

Contextualizing ASR Lattice Rescoring with Hybrid Pointer Network Language Model

Da-Rong Liu[†], Chunxi Liu^{*}, Frank Zhang^{*}, Gabriel Synnaeve^{*}, Yatharth Saraf^{*}, Geoffrey Zweig^{*}

[†]National Taiwan University

^{*}Facebook AI, USA

f07942148@ntu.edu.tw {chunxiliu, frankz, gab, ysaraf, gzweig}@fb.com

Abstract

Videos uploaded on social media are often accompanied with textual descriptions. In building automatic speech recognition (ASR) systems for videos, we can exploit the contextual information provided by such video metadata. In this paper, we explore ASR lattice rescoring by selectively attending to the video descriptions. We first use an attention based method to extract contextual vector representations of video metadata, and use these representations as part of the inputs to a neural language model during lattice rescoring. Secondly, we propose a hybrid pointer network approach to explicitly interpolate the word probabilities of the word occurrences in metadata. We perform experimental evaluations on both language modeling and ASR tasks, and demonstrate that both proposed methods provide performance improvements by selectively leveraging the video metadata.

Index Terms: speech recognition, pointer network, video metadata, lattice rescoring

1. Introduction

Personalized or contextual automatic speech recognition, which aims to improve accuracy by leveraging additional information or external knowledge, has been an important research topic [1, 2, 3, 4, 5]. These prior works usually assume that a set of word-level biasing phrases are known ahead of time, e.g. a user’s personal contact list, which are used to nudge the ASR model towards outputting these particular phrases. In the conventional hybrid ASR framework, bias phrases can be compiled into a weighted finite state transducer (WFST), with vocabulary injection, and on-the-fly language model biasing techniques [1, 2] have shown significant performance gains. Similarly, for end-to-end ASR architectures like Listen, Attend and Spell (LAS) [6], the WFST representation of context n -grams is traversed along with the outputs from the LAS network, and the beam search decoding is biased either in first pass decoding [5] or rescoring [3]. Alternatively, each context n -gram can also be embedded into a fixed dimensional representation, and such contextual information is summarized by an attention mechanism and further fed as an additional input to the decoder [4, 7].

One main distinction in our work is that rather than using a list of named entities, such as user’s contact lists or song names as in many prior approaches, here we aim to exploit contextual information from word sequences or paragraphs, as illustrated in Figure 1. Henceforth, we refer to such textual content as *video metadata*. Utilizing the video metadata effectively can be challenging, since it not only contains potentially relevant information, but also irrelevant text. To address this challenge,

[†] Work was done when Da-Rong was an intern at Facebook.



Figure 1: **Video metadata:** Social media videos are often associated with surrounding text, descriptions or titles, as denoted in the red box.

our neural LM training is explicitly conditioned on the video metadata, and selectively attends to the metadata via an attention mechanism. The resulting contextual neural LM is used to rescore the lattices generated from the first-pass hybrid ASR decoding.

Therefore, our proposed methods are similar to [4] with three important distinctions. First, we produce the ASR lattice via a conventional WFST-based hybrid ASR model, and contextual biasing is performed by jointly rescoring the lattice and attending to the metadata. Second, rather than tying the contextual biasing with the end-to-end LAS training, we build the contextual LM separately from the acoustic model training, which effectively allows for modular evaluation and improvements of the contextual LM component. In common with earlier generations of technology, language model changes can be made independently of the acoustic model. Third, based on [4] that utilized contextual information by an attention mechanism, we further propose the hybrid pointer network (which will be introduced in Section 2.2).

While conventional sequence-to-sequence (seq2seq) models typically generate tokens from a predefined vocabulary [8, 9], pointer networks [10] can be used to explicitly select and output tokens from the input (source) sequence. Recently, hybrid pointer-generator networks (PGN), combining seq2seq models with pointer networks, have been proposed and used in summarization tasks [11]. In such models, an additional scalar variable is generated at each time step and serves as a soft switch to choose between generating the token from a predefined vocabulary or selecting from the input sequence. This has been shown to be particularly effective in generating rare words that have very few occurrences in training data [11].

Our main contributions can be summarized into three categories:

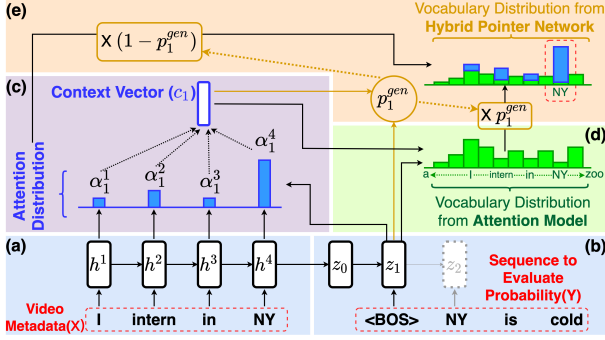


Figure 2: **Contextual Language Model Overview:** In this figure, video metadata is set to “I intern in NY” and the sequence to evaluate probability is set to “NY is cold”. At time-step one, the target of model is predicting $P(“NY”)$. (a) The encoder, which encodes the video metadata into a sequence of hidden vectors h^i . (b) The decoder. (c) The attention mechanism, which generates the attention distribution over encoder hidden vectors, and summarizes them into a context vector. (d) The context vector and the decoder hidden state are used to generate the vocabulary distribution. (e) We can further interpolate the vocabulary distribution with attention distribution to boost the probability of rare word. In general, (a)+(b)+(c)+(d) form the attention model depicted in Section 2.1, and (a)+(b)+(c)+(d)+(e) form the hybrid pointer network depicted in Section 2.2.

- We build a contextual language model that conditions on the video metadata. We compare various alternatives for such a language model and demonstrate that a hybrid pointer network substantially outperforms all competing baselines in perplexities.
- We then use this language model to rescore the lattice generated from the first pass ASR decoding. We employed the pruned lattice rescoring algorithm [12], and show that after rescoring, our contextual LM performs better than all other baseline LMs in word error rate (WER).
- We further perform analysis on how the quality of the video metadata affects the ASR performance.

2. Contextual Language Model

A language model (LM) represents a probability distribution over sequences of N tokens $Y = (y_1, y_2, \dots, y_N)$. Given such a sequence, a LM assigns a probability to the sequence by modeling the probability of token y_k conditioned on its history $\{y_1, y_2, \dots, y_{k-1}\}$. The probability of the whole sequence can be decomposed as:

$$P(y_1, y_2, \dots, y_N) = \prod_{k=1}^N P(y_k | y_1, y_2, \dots, y_{k-1}) \quad (1)$$

In this paper, our language model can be conditioning on the video metadata $\mathbf{X} = \{x_1, x_2, \dots, x_M\}$, where M is the metadata sequence length. Then, the conditional probability can be represented and decomposed as:

$$P(y_1, y_2, \dots, y_N | \mathbf{X}) = \prod_{k=1}^N P(y_k | y_1, y_2, \dots, y_{k-1}, \mathbf{X}) \quad (2)$$

This is trained by minimizing the negative log-likelihood:

$$\sum_{k=1}^N -\log P(y_k | y_1, y_2, \dots, y_{k-1}, \mathbf{X}; \theta_{model}) \quad (3)$$

where θ_{model} is the set of trainable parameters. For modeling $P(y_k | y_1, y_2, \dots, y_{k-1}, \mathbf{X})$, we first describe the attention model in Section 2.1, and then we will show how to adapt the attention model to the hybrid pointer network in Section 2.2.

2.1. Attention Model

The attention model, which is depicted as a component of the hybrid pointer network in Figure 2, is similar to the attention model proposed in [13] with only one difference. The model proposed in [13] is targeted to the translation task, which will auto-regressively generate words at the decoder. The auto-regressive generation means we will generate one word at one time-step, and feed the word as input to the next time-step. However, as our model is a language model [14, 15], the input to the decoder is the word sequence of which we evaluate the probabilities. In this work, these word sequences will be the ASR hypotheses from the lattice. We describe the detail formulation below.

The tokens of the metadata x_i are fed one-by-one into the encoder (multi-layer LSTM), producing a sequence of hidden states h^i . At each time step t , the inputs to the decoder are the word embedding of y_t and the previous hidden state z_{t-1} , generating the current hidden state z_t . The attention distribution is computed as:

$$\hat{\alpha}_t^i = (W_z z_t + b_z)^T h^i \quad (4)$$

$$\alpha_t = \text{softmax}(\hat{\alpha}_t) \quad (5)$$

Next, the attention distribution is used to produce a weighted sum of the encoder hidden states, known as the context vector:

$$c_t = \sum_i \alpha_t^i h^i \quad (6)$$

The context vector can be viewed as the summary of the encoder information, which is then concatenated with decoder hidden state z_t and passed through two linear layers to determine the probability of the next word y_{t+1} :

$$P_{vocab}(w) = \text{softmax}(W'(W[z_t; c_t] + b) + b') \quad (7)$$

$$\begin{aligned} P(y_{t+1} | y_1, y_2, \dots, y_t, \mathbf{X}) \\ = P(y_{t+1} | z_t, c_t) = P_{vocab}(y_{t+1}) \end{aligned} \quad (8)$$

where W, W', b, b' are learnable parameters. P_{vocab} is the distribution over the whole vocabulary.

2.2. Hybrid pointer network

In the attention model, we have incorporated video metadata information while predicting the probabilities. However, it will be difficult for the model to predict words of rare occurrences in the training data. Take Figure 2 as an example. We assume the word “NY” is rare in training data and the attention model is perfect in predicting $p(NY)$, which means:

$$\begin{aligned} \alpha_1^4 &= 1 \\ \alpha_1^1 &= \alpha_1^2 = \alpha_1^3 = 0 \end{aligned}$$

In this way, we can get $c_1 = h^4$. Because “NY” is rare in training data, there will be two consequences. First, the word

Table 1: Dataset sizes in hours

	Train	Valid	Test		
			clean	noisy	extreme
English	14k (hrs)	9.7	20.2	18.6	49.1
Spanish	7.5k	9.9	17.2	19.5	46.1

embedding of “NY”, i.e. h^4 , may not be well trained. Second, even if h^4 does contain the information of the word “NY”, according to equation (8):

$$P(NY) = P(NY|z_1, c_1) = P(NY|z_1, h^4) \quad (9)$$

it can be difficult for the model to map h^4 back to “NY”, because again it is rare in training data.

To address this problem, we propose to use a hybrid pointer network similar to [11]. In a hybrid pointer network, another random variable $P_t^{gen} \in (0, 1)$ is introduced:

$$P_t^{gen} = \sigma(W^{gen}[c_t; z_t; y_t] + b) \quad (10)$$

This P_t^{gen} is used as a soft switch to select between the distribution P_{vocab} generated in equation (7) or sample from the attention distribution $\hat{\alpha}_t$:

$$P_{vocab}^{ptr}(w) = p_t^{gen} P_{vocab}(w) + (1 - p_t^{gen}) \sum_{i:w_i=w} \alpha_t^i \quad (11)$$

When p_t^{gen} is close to one, it means the language model has high confidence to directly generate the distribution. In contrast, if p_t^{gen} is close to zero, it means the model has low confidence, and turns to utilizing the information from metadata. Finally, the probability of the next word y_{t+1} will be:

$$P(y_{t+1}|y_1, y_2, \dots, y_t, X) = P_{vocab}^{ptr}(y_{t+1}) \quad (12)$$

3. Evaluation

We evaluate the effectiveness of our proposed approaches on our in-house English (EN) and Spanish (ES) video datasets, which are sampled from public social media videos and de-identified before transcription. These videos contain a diverse range of speakers, accents, topics, and acoustic conditions making automatic recognition difficult. Each data instance consists of the audio, reference transcription and the corresponding video metadata text. The test sets for each language are composed of `clean`, `noisy` and `extreme` categories, with `extreme` being more acoustically challenging than `clean` and `noisy`. The dataset sizes are shown in Table 1.

We first evaluate the language model perplexities in Section 3.1. Five different language models are compared:

- 5-gram: the language model with Kneser-Ney smoothing and used in the first-pass ASR decoding.
- LSTM: a multi-layer LSTM.
- cache-LSTM: this is a simple way for the LSTM LM to leverage video metadata. We interpolate the output distribution of LSTM LM with the unigram probability distribution of video metadata. The interpolation weight is a tunable hyperparameter.
- attention model: the model described in Section 2.1
- hybrid pointer network: the model described in Section 2.2

Table 2: Perplexities of each language model on the test sets. The values in the parenthesis for (c) and (d) denote the interpolation weights.

	English		
	clean	noisy	extreme
(a) 5-gram	129.9	150.1	150.4
(b) LSTM	109.6	114.9	119.6
(c) cache-LSTM (0.1)	105.1	115.35	119.8
(d) cache-LSTM (0.2)	113.0	125.7	130.8
(e) attention model	99.1	106.2	110.2
(f) hybrid pointer network	76.9	91.0	95.2
	Spanish		
	clean	noisy	extreme
(a) 5-gram	176.4	194.0	209.6
(b) LSTM	119.0	130.3	151.6
(c) cache-LSTM (0.1)	118.0	137.2	160.4
(d) cache-LSTM (0.2)	127.6	151.7	177.8
(e) attention model	107.4	118.9	139.7
(f) hybrid pointer network	84.4	101.4	121.2

Specifically, the 5-gram and LSTM LMs do not use the information from the video metadata, while the remaining do. A two-layer LSTM with 512 hidden units and 0.1 dropout are used as the recurrent part of all recurrent models. Adaptive softmax [16] is used for efficient training. We implement all our recurrent LMs based on the fairseq toolkit [17]. Kaldi decoder [18] is used to produce ASR lattices. All language models, except the 5-gram model, are then used to perform lattice rescoring. The WER comparisons are shown in Section 3.2.

3.1. Language Modeling

We first evaluate the effectiveness of each LM in terms of perplexity. LMs are trained using both transcriptions and video metadata from the `train` set. We build our vocabulary as all words seen in the training data, while leaving remaining words as OOVs. We use cosine learning rate (LR) scheduler [19] and NAG optimizer with initial LR 0.001. The results are shown in Table 2.

We can see there are significant improvements from the 5-gram LM to each of the recurrent neural LMs, i.e., (a) vs (b), (c), (d), (e), and (f). Comparing the cache-LSTM with the LSTM ((b) vs (c), (d)), we see that even though the cache-LSTM does leverage the video metadata, it does not perform better than the LSTM trained only on transcripts. This indicates that naively interpolating with the unigram distribution of the metadata may not be helpful. The attention model performs better than LSTM and cache-LSTM ((e) vs (b), (c), (d)), because in attention LSTM, the model automatically learns how to leverage the video metadata. Finally, the hybrid pointer network performs best, as it can overcome the shortcomings of the attention model as described in Section 2.2.

3.2. ASR performance

We now evaluate the effectiveness of the proposed method on the ASR task. We first produce lattices via first-pass decoding with a graphemic hybrid ASR system [20]. For acoustic modeling, we utilize a hybrid ASR model with a graphemic lexicon trained with the lattice-free MMI criterion [21]. In the first-pass decoding, we use Kaldi decoder with 5-gram LM (from Table2 (a)) to generate lattices, with lattice beam 8. The lattices

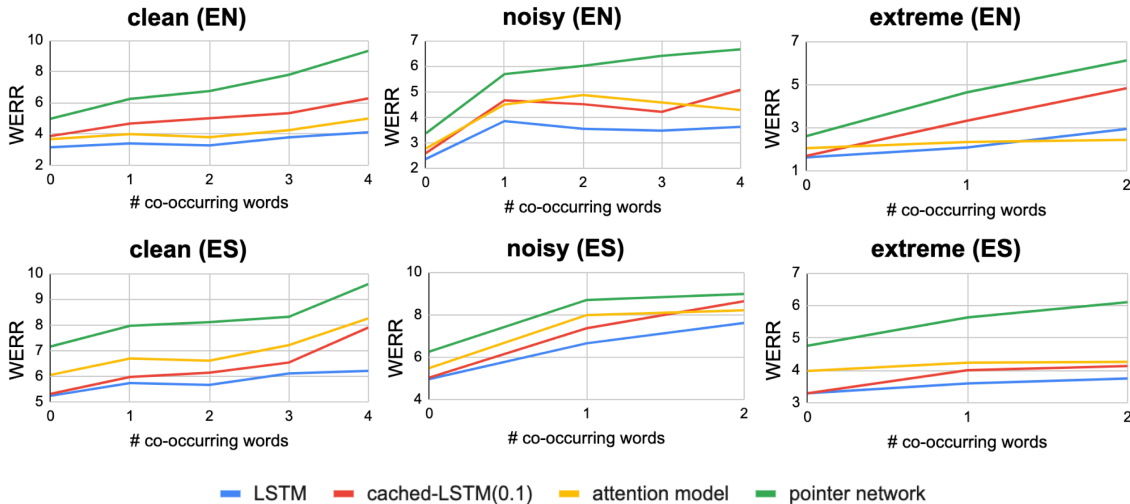


Figure 3: Analysis: in this figure, we report the WERR as the number of co-occurring words between video metadata and transcription varies. WERR denotes the relative WER reduction compared to the WER of first-pass decoding. We only report the number if the remaining test data size is more than half an hour, since there could be too much variance if the remaining test data size is too small.

are then rescored with the neural LMs by the pruned lattice algorithm [12]. A 5-gram approximation is adopted [12, 22] to reduce the search space, i.e., we merge search paths containing the same last 5 history words.

The results are shown in Table 3. As expected, neural lattice rescoring improves the first-pass decoding results ((a) vs (b)). The cache-LSTM does marginally better than LSTM ((b) vs (c)), but the improvement is small and unstable. While attention model can further boost the performance, the hybrid pointer network achieves the best performance.

Table 3: ASR results in WER. In rescoring with neural LMs, we interpolate the neural LM and n -gram LM scores as in [12]. The weight is tuned on valid set, and is shown in parenthesis for each language.

	English (0.6)		
	clean	noisy	extreme
(a) first-pass	15.3	22.0	28.3
(b) LSTM	14.81	21.47	27.8
(c) cache-LSTM (0.1)	14.7	21.4	27.8
(d) attention model	14.7	21.4	27.7
(e) hybrid pointer network	14.5	21.3	27.6
	Spanish (0.7)		
	clean	noisy	extreme
(a) first-pass	13.6	15.5	21.9
(b) LSTM	12.8	14.7	21.1
(c) cache-LSTM (0.1)	12.8	14.7	21.1
(d) attention model	12.7	14.7	21.0
(e) hybrid pointer network	12.6	14.5	20.8

Table 4: The proportion of test data that has none of video metadata in each test category.

	clean	noisy	extreme
English	78 / 1203	296 / 903	3214 / 7092
Spanish	53 / 732	452 / 821	971 / 2026

3.3. Analysis

Although we have shown the feasibility of hybrid pointer network in Section 3.2, we can see that the relative WER reduction (WERR) from LSTM to pointer LSTM is stable, but up to 2% (Table 3 (b) vs (e)). The improvements can be correlated with the video metadata quality. Table 4 shows there is a large portion of dataset where the video metadata is absent. While the video metadata text could be irrelevant to the video transcriptions, being null will certainly limit its effectiveness. To understand the importance of the available metadata size, Figure 3 shows the results as the number of co-occurring words in both metadata and reference transcription varies. In our analysis, we select the test data instances with 1, 2, 3 or 4 co-occurring words between transcription and video metadata. The 3000 most frequent words are not counted as co-occurring words, because hypothetically a co-occurring high frequency word, like ‘the’, may not indicate the video metadata quality.

The results are shown in Figure 3. As the number of co-occurring words increases, we first observe that, the WERR given by hybrid pointer network is more significant than the gain seen in the overall test set. Although cache-LSTM and attention-LSTM also utilize the video metadata information, the WERR is not as substantial as in the pointer network. Also, comparing the pointer network and the LSTM models, the gap between the WERR curves increase in most cases as the number of co-occurring words grows, which indicates the effectiveness of our proposed method given a reasonable video metadata quality.

4. Conclusions

In this work, we propose the use of a hybrid pointer network LM for lattice rescoring, thus making use of text metadata accompanying social media videos. We analyze the conditions of its effectiveness, and demonstrate that it can provide improvements in both LM perplexity and ASR WER. Also, in the hybrid pointer network framework, we can replace the recurrent components of multi-layer LSTM with other neural models, such as neural transformers [23].

5. References

- [1] K. Hall, E. Cho, C. Allauzen, F. Beaufays, N. Coccaro, K. Nakajima, M. Riley, B. Roark, D. Rybach, and L. Zhang, "Composition-based on-the-fly rescoring for salient n-gram biasing," 2015.
- [2] I. McGraw, R. Prabhavalkar, R. Alvarez, M. G. Arenas, K. Rao, D. Rybach, O. Alsharif, H. Sak, A. Gruenstein, F. Beaufays *et al.*, "Personalized speech recognition on mobile devices," in *Proc. ICASSP*, 2016.
- [3] I. Williams, A. Kannan, P. S. Aleksic, D. Rybach, and T. N. Sainath, "Contextual speech recognition in end-to-end neural network systems using beam search."
- [4] G. Pundak, T. N. Sainath, R. Prabhavalkar, A. Kannan, and D. Zhao, "Deep context: end-to-end contextual speech recognition," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 418–425.
- [5] Z. Chen, M. Jain, Y. Wang, M. L. Seltzer, and C. Fuegen, "End-to-end contextual speech recognition using class language models and a token passing decoder," in *Proc. ICASSP*, 2019.
- [6] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell," *arXiv preprint arXiv:1508.01211*, 2015.
- [7] Z. Chen, M. Jain, Y. Wang, M. L. Seltzer, and C. Fuegen, "Joint grapheme and phoneme embeddings for contextual end-to-end ASR," in *Proc. Interspeech 2019*, 2019.
- [8] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [9] R. Nallapati, B. Zhou, C. Gulcehre, B. Xiang *et al.*, "Abstractive text summarization using sequence-to-sequence rnns and beyond," *arXiv preprint arXiv:1602.06023*, 2016.
- [10] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer networks," in *Advances in Neural Information Processing Systems*, 2015, pp. 2692–2700.
- [11] A. See, P. J. Liu, and C. D. Manning, "Get to the point: Summarization with pointer-generator networks," *arXiv preprint arXiv:1704.04368*, 2017.
- [12] H. Xu, T. Chen, D. Gao, Y. Wang, K. Li, N. Goel, Y. Carmiel, D. Povey, and S. Khudanpur, "A pruned rnnlm lattice-rescoring algorithm for automatic speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5929–5933.
- [13] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," *arXiv preprint arXiv:1508.04025*, 2015.
- [14] E. Grave, A. Joulin, and N. Usunier, "Improving neural language models with a continuous cache," *arXiv preprint arXiv:1612.04426*, 2016.
- [15] S. Merity, C. Xiong, J. Bradbury, and R. Socher, "Pointer sentinel mixture models," *arXiv preprint arXiv:1609.07843*, 2016.
- [16] E. Grave, A. Joulin, M. Cissé, H. Jégou *et al.*, "Efficient softmax approximation for gpus," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 2017, pp. 1302–1310.
- [17] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [18] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The kaldi speech recognition toolkit," in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society, Dec. 2011, iEEE Catalog No.: CFP11SRW-USB.
- [19] I. Loshchilov and F. Hutter, "Sgdr: Stochastic gradient descent with warm restarts," *arXiv preprint arXiv:1608.03983*, 2016.
- [20] D. Le, X. Zhang, W. Zheng, C. Fügen, G. Zweig, and M. L. Seltzer, "From senones to chenones: Tied context-dependent graphemes for hybrid speech recognition," *Proc. ASRU*, 2019.
- [21] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi." in *Proc. Interspeech*, 2016.
- [22] X. Liu, X. Chen, Y. Wang, M. J. Gales, and P. C. Woodland, "Two efficient lattice rescoring methods using recurrent neural network language models," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 8, pp. 1438–1449, 2016.
- [23] K. Irie, A. Zeyer, R. Schlüter, and H. Ney, "Language modeling with deep transformers," *Proc. Interspeech*, 2019.