

# Improving RNN-Transducers with Acoustic LOOKAHEAD

Vinit S. Unni<sup>1</sup>, Ashish Mittal<sup>1,2</sup>, Preethi Jyothi<sup>1</sup>, Sunita Sarawagi<sup>1</sup>

<sup>1</sup>Indian Institute of Technology Bombay, India

<sup>2</sup>IBM Research, India

vinit@cse.iitb.ac.in, arakeshk@in.ibm.com, pjyothi@cse.iitb.ac.in, sunita@cse.iitb.ac.in

## Abstract

RNN-Transducers (RNN-Ts) have gained widespread acceptance as an end-to-end model for speech to text conversion because of their high accuracy and streaming capabilities. A typical RNN-T independently encodes the input audio and the text context, and combines the two encodings by a thin joint network. While this architecture provides SOTA streaming accuracy, it also makes the model vulnerable to strong LM biasing which manifests as multi-step hallucination of text without acoustic evidence. In this paper we propose LOOKAHEAD that makes text representations more acoustically grounded by looking ahead into the future within the audio input. This technique yields a significant 5% – 20% relative reduction in word error rate on both in-domain and out-of-domain evaluation sets.

**Index Terms:** speech recognition, RNN transducer, acoustic hallucinations

## 1. Introduction

RNN-Transducers (RNN-Ts) [1] are the predominant choice for end-to-end automatic speech recognition (ASR) offering both high accuracy and streaming capabilities [2, 3]. They comprise a speech encoder that can process speech to generate an acoustic representation and a text encoder that is conditioned on label outputs from previous time-steps to generate a textual representation. Both the acoustic and textual representations are further combined by a simple joint network to predict the final output sequence. Apart from making the model streaming-friendly, separate speech and text modules in RNN-Ts also allow for text-only data to be used in training the text encoder [4].

While having separate speech and text modules in RNN-Ts has its benefits, it also makes the model vulnerable to strong biases from the language model. Driven by strong textual priors, the representation from the text encoder could be very biased towards an output unit that is eventually adopted by the joint network but does not have any acoustic correlates in the speech input. Such outputs could be considered hallucinations that arise due to the overconfidence of the language model in the RNN-T [5]. This problem is more severe when the RNN-T is used to decode out-of-domain utterances. Apart from the more egregious hallucination errors, we also find that language model biases in RNN-Ts lead to word-boundary errors. For example, “villeroy took” is mispredicted by an RNN-T baseline as “villar I took”.

Hallucinated outputs have been studied a lot more in the context of neural machine translation where the decoder language model hallucinates content that is not aligned to the source sentence [6]. In contrast, the problem of hallucination in RNN-Ts, that stems from its very design, has been far less studied and demands more attention.

In this work, we propose LOOKAHEAD as a fix for the problem of hallucinations in RNN-Ts. LOOKAHEAD aims to make the textual representations more acoustically grounded by *looking ahead* into the future within the speech signal. To achieve such a lookahead without interfering with the RNN-T’s online decoding capabilities, we extract a limited number of lookahead output tokens for each frame of the input speech using only the acoustic encoder and further use these extracted tokens to modify the textual representation. This technique yields significant reductions in word error rates (WERs) on the established Librispeech benchmark and a variety of out-of-domain evaluation sets. We also show that beyond improving WERs, LOOKAHEAD results in predictions that are more acoustically faithful to the speech input. For example, the reference “la valliere” is misrecognized as “the valet” by an RNN-T baseline, while an RNN-T baseline with LOOKAHEAD predicts “lavalier”.

**Contributions:** Thus, overall the contributions of this paper are as follows: (1) We highlight the problem of hallucination in SOTA online ASR models and attribute it to the speech independent encoding of text representations. (2) We propose a fix based on enriching text representations with a lookahead of future tokens extracted from the audio. (3) We present a simple extension of the RNN-T architecture called LOOKAHEAD with very modest computational overheads. (4) We present an evaluation on three benchmarks on various settings of model sizes and show that our proposal improves WER and reduces hallucination significantly.

## 2. Background: RNN Transducer

Let the input audio be denoted by  $\mathbf{x} = \{x_1, x_2, x_3, \dots, x_T\}$  where each  $x_t$  represents the acoustic features at time  $t$ . Let the corresponding text transcript be  $\mathbf{y} = \{y_1, y_2, y_3, \dots, y_U\}$  where  $y_u \in \mathcal{V}$  denotes the  $u^{th}$  output token drawn from a vocabulary  $\mathcal{V}$ . One of the most distinguishable features of an RNN-T is the presence of two separate encoders for text and acoustic signals respectively. The acoustic encoder (AE) takes as input  $\mathbf{x} = \{x_1, x_2, x_3, \dots, x_T\}$  and generates the acoustic encoded representation  $\mathbf{h} = \{h_1, h_2, h_3, \dots, h_T\}$ . The text or language encoder (LE) generates representation  $g_u$  appropriate for the next output token as a function of previous tokens  $\mathbf{y}_{<u} = y_1 \dots, y_{u-1}$

$$g_u = \text{LE}(\mathbf{y}_{<u}) \quad h_t = \text{AE}(\mathbf{x}, t)$$

A *Joint Network* (JN) combines the two encodings for each  $t \in [1 \dots T]$  and each  $u \in [1, \dots, U]$  to generate a lattice  $S$ . Each cell of the lattice represents a state  $s(t, u)$ , from which we generate a probability distribution of a token belonging to

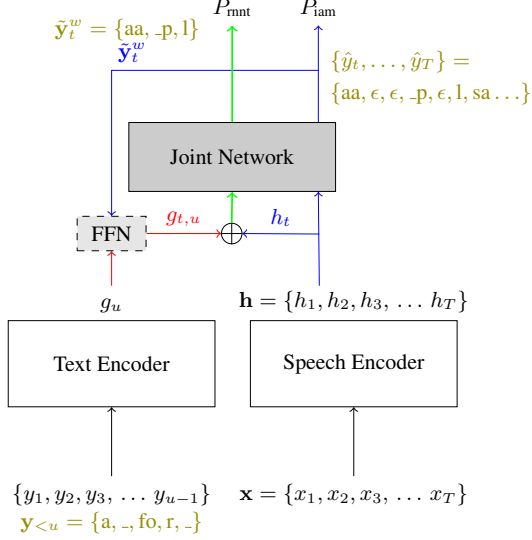


Figure 1: Architecture of LOOKAHEAD. We first calculate  $P_{iam}$  which is used to calculate  $P_{mnt}$ . An illustrative example with  $w = 3$  is provided in olive.  $\hat{y}_t$  is the argmax output associated with frame  $h_t$ .

vocabulary  $\mathcal{V}$

$$P(y_u^t | x_t, \mathbf{y}_{<u}) = \text{softmax}\{\text{JN}(h_t \oplus g_u)\} \quad (1)$$

where  $P(y_u^t)$  is the probability of emitting  $y_u$  at time  $t$ . The vocabulary includes a special blank token  $\epsilon$ .

Using this probability lattice, the RNN-T estimates the conditional distribution  $P(\mathbf{y}|\mathbf{x})$  by marginalizing over all possible monotonic alignments of acoustic frame  $t$  with output token position  $u$  on the lattice. This can be computed in polynomial time using a DP-sum method [1]. During inference, high-probability outputs are generated using beam-search.

### 2.1. Related Work

The inability of RNN-T models to generalize to out-of-domain utterances has been previously well-documented [7]. Previous approaches have adapted RNN-T models to new domains via shallow fusion with an external LM [8, 9, 10, 11, 12] and discounting RNN-T scores using an external LM [13] or an implicit LM derived from the RNN-T [14, 15, 5]. Apart from these inference-time techniques, there have also been other approaches that required training-time interventions such as subword regularization [11], augmenting pronunciation dictionaries [12, 16] and text-only adaptation of the RNN-T’s language encoder [4, 17, 18].

Apart from adaptation-based techniques, prior work has also explored many approaches that aim at improving interactions between the acoustic encoder (AE) and the language encoder (LE) of the RNN-T. These approaches have focused on more complex mechanisms to combine the AE and LE outputs rather than simply adding the respective representations [19, 20], using quantization in the LE [21] and fixing overconfident LE predictions in the RNN-T lattice [22]. Similar in spirit to our technique are non-causal models [23, 24, 25, 26, 27] that refer to future information via a self-attention or convolution mechanism. However, all these methods peek into the future at the

level of AE and thus with increased number of layers, the effective window size to the future increases which is not desirable in streaming applications.

## 3. Our Approach

In the RNN-T architecture the text and the audio are encoded independently and combined using a thin joint layer. The text encoding  $g_u$  captures the next token state conditioned only on tokens generated before it. If the language model priors are strong, it is possible for the text encoding  $g_u$  to strongly prefer a word unit, and to bias the softmax to generate that unit, even if that has little agreement with the acoustic input. This often manifests as hallucinated words that are fluent as per the language model, but which have no support with the acoustic input. In Table 4, we present examples of such hallucinations. We observe that rare words, e.g., word “MAGATAMA” in the first example, gets recognized as “BY THE TOWN” which has little acoustic overlap with the utterance.

Next we present our approach called LOOKAHEAD of fixing this mismatch. We identify that hallucinations are generated by representations  $g_u$  which have been computed independent of the acoustic input. Our key insight is to improve the LE representation  $g_u$  by looking ahead a few tokens into the future from the acoustic input. However, implementing this insight, without adversely impacting the efficient online operation of the RNN-T model is non-trivial. We perform this modification in two steps. First, we extract from the acoustic encoder, a lookahead of  $k$  tokens after each frame  $t$  of the acoustic input. In Section 3.1 we present how we do that. Second, we modify  $g_u$  with the extracted future tokens to get now a time-aware, acoustically informed encoding  $g_{u,t}$ . Section 3.2 presents the details of this step. Overall modified architecture of our approach appears in Figure 1.

### 3.1. Extraction of Lookahead tokens from the Acoustic Encoder

We use the notion of *implicit acoustic model* (IAM) and generate a probability distribution over vocabulary tokens based on only the acoustic signals  $h_t$

$$P_{iam}(y|h_t) = \text{softmax}\{\text{JN}(h_t \oplus \mathbf{0})\} \quad (2)$$

This distribution is supervised using gold transcripts by marginalizing over all non-blank tokens, much like the regular RNN-T loss.

Next, for each frame of  $\mathbf{x}$  we get the most likely output based on acoustic signal  $h_t$ .

$$\hat{y}_t = \underset{y}{\text{argmax}} P_{iam}(y|h_t) \quad \forall t \quad (3)$$

We do this for all input frames. On this set of  $T$  tokens, at each  $t$ , we find the first  $w$  non-blank tokens after  $t$  and call it  $\tilde{\mathbf{y}}_t^w$ . We choose small values of  $w$  (e.g., 2 or 3) to ensure that we still inherit the streaming capabilities of the underlying RNN-T model. The  $\tilde{\mathbf{y}}_t^w$  provides a look ahead into the future tokens after each  $t$  based purely on the acoustic input. Figure 1 shows an example where for  $t = 1, w = 3$  we extract a  $\tilde{\mathbf{y}}_t^w$  comprising of the first three non-blank outputs from the acoustic model.

### 3.2. Modifying Text Encoding with Lookahead Tokens

With the presence of the above future context, we condition the text encoding  $g_u$  at each  $t$  using a simple FFN  $F$  as

$$\hat{g}_{t,u} = F(g_u, \tilde{\mathbf{y}}_t^w) \quad (4)$$

Split \ Accent	US	Eng	Can	Ind	Scot	Phil	HK
Train	55.3 / 42023 / 517	45.0 / 34615 / 1856	35.6 / 24814 / 575	51.2 / 34890 / 982	13.8 / 8671 / 85	2.8 / 2035 / 47	NA
Dev	4.3 / 3089 / 78	4.2 / 3465 / 62	4.0 / 3033 / 68	5.8 / 3673 / 50	2.6 / 2004 / 33	1.5 / 1009 / 39	NA
Test	8.0 / 5363 / 93	6.4 / 5167 / 71	6.6 / 5008 / 140	7.8 / 5094 / 67	2.8 / 2001 / 37	1.6 / 1114 / 22	3.5 / 2750 / 111

Table 1: Statistics of the MCV dataset. Each cell lists duration (in hrs)/# utterances/# of speakers in the subset.

Dataset	Model	Librispeech-Test				MCV-Test						
		dev-clean	test-clean	dev-other	test-other	US	Eng	Can	Ind	Scot	Phil	HK
L100+P	Baseline	7.1	7.6	20.5	20.8	46.5	35.6	27.0	69.2	40.1	46.3	48.1
	LOOKAHEAD	<b>6.4</b>	<b>6.6</b>	<b>18.8</b>	<b>19.1</b>	<b>44.9</b>	<b>33.2</b>	<b>24.8</b>	<b>68.6</b>	<b>37.3</b>	<b>42.6</b>	<b>45.7</b>
L360+P	Baseline	4.2	4.5	13.5	13.1	38.1	23.6	19.5	56.9	30.2	34.4	39.1
	LOOKAHEAD	<b>4.0</b>	<b>4.3</b>	<b>12.8</b>	<b>12.5</b>	<b>36.9</b>	<b>22.5</b>	<b>19.2</b>	<b>56.6</b>	<b>28.9</b>	<b>33.2</b>	<b>37.5</b>
L960	Baseline	<b>3.1</b>	<b>3.2</b>	8.6	8.5	29.2	15.0	16.1	35.5	23.1	23.9	29.7
	LOOKAHEAD	3.2	3.3	<b>8.3</b>	<b>8.4</b>	<b>28.8</b>	<b>14.6</b>	<b>15.9</b>	<b>35.2</b>	<b>23.0</b>	<b>23.8</b>	<b>29.0</b>
MCV	Baseline	18.9	18.8	31.4	31.7	27.9	18.1	19.1	27.4	25.8	27.2	31.9
	LOOKAHEAD	<b>16.1</b>	<b>15.8</b>	<b>27.6</b>	<b>28.2</b>	<b>23.8</b>	<b>14.3</b>	<b>15.3</b>	<b>24.2</b>	<b>21.9</b>	<b>23.9</b>	<b>27.8</b>

Table 2: Main results of baseline vs LOOKAHEAD.

The remainder of the RNN-T pipeline stays the same and we pass along the new LE representation  $\hat{g}_{t,u}$  to the JN for further processing

$$P_{LA}(y_u^t | \mathbf{x}, t, \mathbf{y}_{<u}) = \text{softmax}\{\text{JN}(h_t \oplus \hat{g}_{t,u})\} \quad (5)$$

Thus, our overall training is a sum of losses over the gold transcript on marginalizations of  $P_{LA}$ , and of  $P_{iam}$ .

## 4. Experiments and Results

In this section, we present various results comparing our method to an existing SoTA baseline RNN-T ASR system.

### 4.1. Datasets

We use the popular Librispeech benchmark [28] in three different well-known settings of varying training durations. We also use the Mozilla Common Voice dataset to derive accented out-of-domain speech samples. These are detailed below:

1. L100: This is the well-known Librispeech-100 dataset which is a 100 hour subset of the Librispeech corpus[28].
2. L360: This is a 360 hour subset of the Librispeech corpus.
3. L960: This is the complete 960-hour Librispeech corpus.
4. MCV: This is a dataset obtained from the Mozilla Common Voice [29] Version 7.0 corpus. Using this dataset, we create a 200hr training subset consisting of seven English accents of which one is a test only accent. The accent-wise composition of the MCV dataset can be found in Table 1.

X+P refers to systems with speed-perturbation that augments the dataset X with time-warped copies using factors

	Libri-test						
	dev-clean		test-clean		dev-other		test-other
Baseline	20.0		21.7		48.3		51.6
LOOKAHEAD	20.3		20.1		47.2		50.7
	MCV-test						
	US	Eng	Can	Ind	Scot	Phil	HK
Baseline	61.1	63.4	46.2	85.2	60.7	61.2	70.7
LOOKAHEAD	58.7	59.9	41.5	84.4	58.0	58.3	67.7

Table 3: Rare-WER comparison of baseline vs LOOKAHEAD on L100+P.

0.9, 1.0, 1.1 respectively. In Table 2, we indicate wherever this augmentation is used. We test the above models both on in-domain and out-of-domain test splits, where the latter is obtained by changing the accent from the MCV corpus.

### 4.2. Implementation Details

We use the ESPNet toolkit [30] to implement our model and perform all our experiments. Our base architecture is an RNN-T model with conformers [31] as AE and LSTM [32] as LE. The joint network adds the two encodings and applies a tanh non-linearity. We experiment with two different models corresponding to different encoder sizes.

- SMALL: This model includes 18 AE blocks with internal representation of size 256 and four attention heads. A macaron-style feed forward projecting to size 1024 is used within the conformer block. The LE consists of a single 300 dimensional LSTM cell. Both these encoders project their output to a joint-space of 300 dimensions before feeding it to the joint network. This model was used for performing experiments with the L100+P and MCV datasets. We use a vocabulary size of 300 and 150 sub-words for these experiments, respectively ( $\sim 30M$  params).
- LARGE: This model has 12 AE blocks conformer-encoder of size 512 and a macaron style feed-forward of 2048. The LE block is a 512 unit LSTM. Both representations are projected to a joint-space of 512 dimensions. This model is used for our experiments involving the L960 and L360+P datasets. We use a vocabulary size of 500 for these experiments ( $\sim 80M$  params).

We train for 100 epochs using the *Noam*-optimizer [33, 34] with a Noam-learning rate of 5 and 25000 warm-up steps. For inference, we use ALSD [35] with a beam size of 30. For the purposes of LOOKAHEAD we have used 3 lookahead tokens.

### 4.3. Results

Our primary set of numbers can be seen in Table 2 which shows the capability of LOOKAHEAD against baseline models trained on various training data sizes. Using L100+P, we see statistically significant improvements ( $p < 0.01$  using the MAPSSWE test [36]) in WERs across both training sets for in-domain and out-of-domain evaluation sets. The gains are reduced with using

Reference	Baseline	LOOKAHEAD
shaped like commas or <b>MAGATAMA</b> confession is <b>GOOD FOR</b> the soul for <b>ALL WHO SIN ALL WHO SUFFER</b> you <b>LIFE WAS</b> not <b>EASY</b> do <b>ANDROIDS</b> dream of <b>LA VALLIERE</b> is quite a poetess said	sheep like commas or <b>BY THE TOWN</b> the fashion is <b>GIVEN FORWARD</b> the soul for <b>OLDER SIN OLD RISK</b> you by <b>OR</b> not <b>EVEN</b> do <b>ENJOYS</b> dream of <b>THE VALET</b> is quite poetic said	sheep like commas or <b>MAGATINA</b> confession is <b>GOOD FOR</b> the soul for <b>ALL WHOOSIN OLD WHO SUFFER</b> you <b>LIFE OR</b> not <b>EASY</b> do <b>ANDREWS</b> dream of <b>LAVALIER</b> is quite a poeticus said

Table 4: Anecdotes comparing Baseline with LOOKAHEAD. Baseline hallucination errors are highlighted in blue.

	Model	dev_clean	test_clean	dev_other	test_other	us	can	eng	ind	scot	phil	hk
PER	Baseline	2.52	2.73	9.90	9.81	24.46	11.74	17.34	39.33	19.55	21.35	23.96
	LOOKAHEAD	<b>2.28</b>	<b>2.35</b>	<b>8.99</b>	<b>8.98</b>	<b>23.40</b>	<b>10.42</b>	<b>15.80</b>	<b>38.18</b>	<b>17.55</b>	<b>19.20</b>	<b>22.41</b>
WFED	Baseline	8.31	8.95	31.53	31.25	77.95	39.07	52.89	118.89	63.40	64.02	73.55
	LOOKAHEAD	<b>7.57</b>	<b>8.00</b>	<b>28.84</b>	<b>29.13</b>	<b>74.47</b>	<b>34.20</b>	<b>48.23</b>	<b>115.30</b>	<b>57.12</b>	<b>58.94</b>	<b>69.04</b>
DER	Baseline	1.13	1.22	4.37	4.34	11.48	5.52	7.23	17.50	9.00	8.98	10.23
	LOOKAHEAD	<b>1.03</b>	<b>1.10</b>	<b>4.01</b>	<b>4.03</b>	<b>11.01</b>	<b>4.83</b>	<b>6.58</b>	<b>16.82</b>	<b>8.10</b>	<b>8.27</b>	<b>9.57</b>

Table 5: Comparing acoustically-aware error metrics between Baseline and LOOKAHEAD trained using L100+P.

larger models (L360+P and L960) but still remain consistent, especially for the out-of-domain test sets. We further focus on rare words where the effect of LM biasing is expected to have a greater effect. Table 3 shows WER performance of baseline vs LOOKAHEAD on rare-words on the L100+P dataset. We define any word that occurs lower than 20 times in the dataset as a rare-word. We see consistent improvement in rare-word recognition across both in-domain as well as out-of-domain datasets. Table 4 shows some anecdotes indicating the effectiveness of LOOKAHEAD. Note how the baseline model hallucinates words like TOWN, GIVEN, etc. that have no acoustic overlap with the spoken word. LOOKAHEAD either corrects them or produces a word that is more acoustically similar to the spoken word.

#### 4.4. Error Analysis

Apart from WERs, we provide further analysis in Table 5 to validate our claim that predictions from LOOKAHEAD are indeed acoustically more similar to the ground-truth. We use the following three metrics:

- Phone Error Rate (PER): We use Epitran [37] to transcribe word sequences to phone sequences in IPA and find the IPA-based edit distance to yield PERs, which are a better measure of acoustic discrepancy compared to WERs.
- Weighted Feature-based Edit Distance (WFED): With the PER metric, the distance between two dissimilar phones and two related phones are both 1. WFED is an edit distance metric based on Panphon [38] that helps compute distances between IPA phones based on their articulatory features.
- Dolgopolsky Error Rate (DER): Both PER and WFED rely on an underlying sequence of IPA phones derived from the predicted word sequences using a fixed set of rules. These IPA sequences are not very reliable, especially for the out-of-

Metric	Ref	Hyp	Dist
PER	ball (bɒl)	maul(mɒl)	1
	call (kɒl)	install (mstɒl)	4
WFED	ball (bɒl)	paul (pɒl)	0.25
	ball (bɒl)	call (kɒl)	2.25
DER	pit (pit)	meet (mit)	1
	pit (pit)	beat (bit)	0

Table 6: Illustration of different error rate metrics

w	dev-clean	test-clean	dev-other	test-other
2	6.5	6.9	18.8	19.0
3	<b>6.4</b>	<b>6.6</b>	<b>18.8</b>	19.1
5	6.4	6.8	18.9	<b>18.9</b>

Table 7: Ablation on L100+P by varying w with LOOKAHEAD.

domain MCV speech samples with varying speech accents. To alleviate errors arising from incorrect IPA, we use Dolgopolsky’s IPA clusters [39] to label similar-sounding IPA phones with the same cluster ID. DER is an edit distance using these new cluster IDs that only measures large deviations in phonetic realizations between a reference word and a predicted word and does not penalize small changes in the place and manner of articulation of phones (like PER).

Table 6 lists illustrative examples for each of the three metrics. E.g., WFED imposes a cost of 2.25 if “ball” is misrecognized as “call” (which would have a PER of 1); DER imposes no cost for “pit” being misrecognized as “beat”. Table 5 compares PER, WFED and DER values from the baseline and LOOKAHEAD. DER penalizes only for large acoustic digressions; we see consistent improvements on DER using LOOKAHEAD with larger improvements on the MCV test utterances. We also show ablations on the window size of future tokens considered in Table 7. We observe that using a window size of 3 yields the best results.

## 5. Conclusion

We propose a simple and effective scheme of acoustic LOOKAHEAD for RNN-T models to safeguard against producing acoustic hallucination. From the acoustic encoder, we extract a fixed LOOKAHEAD of tokens and use it to improve the representation from the language encoder to generate a more acoustically aligned output. We obtain significant reductions in WER and various other phonetic error rates across various settings.

## 6. Acknowledgements

The third author gratefully acknowledges financial support from a SERB Core Research Grant, Department of Science and Technology, Govt of India on accented speech processing.

## 7. References

- [1] A. Graves, “Sequence Transduction with Recurrent Neural Networks,” 2012. [Online]. Available: <http://arxiv.org/abs/1211.3711>
- [2] K. Rao, H. Sak, and R. Prabhavalkar, “Exploring architectures, data and units for streaming End-to-End speech recognition with RNN-Transducer,” in *ASRU-2017*.
- [3] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang *et al.*, “Streaming End-to-End speech recognition for mobile devices,” in *ICASSP 2019*.
- [4] J. Pylkkönen, A. Ukkonen, J. Kilpikoski, S. Tamminen, and H. Heikinheimo, “Fast Text-Only Domain Adaptation of RNN-Transducer Prediction Network,” 2021.
- [5] V. Unni, S. Khare, A. Mittal, P. Jyothi, S. Sarawagi, and S. Bharadwaj, “Adaptive discounting of implicit language models in RNN-Transducers,” in *ICASSP*, 2022.
- [6] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. Bang, A. Madotto, and P. Fung, “Survey of hallucination in natural language generation,” in *ACM Computing Surveys*. Association for Computing Machinery (ACM), nov 2022.
- [7] C.-C. Chiu, A. Narayanan, W. Han, R. Prabhavalkar, Y. Zhang, N. Jaitly, R. Pang, T. N. Sainath, P. Nguyen, L. Cao, and Y. Wu, “RNN-T models fail to generalize to out-of-domain audio: Causes and solutions,” 2020. [Online]. Available: <https://arxiv.org/abs/2005.03271>
- [8] C. Gulcehre, O. Firat, K. Xu, K. Cho, L. Barrault, H.-C. Lin, F. Bougares, H. Schwenk, and Y. Bengio, “On using monolingual corpora in neural machine translation,” 2015.
- [9] T. N. Sainath, Y. R. He, A. Narayanan, R. Botros, R. Pang, D. J. Rybach, C. Allauzen, E. Variansi, J. Qin, Q.-N. Le-The, A. Gruenstein, A. Gulati, B. Li, C. Peyser, C.-C. Chiu, D. A. Caseiro, E. Guzman, I. C. McGraw, J. Yu, M. D. Riley, P. Rondon, Q. Liang, S. Mavandadi, S. yiin Chang, T. D. Strohman, W. R. Huang, W. Li, Y. Wu, and Y. Zhang, “An efficient streaming non-recurrent on-device end-to-end model with improvements to rare-word modeling,” 2021.
- [10] V. Ravi, Y. Gu, A. Gandhe, A. Rastrow, L. Liu, D. Filimonov, S. Novotney, and I. Bulkyo, “Improving accuracy of rare words for RNN-Transducer through unigram shallow fusion,” 2020.
- [11] D. Le, G. Keren, J. Chan, J. Mahadeokar, C. Fuegen, and M. L. Seltzer, “Deep shallow fusion for RNN-T personalization,” 2020.
- [12] D. Zhao, T. N. Sainath, D. Rybach, P. Rondon, D. Bhatia, B. Li, and R. Pang, “Shallow-fusion end-to-end contextual biasing,” in *Interspeech-2019*.
- [13] E. McDermott, H. Sak, and E. Variansi, “A Density Ratio approach to language model fusion in End-To-End automatic speech recognition,” 2020.
- [14] E. Variansi, D. Rybach, C. Allauzen, and M. Riley, “Hybrid Autoregressive Transducer (HAT),” 2020.
- [15] Z. M. et al., “Internal Language Model Estimation for Domain-Adaptive End-to-End Speech Recognition,” 2020.
- [16] N. Kitaoka, B. Chen, and Y. Obashi, “Dynamic out-of-vocabulary word registration to language model for speech recognition.”
- [17] X. Chen, Z. Meng, S. Parthasarathy, and J. Li, “Factorized Neural Transducer for efficient language model adaptation,” 2021. [Online]. Available: <https://arxiv.org/abs/2110.01500>
- [18] A. Mittal, S. Sarawagi, and P. Jyothi, “In-situ text-only adaptation of speech models with low-overhead speech imputations,” in *ICLR*, 2023. [Online]. Available: [https://openreview.net/forum?id=T2Ncx\\_PN2K](https://openreview.net/forum?id=T2Ncx_PN2K)
- [19] G. Saon, Z. Tieske, D. Bolanos, and B. Kingsbury, “Advancing RNN Transducer technology for speech recognition,” 2021. [Online]. Available: <https://arxiv.org/abs/2103.09935>
- [20] C. Zhang, B. Li, Z. Lu, T. N. Sainath, and S. yiin Chang, “Improving the fusion of acoustic and text representations in RNN-T,” in *ICASSP-2022*.
- [21] J. Shi, G. Saon, D. Haws, S. Watanabe, and B. Kingsbury, “VQ-T: RNN Transducers using Vector-Quantized prediction network states,” 2022. [Online]. Available: <https://arxiv.org/abs/2208.01818>
- [22] M. Novak and P. Papadopoulos, “RNN-T lattice enhancement by grafting of pruned paths,” in *Interspeech*, 2022, pp. 4960–4964.
- [23] Q. Zhang, H. Lu, H. Sak, A. Tripathi, E. McDermott, S. Koo, and S. Kumar, “Transformer Transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” 2020. [Online]. Available: <https://arxiv.org/abs/2002.02562>
- [24] Y. Shi, Y. Wang, C. Wu, C.-F. Yeh, J. Chan, F. Zhang, D. Le, and M. Seltzer, “Emformer: Efficient memory transformer based acoustic model for low latency streaming speech recognition,” 2020. [Online]. Available: <https://arxiv.org/abs/2010.10759>
- [25] X. Chen, Y. Wu, Z. Wang, S. Liu, and J. Li, “Developing real-time streaming transformer transducer for speech recognition on large-scale dataset,” in *ICASSP*, 2020, pp. 5904–5908.
- [26] W. Han, Z. Zhang, Y. Zhang, J. Yu, C.-C. Chiu, J. Qin, A. Gulati, R. Pang, and Y. Wu, “ContextNet: Improving convolutional neural networks for automatic speech recognition with global context,” *ArXiv*, vol. abs/2005.03191, 2020.
- [27] G. P. Strimel, Y. Xie, B. King, M. Radfar, A. Rastrow, and A. Mouchtaris, “Lookahead when it matters: Adaptive non-causal transformers for streaming neural transducers,” in *ICML*, 2023.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *ICASSP*, 2015, pp. 5206–5210.
- [29] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common Voice: A Massively-Multilingual Speech Corpus,” 2019. [Online]. Available: <https://arxiv.org/abs/1912.06670>
- [30] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, “ESPnet: End-to-End Speech Processing Toolkit,” 2018. [Online]. Available: <https://arxiv.org/abs/1804.00015>
- [31] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented transformer for speech recognition,” 2020.
- [32] S. Hochreiter and J. Schmidhuber, “Long Short-term Memory,” *Neural computation*, vol. 9, pp. 1735–80, 12 1997.
- [33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” 2017.
- [34] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” 2014. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [35] G. Saon, Z. Tieske, and K. Audhkhasi, “Alignment-Length Synchronous Decoding for RNN Transducer,” in *ICASSP-2020*.
- [36] L. Gillick and S. J. Cox, “Some statistical issues in the comparison of speech recognition algorithms,” in *ICASSP*, 1989, pp. 532–535.
- [37] D. R. Mortensen, S. Dalmia, and P. Littell, “Epitran: Precision G2P for many languages,” in *LREC*, May 2018.
- [38] D. R. Mortensen, P. Littell, A. Bharadwaj, K. Goyal, C. Dyer, and L. S. Levin, “Panphon: A resource for mapping IPA segments to articulatory feature vectors,” in *COLING*. ACL, 2016, pp. 3475–3484.
- [39] A. B. Dolgopolsky, “Gipoteza drevnejšego rodstva jazykovych semej severnoj evrazii s verojatnostej točki zrenija.” 1986.