.  $N$  is the size of  $n$ -best list. According to equation 6, the posterior probabilities are normalized over the  $n$ -best list. Therefore, it is not necessary for the RNNLM to produce normalized probabilities for each word of the sequence. In this paper, we explore both a normalized and unnormalized output for the RNNLM.

To stabilize training and achieve better convergence, we interpolate the word error rate losses with cross-entropy loss, weighted by a hyper-parameter  $\lambda$  [9]:

$$\mathcal{L}_{total} = \mathcal{L}_{werr} + \lambda \mathcal{L}_{CE} \quad (8)$$

#### 3.2. Attention to audio

As described in Section 1, unlike conventional ASR systems with a separate acoustic model and language model, end-to-end systems benefit from directly optimizing the  $p(w|O)$ .



**Fig. 1:** LSTM language model with attention to audio embedding from the first pass

Similarly, the input to the standard RNNLM can be modified so that it also models  $p(w|O)$ :

$$P(\vec{w}|O) = \prod_{|\vec{w}|} p(w_i|h_i, O) \quad (9)$$

There is no direct alignment between the words and segments of input audio. The dependence in equation 9 is learnt by using an LSTM with attention to input audio frames. At every step, the probability of next word is computed conditioned on all the previous words as well an attention context over the entire audio sequence:

$$\begin{aligned} c_i &= \text{AttentionContext}(h_{i-1}, O) \\ h_i &= \text{RNN}(h_{i-1}, w_{i-1}, c_{i-1}) \\ p(w_i|w_{<i}, O) &= \text{Softmax}(W_{ho}^T(h_i, c_i)) \end{aligned} \quad (10)$$

The attention context over (frame-by-frame) encoded audio is learnt similar to the way described in [8]. However, unlike end-to-end systems, we make two distinctions:

1. Learning a separate encoder over raw audio features is not a requirement for our model. Instead, we directly take the intermediate layers of the first pass acoustic model as input to our model. This reduces the amount of training data needed to learn a good model. In section 4 we compare using different encoder types on top of the output of first-pass acoustic model.
2. The model is not trained to generate the correct transcription but instead trained only to minimize the word error rate of an n-best list generated by a first pass system. This is an important distinction as the *discriminative* flavor of our model enables it to focus on correcting first-pass errors, as opposed to learn to generate words.

Figure 1 shows the proposed architecture with different ways of adding attention to the LM LSTMs; AM LSTM is the first pass model used for generating the audio embedding, CNN as an optional encoder that can be learnt during model training.

## 4. EXPERIMENTS

In all of the experiments in this paper, we build an ASR system that targets a spoken dialog system used in digital assistants such as Amazon Alexa, Google assistant, Siri, etc. The acoustic model is a low-frame-rate model with 2-layer frequency LSTM [20] followed by a 5-layer time LSTM trained on cross-entropy loss, followed by sMBR loss [21].

All recurrent neural models (NLM) models have a word embedding size of 512, and comprise two LSTM [19] layers, each with 512 hidden layers. During training, the models are trained on either cross-entropy loss (NLM-XENT) or MWER loss (NLM-MWE) as described in section 3.1, with a maximum n-best list size of 64. During inference, the NLMs are used to rescore 10-best hypotheses generated from the first-pass decoding. For NLM models with attention to audio (NLM-AUDIO), we attend to the activation of the last LSTM layer (dimension of 768) of the first pass acoustic model and learn to projection it to a smaller dimension of 200. For all experiments, a CNN-based encoder was used to learn the acoustic embeddings (details in section 4.3).

The first-pass LM is trained on a variety of out-of and in-domain corpora, written text data and transcribed speech data from real user-agent interactions<sup>1</sup> respectively. The written text corpora contain over 50 billion words in total. A Kneser-Ney (KN) [22] smoothed n-gram language model is estimated from each corpus with a vocabulary of 400k, and the final first-pass LM is a linear interpolation of these component LMs. The interpolation weights are estimated by minimizing the perplexity on a development set. The transcribed speech data, used to train the rescoring NLM-XENT comprises of approximately 300 million words of text. A smaller dataset of 5 million tokens is used for training the models that require audio data (NLM-MWE and NLM-AUDIO) in addition to the transcription. From this corpus, we extract the vocabulary of 66k most frequent words. All NLM models use this vocabulary and out of vocabulary tokens are mapped to  $\langle \text{unk} \rangle$ .

We evaluate each model on a 40k utterance test data set and measure the word error rate reduction (WERR) relative to the baseline model. Table 1 shows the result for different rescoring models both in terms of perplexity<sup>2</sup> and word error rate improvements (compared against the baseline 4-gram LM). We see very small improvement when using a cross-entropy trained NLM. However, with MWER rate training with un-normalized output, even on much smaller amount of data, the improvement is 3.6%; the perplexity of the normalized is much higher because its not optimized for maximum likelihood. Further, the model with attention to acoustic embeddings, from the first-pass AM results in about 7.3% improvement. The perplexity of the NLM model with attention

<sup>1</sup>The user data is anonymized and only a subset of data used for our production system

<sup>2</sup>Perplexity is only reported for normalized models.

**Table 1:** Perplexity(PPL) and Relative Word Error Rate Reduction (WERR) for different rescoring models

Model	PPL	WERR
KN-4g	37.2	
NLM-XENT	34.7	1.9%
NLM-MWE-NORM	90.7	2.6%
NLM-MWE	-	3.6%
NLM-MWE-AUDIO-NORM	18.5	7.0%
NLM-MWE-AUDIO	-	7.3%
10-best oracle	-	28%

to audio is significantly low because its able to attend to the audio frames corresponding to the word being predicted. We also see that the un-normalized models perform slightly better than the normalized model because they have less constraints on the output scores.

#### 4.1. Using pre-trained language model

Minimum word error rate training requires audio data to generate the n-best list required for training. However, LMs trained with cross-entropy can be trained on much larger amounts of text-only data. Table 2 shows the result of fine-tuning a pre-trained NLM-XENT (trained on 300 million tokens) with MWER criterion. In case of audio-attention model, it is not possible to pre-train all the layers of model, since it needs an extra context input. Here, we pre-train the two LSTM layers of the model and apply audio context only to the final affine. We see only a minor improvement over the randomly initialized model in all cases. This might be due to the amount of audio training data we have.

**Table 2:** Perplexity(PPL) and Relative Word Error Rate Reduction (WERR) for different attention location

Model	No-pretrain		Pretrain	
	PPL	WERR	PPL	WERR
NLM-XENT	34.7	1.9 %	33.8	1.9%
NLM-MWE	-	3.6%	-	4.5%
NLM-MWE-AUDIO	-	7.3%	-	7.7%

#### 4.2. Adding audio attention to LM

In a standard end-to-end systems, the attention to encoder states is applied at every layer. However, in a language model, the model has to learn both the dependency to acoustics as well as the temporal dependency in the word sequence. Table 3 shows results of different variations of attention location, as shown in figure 1. We see that adding audio context to the last LSTM-LM layer leads to some degradation in performance. However, using audio context at final affine layer or at the first LSTM layer leads to similar results.

**Table 3:** Model size and Relative Word Error Rate Reduction (WERR) for different attention location

Attention location	# params (MM)	WERR
Final Affine (A3)	69	7.3%
First LSTM layer (A1)	71	6.6%
Last LSTM layer (A2)	71	6.1%
First LSTM + Final Affine	87	7.4%

#### 4.3. Learning context encoder for audio embeddings

In addition to attending to hidden layer outputs of the first pass LSTM-AM, we also experimented with applying a learnable context encoders on top of the LSTM AM outputs. Following various papers, we experimented with three different encoder types: a) 2-layer pyramidal LSTM (PyLSTM) [8] b) single layer time-delay Neural Network (TDNN) [23] with context  $\{-1, 2\}$  and c) convolutional neural network (CNN) with filter size 3x3 and non-overlapping max pooling of size 3 [24].

Table 4 shows the result of applying these different context encoders. We see that using Pyramidal LSTM and TDNNs for encoding the input improves the performance of the model. However, both TDNN and Pyramidal LSTM have a significant increase in number of parameters as well as training and inference time. CNN encoders are able to improve the performance over having no encoder. For all our other experiments, we chose to go with CNN encoder because of its fast training and inference speed.

**Table 4:** Model size and Relative Word Error Rate Reduction (WERR) for different attention location

Encoder type	# params (MM)	WERR
No encoder	68	7%
PyLSTM	84	7.9%
TDNN	83	8.0%
CNN	69	7.3%

## 5. CONCLUSION

In this paper, we proposed to use an attention-based language model for second-pass rescoring of n-best lists generated by a conventional ASR system. We show that by training the model with minimum word error rate (MWER) criteria, we can get upto 4.5% word error rate improvement over the the first pass system. Further, we show that the attention-based model can improve the word error rate even when the amount of training data is less and no task-specific audio encoder is learnt. We also show that pre-training the word embeddings and LSTM layers of the model can improve the performance of the model to 8% WERR. Future work will compare other end-to-end systems with this approach as well as extending this work to very low-data regime such as few-shot learning.

## References

- [1] G Hinton, L Deng, D Yu, G Dahl, A.-r Mohamed, N Jaitly, A Senior, V Vanhoucke, P Nguyen, B Kingsbury, et al., “Deep neural networks for acoustic modeling in speech recognition,” *IEEE Signal processing magazine*, vol. 29, 2012.
- [2] T Mikolov, M Karafiát, L Burget, J Černocký, and S Khudanpur, “Recurrent neural network based language model,” in *Eleventh annual conference of the international speech communication association*, 2010.
- [3] M Sundermeyer, R Schlüter, and H Ney, “Lstm neural networks for language modeling,” in *Thirteenth annual conference of the international speech communication association*, 2012.
- [4] X Liu, Y Wang, X Chen, M. J Gales, and P. C Woodland, “Efficient lattice rescoring using recurrent neural network language models,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4908–4912.
- [5] S Kumar, M Nirschl, D Holtmann-Rice, H Liao, A. T Suresh, and F Yu, “Lattice rescoring strategies for long short term memory language models in speech recognition,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 165–172.
- [6] A Raju, D Filimonov, G Tiwari, G Lan, and A Rastrow, “Scalable multi corpora neural language models for asr,” *Interspeech*, 2019.
- [7] A Graves, A.-r Mohamed, and G Hinton, “Speech recognition with deep recurrent neural networks,” in *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2013, pp. 6645–6649.
- [8] W Chan, N Jaitly, Q. V Le, and O Vinyals, “Listen, attend and spell,” *arXiv preprint arXiv:1508.01211*, 2015.
- [9] C.-C Chiu, T. N Sainath, Y Wu, R Prabhavalkar, P Nguyen, Z Chen, A Kannan, R. J Weiss, K Rao, E Gonnina, et al., “State-of-the-art speech recognition with sequence-to-sequence models,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4774–4778.
- [10] R Prabhavalkar, K Rao, T. N Sainath, B Li, L Johnson, and N Jaitly, “A comparison of sequence-to-sequence models for speech recognition.” in *Interspeech*, 2017, pp. 939–943.
- [11] A Sriram, H Jun, S Satheesh, and A Coates, “Cold fusion: Training seq2seq models together with language models,” *CoRR*, 2017.
- [12] A Kannan, Y Wu, P Nguyen, T. N Sainath, Z Chen, and R Prabhavalkar, “An analysis of incorporating an external language model into a sequence-to-sequence model,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 1–5828.
- [13] D Zhao, T. N Sainath, D Rybach, D Bhatia, B Li, and R Pang, “Shallow-fusion end-to-end contextual biasing,” *Submitted to Interspeech*, vol. 2019, 2019.
- [14] G Pundak, T. N Sainath, R Prabhavalkar, A Kannan, and D Zhao, “Deep context: end-to-end contextual speech recognition,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 418–425.
- [15] A Bruguier, G Pundak, R Prabhavalkar, and T Sainath, “Phoebe: Pronunciation-aware contextualization for end-to-end speech recognition,” 2019.
- [16] T Hori, C Hori, S Watanabe, and J. R Hershey, “Minimum word error training of long short-term memory recurrent neural network language models for speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5990–5994.
- [17] J Guo, T. N Sainath, and R. J Weiss, “A spelling correction model for end-to-end speech recognition,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5651–5655.
- [18] C.-C Chiu, D Rybach, I McGraw, M Visontai, Q Liang, R Prabhavalkar, R Pang, T Sainath, T Strohman, W Li, Y. R He, and Y Wu, “Two-pass end-to-end speech recognition,” in *Interspeech*, 2019.
- [19] S Hochreiter and J Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [20] B Li, T. N Sainath, A Narayanan, J Caroselli, M Bacchiani, A Misra, I Shafran, H Sak, G Pundak, K. K Chin, et al., “Acoustic modeling for google home.” 2017.
- [21] K Veselý, A Ghoshal, L Burget, and D Povey, “Sequence-discriminative training of deep neural networks.” 2013.
- [22] R Kneser and H Ney, “Improved backing-off for n-gram language modeling,” in *1995 International Conference on Acoustics, Speech, and Signal Processing*, 1995, vol. 1, pp. 181–184.
- [23] V Peddinti, D Povey, and S Khudanpur, “A time delay neural network architecture for efficient modeling of long temporal contexts,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [24] T. N Sainath, O Vinyals, A Senior, and H Sak, “Convolutional, long short-term memory, fully connected deep neural networks,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4580–4584.