

# INTERNAL LANGUAGE MODEL TRAINING FOR DOMAIN-ADAPTIVE END-TO-END SPEECH RECOGNITION

Zhong Meng, Naoyuki Kanda, Yashesh Gaur, Sarangarajan Parthasarathy, Eric Sun,  
Liang Lu, Xie Chen, Jinyu Li, Yifan Gong

Microsoft Corporation, Redmond, WA, USA

## ABSTRACT

The efficacy of external language model (LM) integration with existing end-to-end (E2E) automatic speech recognition (ASR) systems can be improved significantly using the internal language model estimation (ILME) method [1]. In this method, the internal LM score is subtracted from the score obtained by interpolating the E2E score with the external LM score, during inference. To improve the ILME-based inference, we propose an internal LM training (ILMT) method to minimize an additional internal LM loss by updating only the E2E model components that affect the internal LM estimation. ILMT encourages the E2E model to form a standalone LM inside its existing components, without sacrificing ASR accuracy. After ILMT, the more modular E2E model with matched training and inference criteria enables a more thorough elimination of the source-domain internal LM, and therefore leads to a more effective integration of the target-domain external LM. Experimented with 30K-hour trained recurrent neural network transducer and attention-based encoder-decoder models, ILMT with ILME-based inference achieves up to 31.5% and 11.4% relative word error rate reductions from standard E2E training with Shallow Fusion on out-of-domain LibriSpeech and in-domain Microsoft production test sets, respectively.

**Index Terms**— Speech recognition, language model, recurrent neural network transducer, attention-based encoder-decoder

## 1. INTRODUCTION

End-to-end (E2E) automatic speech recognition (ASR) has achieved state-of-the-art performance by directly mapping speech to word sequences through a single neural network. The most popular E2E models include connectionist temporal classification [2, 3, 4, 5], recurrent neural network transducer (RNN-T) [6, 7, 8, 9], and attention-based encoder-decoder (AED) models [10, 11, 12, 13, 14]. However, E2E models tend to overfit the audio-transcript pairs in source-domain training data, and suffer from performance degradation when evaluated in a mismatched target domain. Numerous ideas have been explored to adapt ASR models, such as regularization method [15, 16, 17, 18, 19], teacher-student learning [20, 21, 22, 23], transformation method [24, 25, 26], and adversarial learning [27, 28, 29, 30]. Nevertheless, all these methods require audio data for adaptation when applied to E2E models [31, 32, 33]. One promising solution without using audio is to train a language model (LM) with a large amount of text readily available in the target domain and fuse it with the E2E model during inference. However, an E2E model does not have a modular LM component as in a traditional hybrid system [34], making external LM integration a challenging task.

Among many approaches proposed for LM integration, Shallow Fusion [3, 35, 36, 37] is a simple yet effective method where the log probabilities of the E2E model and the LM are linearly interpolated during the inference. Towards a better integration, the Density Ratio

method [38, 39] subtracts the source-domain LM score from the interpolated score of Shallow Fusion, and shows improved performance. Further, as a new type of E2E model, the hybrid autoregressive transducer (HAT) was proposed in [40] to preserve the modularity of a traditional hybrid system. HAT allows us to estimate the internal LM scores and subtract them from the Shallow Fusion scores for external LM integration.

More recently, we proposed an internal LM estimation (ILME) method in [1] to facilitate the integration of an external LM for any pre-existing E2E models, including RNN-T and AED models, without any additional training. With ILME-based inference, the internal LM score of an E2E model is estimated by eliminating the contribution of the acoustic encoder, and is then subtracted from the log-linear interpolation between E2E and external LM scores. However, ILME-based inference in [1] is performed with an ASR model trained to optimize a standard E2E loss by updating all model parameters. The accuracy of the internal LM estimation is not guaranteed when the E2E model is not structured in a way that strictly satisfies the conditions of Proposition 1 in [40, Appendix A].

To compensate for the mismatch between the E2E training and the ILME-based inference, we propose an *internal LM training (ILMT)* of the E2E model to minimize an additional internal LM loss by updating only the model components engaged in the prediction of internal LM scores during inference. ILMT facilitates the E2E model to form a standalone LM inside its existing components while maintaining ASR accuracy. ILMT improves the effectiveness of the ILME-based LM integration with a more modular E2E model and a well-aligned training and inference criterion. Evaluated with 30 thousand (K)-hour trained RNN-T and AED models, ILMT with ILME-based inference achieves up to 31.5% and 11.4% relative word error rate (WER) reductions from Shallow Fusion on cross- and intra-domain evaluations, respectively, far outperforming the reductions with standard E2E training.

## 2. RELATED E2E METHODS

An E2E model predicts the conditional distribution  $P(\mathbf{Y}|\mathbf{X}; \theta_{\text{E2E}})$  of token sequences  $\mathbf{Y} = \{y_1, \dots, y_U\}$  given a speech-feature sequence  $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$  as the input, where  $y_u \in \mathcal{V}$  and  $\mathbf{x}_t$  is a feature vector at time  $t$ .  $\mathcal{V}$  is the set of all possible output tokens, i.e., word pieces. We insert a start-of-sentence token  $y_0 = \langle \text{sos} \rangle$  at the beginning of  $\mathbf{Y}$ .

### 2.1. Recurrent Neural Network Transducer (RNN-T)

The RNN-T model [6] comprises an encoder, a prediction network and a joint network. The encoder maps the input speech features  $\mathbf{X}$  to a sequence of hidden states  $\mathbf{H}^{\text{enc}} = \{\mathbf{h}_1^{\text{enc}}, \dots, \mathbf{h}_T^{\text{enc}}\}$ . The prediction network is an RNN that takes the embedding vector  $\mathbf{e}_{u-1}$  of the

previous *non-blank* token  $y_{u-1}$  and generates the hidden state  $\mathbf{h}_u^{\text{pred}}$ , i.e.,  $\mathbf{h}_u^{\text{pred}} = \text{PredictionRNN}(\mathbf{h}_{u-1}^{\text{pred}}, \mathbf{e}_{u-1})$ .

The joint network is a feed-forward network that combines the outputs of the encoder and prediction network to predict the conditional distribution over the next possible token  $\tilde{y}_i \in \mathcal{V} \cup \langle b \rangle$ , i.e.,

$$\mathbf{z}_{t_i, u_i} = W_j \phi(W_e \mathbf{h}_{t_i}^{\text{enc}} + W_p \mathbf{h}_{u_i}^{\text{pred}} + \mathbf{b}_e + \mathbf{b}_p) + \mathbf{b}_j, \quad (1)$$

$$[P(\tilde{y}_i = v | \mathbf{X}_{1:t_i}, \mathbf{Y}_{0:u_{i-1}}; \theta_{\text{RNN-T}})]_{v \in \mathcal{V} \cup \langle b \rangle} = \text{softmax}(\mathbf{z}_{t_i, u_i}), \quad (2)$$

where  $\langle b \rangle$  denotes a blank symbol,  $\phi$  is a non-linear function, e.g., tanh or ReLU.  $W_j, W_e, W_p$  are weight matrices, and  $\mathbf{b}_e, \mathbf{b}_p, \mathbf{b}_j$  are biases.  $\mathbf{z}_{t_i, u_i}$  is a  $|\mathcal{V}| + 1$  dimensional logit vector.  $\tilde{y}_i$  forms a blank-augmented token sequences  $\tilde{\mathbf{Y}} = \{\tilde{y}_1, \dots, \tilde{y}_{T+U}\}$  aligned with the token and feature sequences  $\mathbf{Y}$  and  $\mathbf{X}$  as  $(\tilde{y}_i, y_{u_i}, x_{t_i})_{i=1}^{T+U}$ , i.e., the index  $i$  in  $\tilde{\mathbf{Y}}$  is mapped to the index  $u_i$  in  $\mathbf{Y}$ , and the index  $t_i$  in  $\mathbf{X}$ .

The RNN-T loss is computed by marginalizing over all possible blank-augmented token sequences aligned with each reference  $\mathbf{Y}$ , i.e.,  $\mathcal{A}(\mathbf{X}, \mathbf{Y})$ , on the training corpus.

$$\begin{aligned} \mathcal{L}_{\text{RNN-T}}(\theta_{\text{RNN-T}}) \\ = - \sum_{(\mathbf{X}, \mathbf{Y}) \in \mathcal{D}} \log \sum_{\tilde{\mathbf{Y}} \in \mathcal{A}(\mathbf{X}, \mathbf{Y})} \prod_{i=1}^{T+U} P(\tilde{y}_i | \mathbf{X}_{1:t_i}, \mathbf{Y}_{0:u_{i-1}}; \theta_{\text{RNN-T}}). \end{aligned} \quad (3)$$

## 2.2. Attention-Based Encoder-Decoder (AED)

The AED model [10] consists of an encoder, a decoder and an attention network. The encoder maps a sequence of input speech frames  $\mathbf{X}$  into a sequence of hidden states  $\mathbf{H}^{\text{enc}}$ . The attention network generates an attention weight for  $\mathbf{h}_t^{\text{enc}}$  at each decoder step  $u$ , determining which encoder states should be attended to predict the output label  $y_u$ , i.e.,  $\mathbf{a}_u = \text{AttentionNet}(\mathbf{a}_{u-1}, \mathbf{h}_t^{\text{enc}}, \mathbf{h}_u^{\text{dec}})$ , where  $\mathbf{a}_u$  is a vector of attention weights of dimension  $T$ , and  $\mathbf{h}_u^{\text{dec}}$  is the decoder hidden state. The context vector  $\mathbf{c}_u$  is computed as a linear combination of  $\mathbf{H}^{\text{enc}}$  weighted by the attention, i.e.,  $\mathbf{c}_u = \sum_{t=1}^T a_{u,t} \mathbf{h}_t^{\text{enc}}$ .

At each step  $u$ , the decoder RNN takes the sum of the previous token embedding  $\mathbf{e}_{u-1}$  and the context vector  $\mathbf{c}_{u-1}$  as the input to predict the conditional distribution over  $\mathcal{V} \cup \langle \text{eos} \rangle$ , i.e.,

$$\mathbf{h}_u^{\text{dec}} = \text{DecoderRNN}(\mathbf{h}_{u-1}^{\text{dec}}, \mathbf{e}_{u-1} + \mathbf{c}_{u-1}), \quad (4)$$

$$\mathbf{z}_u = W_d \mathbf{h}_u^{\text{dec}} + \mathbf{b}_d, \quad (5)$$

$$[P(y_u = v | \mathbf{X}, \mathbf{Y}_{0:u-1}; \theta_{\text{AED}})]_{v \in \mathcal{V} \cup \langle \text{eos} \rangle} = \text{softmax}(\mathbf{z}_u), \quad (6)$$

where  $\langle \text{eos} \rangle$  is the end-to-sentence token,  $W_d$  and  $\mathbf{b}_d$  are weight matrix and bias, respectively.

The AED loss is obtained as a summation of token sequence posteriors over the training corpus  $\mathcal{D}$  as follows

$$\mathcal{L}_{\text{AED}}(\theta_{\text{AED}}) = \sum_{(\mathbf{X}, \mathbf{Y}) \in \mathcal{D}} \sum_{u=1}^{U+1} \log P(y_u | \mathbf{X}, \mathbf{Y}_{0:u-1}; \theta_{\text{AED}}). \quad (7)$$

## 3. INTERNAL LM ESTIMATION (ILME)

From audio-transcript training pairs, an E2E model implicitly learns an internal language model (LM) that characterizes the distribution of source-domain training text. The exact computation of internal LM is intractable, but it can be approximated by Proposition 1 in [40, Appendix A] which suggests that the E2E internal LM  $P(y_u | \mathbf{Y}_{0:u-1}; \theta_{\text{E2E}}^S)$  approximately equals to E2E model output  $\text{softmax}[J(\mathbf{g}_u)]$  after zeroing out the acoustic embedding  $\mathbf{f}_t$  if  $P(y_u | \mathbf{X}, \mathbf{Y}_{0:u-1}; \theta_{\text{E2E}}^S) = \text{softmax}[J(\mathbf{f}_t + \mathbf{g}_u)]$  and  $J(\mathbf{f}_t + \mathbf{g}_u) \approx J(\mathbf{f}_t) + J(\mathbf{g}_u)$  are satisfied, where  $\mathbf{g}_u$  is a language embedding.

As shown in [1], the conditional probability of RNN-T internal LM,  $P(y_u | \mathbf{Y}_{0:u-1}; \theta_{\text{RNN-T}})$ , is estimated as a softmax normalization of the non-blank token logits when the hidden states of the encoder are eliminated from the input of the joint network.

$$\mathbf{z}_u^{\text{ILM}} = W_j \phi(W_p \mathbf{h}_u^{\text{pred}} + \mathbf{b}_p) + \mathbf{b}_j, \quad (8)$$

$$P(y_u | \mathbf{Y}_{0:u-1}; \theta_{\text{RNN-T}}) = \text{softmax}(\mathbf{z}_u^{\text{ILM, NB}}), \quad (9)$$

where  $\mathbf{z}_u^{\text{ILM}}$  is a  $(|\mathcal{V}| + 1)$ -dimensional vector with a designated logit for the blank token  $\langle b \rangle$ ,  $\mathbf{z}_u^{\text{ILM, NB}}$  is a logit vector of dimension  $|\mathcal{V}|$  created by taking out the blank logit from  $\mathbf{z}_u^{\text{ILM}}$ . Without the encoder input, the RNN-T is completely driven by the prediction and joint networks with the token sequence  $\mathbf{Y}$  as the only input.

Similarly, [1] has also shown that the conditional probability of the AED internal LM is estimated by the decoder output after zeroing out the context vector, i.e.,

$$\begin{aligned} P(y_u | \mathbf{Y}_{0:u-1}; \theta_{\text{AED}}) \\ = \text{softmax} \left[ W_d \cdot \text{DecoderRNN}(\mathbf{h}_{u-1}^{\text{dec}}, \mathbf{e}_{u-1}) + \mathbf{b}_d \right]. \end{aligned} \quad (10)$$

Without the context vector, AED is entirely driven by the decoder with the token sequence  $\mathbf{Y}$  as the only input, acting exactly the same as an RNN-LM.

During ILME-based inference [1], we subtract the log of the internal LM probability  $P(\mathbf{Y}; \theta_{\text{E2E}})$  from the log-linear combination between the conditional probability of E2E model and the external LM probability  $P(\mathbf{Y}; \theta_{\text{LM}})$ , and search for the optimal token sequence  $\hat{\mathbf{Y}}$  as follows via a left-to-right beam search.

$$\begin{aligned} \hat{\mathbf{Y}} = \arg \max_{\mathbf{Y}} [\log P(\mathbf{Y} | \mathbf{X}; \theta_{\text{E2E}}) + \lambda_E \log P(\mathbf{Y}; \theta_{\text{LM}}) \\ - \lambda_I \log P(\mathbf{Y}; \theta_{\text{E2E}})], \end{aligned} \quad (11)$$

where  $\lambda_E$  and  $\lambda_I$  are the weights for external and internal LMs, respectively.

## 4. INTERNAL LM TRAINING OF E2E MODELS

In standard E2E training, an internal LM is implicitly learned to minimize the E2E loss by updating all parameters of the E2E model. However, during ILME-based inference, only a part of the E2E model contributes to the prediction of the internal LM scores. The estimation of the internal LM scores is not accurate when the conditions of Proposition 1 in [40, Appendix A] is not strictly satisfied by the E2E model.

In this work, we propose an *internal LM training* of the E2E model to mitigate the mismatch between the E2E training and the ILME-based inference. Through the standard E2E training, the decoder of an AED or the prediction and joint networks of an RNN-T acts as an acoustically-conditioned LM that takes both the token and acoustic embeddings as the input to predict the conditional probability of the next token. From Eqs. (9) and (10), the internal LM scores are estimated entirely by the acoustically-conditioned LM of an E2E model during ILME-based inference. Therefore, the goal of ILMT is to encourage the acoustically-conditioned LM of an E2E model to also behave like a standalone internal LM, without sacrificing ASR accuracy. To achieve that, we jointly minimize an internal LM loss together with the standard E2E loss during ILMT.

The internal LM loss of an RNN-T model is obtained by summing up the negative log probabilities of the internal LM over the training corpus  $\mathcal{D}$  as follows

$$\mathcal{L}_{\text{ILM}}(\theta_{\text{pred}}, \theta_{\text{joint}}) = - \sum_{\mathbf{Y} \in \mathcal{D}} \sum_{u=1}^U \log P(y_u | \mathbf{Y}_{0:u-1}; \theta_{\text{pred}}, \theta_{\text{joint}}). \quad (12)$$

Note that, from Eqs. (8) and (9), the RNN-T internal LM loss is conditioned only on the parameters of the prediction and joint networks,  $\theta_{\text{pred}}$  and  $\theta_{\text{joint}}$ . For RNN-T, the ILMT loss is constructed as a weighted sum of the RNN-T loss in Eq. (3) and the ILM loss below

$$\mathcal{L}_{\text{ILMT}}(\theta_{\text{RNN-T}}) = \mathcal{L}_{\text{RNN-T}}(\theta_{\text{RNN-T}}) + \alpha \mathcal{L}_{\text{ILM}}(\theta_{\text{pred}}, \theta_{\text{joint}}), \quad (13)$$

where  $\alpha$  is the weight of the internal LM loss. By minimizing the RNN-T ILMT loss, we maximize the internal LM probability of the E2E training transcripts by updating only the prediction and joint networks while maximizing the conditional probability of the training transcripts given input speech by updating the entire RNN-T.

The internal LM loss of AED is formulated as a summation of negative log probabilities of the internal LM over training corpus  $\mathcal{D}$

$$\mathcal{L}_{\text{ILM}}(\theta_{\text{dec}}) = - \sum_{\mathbf{Y} \in \mathcal{D}} \sum_{u=1}^{U+1} \log P(y_u | \mathbf{Y}_{0:u-1}; \theta_{\text{dec}}). \quad (14)$$

Note that, from Eq. (10), the AED internal LM loss is conditioned only on the parameters of the decoder  $\theta_{\text{dec}}$ . For AED, the ILMT loss is computed as a weighted sum of the AED loss and the ILM loss below

$$\mathcal{L}_{\text{ILMT}}(\theta_{\text{AED}}) = \mathcal{L}_{\text{AED}}(\theta_{\text{AED}}) + \alpha \mathcal{L}_{\text{ILM}}(\theta_{\text{dec}}) \quad (15)$$

By minimizing the AED ILMT loss, we maximize the internal LM probability of the E2E training transcripts by updating only the AED decoder while maximizing the conditional probability of the training transcripts given input speech by updating the entire AED model.

The procedure of ILMT with the ILME-based inference for the LM integration with an E2E model is the following

1. Train an E2E model with source-domain audio-transcript pairs to minimize the ILMT loss in Eq. (13) for RNN-T or in Eq. (15) for AED.
2. Train an external LM with target-domain text-only data.
3. Integrate the ILMT E2E model in Step 1 with the external LM in Step 2 by performing ILME-based inference in Section 3.

With ILMT, a standalone internal LM with a significantly lower perplexity is learned only by the E2E components used to compute the internal LM scores during the ILME-based inference. With increased modularity, the E2E model is more adaptable to the target domain with its increased flexibility to eradicate the effect of the source-domain internal LM through the ILME-based inference.

## 5. EXPERIMENT

In this work, we perform ILMT of RNN-T and AED models and integrate them with external long short-term memory (LSTM) [41, 42, 43] LMs using different methods. We conduct both cross-domain and intra-domain evaluations to investigate the effectiveness of ILMT. Same as [1], we perform beam search inference with a beam size of 25 for all evaluations, and use the 3999 word-piece units generated by byte-pair encoding [44] as  $\mathcal{V}$  for both E2E models and LSTM-LMs.

### 5.1. Internal LM Training

We perform ILMT of the E2E models with the same 30K hours of anonymized and transcribed data as in [1] collected from Microsoft services, including voice search, short message dictation, command and control, and conversations recorded in various conditions.

The RNN-T model is initialized with the parameters of the RNN-T in [1] which was well-trained until full convergence with the

30K-hour data. The encoder and prediction networks are both uni-directional LSTMs with 6 and 2 hidden layers, respectively, and 1024 hidden units in each layer. The joint network has 4000-dimensional output units. The RNN-T has 76M parameters. During ILMT, the weight of the internal LM loss is set to 0.4. The internal LM perplexities of ILMT RNN-T and the standard RNN-T in [1] are 52.0 and 99.4, respectively, on the validation set of 30K-hour data.

The AED model [10, 45, 46] is randomly initialized and shares the same architecture as the one in [1]. The encoder is a bi-directional LSTM with 6 hidden layers and 780 hidden units in each layer. The decoder is a uni-directional LSTM with 2 hidden layers, each with 1280 hidden units. The decoder has 4000-dimensional output units. The AED model has 97M parameters. During ILMT, the weight of the internal LM loss is set to 1.0. The internal LM perplexities of ILMT AED and the standard AED in [1] are 46.1 and 796.7, respectively, on the validation set of the 30K-hour training data.

### 5.2. Cross-Domain Evaluation

We evaluate a 30K-hour E2E model on the LibriSpeech clean test set by integrating an LSTM-LM trained with LibriSpeech text. Collected from read English based on audio book, the LibriSpeech corpus [47] is outside the domains covered by the 30K-hour training speech. The test-clean and dev-clean sets consist of 2620 and 2703 utterances, respectively. We tune the LM weights on dev-clean. We use the same external LSTM-LM as in [1] trained with the transcript of 960K-hour training speech and the additional 813M-word text from LibriSpeech corpus. With 58M parameters, the LSTM-LM has 2 hidden layers with 2048 hidden units for each layer. For Density Ratio, we use the same source-domain LSTM-LM with 2 hidden layers, 2048 hidden units, and 57M parameters as in [1] trained using the transcript of 30K-hour speech.

We list the results of RNN-T in Table 1 with an excerpt of standard E2E training results from [1]. With ILMT, all three LM integration methods show 27.9%-40.9% relative WER reductions from the baseline with standard RNN-T training and inference, significantly larger than the corresponding reductions without ILMT in the range of 16.1%-29.1%. ILMT with ILME inference performs the best achieving 29.6% and 16.6% relative WER reduction from the standard RNN-T with Shallow Fusion and ILME inference, respectively. As in Table 2, the AED results are similar to RNN-T. ILMT with ILME inference performs the best, achieving 57.6%, 31.5% and 25.1% relative WER reductions from the standard AED training with AED inference, Shallow Fusion and ILME inference, respectively.

### 5.3. Intra-Domain Evaluation

#### 5.3.1. Dictation Test Set

We evaluate a 30K-hour E2E model on a dictation test set by integrating a strong external LSTM-LM trained with a large amount of multi-domain text. As in [1], we use the same 2K in-house dictation utterances collected from the keyboard input as the test set, and the same 442 email dictation utterances as the validation set. The test set has a similar style as the dictation data in 30K-hour corpus and is thus considered as in-domain evaluation. We use the same multi-domain LSTM-LM as in [1] trained with 2 billion (B) words of text comprising short message dictation and conversational data such as talks, interviews, and meeting transcripts.

Table 1 lists the RNN-T results with an excerpt of standard E2E training results from [1]. With ILMT, all three LM integration methods show 6.9%-13.6% relative WER reductions from the baseline with standard RNN-T training and inference, significantly larger than

Train Loss	Evaluation Method	Model Params	LibriSpeech			Dictation			Conversation		
			Dev WER	Test WER	Test WERR	Dev WER	Test WER	Test WERR	Dev WER	Test WER	Test WERR
$\mathcal{L}_{\text{RNN-T}}$	No LM	76M	9.27	8.97	-	23.40	16.16	-	14.92	14.26	-
	Shallow Fusion	134M	7.44	7.53	16.1	22.19	15.77	2.4	14.88	14.08	1.3
	Density Ratio	191M	6.80	6.74	24.9	21.54	15.64	3.2	14.76	14.20	0.4
	ILME	134M	6.41	6.36	29.1	21.04	14.70	9.0	14.61	14.03	1.6
$\mathcal{L}_{\text{ILMT}}$	No LM	76M	8.58	8.37	6.7	22.61	15.73	2.7	14.00	13.59	4.7
	Shallow Fusion	134M	6.60	6.47	27.9	21.31	15.04	6.9	13.83	13.29	6.8
	Density Ratio	191M	5.86	5.61	37.5	20.61	14.76	8.7	13.71	13.29	6.8
	ILME	134M	<b>5.57</b>	<b>5.30</b>	<b>40.9</b>	<b>19.94</b>	<b>13.97</b>	<b>13.6</b>	<b>13.32</b>	<b>12.96</b>	<b>9.1</b>

**Table 1.** WERs (%) of 30k-hour **RNN-T models** trained with RNN-T or ILMT loss, and evaluated with different LM integration methods on **out-of-domain** LibriSpeech, **in-domain** dictation, and **in-domain** conversation dev/test sets. WERR is relative WER reduction.

Train Loss	Evaluation Method	Model Params	LibriSpeech			Dictation			Conversation		
			Dev WER	Test WER	Test WERR	Dev WER	Test WER	Test WERR	Dev WER	Test WER	Test WERR
$\mathcal{L}_{\text{AED}}$	No LM	97M	8.56	8.61	-	20.17	14.08	-	14.05	13.43	-
	Shallow Fusion	155M	5.00	5.33	38.1	18.55	12.96	8.0	13.45	12.95	3.6
	Density Ratio	212M	4.74	5.09	40.9	18.76	12.89	8.5	13.55	12.95	3.6
	ILME	155M	4.42	4.87	43.4	18.26	12.36	12.2	13.33	12.67	5.7
$\mathcal{L}_{\text{ILMT}}$	No LM	97M	7.31	7.47	13.2	21.06	13.72	2.6	12.60	12.19	9.2
	Shallow Fusion	155M	6.54	6.61	23.2	19.09	12.32	12.5	12.42	11.90	11.4
	Density Ratio	212M	4.28	4.85	43.7	18.30	12.23	13.1	12.23	11.85	11.8
	ILME	155M	<b>3.30</b>	<b>3.65</b>	<b>57.6</b>	<b>17.00</b>	<b>11.60</b>	<b>17.6</b>	<b>12.11</b>	<b>11.58</b>	<b>13.8</b>

**Table 2.** WERs (%) of 30k-hour **AED models** trained with AED or ILMT loss, and evaluated with different LM integration methods on **out-of-domain** LibriSpeech, **in-domain** dictation, and **in-domain** conversation dev/test sets. WERR is relative WER reduction.

the corresponding reductions without ILMT in the range of 2.4%-9.0%. ILMT with ILME inference performs the best achieving 11.4% and 5.0% relative WER reduction from the standard RNN-T training with Shallow Fusion and ILME inference, respectively. As in Table 2, the AED results are similar to RNN-T. ILMT with ILME inference performs the best, achieving 17.6%, 10.5% and 6.1% relative WER reductions from the standard AED training with AED inference, Shallow Fusion and ILME inference, respectively.

### 5.3.2. Conversation Test Set

We evaluate a 30K-hour E2E model on an conversation test set by integrating a strong multi-domain external LSTM-LM. From the Microsoft telecommunication applications, we collect 2560 in-house conversational utterances as the test set, and another 1280 conversational utterances as the validation set. The test set has a similar style as the conversational data in 30K-hour corpus and is thus considered as in-domain evaluation. For the external LM, we use the same 2B-word LSTM-LM as in Section 5.3.1.

As shown in Table 1, with ILMT RNN-T, all three LM integration methods show 6.8%-9.1% relative WER reductions from the baseline with standard RNN-T training and inference, significantly larger than the corresponding reductions without ILMT in the range of 0.4%-1.6%. ILMT with ILME inference performs the best achieving 8.0% and 7.6% relative WER reduction from the standard RNN-T training with Shallow Fusion and ILME inference, respectively. As shown in Table 2, the AED results are similar to RNN-T. ILMT with ILME inference performs the best, achieving 13.8%, 10.6% and 8.6% relative WER reductions from the standard AED training with AED inference, Shallow Fusion and ILME inference, respectively.

## 5.4. Result Analysis

From the results, we have the following observations for both RNN-T and AED models, and for both cross- and intra-domain evalua-

tions. All LM integration methods consistently achieve remarkably lower WERs with ILMT than with standard E2E training. Among all methods, ILMT with ILME inference consistently performs the best, with 29.6%-31.5% and 8.0%-11.4% relative WER reductions from standard E2E training with Shallow Fusion for cross-domain and intra-domain evaluations, respectively. ILME inference consistently outperforms Density Ratio in terms of lower WER with ILMT or standard E2E training despite having 26.8%-29.8% fewer model parameters. With ILME inference, ILMT achieves 16.6%-25.1% and 5.0%-8.6% relative WER reductions from the standard E2E training for cross-domain and intra-domain evaluations, respectively. All of these manifest the advantage of ILMT over standard E2E training for ILME inference and other LM integration methods.

Note that, with or without ILMT, ILME inference is effective even for intra-domain evaluation because it replaces the weak E2E internal LM with a powerful external LM trained with orders of magnitude more multi-domain text than the E2E training transcript. All three LM fusion methods perform better for AED than RNN-T due to larger relative WER reductions from a stronger baseline. The internal LM perplexity of an E2E model is remarkably reduced by ILMT.

## 6. CONCLUSION

We propose an internal LM training of the E2E model which minimizes an internal LM loss in addition to the standard E2E loss to improve the effectiveness of ILME external LM integration. With ILMT, ILME inference achieves 29.6%-31.5% and 8.0%-11.4% relative WER reductions from the standard E2E training with Shallow Fusion for cross-domain and intra-domain evaluations, respectively. With ILME inference, ILMT outperforms the standard E2E training by 16.6%-25.1% and 5.0%-8.6% relatively in terms of lower WER for cross-domain and intra-domain evaluations, respectively.

## 7. REFERENCES

- [1] Z. Meng, S. Parthasarathy, E. Sun, et al., “Internal language model estimation for domain-adaptive end-to-end speech recognition,” in *Proc. SLT*. IEEE, 2021.
- [2] A. Graves, S. Fernández, F. Gomez, et al., “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *ICML*. ACM, 2006.
- [3] A. Hannun, C. Case, et al., “Deep speech: Scaling up end-to-end speech recognition,” *arXiv preprint arXiv:1412.5567*, 2014.
- [4] H. Soltau, H. Liao, and H. Sak, “Neural speech recognizer: Acoustic-to-word LSTM model for large vocabulary speech recognition,” *Proc. Interspeech*, 2016.
- [5] J. Li, G. Ye, A. Das, et al., “Advancing acoustic-to-word CTC model,” in *Proc. ICASSP*, 2018.
- [6] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [7] M. Jain, K. Schubert, J. Mahadeokar, et al., “RNN-T for latency controlled asr with improved beam search,” *arXiv preprint arXiv:1911.01629*, 2019.
- [8] T. Sainath, Y. He, B. Li, et al., “A streaming on-device end-to-end model surpassing server-side conventional model quality and latency,” in *Proc. ICASSP*, 2020, pp. 6059–6063.
- [9] J. Li, R. Zhao, Z. Meng, et al., “Developing RNN-T models surpassing high-performance hybrid models with customization capability,” in *Interspeech*, 2020.
- [10] J. K Chorowski, D. Bahdanau, D. Serdyuk, et al., “Attention-based models for speech recognition,” in *NIPS*, 2015.
- [11] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. ICASSP*. IEEE, 2016.
- [12] C.-C. Chiu, T. N Sainath, Y. Wu, et al., “State-of-the-art speech recognition with sequence-to-sequence models,” in *Proc. ICASSP*. IEEE, 2018, pp. 4774–4778.
- [13] S. Karita, N. Chen, et al., “A comparative study on transformer vs RNN in speech applications,” in *Proc. ASRU*, 2019.
- [14] J. Li, Y. Wu, Y. Gaur, et al., “On the comparison of popular end-to-end models for large scale speech recognition,” in *Proc. Interspeech*, 2020.
- [15] D. Yu, K. Yao, H. Su, et al., “Kl-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition,” in *Proc. ICASSP*, May 2013.
- [16] Z. Meng, J. Li, and Y. Gong, “Adversarial speaker adaptation,” in *Proc. ICASSP*, 2019.
- [17] H. Liao, “Speaker adaptation of context dependent deep neural networks,” in *Proc. ICASSP*, May 2013.
- [18] Z. Meng, H. Hu, J. Li, et al., “L-vector: Neural label embedding for domain adaptation,” in *Proc. ICASSP*. IEEE, 2020.
- [19] P. Swietojanski and S. Renals, “Learning hidden unit contributions for unsupervised speaker adaptation of neural network acoustic models,” in *Proc. SLT*. IEEE, 2014, pp. 171–176.
- [20] J. Li, R. Zhao, et al., “Learning small-size DNN with output-distribution-based criteria,” in *INTERSPEECH*, 2014.
- [21] Z. Meng, J. Li, Y. Gong, et al., “Adversarial teacher-student learning for unsupervised domain adaptation,” in *Proc. ICASSP*, 2018.
- [22] V. Manohar, P. Ghahremani, D. Povey, et al., “A teacher-student learning approach for unsupervised domain adaptation of sequence-trained asr models,” in *Proc. SLT*. IEEE, 2018.
- [23] Z. Meng, J. Li, Y. Zhao, and Y. Gong, “Conditional teacher-student learning,” in *Proc. ICASSP*, 2019.
- [24] R. Gemello, F. Mana, S. Scanzio, et al., “Linear hidden transformations for adaptation of hybrid ann/hmm models,” *Speech Communication*, vol. 49, no. 10, pp. 827 – 835, 2007.
- [25] T. Tan, Y. Qian, M. Yin, et al., “Cluster adaptive training for deep neural network,” in *Proc. ICASSP*. IEEE, 2015.
- [26] O. Abdel-Hamid and H. Jiang, “Fast speaker adaptation of hybrid nn/hmm model for speech recognition based on discriminative learning of speaker code,” in *Proc. ICASSP*, May 2013.
- [27] Y. Shinohara, “Adversarial multi-task learning of deep neural networks for robust speech recognition,” in *INTERSPEECH*, 2016.
- [28] Z. Meng, J. Li, Z. Chen, et al., “Speaker-invariant training via adversarial learning,” in *Proc. ICASSP*, 2018.
- [29] D. Serdyuk, K. Audhkhasi, et al., “Invariant representations for noisy speech recognition,” in *NIPS Workshop*, 2016.
- [30] Z. Meng, Z. Chen, V. Mazalov, J. Li, and Y. Gong, “Unsupervised adaptation with domain separation networks for robust speech recognition,” in *Proc. ASRU*, 2017.
- [31] T. Ochiai, S. Watanabe, et al., “Speaker adaptation for multi-channel end-to-end speech recognition,” in *Proc. ICASSP*, 2018.
- [32] Z. Meng, Y. Gaur, J. Li, et al., “Speaker adaptation for attention-based end-to-end speech recognition,” *Proc. Interspeech*, 2019.
- [33] Z. Meng, J. Li, Y. Gaur, and Y. Gong, “Domain adaptation via teacher-student learning for end-to-end speech recognition,” in *Proc. ASRU*. IEEE, 2019.
- [34] G. Hinton, L. Deng, D. Yu, et al., “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal Processing Magazine*, 2012.
- [35] C. Gulcehre, O. Firat, K. Xu, et al., “On using monolingual corpora in neural machine translation,” *CoRR*, 2015.
- [36] J. Chorowski and N. Jaitly, “Towards better decoding and language model integration in sequence to sequence models,” *CoRR*, vol. abs/1613.02695, 2016.
- [37] S. Kim, Y. Shangguan, J. Mahadeokar, et al., “Improved neural language model fusion for streaming recurrent neural network transducer,” *arXiv preprint arXiv:2010.13878*, 2020.
- [38] E. McDermott, H. Sak, and E. Variani, “A density ratio approach to language model fusion in end-to-end automatic speech recognition,” in *Proc. ASRU*. IEEE, 2019, pp. 434–441.
- [39] N. Kanda, X. Lu, and H. Kawai, “Maximum a posteriori based decoding for CTC acoustic models,” in *Interspeech*, 2016.
- [40] E. Variani, D. Rybach, C. Allauzen, et al., “Hybrid autoregressive transducer (hat),” *arXiv preprint arXiv:2003.07705*, 2020.
- [41] A. Senior H. Sak and F. Beaufays, “Long short-term memory recurrent neural network architectures for large scale acoustic modeling,” in *Interspeech*, 2014.
- [42] Z. Meng, S. Watanabe, J. R. Hershey, et al., “Deep long short-term memory adaptive beamforming networks for multichannel robust speech recognition,” in *ICASSP*. IEEE, 2017.
- [43] H. Erdogan, T. Hayashi, J. R. Hershey, et al., “Multi-channel speech recognition: LSTMs all the way through,” in *ChiME-4 workshop*, 2016, pp. 1–4.
- [44] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” *arXiv preprint arXiv:1508.07909*, 2015.
- [45] Z. Meng, Y. Gaur, J. Li, et al., “Character-aware attention-based end-to-end speech recognition,” in *Proc. ASRU*, 2019.
- [46] Y. Gaur, J. Li, Z. Meng, et al., “Acoustic-to-phrase end-to-end speech recognition,” in *INTERSPEECH*, 2019.
- [47] V. Panayotov, G. Chen, et al., “Librispeech: an asr corpus based on public domain audio books,” in *Proc. ICASSP*, 2015.