THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Investigating gated recurrent neural networks for speech synthesis

Citation for published version:

Wu, Z & King, S 2016, Investigating gated recurrent neural networks for speech synthesis. in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Institute of Electrical and Electronics Engineers (IEEE), pp. 1-5, 41st IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016, Shanghai, China, 20/03/16. <https://doi.org/10.1109/ICASSP.2016.7472657>

Digital Object Identifier (DOI):

[10.1109/ICASSP.2016.7472657](https://doi.org/10.1109/ICASSP.2016.7472657)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

INVESTIGATING GATED RECURRENT NETWORKS FOR SPEECH SYNTHESIS

Zhizheng Wu Simon King

The Centre for Speech Technology Research, University of Edinburgh, United Kingdom

zhizheng.wu@ed.ac.uk

ABSTRACT

Recently, recurrent neural networks (RNNs) as powerful sequence models have re-emerged as a potential acoustic model for statistical parametric speech synthesis (SPSS). The long short-term memory (LSTM) architecture is particularly attractive because it addresses the vanishing gradient problem in standard RNNs, making them easier to train. Although recent studies have demonstrated that LSTMs can achieve significantly better performance on SPSS than deep feed-forward neural networks, little is known about why. Here we attempt to answer two questions: a) why do LSTMs work well as a sequence model for SPSS; b) which component (e.g., input gate, output gate, forget gate) is most important. We present a visual analysis alongside a series of experiments, resulting in a proposal for a simplified architecture. The simplified architecture has significantly fewer parameters than an LSTM, thus reducing generation complexity considerably without degrading quality.

Index Terms— Speech synthesis, acoustic modelling, recurrent network network, gated recurrent network, long short-term memory

1. INTRODUCTION

Statistical parametric speech synthesis (SPSS) has quite steadily advanced in naturalness in the past decade, as witnessed by the series of Blizzard Challenges [1]. However, the quality of synthetic speech produced by SPSS is still far below that of the natural human speech, and cannot compete with the best unit selection systems, which concatenate waveforms [2]. As suggested in [3], acoustic modelling, which captures the complex relationship between linguistic and acoustic representations, is a key limiting factor and is the focus of this work.

1.1. Relation to prior work

Neural networks have re-emerged as a potential powerful acoustic model for SPSS. In [4, 5, 6, 7, 8], feed-forward neural networks are employed to map a linguistic representation derived from input text directly to acoustic features. In [9], a deep belief network (DBN) was used to model the relationship between linguistic and acoustic representations jointly.

In [10] and [11], mixture density networks (MDNs) and real-valued neural autoregressive density estimators (RNADEs) were proposed, respectively, to predict acoustic feature distributions given input linguistic features. These various implementations can be viewed as a replacement of the decision tree in HMM-based speech synthesis; they map linguistic features to acoustic features frame by frame through multiple hidden layers. However, the temporal sequence nature of speech is not explicitly modelled in the network architectures.

To include temporal constraints, we proposed to include contextual information by stacking low-dimensional bottleneck features from multiple consecutive frames [12]. Still in the DNN framework, minimum trajectory error training [13] or sequence error training criterion [14] have been proposed to minimise the utterance-level trajectory error rather than the frame-by-frame error. On the other hand, recurrent neural networks (RNNs) directly and elegantly include temporal information in the network architecture, making them attractive for modelling speech parameter trajectories. In [15], a standard RNN was employed to predict prosodic information for speech synthesis. In [16], two variants on standard RNNs, the Elman RNN and clockwork RNN, were investigated for speech synthesis.

The most widely used recurrent network in speech processing applications is the long short-term memory (LSTM) architecture. Because the LSTM addresses the vanishing gradient problem of the standard RNN, it is easier to train. In [17], an LSTM was employed to model the F0 contour. In [18], a bidirectional LSTM was employed to map a sequence of linguistic features to the corresponding sequence of acoustic features. In [19], an LSTM with a recurrent output layer was proposed to perform sequence mapping from linguistic to acoustic representations. These studies all formulate SPSS as sequence-to-sequence mapping and all demonstrate the effectiveness of LSTMs. However, LSTM architecture seems rather ad-hoc and it is not obvious what its various components are actually contributing to performance.

This raises at least two questions that have not been answered in previous studies: a) *how exactly does the LSTM architecture model a speech parameter sequence*; b) *which components of the LSTM architecture are important, and which could be discarded*. Answers to these questions may suggest better and perhaps simpler recurrent network architectures.

1.2. The novelty of this work

We attempt to reach a better understanding of the “black-box” LSTM architecture and our findings lead us to propose a simplified architecture for speech modelling.

First, we give an analysis of the forget gate and memory cell in the LSTM architecture. Specifically, we visualise the activation of the forget gate to understand when the forget gate resets the memory cell state, and how the forget gate relates to speech structure. We analyse how the cell state correlates with the trajectory to be predicted. These visualisations enable us to understand how LSTMs model the temporal structure in speech synthesis. *To the best of our knowledge, this is the first attempt to visually analyse the LSTM architecture in predicting a speech parameter sequence.*

Second, we analyse the importance of each LSTM component for speech synthesis and propose a simplified architecture. The analysis is done empirically with several variants of the LSTM. Each removes a different component of the vanilla LSTM. The analysis was inspired by the studies in [20, 21], and we focus on the speech synthesis application. Based on this analysis, we present a simplified architecture, which only has the forget gate. The simplified architecture has significantly fewer parameters than the vanilla LSTM, and so reduces the computational cost of generation considerably without degrading the quality of the synthesised speech.

2. LONG SHORT-TERM MEMORY

Standard RNNs are hard to train due to the well-known vanishing or exploding gradient problems [22, 23]. To address the vanishing gradient problem, the LSTM architecture was proposed, the basic idea of which was presented in [24]. The most commonly used architecture was described in [25], and is formulated as,

$$\begin{aligned} \mathbf{i}_t &= \delta(\mathbf{W}^i \mathbf{x}_t + \mathbf{R}^i \mathbf{h}_{t-1} + \mathbf{p}^i \odot \mathbf{c}_{t-1} + \mathbf{b}^i) \\ \mathbf{f}_t &= \delta(\mathbf{W}^f \mathbf{x}_t + \mathbf{R}^f \mathbf{h}_{t-1} + \mathbf{p}^f \odot \mathbf{c}_{t-1} + \mathbf{b}^f) \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot g(\mathbf{W}^c \mathbf{x}_t + \mathbf{R}^c \mathbf{h}_{t-1} + \mathbf{b}^c) \\ \mathbf{o}_t &= \delta(\mathbf{W}^o \mathbf{x}_t + \mathbf{R}^o \mathbf{h}_{t-1} + \mathbf{p}^o \odot \mathbf{c}_t + \mathbf{b}^o) \\ \mathbf{h}_t &= \mathbf{o}_t \odot g(\mathbf{c}_t) \end{aligned}$$

In these formulations, $\mathbf{i}_t$, $\mathbf{f}_t$, $\mathbf{c}_t$, $\mathbf{o}_t$, and $\mathbf{h}_t$ are the input gate, forget gate, cell state, output gate and block output at time instance t , respectively; $\delta(\cdot)$ and $g(\cdot)$ are the sigmoid and tangent activation functions, respectively; $\mathbf{x}_t$ is the input at time t ; $\mathbf{W}^*$, and $\mathbf{R}^*$ are the weight matrices applied on input and recurrent hidden units, respectively; $\mathbf{p}^*$ and $\mathbf{b}^*$ are the peep-hole connections and biases, respectively; and $\odot$ means element-wise product. We will call this the *vanilla LSTM*.

The central idea of the LSTM is the so-called memory cell $\mathbf{c}$ which maintains its state over time, and the gating units which are used to regulate the information flow into and out of the memory cell [20]. More specifically, the input gate can allow the input signal to adjust the cell state or prevent that

(e.g., setting the input gate to zero); the output gate can allow the cell state to affect other neurons or block that; and the forget gate enables the cell to remember or forget its previous state. However, as discussed in [20, 21], the architecture might not be optimal for all the tasks, and the relative importance of each component is not at all clear.

3. GATED RECURRENT NEURAL NETWORKS

In this section, we present several variants of the LSTM and propose a simplified version that only has the forget gate; it therefore has significantly fewer parameters and lower computational cost. As these variants all share with the LSTM the concept of a memory cell with gates, we will call them *gated recurrent neural networks*.

3.1. Four variants on the LSTM

To assess the importance of each component, we start with four variants of the LSTM architecture. Each removes one component from the LSTM architecture, so we can understand how much each component contributes to performance. The differences with the vanilla LSTM are:

- No Peep-holes (NPH): Set $\mathbf{p}^i$, $\mathbf{p}^f$, $\mathbf{p}^o$ to zero
- No input gate (NIG): $\mathbf{i}_t = 1$
- No forget gate (NFG): $\mathbf{f}_t = 1$
- No output gate (NOG): $\mathbf{o}_t = 1$

In the NFG variant, the past cell state will still contribute to the current cell state but without any controlling or scaling by the forget gate. Note that, when removing the input, forget or output gates, the number of parameters is reduced.

3.2. Gated Recurrent Unit (GRU)

As an alternative to the LSTM, the Gated Recurrent Unit (GRU) architecture was proposed in [26]. In [27], the GRU was found to achieve better performance than the LSTM on some tasks. The GRU is formulated as:

$$\begin{aligned} \mathbf{r}_t &= \delta(\mathbf{W}^r \mathbf{x}_t + \mathbf{R}^r \mathbf{h}_{t-1} + \mathbf{b}^r) \\ \mathbf{z}_t &= \delta(\mathbf{W}^z \mathbf{x}_t + \mathbf{R}^z \mathbf{h}_{t-1} + \mathbf{b}^z) \\ \tilde{\mathbf{h}}_t &= g(\mathbf{W}^h \mathbf{x}_t + \mathbf{r}_t \odot (\mathbf{R}^h \mathbf{h}_{t-1}) + \mathbf{b}^h) \\ \mathbf{h}_t &= \mathbf{z}_t \odot \mathbf{h}_{t-1} + (1 - \mathbf{z}_t) \odot \tilde{\mathbf{h}}_t \end{aligned}$$

From these formulae, we can observed that the GRU architecture is similar to LSTM but without a separate memory cell. The GRU does not use peep-hole connections and output activation functions, and combines the input and forget gates into an update gate $\mathbf{z}_t$ to balance between previous activation $\mathbf{h}_{t-1}$ and the candidate activation $\tilde{\mathbf{h}}_t$. The reset gate $\mathbf{r}_t$ allows it to forget the previous state.

3.3. Simplified LSTM (S-LSTM)

As we will see in the experiments reported in the next section, the input gate, output gate and peep-hole connections

Fig. 1. Averaged activations of all 256 forget gates as a function of time (in frames) are shown by the solid red line, with phoneme boundaries marked as dashed vertical blue lines.

can be removed without degrading speech synthesis performance significantly. Hence, we can propose an even simpler variant, that removes output gates and peep-hole connections, and replaces the input gate by the forget gate in the form of $1 - f_t$. In this way, only the forget gate is retained. This simplest variant can be written as:

$$\begin{aligned} f_t &= \delta(\mathbf{W}^f \mathbf{x}_t + \mathbf{R}^f \mathbf{h}_{t-1} + \mathbf{b}^f) \\ \mathbf{c}_t &= f_t \odot \mathbf{c}_{t-1} + (1 - f_t) \odot g(\mathbf{W}^c \mathbf{x}_t + \mathbf{R}^c \mathbf{h}_{t-1} + \mathbf{b}^c) \\ \mathbf{h}_t &= g(\mathbf{c}_t) \end{aligned}$$

The simplified architecture is similar to the GRU, except that it uses a memory cell state. The cell state is controlled by the forget gate only, which trades off between past cell state and current block input. When the activation of forget gate is small, the cell state will mainly depend on the block input, otherwise it will mainly copy the past cell state.

4. EXPERIMENTS

4.1. Experimental setup

A corpus from a British male speaker was employed in our experiments, divided into three subsets: training, development and testing (2400, 70 and 72 utterances). The sampling rate was 48 kHz, and we used the STRAIGHT vocoder [28] to extract 60-dimensional Mel-Cepstral Coefficients (MCCs), 25 band aperiodicities (BAPs), and fundamental frequency (F_0) on log-scale, all at 5-ms frame step. All systems used the same acoustic features. F_0 was linearly interpolated before modelling and a binary voiced/unvoiced feature was used to record voicing information. Dynamic features for MCCs, BAPs and F_0 were also computed. The acoustic features were mean-variance normalised before modelling, and the mean and variance was restored at the generation time. At generation time, maximum likelihood parameter generation algorithm [29] was applied to smooth parameter trajectories.

All systems used the same input linguistic features comprising 601 features. 592 of these are binary features derived from linguistic context, such as quin-phone identities, part-of-speech, positional information of phoneme, syllable, word and phrase, and the number of syllables, words and phrases,

Fig. 2. Comparison between the first MCC trajectory and the 120th cell state (vertically offset for clarity).

etc. The remaining 9 numerical features capture frame position information, e.g., frame position in HMM state and phoneme. Linguistic features were normalised to [0.01 0.99] before modelling.

In all RNNs, we employed a three-layer feed-forward neural network at the bottom. On top of the feed-forward layers, we used the gated recurrent neural networks. The bottom feed-forward layers were intended to act as feature extraction layers, with 512 hidden units using tangent activation function in each layer. All RNN implementations used 256 units (e.g., LSTM blocks) in the recurrent layer. Hyperparameters for each system were optimised on the development set. We fixed the momentum, and only tuned learning rates.

4.2. Analysis of LSTM

We first visualised the forget gate and cell state, which are thought to be the two most important components in modelling long-term temporal structure. The averaged activations (over the 256 units) of the forget gate as a function of the frame index is presented in Fig. 1. The red solid line is the forget gates averaged activations; blue dashed lines show phoneme boundaries. It is clear that the peaks of the forget gate activation trajectory have a strong correspondence with the phoneme boundaries; within a phoneme, the contribution of past cell state decays linearly. The forget gate is capturing some important temporal structure of speech; this is not surprising, since the phoneme boundaries are explicitly represented in the input linguistic features.

The memory cell should maintains its state over time [20] and so could store the trend of the trajectory to be predicted. To analyse the relationship between the cell states and the MCC trajectories, we computed the correlation between the cell states and the first MCC trajectory, and found that the 120th cell state has the highest correlation with the first MCC trajectory. The correlation is as high as 0.9. A comparison between them is presented in Fig. 2, which shows that the cell state tracks the shape of the MCC trajectory.

4.3. Objective results

Even though objective measures might not always correlate with human perception, they offer a way to tune the systems

Table 1. Subjective preference scores (%). $p < 0.01$ indicates significant difference between the two systems.

	LSTM	NIG	NOG	NFG	NPH	GRU	S-LSTM	Neutral	p-value
1	46.7	42.3	-	-	-	-	-	11.0	0.1582
2	40.7	-	47.8	-	-	-	-	11.5	0.0649
3	74.3	-	-	20.3	-	-	-	5.4	$< 10^{-8}$
4	44.0	-	-	-	45.3	-	-	10.7	0.7820
5	47.8	-	-	-	-	42.7	-	9.5	0.1086
6	46.5	-	-	-	-	-	45.3	8.2	0.7948
7	-	46.0	-	-	-	-	46.8	7.2	0.8305
8	-	-	48.0	-	-	-	44.3	7.7	0.3828
9	-	-	-	-	51.7	-	40.5	7.8	0.0018
10	-	-	-	-	-	48.0	45.0	7.0	0.4703

Table 2. Objective measures. MCD: Mel-Cepstral Distortion. BAP: distortion of band aperiodicities. F0 RMSE is calculated on a linear scale. V/UV: voiced/unvoiced error. Note that the number of parameters listed is for the recurrent layer only and does not include the bottom three feed-forward layers, which are the same size across all systems. The generation time is to generate all the 142 utterances in both development and testing sets.

	MCD (dB)	BAP (dB)	F0 RMSE (Hz)	V/UV (%)	# parameters	generation time (s)
LSTM	4.14	1.95	8.96	4.15	788, 224	214
NIG	4.18	1.95	9.10	4.15	591, 104	180
NOG	4.14	1.94	8.84	4.29	591, 104	167
NFG	4.68	1.99	9.69	4.41	591, 104	174
NPH	4.14	1.95	9.02	4.20	787, 456	180
GRU	4.17	1.95	9.00	4.22	590, 592	159
S-LSTM	4.19	1.95	8.87	4.14	393, 728	154

and roughly predict model performance. The objective results are in Table 2. Compared to LSTM, NIG, NOG and NPH all achieve similar objective distortion, with considerably fewer parameters and lower generation time: the input gate, output gate and peep-hole connections are not necessary. The NFG system increases distortion considerably: the forget gate is important. This finding is consistent with [20].

The GRU system achieves similar performance to the LSTM system: even though it has even fewer parameters, it performs as well as NIG, NOG or NPH. This is also consistent with studies on other tasks [27, 21]. Although S-LSTM slightly increases MCD distortion from 4.14 dB to 4.19 dB compared to LSTM, it achieves similar performance on the other measures. The S-LSTM has about half the number of parameters in its recurrent layer compared to the LSTM, and reduces generation time from 214 seconds to 154 seconds. The generation time is the total time to generate all the 142 utterances in both development and testing sets.

In summary, the S-LSTM has the smallest number of parameters and achieves the fastest generation, whilst achieving similar objective results to the LSTM and GRU architectures.

4.4. Subjective results

Subjective preference tests were conducted using 30 paid native English speakers. Each listener was asked to listen 20 pairs of synthesised utterances. The sentence was the same

in both items within a pair, and was randomly selected from the 72 test sentences¹. For each pair, the listener was asked to decide which one sounded more natural; a “neutral” option was allowed if the listener had no preference.

Preference results are in Table 1. Comparing against the LSTM system, all the systems except NFG show no significant difference in preference.

The NFG system achieves only a 20.3% preference score when paired against the LSTM which is preferred 74.3% of the time. As with the objective results in Table 2, we conclude that the forget gate is the only critical component in the LSTM architecture; the input gate, output gate and peep-hole connections can be omitted.

We also compares the proposed S-LSTM system against with all other systems (except NFG, since it is worse than LSTM). Consistent with the objective results, the subjective results also demonstrate that S-LSTM is as good as any other systems.

5. CONCLUSIONS

We have analysed the forget gate and cell state of the LSTM architecture, and examined the performance of several variants of LSTM. We conclude that:

- The forget gate can learn the temporal structure of speech; its activations have a high correspondence with phone boundaries.
- The memory cell maintains a state over time, which matched the shape of the trajectory to be predicted.
- For this task, the forget gate is the only critical component of the LSTM; other components can be omitted with no reduction in naturalness.

From these results, we propose a simplified LSTM architecture that only uses the critical forget gate. The simplified LSTM has significantly fewer parameters than the vanilla LSTM, but achieves similar performance in both objective and subjective evaluations.

Acknowledgements: This research was supported by EPSRC Programme Grant EP/I031022/1, Natural Speech Technology (NST). The NST research data collection may be accessed at <http://dashare.is.ed.ac.uk/handle/10283/786>.

¹Samples are available at: <http://homepages.inf.ed.ac.uk/zwu2/demo/icassp16/lstm.html>

6. REFERENCES

- [1] Simon King, “Measuring a decade of progress in text-to-speech,” *Loquens*, vol. 1, no. 1, 2014.
- [2] Andrew J Hunt and Alan W Black, “Unit selection in a concatenative speech synthesis system using a large speech database,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 1996.
- [3] Heiga Zen, Keiichi Tokuda, and Alan W Black, “Statistical parametric speech synthesis,” *Speech Communication*, vol. 51, no. 11, pp. 1039–1064, 2009.
- [4] Heiga Zen, Andrew Senior, and Mike Schuster, “Statistical parametric speech synthesis using deep neural networks,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2013.
- [5] Heng Lu, Simon King, and Oliver Watts, “Combining a vector space representation of linguistic context with a deep neural network for text-to-speech synthesis,” *Proc. the 8th ISCA Speech Synthesis Workshop (SSW)*, 2013.
- [6] Yao Qian, Yuchen Fan, Wenping Hu, and Frank K Soong, “On the training aspects of deep neural network (DNN) for parametric TTS synthesis,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2014.
- [7] Qiong Hu, Zhizheng Wu, Korin Richmond, Junichi Yamagishi, Yannis Stylianou, and Rannieri Maia, “Fusion of multiple parameterisations for DNN-based sinusoidal speech synthesis with multi-task learning,” in *Proc. Interspeech*, 2015.
- [8] Cassia Valentini-Botinhao, Zhizheng Wu, and Simon King, “Towards minimum perceptual error training for DNN-based speech synthesis,” in *Proc. Interspeech*, 2015.
- [9] Shiyin Kang, Xiaojun Qian, and Helen Meng, “Multi-distribution deep belief network for speech synthesis,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2013.
- [10] Heiga Zen and Andrew Senior, “Deep mixture density networks for acoustic modeling in statistical parametric speech synthesis,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2014.
- [11] Benigno Urias, Iain Murray, Steve Renals, and Cassia Valentini, “Modelling acoustic feature dependencies with artificial neural networks: Trajectory-rnade,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2015.
- [12] Zhizheng Wu, Cassia Valentini-Botinhao, Oliver Watts, and Simon King, “Deep neural networks employing multi-task learning and stacked bottleneck features for speech synthesis,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2015.
- [13] Zhizheng Wu and Simon King, “Minimum trajectory error training for deep neural networks, combined with stacked bottleneck features,” in *Proc. Interspeech*, 2015.
- [14] Yuchen Fan, Yao Qian, Frank K. Soong, and Lei He, “Sequence generation error (SGE) minimization based deep neural networks training for text-to-speech synthesis,” in *Proc. Interspeech*, 2015.
- [15] Sin-Horng Chen, Shaw-Hwa Hwang, and Yih-Ru Wang, “An RNN-based prosodic information synthesizer for Mandarin text-to-speech,” *IEEE Transactions on Speech and Audio Processing*, vol. 6, no. 3, pp. 226–239, 1998.
- [16] Sivanand Achanta, Tejas Godambe, and Suryakanth V. Gan-gashetty, “An investigation of recurrent neural network architectures for statistical parametric speech synthesis,” in *Proc. Interspeech*, 2015.
- [17] Raul Fernandez, Asaf Rendel, Bhuvana Ramabhadran, and Ron Hoory, “Prosody contour prediction with long short-term memory, bi-directional, deep recurrent neural networks,” in *Proc. Interspeech*, 2014.
- [18] Yuchen Fan, Yao Qian, Fenglong Xie, and Frank K. Soong, “TTS synthesis with bidirectional LSTM based recurrent neural networks,” in *Proc. Interspeech*, 2014.
- [19] Heiga Zen and Hasim Sak, “Unidirectional long short-term memory recurrent neural network with recurrent output layer for low-latency speech synthesis,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2015.
- [20] Klaus Greff, Rupesh Kumar Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber, “Lstm: A search space odyssey,” *arXiv preprint arXiv:1503.04069*, 2015.
- [21] Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever, “An empirical exploration of recurrent network architectures,” in *Proc. IEEE Int. Conf. on Machine Learning (ICML)*, 2015.
- [22] Yoshua Bengio, Patrice Simard, and Paolo Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.
- [23] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio, “On the difficulty of training recurrent neural networks,” *arXiv preprint arXiv:1211.5063*, 2012.
- [24] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] Alex Graves and Jürgen Schmidhuber, “Framewise phoneme classification with bidirectional lstm and other neural network architectures,” *Neural Networks*, vol. 18, no. 5, pp. 602–610, 2005.
- [26] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” *arXiv preprint arXiv:1406.1078*, 2014.
- [27] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv preprint arXiv:1412.3555*, 2014.
- [28] Hideki Kawahara, Ikuyo Masuda-Katsuse, and Alain de Cheveigné, “Restructuring speech representations using a pitch-adaptive time–frequency smoothing and an instantaneous-frequency-based F0 extraction: Possible role of a repetitive structure in sounds,” *Speech communication*, vol. 27, no. 3, pp. 187–207, 1999.
- [29] Keiichi Tokuda, Takayoshi Yoshimura, Takashi Masuko, Takao Kobayashi, and Tadashi Kitamura, “Speech parameter generation algorithms for HMM-based speech synthesis,” in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2000.