

ON USING 2D SEQUENCE-TO-SEQUENCE MODELS FOR SPEECH RECOGNITION

Parnia Bahar^{1,2}, Albert Zeyer^{1,2}, Ralf Schlüter¹, Hermann Ney^{1,2}

Human Language Technology and Pattern Recognition Group

¹Computer Science Department, RWTH Aachen University, D-52056 Aachen, Germany

²AppTek, McLean, USA, <http://www.apptek.com/>

{bahar, zeyer, schluter, ney}@cs.rwth-aachen.de

ABSTRACT

Attention-based sequence-to-sequence models have shown promising results in automatic speech recognition. Using these architectures, one-dimensional input and output sequences are related by an attention approach, thereby replacing more explicit alignment processes, like in classical HMM-based modeling. In contrast, here we apply a novel two-dimensional long short-term memory (2DLSTM) architecture to directly model the input/output relation between audio/feature vector sequences and word sequences. The proposed model is an alternative model such that instead of using any type of attention components, we apply a 2DLSTM layer to assimilate the context from both input observations and output transcriptions. The experimental evaluation on the Switchboard 300h automatic speech recognition task shows word error rates for the 2DLSTM model that are competitive to end-to-end attention-based model.

Index Terms— 2D sequence-to-sequence model, end-to-end, speech recognition, multi-dimensional LSTM

1. INTRODUCTION

Conventional automatic speech recognition (ASR) systems using Gaussian mixture model (GMM) and/or hybrid deep neural network (DNN) hidden Markov models (HMM) consist of several components that are trained separately, depend on pretrained alignments and require a complex search [1, 2, 3, 4]. Unlike the conventional approaches, attention-based sequence-to-sequence models propose a standalone and single neural network that trains end-to-end, does not need explicit alignments or context-dependent phonetic labels as in HMM and simplify the inference. In these models, an implicit probabilistic notion of alignment is used as part of a neural network. However, it does not work the same way as its analogy of alignment models in the conventional methods.

The widely used attention-based sequence-to-sequence systems are based on an encoder-decoder architecture, where one or more long short-term memory (LSTM) layers read the observation sequence and another LSTM decodes it to a variable length output sequence of characters or words. In such architectures, both input and output sequences are

separately handled as a one-dimensional sequence over time. An attention mechanism is then added into the architecture to combine the encoder and the decoder by allowing the decoder to selectively focus on individual parts of the encoder state sequences [5, 6, 7, 8, 9].

The LSTM [10] is well suited for sequence modeling, where the sequence is strongly correlated along a one-dimensional time axis. Handling dynamic length, encoding positional information, the ability to make use of the previous context and tracking long-term dependencies by the gating strategy are some of the properties which make LSTM appropriate for the sequence to sequence modeling. Although an LSTM processes essentially one-dimensionally, it can be extended for the processing of multi-dimensional data such as an image or a video [11].

In this work, we investigate the use of two-dimensional LSTM (2DLSTM) [11, 12] in sequence-to-sequence modeling as an alternative model for the attention component. In this architecture, we apply a 2DLSTM on top of a deep bidirectional encoder to relate input and output representations in a 2D space. One dimension of the 2DLSTM processes the input sequence, and another dimension predicts the output (sub)words. In contrast to the attention-based sequence-to-sequence model, where the encoder states are not updated and the model is not able to re-interpret the encoder states while decoding, this model enables the computation of the encoding of the observation sequence as a function of the previously generated transcribed words. Our model is similar to an architecture used in machine translation described in [13]. We believe that the 2DLSTM is able to capture necessary monotonic alignments as well as retrieve coverage concepts internally by its cell states. Experimental results on the 300h-Switchboard task show competitive performance compared to an attention-based sequence-to-sequence system.

2. RELATED WORKS

A way of building multidimensional context into recurrent networks is provided by a strategy that is based on networks with tree-structured update graphs. In handwriting recognition (HWR), 2DLSTM has shown successful results in auto-

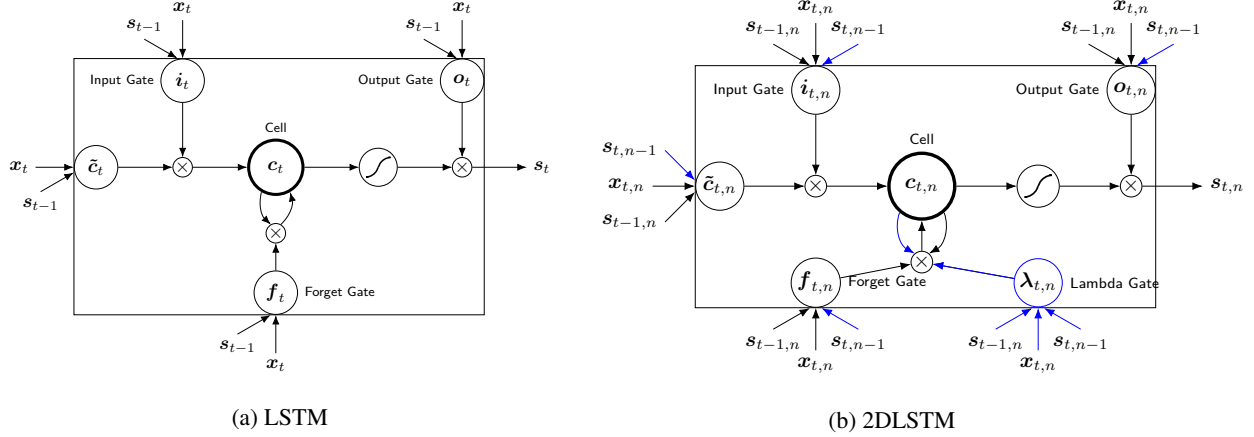


Fig. 1: The internal architecture of the standard and the 2DLSTM. The additional connections are marked in blue [13].

matic extraction of features from raw 2D-images over convolutional neural networks (CNNs) [14]. In order to investigate deeper and larger models using 2DLSTM, an algorithm to use the GPU power has been implemented [15].

Different neural networks have been proposed in automatic speech recognition (ASR) to model 2D correlations in the input signal. One of them is a 2DLSTM layer which scans the input over both time and frequency jointly for spatio-temporal modeling and aggregates more variations [16]. Moreover, various architectures to model time-frequency patterns based on deep DNN, CNN, RNN and 2DLSTM layers are compared for large vocabulary ASR [17].

As an alternative method to the concept of the 2DLSTM, a network of one-dimensional LSTM cells arranged in a multidimensional grid has been introduced [18]. In this topology, the LSTM cells communicate not only along time sequence but also between the layers. The grid LSTM network is also applied for the endpoint detection task in ASR to model both spectral and temporal variations [19]. A 2D attention matrix is also applied in a neural pitch accent recognition model [20], in which graphemes are encoded in one dimension and audio frames are encoded in the other.

Recently, the 2DLSTM layer also has been used for sequence-to-sequence modeling in machine translation [13] where it implicitly updates the source representation conditioned on the generated target words. In a similar direction, a 2D CNN-based network has been proposed where the positions of the source and the target words define the 2D grid for translation modeling [21].

Similar to [13], we apply a 2DLSTM layer to combine the acoustic model (the LSTM encoder) and the language model (the decoder) without any attention components. The 2DLSTM reconciles the context from both the input and the output sequences and re-interprets the encoder states while a new word has been predicted. Compared to [13], our model is much deeper. We use max-pooling to select the most relevant encoder state whereas [13] uses the last horizontal state of the 2DLSTM. Furthermore, we utilize the same pretraining

scheme explained in [9] during training and a faster decoding.

3. 2D LONG SHORT-TERM MEMORY

The 2DLSTM is characterized as a general form of the standard LSTM [11, 22]. It has been proposed to process inherent 2D data of arbitrary lengths, T and N . Therefore, it uses both horizontal and vertical recurrences. The building block of both the LSTM and the 2DLSTM are shown in Figure 1. At time step (t, n) , it gets an input $x_{t,n}$, and its computation relies on both the vertical $s_{t,n-1}$ and the horizontal hidden states $s_{t-1,n}$. Besides the input $i_{t,n}$, the forget $f_{t,n}$ and the output $o_{t,n}$ gates that are similar to those in the LSTM, the 2DLSTM employs an additional lambda gate. As written in Equation 5, its activation is computed analogously to the other gates [13, 11].

$$i_{t,n} = \sigma(W_1 x_{t,n} + U_1 s_{t-1,n} + V_1 s_{t,n-1}) \quad (1)$$

$$f_{t,n} = \sigma(W_2 x_{t,n} + U_2 s_{t-1,n} + V_2 s_{t,n-1}) \quad (2)$$

$$o_{t,n} = \sigma(W_3 x_{t,n} + U_3 s_{t-1,n} + V_3 s_{t,n-1}) \quad (3)$$

$$\tilde{c}_{t,n} = g(W_4 x_{t,n} + U_4 s_{t-1,n} + V_4 s_{t,n-1}) \quad (4)$$

$$\lambda_{t,n} = \sigma(W_5 x_{t,n} + U_5 s_{t-1,n} + V_5 s_{t,n-1}) \quad (5)$$

$$c_{t,n} = f_{t,n} \circ [\lambda_{t,n} \circ c_{t-1,n} + (1 - \lambda_{t,n}) \circ c_{t,n-1}] + \tilde{c}_{t,n} \circ i_{t,n} \quad (6)$$

$$s_{t,n} = g(c_{t,n}) \circ o_{t,n} \quad (7)$$

The internal cell state $c_{t,n}$, is computed based on the sum of the two previous cell's states $c_{t-1,n}$ and $c_{t,n-1}$, weighted by the lambda gate $\lambda_{t,n}$ and its complement (see Equation 6). Similar to the LSTM, the internal cell $c_{t,n}$ is combined with the output gate to yield the hidden state. g and σ are the hyperbolic tangent and the sigmoid functions. V_i , W_i and U_i , are the weight matrices. For notational simplicity, we omit the bias vectors.

We process the 2D data in a forward pass from the time step $(1, 1)$ to (T, N) and thus the gradient is passed backwards in an opposite direction from the time step (T, N) to $(1, 1)$. Training a 2DLSTM unit involves back-propagation through two dimensions. For more details, We refer the readers to [11, 22].

4. 2D SEQUENCE-TO-SEQUENCE MODEL

Bayes decision rule requires maximization of the class posterior given an input observation. In ASR, classes are discrete label sequences of unknown length N (e.g. word, subword, character) sequences, denoted as $w_1^N = w_1, \dots, w_N$. Given an input observation $x_1^T = x_1, \dots, x_T$ of variable length T where usually $T > N$, the posterior probability of a label sequence w_1^N is defined as $p(w_1^N | x_1^T)$. This conditional distribution usually covers the alignment information between the input observation sequence and the output word sequence either implicitly or explicitly.

In the attention-based sequence-to-sequence approach, the attention weights serve as the implicit probabilistic notion of alignments aligning output labels to encoder states. The freedom of the attention model to focus on the entire input sequence might contradict monotonicity in ASR. In this work, we remove the attention component and intend to investigate whether the 2D sequence-to-sequence modeling is able to properly capture the input-output monotonic relation.

As shown in Figure 2, we apply a deep bidirectional LSTM encoder ($L = 6$) to scan an observation sequence. On top of each bidirectional LSTM layer, we conduct max-pooling over the time dimension to reduce the observation length. Hence, the encoder states are formulated as follows:

$$h_1^{T'} = \text{biLSTM}_L \circ \dots \circ \text{max-pool}_1 \circ \text{biLSTM}_1(x_1^T) \quad (8)$$

where T' is the reduced length by a reduction factor. Similar to [13], we then equip the network by a 2DLSTM layer to relate the encoder and the decoder states. At time step (t', n) , the 2DLSTM receives both the encoder state $h_{t'}$, and the last target embedding vector w_{n-1} , as inputs. One dimension of the 2DLSTM (horizontal-axis in the figure) sequentially reads the encoder states and another (vertical axis) plays the role of the decoder. Therefore, there is no additional decoder LSTM. Unlike the attention-based sequence-to-sequence model, where the encoder states are obtained once at the beginning, our model repeatedly updates the encoder representations $h_1^{T'}$, while generating a new output word w_n . We note that in this model, we do not use any attention component. The state of the 2DLSTM is derived as follows:

$$s_{t',n} = \text{2DLSTM}\left([h_{t'}; w_{n-1}], s_{t'-1,n}, s_{t',n-1}\right) \quad (9)$$

It is significant to note that the 2DLSTM state for a label/word step n only have a dependence on the preceding

word sequence w_1^{n-1} , while it takes into account the whole temporal context of the input observation sequence.

At each decoder step, once the whole input sequence is processed from 1 to T' , we do max-pooling over all horizontal states to obtain the context vector. We have also tried average-pooling or the last horizontal state instead of max-pooling, but none is better in this case. In order to generate a next output word, w_n , a transformation followed by a softmax operation is applied. Therefore:

$$p(w_n | w_1^{n-1}, x_1^T) = \text{softmax}\left(\tanh\left(\text{max-pool}(s_{1,n-1}^{T',n-1})\right)\right)_{|w_n} \quad (10)$$

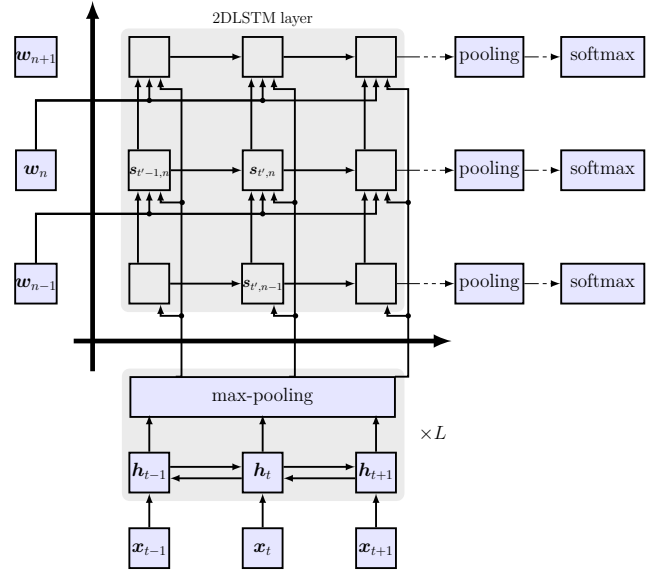


Fig. 2: The 2D seq2seq architecture using the 2DLSTM layer on top of L -layer of encoder. Neither attention components nor explicit LSTM decoders are used. Inspired by [13].

5. EXPERIMENTS

We have conducted experiments on the Switchboard 300h task. We apply 40-dimensional Gammatone features [23] using the RASR feature extractor [24]. We use the full Hub5'00 including Switchboard (SWB) and Callhome (CH) as the development set and the Hub5'01 as a test set. In order to enable an open-vocabulary system, we use byte-pair-encoding (BPE) [25] with 1k merge operations.

As our baseline, we utilize the attention-based sequence-to-sequence architecture similar to that described in [9] with the exact pretraining scheme and the same reduction factor. The baseline model includes a one-layer LSTM decoder with additive attention equipped with fertility feedback.

The feature vectors are passed into a stack of 6 bidirectional LSTM layers of size 1000 in each direction followed by the max-pooling operation. We downsample the input sequence by factor of 8 in total as described in [9]. The

Table 1: Total number of parameters, perplexity and FER^[%] on the development set.

model	# params	perplexity	FER
baseline [9]	157M	1.56	10.9
this work	160M	1.53	10.6

Table 2: WER^[%] on Switchboard 300h. [†]average of 3 runs.

model	LM	Hub5'00 (dev)		Hub5'01 (test)
		SWB	CH	
prior works				
hybrid	LSTM	8.3	17.3	12.9
CTC [30]	RNN	14.0	25.3	-
attention [31]	-	23.1	40.8	-
baseline				
attention [9] [†]	-	13.0	26.2	19.4
this work [†]	-	12.9	26.4	19.0

2DLSTM layer is equipped with 1000 nodes and the output subwords are projected into a 620-dimensional embedding space. The models are trained end to end using the Adam optimizer [26], dropout of 30% [27], label smoothing of 0.1 [28] and warmup technique. We reduce the learning rate by a factor of 0.7 following a variant of the Newbob scheme based on the perplexity on the development set for a few checkpoints.

In our training, we use layer-wise pretraining for the encoder, where we start with two encoder layers and a single max-pool in between with the same multiple-step reduction factor similar to [9]. Decoding is performed using beam search with a beam size of 12 and the subwords are merged into words. We do not utilize any language model (LM) neither in the baseline system nor in the 2D sequence-to-sequence model. The model is built using our in-house CUDA implementation of 2DLSTM [15] utilizing optimal speedups in RETURNN [29]. The code is open source and the configuration of the setups are available online¹.

Table 1 compares the total number of parameters, perplexity and frame error rate (FER) on the development set between our model and the attention baseline. Both models have the same vocabulary size of almost 1K. Our model has 3M more parameters. The perplexity and the FER are comparable. We also compare our model over prior works based on the WER listed in Table 2. As a simple significance test, the reported WERs are averaged over 3 runs. Although our 2D sequence-to-sequence model is still behind the hybrid methods, it leads to competitive results over the attention baseline. We observe that our model outperforms the baseline on the Hub5'01 subset by 0.4% absolute. Including a separate LM during the search, we expect to obtain improvements.

We also compare our model and the attention-based sequence-to-sequence model in terms of decoding speed.

¹<https://github.com/rwth-i6/returnn>

Table 3: Decoding speed measured in minutes on the entire development set.

model	decoding speed (mins)
baseline [9]	4
this work	26

Based on the fact that the whole output label sequence is known during the training, the entire 2DLSTM states can be computed once and at each time step, one row of it is taken. This computation cannot be done as a single operation in the search since the output sequence has to be predicted; therefore, during the decoding, we need to compute the states of the 2DLSTM row-wise which slows down the search procedure. This algorithm is faster than [13], where at each output step, they recompute all previous states of 2DLSTM from scratch which are not required. Table 3 lists the decoding speed of the models to decode the entire development set using a single GPU. In general, the decoding speed of our model is about 6 times slower than that of a standard attention-based model.

6. CONCLUSION

We have applied a simple 2D sequence-to-sequence model as an alternative to the attention-based model. In our model, a 2DLSTM layer has been utilized to jointly combine the input and the output representations. It processes the observation sequence via the horizontal dimension and generates the output (sub)word sequence through the vertical axis. It does not have any additional LSTM decoder and does not benefit from any attention components. Contrary to the attention-based sequence-to-sequence model, it repeatedly re-encodes the encoder representation when a new output (sub)word is generated. The experimental results are competitive with the baseline on the 300h-Switchboard Hub'00 and show 0.4% improvements on the Hub'01. Our future goal is to develop a bidirectional 2DLSTM to model completely independent of the standard LSTM layers as well as run more experiments on various speech tasks.

7. ACKNOWLEDGEMENTS



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.

8. REFERENCES

- [1] Hans-Peter Hutter, "Comparison of a new hybrid connectionist schmm approach with other hybrid approaches for speech recognition," in *IEEE Inter. Conf. ICASSP, Detroit, Michigan, USA, May 08-12, 1995*, pp. 3311–3314.
- [2] Anthony J. Robinson, "An application of recurrent nets to phone probability estimation," *IEEE Trans. Neural Networks*, vol. 5, no. 2, pp. 298–305, 1994.
- [3] Herve A Bourlard and Nelson Morgan, *Connectionist speech recognition: a hybrid approach*, vol. 247, Springer Science & Business Media, 2012.
- [4] Albert Zeyer, Patrick Doetsch, Paul Voigtlaender, Ralf Schlüter, and Hermann Ney, "A comprehensive study of deep bidirectional LSTM RNNs for acoustic modeling in speech recognition," in *IEEE Inter. Conf. ICASSP, New Orleans, LA, USA, Mar. 5-9, 2017*, 2017, pp. 2462–2466.
- [5] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, "Sequence to sequence learning with neural networks," in *Annual Conf. NIPS, Dec. 8-13, Montreal, Quebec, Canada, 2014*, pp. 3104–3112.
- [6] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, "Neural machine translation by jointly learning to align and translate," *CoRR*, vol. abs/1409.0473, 2015.
- [7] Jan Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, "Attention-based models for speech recognition," in *Annual Conf. NIPS, Dec. 7-12, 2015, Montreal, Quebec, Canada, 2015*, pp. 577–585.
- [8] Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio, "End-to-end attention-based large vocabulary speech recognition," in *IEEE Inter. Conf. ICASSP, Shanghai, China, Mar. 20-25, 2016*, 2016, pp. 4945–4949.
- [9] Albert Zeyer, Kazuki Irie, Ralf Schlüter, and Hermann Ney, "Improved training of end-to-end attention models for speech recognition," in *19th Annual Conf. Interspeech, Hyderabad, India, 2-6 Sep. 2018.*, 2018, pp. 7–11.
- [10] Sepp Hochreiter and Jürgen Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [11] Alex Graves, *Supervised sequence labelling with recurrent neural networks*, Ph.D. thesis, Technical University Munich, 2008.
- [12] Alex Graves, Santiago Fernández, and Jürgen Schmidhuber, "Multi-dimensional recurrent neural networks," *CoRR*, vol. abs/0705.2011, 2007.
- [13] Parnia Bahar, Christopher Brix, and Hermann Ney, "Towards two-dimensional sequence to sequence model in neural machine translation," *arXiv preprint arXiv:1810.03975*, 2018.
- [14] Gundram Leifert, Tobias Strauß, Tobias Grüning, and Roger Labahn, "Citlab ARGUS for historical handwritten documents," *CoRR*, vol. abs/1605.08412, 2016.
- [15] Paul Voigtlaender, Patrick Doetsch, and Hermann Ney, "Handwriting recognition with large multidimensional long short-term memory recurrent neural networks," in *15th Intern. Conf. ICFHR, Shenzhen, China, Oct. 23-26, 2016*, pp. 228–233.
- [16] Jinyu Li, Abdelrahman Mohamed, Geoffrey Zweig, and Yifan Gong, "Exploring multidimensional lstms for large vocabulary ASR," in *2016 IEEE Intern. Conf. ICASSP, Shanghai, China, Mar. 20-25, 2016*, 2016, pp. 4940–4944.
- [17] Tara N. Sainath and Bo Li, "Modeling time-frequency patterns with LSTM vs. convolutional architectures for LVCSR tasks," in *17th Annual Conf. of the International Speech, Interspeech, San Francisco, CA, USA, Sep. 8-12, 2016*, 2016, pp. 813–817.
- [18] Nal Kalchbrenner, Ivo Danihelka, and Alex Graves, "Grid long short-term memory," *CoRR*, vol. abs/1507.01526, 2015.
- [19] Bo Li, Carolina Parada, Gabor Simko, Shuo yiin Chang, and Tara Sainath, "Endpoint detection using grid long short-term memory networks for streaming speech recognition," in *In Proc. Interspeech 2017*, 2017.
- [20] Antoine Bruguier, Heiga Zen, and Arkady Arkhangorodsky, "Sequence-to-sequence neural network model with 2d attention for learning japanese pitch accents," in *19th Annual Conf. of Interspeech, Hyderabad, India, 2-6 Sep. 2018.*, 2018, pp. 1284–1287.
- [21] Maha Elbayad, Laurent Besacier, and Jakob Verbeek, "Pervasive attention: 2d convolutional neural networks for sequence-to-sequence prediction," *CoRR*, vol. abs/1808.03867, 2018.
- [22] Gundram Leifert, Tobias Strauß, Tobias Grüning, Welf Wustlich, and Roger Labahn, "Cells in multidimensional recurrent neural networks," *The Journal of Machine Learning Research*, vol. 17, pp. 3313–3349, 2016.
- [23] Ralf Schlüter, Ilja Bezrukov, Hermann Wagner, and Hermann Ney, "Gammatone features and feature combination for large vocabulary speech recognition," in *IEEE Inter. Conf. ICASSP, Honolulu, Hawaii, USA, Apr. 15-20, 2007*, pp. 649–652.
- [24] Simon Wiesler, Alexander Richard, Pavel Golik, Ralf Schlüter, and Hermann Ney, "RASR/NN: the RWTH neural network toolkit for speech recognition," in *IEEE Inter. Conf. ICASSP, Florence, Italy, May 4-9, 2014*, 2014, pp. 3281–3285.
- [25] Rico Sennrich, Barry Haddow, and Alexandra Birch, "Neural machine translation of rare words with subword units," in *Proc. of the 54th ACL, Aug. 7-12, Berlin, Germany, Volume 1*, 2016.
- [26] Diederik P. Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [27] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [28] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna, "Rethinking the inception architecture for computer vision," in *IEEE Conf. CVPR, Las Vegas, NV, USA, Jun. 27-30, 2016*, 2016, pp. 2818–2826.
- [29] Albert Zeyer, Tamer Alkhouli, and Hermann Ney, "RETURNN as a generic flexible neural toolkit with application to translation and speech recognition," in *In Proc. of ACL, Melbourne, Australia, Jul. 15-20, 2018, System Demonstrations*, 2018, pp. 128–133.
- [30] Geoffrey Zweig, Chengzhu Yu, Jasha Droppo, and Andreas Stolcke, "Advances in all-neural speech recognition," in *IEEE Inter. Conf. ICASSP, New Orleans, LA, USA, Mar. 5-9, 2017*, pp. 4805–4809.
- [31] Shubham Toshniwal, Hao Tang, Liang Lu, and Karen Livescu, "Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition," in *18th Annual Conf. of Interspeech, Stockholm, Sweden, Aug. 20-24, 2017*, 2017, pp. 3532–3536.