



Contextual Density Ratio for Language Model Biasing of Sequence to Sequence ASR Systems

Jesús Andrés-Ferrer^{1*}, Dario Albesano^{2*}, Puming Zhan³, Paul Vozila³

¹Nuance Communications, Valencia, Spain

²Nuance Communications, Torino, Italy

³Nuance Communications, Burlington, MA, USA

{jesusandres.ferrer, dario.albesano, puming.zhan, paul.vozila}@nuance.com

Abstract

End-2-end (E2E) models have become increasingly popular in some ASR tasks because of their performance and advantages. These E2E models directly approximate the posterior distribution of tokens given the acoustic inputs. Consequently, the E2E systems implicitly define a language model (LM) over the output tokens, which makes the exploitation of independently trained language models less straightforward than in conventional ASR systems. This makes it difficult to dynamically adapt E2E ASR system to contextual profiles for better recognizing special words such as named entities. In this work, we propose a contextual density ratio approach for both training a contextual aware E2E model and adapting the language model to named entities. We apply the aforementioned technique to an E2E ASR system, which transcribes doctor and patient conversations, for better adapting the E2E system to the names in the conversations. Our proposed technique achieves a relative improvement of up to 46.5% on the names over an E2E baseline without degrading the overall recognition accuracy of the whole test set. Moreover, it also surpasses a contextual shallow fusion baseline by 22.1 % relative.

Index Terms: speech recognition, end-to-end, sequence-to-sequence, language model, shallow fusion, density ratio

1. Introduction

In recent years, End-to-end (E2E) systems [1, 2, 3] have been greatly adopted for ASR because of their performance and simplicity. In conventional ASR systems, several components, such as the pronunciation lexicon, the language models, and the acoustic model, are optimized independently. In contrast, E2E systems integrate all components of conventional ASR systems in a single neural network by directly approximating the posterior probability given the acoustic features.

Despite the direct posterior approximation simplifying many aspects such as training, decoding and deployment, these advantages do not come without drawbacks. Specifically, one of the disadvantages is the difficulty of exploiting external language models. Since the E2E models directly approximate the posterior token distribution, there is no straightforward way to integrate them, which generates recognition problems for rare events among other problems. These tail probability events cannot be expected to be comprehensively observed in training. The selection of the output units such as word pieces [4] or characters instead of full words could theoretically mitigate these problems, but in practice this is not the case [5].

Several early attempts were proposed to integrate external language models (LM) into E2E systems, ranging from shallow or deep fusion [6, 7], to cold fusion [8] among many others [9, 10]. In the ASR domain, [7] shown that shallow fusion

with same output units in both LM and E2E worked best to incorporate an external language model in some tasks. However, one of the disadvantages of shallow fusion comes from the fact that most of the E2E systems [1, 2, 11] arguably already incorporate an internal LM [12, 13, 14]. Recently, a *density ratio* approach [15] was proposed to better integrate external LM by removing the internal LM contribution of E2E models while decoding. In this initial work [15], authors reported improvements over shallow fusion for domain adaptation tasks. Further works modified the architecture [14] or made some assumptions [16] to better approximate the implicit LM of E2E systems in combination with density ratio.

In many tasks, ASR systems are expected to recognize singletons or infrequent words such as contact list names or other named entities. These entities may not be seen during training, and even for those which are observed, the training data is clearly neither representative of the particular distribution nor of the relevant set associated with the application context for a given utterance. Consequently, even if the E2E model is using BPE [4] or graphemes as output units, the system struggles to recognize them. Several approaches have been introduced in the literature to contextualize E2E models to named entities such as song names or contact lists. In [17] and [18], an additional attention set input is proposed for both attention and recurrent neural transducer (RNN-T) models respectively. Shallow fusion approach [19] was applied surpassing state-of-the-art conventional models. More recently [5], shallow fusion was used in combination with special tokens to delimit class-based entity names together with a mapping to transform rare words to common words through pronunciation.

In this paper, we propose *contextual density ratio* for contextualizing E2E models so that the internal E2E LM is dynamically adapted to a priori known named entities. The proposed technique builds upon both density ratio [15] and class-based LM tags [5] to contextualize the E2E models. During training, we introduce special tokens to enclose known named entities so that the E2E system learns to predict when a named entity is spoken by the different statistical clues obtained from both the acoustics and the internal LM. We also approximate an internal E2E language model by training an independent LM with the transcriptions on which the E2E system was trained. During decoding, we apply density ratio within the named entities segments identified by the E2E model to dynamically adapt and contextualize the system's LM. We show that the proposed contextual density ratio (CDR) reduces names recognition errors over a E2E baseline with and without contextual shallow fusion [19] by 46.5% and 22.1 % respectively.

The paper is organized as follows. We first briefly review density ratio in section 2.1 to pave the way for the following section 2.2 where the proposed contextual density ratio is detailed. In section 3, we apply the proposed technique to a doctor-to-patient conversation task for better recognizing both doctors and

* equal contribution

patient names. Finally, in section 4, we conclude with some reflections.

2. Proposed approach

In this section we briefly summarize the *contextual density ratio* approach. We first review the density ratio [15], and then we propose our extension to the contextual biasing scenario.

2.1. Density ratio (DR)

Neural network E2E models such as [1, 20], directly approximate the probability of a target token sequence $y = y_1^J$, given acoustic features $\mathbf{x} = x_1^T$, as follows:

$$p(\mathbf{y}|\mathbf{x}) \approx p_{e2e}(\mathbf{y}|\mathbf{x}) \quad (1)$$

where I is the length of the input acoustic features along time and J the length of a possible token sequence; and where p denotes the actual probability of the token sequence given the input acoustic features and p_{e2e} the approximated model distribution.

In density ratio, we wish to adapt the posterior in-domain (ID) distribution $p(\mathbf{y}|\mathbf{x})$ to a new out-of-domain (OOD) posterior, $q(\mathbf{y}|\mathbf{x})$. Via standard noisy channel decomposition (Bayes law) and under assumption that acoustics of both domains are similar, $p(\mathbf{x}|\mathbf{y}) \approx q(\mathbf{x}|\mathbf{y})$, we decompose the OOD posterior as:

$$q(\mathbf{y}|\mathbf{x}) = \frac{q(\mathbf{x}|\mathbf{y})q(\mathbf{y})}{q(\mathbf{x})} = \frac{p(\mathbf{x}|\mathbf{y})q(\mathbf{y})}{q(\mathbf{x})} \quad (2)$$

Utilizing the same noisy channel decomposition for ID distribution $p(\mathbf{x}|\mathbf{y})$ and rearranging terms, we obtain:

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} \quad (3)$$

Plugging eq (3) into eq. (2) yields:

$$q(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{y})q(\mathbf{y})}{q(\mathbf{x})} = p(\mathbf{y}|\mathbf{x}) \frac{q(\mathbf{y})}{p(\mathbf{y})} \frac{p(\mathbf{x})}{q(\mathbf{x})} \quad (4)$$

Note that, eq. (4) has a marginal LM ratio $q(\mathbf{y})/p(\mathbf{y})$ and an acoustics density ratio $p(\mathbf{x})/q(\mathbf{x})$, which does not modify token sequences scores. We obtain the final density ratio (DR) approach by approximating distributions with models:

$$q(\mathbf{y}|\mathbf{x}) = p_{e2e}(\mathbf{y}|\mathbf{x}) \frac{q_{\text{OOD}}(\mathbf{y})}{p_{\text{ID}}(\mathbf{y})} \frac{p(\mathbf{x})}{q(\mathbf{x})} \quad (5)$$

which approximates the acoustic distribution with both an ID language and E2E models, p_{ID} and p_{e2e} ; and shifts the posterior with an OOD language model, q_{OOD} , up to a constant ratio on the acoustics. Note that p_{ID} corresponds to the marginal distribution obtained from E2E posterior, p_{e2e} ; in practice we take the E2E model training transcriptions as approximately sampling from this distribution.

Finally, when performing the beam search we apply weights to balance out the importance of each score, yielding the search score:

$$\text{score}(\mathbf{y}; \mathbf{x}) = \log p_{e2e}(\mathbf{y}|\mathbf{x}) - \alpha \cdot \log p_{\text{ID}}(\mathbf{y}) + \beta \cdot \log q_{\text{OOD}}(\mathbf{y}) \quad (6)$$

2.2. Contextual density ratio (CDR)

Density ratio (DR) [15] can bias the distribution for the full utterance, but we would like to adapt the language model of E2E systems to be able to recognize named entities such as contact list names, or client names. These named entities may occur at any position inside a utterance and often only span a few words.

In order to do so, we extend the token vocabulary with special words for both start and end of entities, $\langle ne \rangle$ and $\langle /ne \rangle$.

During training, our E2E model inserts those tags in between known named entities, for instance, `the_next_patient_<ne> be_e th o ve n. </ne> is_losing_his_hear ing_ .` The system learns to predict when those named entities will occur, and during beam search or decoding, the system hypothesizes those special tokens as if they were standard tokens. Although we focus on a single tag for simplicity in this work, we could add as many specialized tags as desired,

In the previous section, we reviewed the DR approach for full utterance biasing. However, in practice, each token distribution on eq. (5) is approximated autoregressively as follows:

$$q(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x})}{q(\mathbf{x})} \prod_t p_{e2e}(y_t|\mathbf{x}, \mathbf{y}_{<t}) \frac{q_{\text{OOD}}(y_t|\mathbf{y}_{<t})}{p_{\text{ID}}(y_t|\mathbf{y}_{<t})} \quad (7)$$

Because of the new tokens added to the E2E vocabulary, we know the positions where the named entities will occur. For simplicity, we assume a single named entity occurrence at position (t_b, t_e) . Then, the token posterior probability of an utterance with a named entity is decomposed by:

$$q_{\text{ne}}(\mathbf{y}|\mathbf{x}) = \left(\prod_{t < t_b} p_{e2e}(y_t|\mathbf{x}, \mathbf{y}_{<t}) \right) \cdot p_{e2e}(\langle ne \rangle|\mathbf{x}, \mathbf{y}_{<t_b}) \cdot \prod_{t=t_b+1}^{t_e} p_{e2e}(y_t|\mathbf{y}_{<t}, \mathbf{x}) \frac{q_{\text{NE}}(y_t|\mathbf{y}_{t_b}^{t-1})}{p_{\text{ID}}(y_t|\mathbf{y}_{<t})} \frac{p(\mathbf{x})}{q(\mathbf{x})} \cdot \prod_{t > t_e} p_{e2e}(y_t|\mathbf{x}, \mathbf{y}_{<t}) \quad (8)$$

where recall $y_{t_e} = \langle /ne \rangle$. Note that the in-domain (ID) and named entity (NE) LMs condition on different context. It is worth mentioning that the acoustic density ratio, $p(\mathbf{x})/q(\mathbf{x})$, may introduce a bias when comparing hypotheses with different number of named entities. However, assuming acoustics are similar, this ratio should approach 1 and we verified in the next section that this is not a problem in practice.

During beam search or decoding, each token is scored with $\text{score}(y_t) := \text{score}(y_t; \mathbf{y}_{<t}, \mathbf{x})$ computed as:

$$\text{score}(y_t) = \begin{cases} \log p_{e2e}(y_t|\dots) + \log \frac{q_{\text{NE}}(y_t|\mathbf{y}_{t_b}^{t-1})^\beta}{p_{\text{ID}}(y_t|\mathbf{y}_{<t})^\alpha}, & \text{ne}(\mathbf{y}_{<t}) \\ \log p_{e2e}(y_t|\dots), & \text{o.w.} \end{cases} \quad (9)$$

where we omit $\mathbf{y}_{<t}, \mathbf{x}$ from $p_{e2e}(y_t|\dots)$. The notation $\text{ne}(\mathbf{y}_{<t})$ indicates whether there is an open named entity tag in the previous tokens that has not yet been closed and in such a case, $t_b := t_b(\mathbf{y}_{<t})$ is the position of the start tag token, $\langle ne \rangle$. In practice, LM distributions are weighted by α and β to smooth them.

Figure 1 depicts a small segment of an utterance with a named entity. The E2E model scores each token independently of whether it is within a named entity tag, in contrast to both

Model	mr_	<ne>	moz	art_	</ne>
E2E	$p(\text{mr}_- \text{hello}_-, \mathbf{x})$	$p(\langle \text{ne} \rangle \dots, \text{mr}_-, \mathbf{x})$	$p(\text{moz} \dots, \text{mr}_-, \langle \text{ne} \rangle, \mathbf{x})$	$p(\text{art}_- \dots, \text{moz}, \mathbf{x})$	$p(\langle / \text{ne} \rangle \dots, \text{art}_-, \mathbf{x})$
ID	1	1	$p(\text{moz} \dots, \text{mr}_-, \langle \text{ne} \rangle)$	$p(\text{art}_- \dots, \text{moz})$	$p(\langle / \text{ne} \rangle \dots, \text{art}_-)$
NE	1	1	$q(\text{moz} \langle \text{ne} \rangle)$	$q(\text{art}_- \langle \text{ne} \rangle, \text{moz})$	$q(\langle / \text{ne} \rangle \langle \text{ne} \rangle, \text{moz}, \text{art}_-)$

Figure 1: An example of how the different components of contextual density ratio score a hypothesis. The final hypothesis score is computed as $\log(e2e) - \alpha \cdot \log(\text{ID}) + \beta \cdot \log(\text{NE})$. Note how the different language model components track the context differently according to their goal. The ID LM, is tracking full context from the beginning of the utterance in contrast with the NE LM which only tracks the LM from the most recent named entity start tag, $\langle \text{ne} \rangle$.

in-domain (ID) and LM named-entity (NE) LM. Since the ID LM is subtracted to counteract the internal E2E LM, it is tracking the context of the full utterance, whereas the NE LM, only tracks the context within the contextual named entity tags.

3. Experiments

3.1. Data set

We apply the proposed contextual density ratio to a conversational speech transcription task (doctor-patient conversations). All speech data are field data, from which we computed 80 filter bank features from the input audio. For this study, the train and test set consisted of 1.5K hours and 26.8 hours respectively from several medical domains.

We focused on person names as named entities and automatically marked them in both training and testing (references) by intersecting the transcriptions with a list of 173K English names. During training we enclosed occurrences of such names between tags ($\langle \text{ne} \rangle$ and $\langle / \text{ne} \rangle$) and during testing we extracted for each conversation both doctor and patient names from conversation utterances. In total 5% of the utterances in the test set have names.

3.2. Models

The E2E model is a fairly standard E2E attention model similar to LAS [1]. Specifically, our model follows the same architecture of the speaker independent system from LASA [20]. Our model has a set of 3 convolutional neural network (CNN) layers which are configured to produce a 768-dimensional embedding for each input frame. The 3 CNN layers have a kernel size of 3 and do not perform any time reduction. Subsequently, there is a pyramid LAS-like encoder composed of 6 bidirectional Long Short-Term Memory (bLSTM) [21] layers with 768 hidden units per layer and direction. The encoder performs input decimation by a factor of 2 after every other layer. The total time reduction is 8x with respect to the filter banks inputs. The decoder component uses 2 unidirectional LSTM [22] layers of 1536 units per layer. We use Bahdanau attention in an architecture similar to Luong [23, 20] on the attention component, with a single head with 1536 nodes. We apply dropout [24] to regularize the layers and SpecAugment [25] to both regularize the model and as data augmentation. We use a single time band of length 40 frames and a frequency band of up to 20 filters. The final softmax layer predicts the posterior probability of 2.5K word-pieces [4] computed on the training set. The model accounts for a total of 143.7M parameters.

In order to approximate the internal E2E language model for the contextual density ratio technique, we use the exact transcriptions of the utterances on which the E2E model was trained, together with the very same 2.5K word-pieces lexicon. We trained the language model with truncated backpropagation through time (BPTT) [26] to better approximate the internal

E2E language model. We settled for a language model with 512 embedding projection followed by 2 unidirectional LSTM layers of 1536 each and a final down projection back to 512 logits followed by the softmax layer of 2.5K outputs. The internal LM accounts for a total of 31.99M parameters. We regularized the model with small weight-decay and dropout. For decoding, both the α and β parameters were tuned on a different domain adaptation density ratio setup and both equal to 0.1. This ensures that search parameters are not over-trained to the specifics of the contextual density ratio task or the test set. In addition, because of the attention models, we used a coverage score [27] of 0.3 optimized in conjunction with α and β .

For the contextual LMs, we initialized them with the same in-domain LSTM based language model and trained them until convergence on the specific list of names extracted per each conversation.

3.3. Results

In order to measure the different systems, we report the word error rate (WER), as well as the WER within the name tags (WERT)¹. This allow us to evaluate both overall performance of the system as well as the specific improvements on name recognition.

First, we trained both an E2E baseline model without tags and an E2E model with the name tags ($\langle \text{ne} \rangle$ and $\langle / \text{ne} \rangle$). Second line of Table 1 shows that adding the tags slightly hurts the model performance, compared to the first line, but the difference is small and within training repetition variance (± 0.1). Then, we added contextual biasing by either shallow fusion [5] or contextual density ratio in the realistic scenario, in which we build a context biasing language model per conversation. As reported in the last 2 lines of Table 1, in both cases, we recover part of the small drop in accuracy while significantly improving name performance. Adding a contextual biasing shallow fusion component reduces name errors by 31.3 % relative. Contextual density ratio biasing reduces name errors by 46.5 % relative with respect to the E2E baseline and 22.1 % compared to shallow fusion. Note that if all names were wrongly recognized by substitutions, this will account for only 0.43% of the whole test set WER.

Table 2, compares the precision and the recall on the contextual tags themselves for the 3 systems with tags from Table 1. Both contextual density ratio and contextual shallow fusion improve the tagging precision and recall over the baseline, with contextual density ratio being better than contextual shallow fusion. The improvements in terms of precision and recall are small compared to the WERT improvements from Table 1, and consequently the name WERT improvement cannot be attributed as a side effect of better tagging. We hypothesize

¹We aligned hypotheses and references with tags and computed the errors aligned with the references tags.

Table 1: Word Error Rate (WER) and Word Error Rate within the named entities Tags (WERT) for several systems. Both contextual systems based on either SF or DR, extend the E2E model with contextual tags and are based on full conversations.

System	WER	WERT
E2E without contextual tags	14.94	n.a.
E2E with contextual tags	15.10	44.1
+ contextual SF	15.04	30.3
+ contextual DR (CDR)	15.00	23.6

Table 2: Precision and recall computed on the contextual tags themselves for the different contextual systems on table 1.

System	Precision	Recall
E2E with contextual tags	78.5	79.4
+ contextual SF	80.7	82.6
+ contextual DR (CDR)	82.0	83.5

instead that the a priori known entity values for the conversation, allow the contextually biased techniques to better detect the span for such entities, in particular when linguistic contextual cues are insufficient.

3.4. Oracle and noisy names list experiments

We studied the proposed contextual DR (CDR) in several scenarios. Initially, we evaluated the oracle best case where we know whether a given utterance has a name and the name itself in such a case. For those utterances with names, a contextual name LM is trained per utterance.

This oracle setup bounds the potential improvements with respect to the more realistic experiments from the previous section, and measures the effect of the acoustic ratio in Eq. (8) because no utterance without names is decoded with the biasing component. CDR applied to full conversations, is close to the oracle case as shown in first row of Table 3. On the other extreme, a single contextual name LM was trained on all names in the full test set as if the conversation level information was lost. The second row of Table 3 shows that using a single LM has a larger negative effect than using conversation-specific LMs, yet it is still better than the E2E baseline

Following the worst case experiment from Table 3, we studied the effect of extending the entity (or name) lists per each conversation, by adding an increasing number of randomly sampled distracting names (from 1 to 256) unrelated to the conversation. As shown in Figure 2, CDR is relatively robust to noisy names, only starting to degrade to 26.0% WERT at 256 distracting names. Specifically, for 16 distracting names we obtain a WERT of 23.9 which is very close to 23.6 of the conversation level LMs in Table 1.

Finally, the entity list per conversation, which is used to train the CDR contextual LMs, was extended with 16 adversarial names so as to further investigate the robustness of CDR. These names are randomly selected from a pool of 170K names (and their combinations) according to their character lexicographical Levensthein distance, in $\{1,2,4\}$, from any of the actual conversation names. Note that for Levensthein distance 4, we can drop part of the name (e.g. in case of John Smith the resulting adversarial name can be Smith). Table 4 assesses how adversarial examples degrade the performance of the CDR biasing; it is still relatively robust and much better than that of the baseline system. Note that the performance drop for edit

Table 3: WERT precision and recall for a single LM per utterance (oracle case) and a single LM for the full test set dropping conversation level information.

System	WERT	Precision	Recall
1 LM per utterance (Oracle)	20.1	82.6	85.1
1 LM for full test set	37.3	78.5	79.7

Table 4: WERT for 16 names with different Levensthein distances, applied to per-conversation CDR case.

Levensthein distance	random	4	2	1
WERT	23.9	28.2	30.5	28.3

distance 4 may seem unexpected, but it is in part a side effect of random sampling being able to drop full short names or surnames via deletions, in contrast to 2 or 1, where names are rarely dropped if at all. We intentionally allowed for full name deletions so as to make the list insidiously difficult.

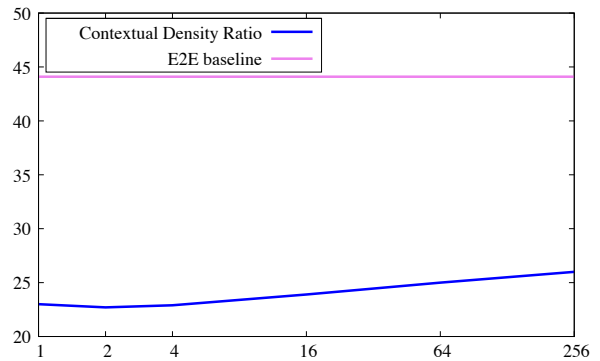


Figure 2: WERT (y-axis) at increasing number of per-conversation randomly sampled distracting names (log-scale).

4. Conclusions

In this paper, a contextual density ratio for contextual language model biasing was proposed. This technique was applied to the task of name recognition in doctor-patient conversations. The proposed approach improves name recognition up to 46.5% with respect to a standard E2E system or 22.1% relative with respect to contextual shallow fusion. Moreover, the technique does not degrade the system performance on utterances without names significantly and has no major side effects.

The behaviour of the proposed technique was studied by adding random distracting and adversarial names to the biasing name list. Contextual density ratio is robust to noise, being more sensitive to similar names. As future work, we want to improve the technique in those adversarial conditions, for instance by also taking into account the phonetics.

5. References

- [1] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. of ICASSP*. Shanghai, China: IEEE, 2016, pp. 4960–4964.
- [2] A. Graves, A. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE International Con-*

- ference on Acoustics, Speech and Signal Processing, May 2013, pp. 6645–6649.
- [3] E. Battenberg, J. Chen, R. Child, A. Coates, Y. G. Y. Li, H. Liu, S. Satheesh, A. Sriram, and Z. Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *Proc. of IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. Okinawa, Japan: IEEE, 2017, pp. 206–213.
 - [4] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 1715–1725. [Online]. Available: <https://www.aclweb.org/anthology/P16-1162>
 - [5] R. Huang, O. Abdel-Hamid, X. Li, and G. Evermann, “Class LM and word mapping for contextual biasing in end-to-end ASR,” in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 4348–4351. [Online]. Available: <https://doi.org/10.21437/Interspeech.2020-1787>
 - [6] Ç. Gülçehre, 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,” *arXiv e-prints*, vol. abs/1503.03535, Mar. 2015. [Online]. Available: <https://arxiv.org/abs/1503.03535>
 - [7] 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)*, 2018, pp. 1–5828.
 - [8] A. Sriram, H. Jun, S. Satheesh, and A. Coates, “Cold fusion: Training seq2seq models together with language models,” in *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, B. Yegnanarayana, Ed. ISCA, 2018, pp. 387–391. [Online]. Available: <https://doi.org/10.21437/Interspeech.2018-1392>
 - [9] C. Shan, C. Weng, G. Wang, D. Su, M. Luo, D. Yu, and L. Xie, “Component fusion: Learning replaceable language model component for end-to-end speech recognition system,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 5361–5365.
 - [10] S. Kim, Y. Shangguan, J. Mahadeokar, A. Bruguier, C. Fuegen, M. L. Seltzer, and D. Le, “Improved neural language model fusion for streaming recurrent neural network transducer,” *CoRR*, vol. abs/2010.13878, 2020. [Online]. Available: <https://arxiv.org/abs/2010.13878>
 - [11] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, “End-to-end attention-based large vocabulary speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016, Shanghai, China, March 20-25, 2016*. IEEE, 2016, pp. 4945–4949. [Online]. Available: <https://doi.org/10.1109/ICASSP.2016.7472618>
 - [12] M. Ghodsi, X. Liu, J. Apfel, R. Cabrera, and E. Weinstein, “Rnn-transducer with stateless prediction network,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7049–7053.
 - [13] R. Prabhavalkar, Y. He, D. Rybach, S. Campbell, A. Narayanan, T. Strohmaier, and T. N. Sainath, “Less is more: Improved RNN-T decoding using limited label context and path merging,” *CoRR*, vol. abs/2012.06749, 2020. [Online]. Available: <https://arxiv.org/abs/2012.06749>
 - [14] E. Variani, D. Rybach, C. Allauzen, and M. Riley, “Hybrid autoregressive transducer (HAT),” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*. IEEE, 2020, pp. 6139–6143. [Online]. Available: <https://doi.org/10.1109/ICASSP40776.2020.9053600>
 - [15] E. McDermott, H. Sak, and E. Variani, “A density ratio approach to language model fusion in end-to-end automatic speech recognition,” *CoRR*, vol. abs/2002.11268, 2020. [Online]. Available: <https://arxiv.org/abs/2002.11268>
 - [16] Z. Meng, S. Parthasarathy, E. Sun, Y. Gaur, N. Kanda, L. Lu, X. Chen, R. Zhao, J. Li, and Y. Gong, “Internal language model estimation for domain-adaptive end-to-end speech recognition,” in *SLT 2021*. IEEE, January 2021.
 - [17] 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)*, 2018, pp. 418–425.
 - [18] M. Jain, G. Keren, J. Mahadeokar, G. Zweig, F. Metzger, and Y. Saraf, “Contextual RNN-T for Open Domain ASR,” in *Proc. Interspeech 2020*, 2020, pp. 11–15. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-2986>
 - [19] D. Zhao, T. N. Sainath, D. Rybach, P. Rondon, D. Bhatia, B. Li, and R. Pang, “Shallow-Fusion End-to-End Contextual Biasing,” in *Proc. Interspeech 2019*, 2019, pp. 1418–1422. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-1209>
 - [20] F. Weninger, F. Mana, R. Gemello, J. Andrés-Ferrer, and P. Zhan, “Semi-supervised learning with data augmentation for end-to-end ASR,” in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 2802–2806.
 - [21] A. Graves, N. Jaitly, and A. Mohamed, “Hybrid speech recognition with deep bidirectional lstm,” in *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, 2013, pp. 273–278.
 - [22] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
 - [23] T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 1412–1421. [Online]. Available: <https://www.aclweb.org/anthology/D15-1166>
 - [24] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
 - [25] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 2613–2617. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2680>
 - [26] P. J. Werbos, “Backpropagation through time: what it does and how to do it,” *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
 - [27] J. Chorowski and N. Jaitly, “Towards better decoding and language model integration in sequence to sequence models,” in *Proc. of INTERSPEECH*. Stockholm, Sweden: ISCA, 2017, pp. 523–527.