



Early Stage LM Integration Using Local and Global Log-Linear Combination

Wilfried Michel, Ralf Schlüter, Hermann Ney

Human Language Technology and Pattern Recognition, Computer Science Department,
RWTH Aachen University, 52056 Aachen, Germany
AppTek GmbH, 52062 Aachen, Germany
{michel, schluter, ney}@cs.rwth-aachen.de

Abstract

Sequence-to-sequence models with an implicit alignment mechanism (e.g. attention) are closing the performance gap towards traditional hybrid hidden Markov models (HMM) for the task of automatic speech recognition. One important factor to improve word error rate in both cases is the use of an external language model (LM) trained on large text-only corpora. Language model integration is straightforward with the clear separation of acoustic model and language model in classical HMM-based modeling. In contrast, multiple integration schemes have been proposed for attention models.

In this work, we present a novel method for language model integration into implicit-alignment based sequence-to-sequence models. Log-linear model combination of acoustic and language model is performed with a per-token renormalization. This allows us to compute the full normalization term efficiently both in training and in testing.

This is compared to a global renormalization scheme which is equivalent to applying shallow fusion in training.

The proposed methods show good improvements over standard model combination (shallow fusion) on our state-of-the-art LibriSpeech system. Furthermore, the improvements are persistent even if the LM is exchanged for a more powerful one after training.

Index Terms: sequence-to-sequence, language model integration, deep fusion, shallow fusion

1. Introduction

For the task of automatic speech recognition (ASR) state-of-the-art approaches rely on statistical models which are estimated on data. The most valuable kind of data is parallel data with audio signals and the corresponding transcriptions. We need this kind of data to train our acoustic models (AM). But unfortunately the amount of parallel data is usually quite limited and it is the most expensive data to obtain.

The amount of available uni-modal data (e.g. audio only or text only) is often several orders of magnitude larger than the amount of parallel data and significant improvements can be obtained by utilizing this additional data. For audio only data [1] and monolingual data in the translation setting [2] the missing labels are created using a pre-trained model and the data can then be added to the available parallel data. The same can be done for text only data in the ASR setting [3], but this generally involves a lot of computational effort and constitutes an ill posed problem.

The standard approach for using text only data is to estimate language models, which are then integrated into the speech recognition system. There are several methods to integrate language models depending on the structure of the acoustic model and the solution to the alignment problem.

1.1. Language model integration for HMM based models

The most straightforward approach exists for hidden Markov model (HMM) based models where the need for a language model arises naturally from the Bayes decomposition $P(W|X) = P(X|W) \cdot P(W)/P(X)$. During the training of the models the denominator term $P(X)$ usually is ignored, which leads to a decomposition of the estimation problem into two separate instances. Language model and acoustic model are then estimated independent of each other.

Further improvements can be obtained by including a pre-trained language model into the acoustic model training. This can be achieved by including $P(X)$ in the training criterion, which leads to the maximum mutual information (MMI) [4] training criterion, or by switching to a Bayes risk based training criterion such as minimum phoneme error (MPE) [5] or state-level minimum Bayes risk (sMBR) [6].

1.2. Language model integration for attention models

For models with an implicit alignment model such as attention, no language model is needed in principle as these models directly output the word posterior probability $P(W|X)$. However, it has been shown that including additional text only data is helpful independent of the method used [7]. Commonly used LM integration techniques include *shallow fusion* [8], *deep fusion* [8], and *cold fusion* [9].

Shallow fusion is most similar to the combination of HMM based models as it uses log-linear combination of the word probabilities $P(W|X) = P_{AM}^\alpha(W|X) \cdot P_{LM}^\beta(W)/Z$ in place of the Bayes decomposition. The normalization term Z is usually ignored which again leads to a decomposition of the training problem.

Deep and cold fusion both integrate an explicit language model into the acoustic model in using the LM hidden states or outputs as additional features for the acoustic model decoder. These approaches also include the LM in the training process, but no significant improvements over shallow fusion have been observed in [7]. The reader is referred to [7] for an in-depth comparison of these combination approaches.

In this work we propose a novel combination technique that combines the simplicity and mathematical elegance of log-linear model combination with the early training integration of cold or deep fusion.

2. Local and global renormalization

In this work we always use the posterior log probability $p(\tilde{w}_1^N | x_1^T)$ of the reference sentence \tilde{w}^N given the acoustic features x_1^T as the starting point for training criterion:

$$F = \log p(\tilde{w}_1^N | x_1^T) \quad (1)$$

Different training procedures will arise by how this probability is being modeled.

2.1. Standard cross entropy

In most of the current systems the sentence posterior is decomposed over word positions and then directly modeled by the softmax output of a recurrent decoder network.

$$F_{\text{CE}} = \log p(\tilde{w}_1^N | x_1^T) \quad (2)$$

$$= \log \prod_{n=1}^N p_n(\tilde{w}_n | \tilde{w}_1^{n-1}, x_1^T) \quad (3)$$

$$= \sum_{n=1}^N \log q_{\text{AM},n}(\tilde{w}_n | \tilde{w}_1^{n-1}, x_1^T) \quad (4)$$

This is often referred to as *cross entropy* training criterion. In this case no external language model is present in training and no text-only data can be used.

2.2. Maximum mutual information (MMI)

A straightforward way to include an external LM is by linear model combination. This is usually done only during decoding and then called *shallow fusion*. In this work, we propose to include the LM also during the training of the AM. This is similar to what has been proposed for RNN-T models in [10].

$$F_{\text{MMI}} = \log p(\tilde{w}_1^N | x_1^T) \quad (5)$$

$$= \log \frac{q_{\text{AM}}^\alpha(\tilde{w}_1^N | x_1^T) \cdot q_{\text{LM}}^\beta(\tilde{w}_1^N)}{\sum_{w_1^N} q_{\text{AM}}^\alpha(w_1^N | x_1^T) \cdot q_{\text{LM}}^\beta(w_1^N)} \quad (6)$$

$$\begin{aligned} &= \alpha \sum_{n=1}^N \log q_{\text{AM},n}(\tilde{w}_n | \tilde{w}_1^{n-1}, x_1^T) \\ &+ \beta \sum_{n=1}^N \log q_{\text{LM},n}(\tilde{w}_n | \tilde{w}_1^{n-1}) \\ &- \log \sum_{w_1^N} \prod_{n=1}^N q_{\text{AM},n}^\alpha(w_n | w_1^{n-1}, x_1^T) \cdot q_{\text{LM},n}^\beta(w_n | w_1^{n-1}) \end{aligned} \quad (7)$$

In decoding we are only interested in the word sequence that maximizes the probability so the denominator is omitted as it does not contribute to the decision. During training we need the full probability. The denominator is especially important for training as this is the only place where the external language model enters the gradient of the acoustic model.

The denominator contains a sum over all possible word sequences which is infeasible to compute in practice. In our experiments, this sum will be approximated by an n -best list.

In Equation 6 two scales (α : AM-scale, β : LM-scale) are introduced. Unlike in decoding, where only the ratio between these scales is important, here, because of the denominator, also the absolute magnitude matters.

This training criterion looks very similar to the *Maximum Mutual Information* (MMI) training criterion used for sequence discriminative training of hybrid-HMM models [4]. Hence we will call it MMI criterion in the following.

2.3. Local fusion

In the previous section we first decomposed the posterior (1) into acoustic and language model and then into different word

positions. Now we do it in the opposite order:

$$F_{\text{local}} = \log p(\tilde{w}_1^N | x_1^T) \quad (8)$$

$$= \log \prod_{n=1}^N p_n(\tilde{w}_n | \tilde{w}_1^{n-1}, x_1^T) \quad (9)$$

$$= \sum_{n=1}^N \log \frac{q_{\text{AM},n}^\alpha(\tilde{w}_n | \tilde{w}_1^{n-1}, x_1^T) \cdot q_{\text{LM},n}^\beta(\tilde{w}_n | \tilde{w}_1^{n-1})}{\sum_w q_{\text{AM},n}^\alpha(w | \tilde{w}_1^{n-1}, x_1^T) \cdot q_{\text{LM},n}^\beta(w | \tilde{w}_1^{n-1})} \quad (10)$$

Here we see, that instead of one sum over *all word sequences* which is in $\mathcal{O}(V^N)$ we have N sums over the vocabulary which is in $\mathcal{O}(V \cdot N)$. This makes it much more tractable and we can calculate all sums exactly. We denote this criterion *local fusion*.

We have again two scales (α , β) which are both relevant due to the sum in the denominator.

Also note that in Equation 10 the probabilities in the denominator are always conditioned on the reference history \tilde{w}_1^{n-1} and not on a separate history w_1^{n-1} as in the MMI criterion (Equation 6). This implies that the denominator cannot be dropped at decoding time since it will depend on the word sequence we are optimizing over. It also means that while MMI is a sequence discriminative criterion, local fusion still only is token-level discriminative.

3. Experimental setup

We will investigate the proposed methods on the full LibriSpeech corpus (1000h) [11].

The same acoustic model architecture as in our previous works [12, 3] will be used for all experiments. The encoder consists of a combination of CNN followed by 6 BLSTM layers with time sub-sampling by max-pooling. We use MLP-style attention with weight feedback and a decoder with a single LSTM layer. As input to the encoder we use 40 dimensional MFCC features which are perturbed by a variant of SpecAugment [13, 14]. The output consists of 10k grapheme BPE units.

The training of the network involves an intricate pre-training scheme with gradually increasing layer size and an additional ctc-loss on the encoder outputs. The complete scheme can be found in [12]. Just like in [3] we will store one checkpoint after the pre-training phase but well before convergence and if not noted otherwise use this as initialization for all further experiments. For some experiments (notably all MMI related experiments) we continued training from this checkpoint until convergence with the CE criterion and use this converged CE model as initialization for further training.

For our experiments we use one out of two neural language models. The *LSTM* LM consists of 4 layers with 2k units each. The *Transformer* LM consists of 24 layers with dimension 1k and uses multiheaded attention (8 heads) as described in [15]. Both LMs operate on the same 10k BPE vocabulary as the AM. If not noted otherwise, the LSTM LM has been used.

In all experiments where the acoustic model is augmented by an LM, we have the choice to do local or global renormalization. For the training this is denoted as MMI in the global case or local fusion in the local case. For decoding we again have the same choice. We tried exchanging the normalization techniques but found that using the matching technique gave slight improvements in some cases. Therefore, we use shallow fusion to decode MMI and CE models and local fusion otherwise.

All Experiments have been performed using our speech recognition toolkits RASR [16] and RETURNN [17], config files are available at our GitHub repository [18].

3.1. Practical guidelines for the MMI criterion

While the local fusion criterion worked out of the box, there were some subtleties for the MMI criterion we would like to point out.

Replacing the sum over all word sequences by an n -best list is an approximation whose accuracy strongly depends on an appropriate choice of n . A larger n , however, requires increased memory and computation time. Here we chose $n = 8$, which was the maximum we could fit in our current GPU memory given our current decoder. Some works suggest sequence training to be possible with even lower beam sizes [19].

Similar to what has been found for hybrid HMM sequence training [20], we found it imperative to always include the reference transcript in the n -best list. If the reference transcript is absent from the top n results, the hypothesis with lowest probability is replaced by the reference. Otherwise the beam is kept intact.

The decoding of sequence to sequence models suffers a length bias problem [21]. While there are some ways around it without modification of the hypothesis scores [22], the prevalent solution is applying a length normalization [23] or silence penalty to the decoding scores [24]. Naturally we must not distort the scores for the MMI criterion. All experiments are performed with the vanilla AM and LM scores.

4. Results

4.1. Runtime comparison

It is expected that the more complex training criteria which also involve forwarding through a language model and the computation of the normalization term will take more time per training epoch than the cross entropy criterion. In Table 1 we compare the average training time per sub-epoch (50h) using a single GPU.

For the training using the local fusion criterion we always compute the full normalization. For the MMI training we approximate the sum by an n -best list with $n = 8$. Also the batch size had to be reduced greatly for the MMI training which results in longer training times.

Table 1: Training time in minutes per sub epoch of 50h for different training criteria and language models used in training.

Criterion	LM	time [min]	slowdown factor
CE	-	48	1.00
local	LSTM	51	1.06
fusion	TRAFO	49	1.04
MMI	LSTM	249	5.26

We see that the local LM combination does not slow down the training significantly. The MMI training increases training time by a factor of 5. We therefore decided to start all MMI related experiments from a fully converged CE model as initialization as it is usually the case for the MMI training of hybrid HMM models [20].

4.2. Local fusion training criterion

4.2.1. AM and LM scales

From Equation 10 we observe that for local fusion both relative (γ_{rel}) and absolute scale (γ_{abs}) of the acoustic and language model influence the result. We define these scales as: $\gamma_{abs} = \alpha$

and $\gamma_{rel} = \beta/\gamma_{abs}$. As a starting point of our tuning experiments we use $\gamma_{rel} = 0.35$ which we estimated from decoding experiments.

Table 2: Comparing performance of models trained with local fusion training criterion with LSTM LM and different scales.

absolute scale γ_{abs}	relative scale γ_{rel}	dev WER [%]	
		clean	other
1.0	0.0	2.8	7.9
5.0	0.35	3.6	10.4
3.0	0.35	2.6	7.6
2.0	0.35	2.6	7.3
1.5	0.35	2.7	7.3
1.2	0.35	2.7	7.4
1.0	0.35	2.9	8.1
0.5	0.35	3.4	9.9
3.0	0.5	2.8	7.7
3.0	0.4	2.7	7.7
3.0	0.35	2.6	7.6
3.0	0.3	2.7	7.5
3.0	0.25	2.7	8.6
2.0	0.5	2.6	7.4
2.0	0.35	2.6	7.3
2.0	0.25	2.8	7.6

In Table 2 we report the result of the scale tuning. First we notice that the decoding optimum is also a good choice for training. For the absolute scale we find an optimum at around 2 – 3. This is in accordance to the “old rule of thumb” we know from sequence training of hybrid HMM models, that the training LM scale should be set to 1 and AM scale should be the inverse of the optimum decoding LM scale. [25]

4.2.2. Interchangeability of LM

When we train the AM together with an LM then it learns to work together with the LM at decoding. Now we wonder if the AM learns to make use of LMs in general or if it is adapted to the specific LM we used in training. We therefore train our AM together with the LSTM LM and use it together with the transformer LM in decoding. The results are shown in Table 3.

Table 3: Comparing the effect of using a different LM for training and testing. Models trained with local fusion criterion except for the first block, which is CE.

training LM	testing LM	dev WER [%]	
		clean	other
-	-	3.9	10.6
	LSTM	2.8	7.9
	TRAFO	2.6	7.3
LSTM	-	4.9	12.7
	LSTM	2.5	6.9
	TRAFO	2.3	6.6
TRAFO	TRAFO	2.4	6.5

The first observation of Table 3 is that the transformer LM improves WER compared to the LSTM LM in any case and using the local fusion criterion always improves significantly over the CE baseline. Using the transformer LM with the local fusion criterion nicely stacks both improvements.

In this simple setting of matched LMs it is indeed possible to exchange the decoding LM without the need for retraining the

AM. We also notice that decoding a local fusion model without LM degrades performance compared to the CE baseline.

4.2.3. Joint training of AM and LM

Until now we have only considered the training of the acoustic model while keeping the language model as is. Within the framework set by the local fusion training criterion it is also possible to jointly train acoustic and language model. For this we initialized the AM with the converged CE model and started a joint training with the LSTM LM and local fusion criterion.

The result (WER: 3.2% dev-clean, 8.3% dev-other) was much worse than even using the initial CE model together with the initial language model so that we quickly abandoned this approach. We attribute the degradation to the fact that the rich amount of text only data used to train the initial model is absent during this training stage so that *catastrophic forgetting* [26] happens.

4.3. MMI training criterion

4.3.1. Scale tuning

Also when using the MMI training criterion the absolute magnitude of AM and LM scale matters. We would expect similar results as with the local fusion criterion for attention and the MMI criterion for HMM based models.

Table 4: Comparing performance of models trained with MMI training criterion with LSTM LM.

absolute scale γ_{abs}	relative scale γ_{rel}	dev WER [%]	
		clean	other
1.0	0.0	2.8	7.9
5.0	0.35	2.6	7.3
1.0	0.35	2.5	7.1
0.5	0.35	2.5	6.9
0.1	0.35	2.4	6.6
0.01	0.35	2.4	6.6

In Table 4 we see that at high absolute scales the performance of MMI trained models is comparable to the performance of local fusion trained models. With reduced absolute scale, however, the performance of MMI models increases while local fusion models suffer from reduced performance.

4.3.2. Cross entropy smoothing

Now we investigate if cross entropy smoothing, an invaluable heuristic for hybrid HMM based sequence training [20], will be useful in this setting. For cross entropy smoothing the MMI objective function is linearly interpolated with the cross entropy objective function. Usually only a small weight (e.g. 10%) is assigned to the CE criterion.

Here we introduce an additional scale γ_{den} as a prefactor to the third line of Equation 7. A scale of $\gamma_{den} = 1$ then corresponds to pure MMI while $\gamma_{den} = 0$ would indicate a pure cross entropy criterion.

As we can see in Table 5 any value lower than 1 reduces the model performance. We tried values larger than 1, but found that the models diverge quickly. We therefore conclude that CE smoothing is not useful for attention based models.

4.4. Comparison of training criteria

In Table 6 we compare the best models we could get with each training criterion. In this case we also give the word error rate

Table 5: Comparing performance of models trained with MMI training criterion with LSTM LM and different cross entropy smoothing settings. Scales were set to $\gamma_{abs} = 0.3$, $\gamma_{rel} = 0.35$

denominator scale γ_{den}	dev WER [%]	
	clean	other
0.0	2.8	7.9
0.1	2.7	7.7
0.3	2.7	7.5
0.6	2.6	7.3
0.9	2.5	7.0
1.0	2.4	6.8

on the test sets. The MMI criterion clearly outperforms the local fusion criterion, which in turn improves over the CE baseline.

Table 6: Comparing performance of different training criteria. All hyperparameters were optimized individually on the respective dev sets. Numbers are WER[%]

init. model	train. crit.	testing LM	dev WER[%]		test WER[%]	
			clean	other	clean	other
pre-	CE	LSTM	2.8	7.9	3.0	8.4
		TRAFO	2.6	7.3	2.8	7.8
train	local	LSTM	2.6	7.3	2.9	7.8
		TRAFO	2.3	6.4	2.6	6.9
con-	local	LSTM	2.5	6.9	2.8	7.6
		TRAFO	2.3	6.5	2.6	6.9
verged	MMI	LSTM	2.4	6.6	2.6	7.0
		TRAFO	2.2	6.1	2.3	6.4

5. Conclusion

In this work we proposed two methods to include additional text only data into the training of attention based implicit alignment models in form of an external trained language model. Both methods improve over the cross entropy baseline with shallow fusion at decoding time.

The *MMI* criterion leads to the largest improvement of 18% relative on test-other, but at the cost of a five-fold increase in training time. This criterion can be applied in a late training stage for fine-tuning of a pretrained model.

The *local fusion* criterion improves over the baseline by 10% rel. on test-other without requiring additional effort or resources. We see no reason to not adapt to this training criterion immediately.

6. Acknowledgments

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 694537, project ”SEQCLAS”) and from a Google Focused Award. The work reflects only the authors’ views and none of the funding parties is responsible for any use that may be made of the information it contains. The GPU cluster used for the experiments was partially funded by Deutsche Forschungsgemeinschaft (DFG) Grant INST 222/1168-1. Simulations were partially performed with computing resources granted by RWTH Aachen University under project nova0003.

The authors would like to thank Albert Zeyer and Wei Zhou for many helpful discussions and Kazuki Irie for providing the pretrained language models.

7. References

- [1] Y. Long, Y. Li, S. Wei, Q. Zhang, and C. Yang, "Large-scale semi-supervised training in deep learning acoustic model for ASR," *IEEE Access*, vol. 7, pp. 133 615–133 627, 2019.
- [2] M. Przystupa and M. Abdul-Mageed, "Neural machine translation of low-resource and similar languages with backtranslation," in *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 3: Shared Task Papers, Day 2*, 2019, pp. 224–235.
- [3] N. Rossenbach, A. Zeyer, R. Schlüter, and H. Ney, "Generating synthetic audio data for attention-based speech recognition systems," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Barcelona, Spain, May 2020.
- [4] Y. Normandin, "Maximum mutual information estimation of hidden markov models," in *Automatic Speech and Speaker Recognition*, vol. 355, pp. 57–81, 01 1996.
- [5] D. Povey and P. C. Woodland, "Minimum phone error and i-smoothing for improved discriminative training," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2002, May 13-17 2002, Orlando, Florida, USA*. IEEE, 2002, pp. 105–108.
- [6] M. Gibson and T. Hain, "Hypothesis spaces for minimum bayes risk training in large vocabulary speech recognition," in *INTER-SPEECH 2006 - ICSLP, Ninth International Conference on Spoken Language Processing, Pittsburgh, PA, USA, September 17-21, 2006*, 2006.
- [7] S. Toshniwal, A. Kannan, C. Chiu, Y. Wu, T. N. Sainath, and K. Livescu, "A comparison of techniques for language model integration in encoder-decoder speech recognition," in *2018 IEEE Spoken Language Technology Workshop, SLT 2018, Athens, Greece, December 18-21, 2018*, dec 2018, pp. 369–375.
- [8] Ç. Gülgehre, O. Firat, K. Xu, K. Cho, L. Barrault, H. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On using monolingual corpora in neural machine translation," 2015. [Online]. Available: <http://arxiv.org/abs/1503.03535>
- [9] A. Sriram, H. Jun, S. Sathesh, 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*. ISCA, 2018, pp. 387–391.
- [10] C. Weng, C. Yu, J. Cui, C. Zhang, and D. Yu, "Minimum bayes risk training of rnn-transducer for end-to-end speech recognition," *CoRR*, 2019. [Online]. Available: <http://arxiv.org/abs/1911.12487>
- [11] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, 2015, pp. 5206–5210.
- [12] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, "A comparison of transformer and LSTM encoder decoder models for ASR," in *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2019, Singapore, December 14-18, 2019*, 2019, pp. 8–15.
- [13] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, 2019, pp. 2613–2617.
- [14] W. Zhou, W. Michel, K. Irie, M. Kitzka, R. Schlüter, and H. Ney, "The RWTH ASR system for TED-LIUM release 2: Improving hybrid HMM with specaugment," may 2020.
- [15] K. Irie, A. Zeyer, R. Schlüter, and H. Ney, "Language modeling with deep transformers," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, 2019, pp. 3905–3909.
- [16] D. Rybach, S. Hahn, P. Lehnen, D. Nolden, M. Sundermeyer, Z. Tüske, S. Wiesler, R. Schlüter, and H. Ney, "RASR - the RWTH aachen university open source speech recognition toolkit," in *IEEE Automatic Speech Recognition and Understanding Workshop*, Waikoloa, HI, USA, Dec. 2011.
- [17] P. Doetsch, A. Zeyer, P. Voigtlaender, I. Kulikov, R. Schlüter, and H. Ney, "Returnn: The RWTH extensible training framework for universal recurrent neural networks," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2017, New Orleans, LA, USA, March 5-9, 2017*, 2017, pp. 5345–5349.
- [18] Config files for the experiments in this paper. [Online]. Available: <https://github.com/rwth-i6/returnn-experiments/tree/master/2020-lm-integration>
- [19] R. Prabhavalkar, T. N. Sainath, Y. Wu, P. Nguyen, Z. Chen, C. Chiu, and A. Kannan, "Minimum word error rate training for attention-based sequence-to-sequence models," in *IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP*, April 2018, pp. 4839–4843.
- [20] K. Veselý, A. Ghoshal, L. Burget, and D. Povey, "Sequence-discriminative training of deep neural networks," in *INTER-SPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France, August 25-29 2013*, pp. 2345–2349.
- [21] P. Sountsov and S. Sarawagi, "Length Bias in Encoder Decoder Models and a Case for Global Conditioning," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP), Austin, Texas, USA, November 1-4, 2016*, pp. 1516–1525.
- [22] W. Zhou, R. Schlüter, and H. Ney, "Robust beam search for encoder-decoder attention based speech recognition without length bias," in *Submitted to Interspeech 2020, 21th Annual Conference of the International Speech Communication Association, Shanghai, China, 25-29 October 2020*, 2020.
- [23] K. Murray and D. Chiang, "Correcting Length Bias in Neural Machine Translation," in *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pp. 212–223.
- [24] A. Hannun, A. Lee, Q. Xu, and R. Collobert, "Sequence-to-sequence speech recognition with time-depth separable convolutions," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, 2019, pp. 3785–3789.
- [25] R. Schlüter, B. Müller, F. Wessel, and H. Ney, "Interdependence of language models and discriminative training," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, vol. 1, Keystone, CO, Dec 1999, p. 119–122.
- [26] A. V. Robins, "Catastrophic forgetting in neural networks: the role of rehearsal mechanisms," in *First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems, ANNES '93, Dunedin, New Zealand, November 24-26, 1993*, 1993, pp. 65–68.