



Dynamic Gated Recurrent Neural Network for Compute-efficient Speech Enhancement

Longbiao Cheng¹, Ashutosh Pandey², Buye Xu², Tobi Delbruck¹, Shih-Chii Liu¹

¹ Institute of Neuroinformatics, University of Zurich and ETH Zurich

² Reality Labs Research, Meta

{longbiao, tobi, shih}@ini.uzh.ch, {apandey620, xub}@meta.com

Abstract

This paper introduces a new Dynamic Gated Recurrent Neural Network (DG-RNN) for compute-efficient speech enhancement models running on resource-constrained hardware platforms. It leverages the slow evolution characteristic of RNN hidden states over steps, and updates only a selected set of neurons at each step by adding a newly proposed select gate to the RNN model. This select gate allows the computation cost of the conventional RNN to be reduced during network inference. As a realization of the DG-RNN, we further propose the Dynamic Gated Recurrent Unit (D-GRU) which does not require additional parameters. Test results obtained from several state-of-the-art compute-efficient RNN-based speech enhancement architectures using the DNS challenge dataset, show that the D-GRU based model variants maintain similar speech intelligibility and quality metrics comparable to the baseline GRU based models even with an average 50% reduction in GRU computes.

Index Terms: dynamic gated networks, compute-efficient networks, recurrent neural network, speech enhancement

1. Introduction

The presence of ambient noise in real-world environments can diminish significantly the speech quality and intelligibility of a speaker. This degradation adversely affects the listener's experience in various applications involving telecommunication systems, online meeting platforms, and Augmented Reality (AR) / Mixed Reality (MR) devices. Despite the promising advancements made in deep-learning-based speech enhancement algorithms [1], their deployment on resource-constrained platforms and subsequent run-time latency is impacted by incompatible memory and compute requirements needed of these platforms.

U-Net based architectures are one class of architectures used extensively for speech enhancement [2, 3, 4, 5, 6, 7]. At its core, the model typically features a Convolutional Neural Network (CNN)-based encoder and decoder along with a temporal information modeling layer. The U-Net architecture delivers great results in reducing noise while preserving the quality and intelligibility of the original speech, primarily because of the skip connections between the encoder and decoder layers. However these connections require that the intermediate feature maps be stored therefore requiring memory that is not easily available on mobile platforms.

Recent studies show that Recurrent Neural Network (RNN)-based architectures can match the speech enhancement performance of the U-Net architectures while using less hard-

ware resource. For instance, FullSubNet [8, 9] employs a two-layer RNN architecture to first model the full-band global information, which is then used in conjunction with another two layers of RNNs to estimate the masks at each time-frequency bin. Similarly, TaErLite [10] segments the enhancement process into two phases, first, coarse magnitude spectrum enhancement and, then, multi-stage complex spectrum refinement, utilizing RNNs as the core component in both phases. The introduction of dual-path [11, 12] and multi-path [13] strategies in RNN-based models significantly reduce parameter size by reusing smaller RNNs across various dimensions, such as different frames or frequency bins. Additional methods that incorporate band splitting [14, 15] and frame skipping [16] strategies further reduce the computation costs of these models.

This work proposes a new Dynamic Gated Recurrent Neural Network (DG-RNN) framework that can decrease further the computation cost of these RNN-based models during inference. The framework was tested on a network trained for a regression task, in this case, speech enhancement. The contributions of the work are as follows:

- A novel RNN architecture, called Dynamic Gated Recurrent Neural Network (**DG-RNN**), which updates only a selected subset of neurons at each step. DG-RNN incorporates a new select gate to the conventional RNN for determining the dynamic set of neurons to be updated at each step. Neurons that are not selected will skip their update process at that step, resulting in a computation reduction of the conventional RNN.
- The application of DG-RNN's selective update strategy to the conventional Gated Recurrent Unit (**GRU**), leading to the creation of Dynamic Gated Recurrent Unit (**D-GRU**). Notably, the D-GRU utilizes the update gate of conventional GRU to derive the select gate, without the introduction of additional parameters.
- Comprehensive experimental validation of D-GRU in a single-channel speech enhancement task using the Deep Noise Suppression (**DNS**) challenge dataset [17]. Our evaluation encompasses models that solely employ a two-layer GRU, as well as models that use Multi-Path Transformer (**MPT**) blocks [15] or Dual-Path RNN (**DPRNN**) blocks [11]. Results show minimal impact on speech quality and intelligibility with a 50% reduction in GRU computation costs even under -5 dB Signal-to-Noise Ratio (**SNR**) conditions. Audio samples can be found [here](#).

2. Dynamic Gated RNN (DG-RNN)

As shown in previous studies [18], natural inputs to a recurrent neural network tend to have a high degree of temporal auto-correlation, resulting in slowly changing hidden states. From these observations, we propose the DG-RNN framework, de-

This work was supported by a research contract from Meta Reality Labs Research.

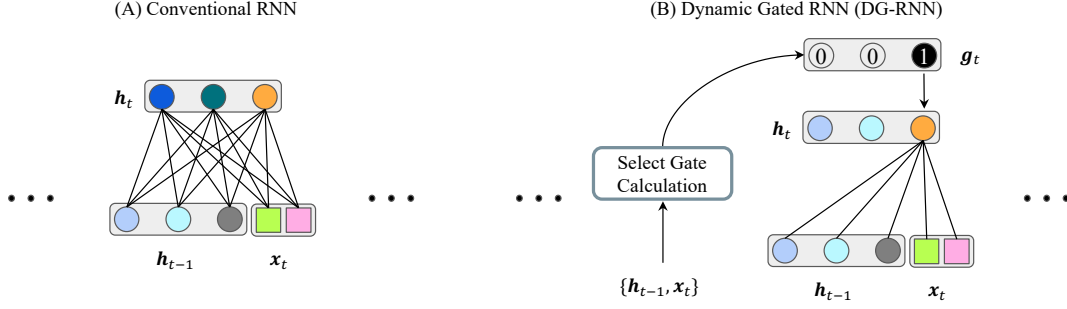


Figure 1: Illustration of the update processes of (A) conventional Recurrent Neural Network (RNN) and (B) Dynamic Gated Recurrent Neural Network (DG-RNN) at step t . (A) For conventional RNN, all neurons in the hidden state are updated at each step. (B) DG-RNN first identifies which neurons need updating, as indicated by a 1 in the proposed select gate g_t . When a neuron is marked with 1, it undergoes the RNN update process. Those marked with 0 retain their values from the previous hidden state.

signed to reduce the computation cost of conventional RNNs. In this framework, only a percentage of neurons is updated at each step, while the remaining neurons retain their previous state value, based on the assumption that their hidden state value underwent minimal change during the current step.

The selective update of the proposed DG-RNN is depicted in Fig. 1. Upon receiving an input x_t at step t , the DG-RNN first determines the set of neurons that will be updated at the current step using our new proposed select gate, g_t . The select gate is computed using x_t and the previous hidden state h_{t-1} :

$$g_t = \mathcal{B}(\mathcal{G}(x_t, h_{t-1})), \quad (1)$$

where $\mathcal{G}(\cdot)$ represents the computation process of a neural network, and $\mathcal{B}(\cdot)$ denotes a binarization operation, yielding g_t as a binary vector of length J , where J is the number of total neurons. Based on the output of the select gate, the hidden state in DG-RNN at step t is updated as follows:

$$h_t^j = \begin{cases} \mathcal{F}^j(x_t, h_{t-1}) & \text{if } g_t^j = 1, \\ h_{t-1}^j & \text{if } g_t^j = 0. \end{cases} \quad (2)$$

where h_t^j indicates the state of the j -th neuron at step t and $\mathcal{F}^j(\cdot)$ is its update function.

We additionally define the **neuron update percentage** \mathcal{P}_t at step t as $\mathcal{P}_t = (\sum_{j=1}^J g_t^j) / J * 100$. Note that \mathcal{P} is the most important parameter in determining cost savings; reducing \mathcal{P} as much as possible without compromising accuracy will lead to lower costs. Moreover, the select gate generation process should be lightweight to not negate the advantages of sparse update. In Section 3, we demonstrate that by applying DG-RNN's selective update strategy to the GRU, the select gate can be derived without additional parameters.

3. Dynamic GRU (D-GRU)

This section begins with a brief summary of the neuron update process in a conventional GRU, followed by the introduction of the D-GRU. Then cost savings of the D-GRU compared to the conventional GRU are presented.

3.1. GRU Update

The update equations for a conventional GRU are:

$$r_t = \sigma(\mathbf{W}_{ir}x_t + \mathbf{b}_{ir} + \mathbf{W}_{hr}h_{(t-1)} + \mathbf{b}_{hr}) \quad (3)$$

$$z_t = \sigma(\mathbf{W}_{iz}x_t + \mathbf{b}_{iz} + \mathbf{W}_{hz}h_{(t-1)} + \mathbf{b}_{hz}) \quad (4)$$

$$c_t = \tanh(\mathbf{W}_{ic}x_t + \mathbf{b}_{ic} + r_t * (\mathbf{W}_{hc}h_{t-1} + \mathbf{b}_{hc})) \quad (5)$$

$$h_t = z_t * c_t + (1 - z_t) * h_{t-1} \quad (6)$$

where h_t is the hidden state at step t , x_t is the input, and r_t , z_t , c_t are the reset, update gates and candidate hidden state, respectively. σ is the sigmoid function and $*$ is the Hadamard product.

3.2. D-GRU Update

As shown in Eq. (6), the new hidden state h_t is the weighted average of the previous hidden state h_{t-1} and the current candidate state c_t . This process is determined by the update gate z_t , in which every element z_t^j has a value between 0 and 1 according to Eq. (4). When z_t^j is close to 1, the hidden state h_t^j is largely replaced by the candidate hidden state c_t^j . Conversely, z_t^j close to 0 means that h_t^j is close to h_{t-1}^j .

In our proposed D-GRU, we only update neurons with the top- A largest values in z_t , i.e., the select gate output is mathematically represented as:

$$g_t^j = \begin{cases} 1 & \text{if } z_t^j \text{ is among the top-}A \text{ largest elements of } z_t \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where $A = (\mathcal{P}/100) * J$ is the number of neurons that need to be updated, with \mathcal{P} being the predefined constant update percentage across steps.

Note that instead of selecting neurons based on the top- A largest z_t^j values, we could also select neurons based on a set threshold, θ^j :

$$g_t^j = \begin{cases} 1 & \text{if } z_t^j > \theta^j \\ 0 & \text{if } z_t^j \leq \theta^j \end{cases} \quad (8)$$

However, this approach results in varying numbers of neurons that require updating in each step, leading to dynamically changing computation loads per step. For practical implementations, the hardware still needs to be capable of doing all the compute for updating of all neurons in order to satisfy the real-time processing constraint. The performance of the threshold based select gate will be discussed in future work.

Following Eq. (2) and Eq. (6), the hidden state of a neuron in the D-GRU is updated as follows:

$$h_t^j = \begin{cases} z_t^j * c_t^j + (1 - z_t^j) * h_{t-1}^j & \text{if } g_t^j = 1 \\ h_{t-1}^j & \text{if } g_t^j = 0 \end{cases} \quad (9)$$

The computations for r_t^j and c_t^j hence can be skipped for neu-

Table 1: Parameter sizes and computation costs of tested models and objective results of them on DNS testset; see Section 5.

| Model | Para. (M) | GRU type | \mathcal{P} (%) | Multiply-Accumulate (MAC) (M/s) | | | PESQ | ESTOI | SISNR | DNS _{SOVAL} |
|-------|-----------|----------|-------------------|---------------------------------|---------|----------------|-------------|-----------|--------------|----------------------|
| | | | | non-GRU | GRU | All layers | | | | |
| noisy | - | - | - | - | - | - | 1.58 | 81 | 9.23 | 2.48 |
| GRU | 1.34 | GRU | 100 | 10.34 | 124.98 | 135.32 (100%) | 2.58 | 90 | 15.47 | 3.22 |
| | | | 75 | | 104.15 | 114.48 (84%) | 2.62 | 90 | 15.68 | 3.22 |
| | | | 50 | | 83.31 | 93.65 (69%) | 2.58 | 90 | 15.59 | 3.20 |
| | | | 25 | | 62.49 | 72.83 (53%) | 2.45 | 89 | 14.82 | 3.15 |
| MPT | 0.28 | GRU | 100 | 71.61 | 210.71 | 282.32 (100%) | 2.66 | 90 | 15.04 | 3.25 |
| | | | 75 | | 175.58 | 247.19 (87%) | 2.65 | 90 | 14.96 | 3.23 |
| | | | 50 | | 140.46 | 212.06 (75%) | 2.63 | 90 | 14.96 | 3.22 |
| | | | 25 | | 105.36 | 176.96 (62%) | 2.57 | 90 | 14.75 | 3.20 |
| DPRNN | 0.35 | GRU | 100 | 811.02 | 4878.35 | 5689.37 (100%) | 2.85 | 92 | 18.08 | 3.29 |
| | | | 75 | | 4065.13 | 4876.14 (85%) | 2.88 | 92 | 17.71 | 3.26 |
| | | | 50 | | 3251.91 | 4062.92 (71%) | 2.82 | 92 | 17.74 | 3.27 |
| | | | 25 | | 2439.18 | 3250.19 (57%) | 2.72 | 91 | 17.28 | 3.24 |

rons whose $g_t^j = 0$:

$$r_t^j = \begin{cases} \sigma(\mathbf{w}_{ir}^j \mathbf{x}_t + b_{ir}^j + \mathbf{w}_{hr}^j \mathbf{h}_{t-1} + b_{hr}^j) & \text{if } g_t^j = 1 \\ \text{(skipped)} & \text{if } g_t^j = 0 \end{cases} \quad (10)$$

$$c_t^j = \begin{cases} \tanh(\mathbf{w}_{ic}^j \mathbf{x}_t + b_{ic}^j + r_t^j (\mathbf{w}_{hc}^j \mathbf{h}_{t-1} + b_{hc}^j)) & \text{if } g_t^j = 1 \\ \text{(skipped)} & \text{if } g_t^j = 0 \end{cases} \quad (11)$$

3.3. Cost Savings

There is an additional cost that is incurred for determining the neurons to update following Eq. (7). However, the computational complexity of this process is $\mathcal{O}(J)$ [19] and is negligible compared to other costs of updating the GRU equations. Following Eqs. (10, 11), we require \mathcal{P} percent of the original computation costs associated with the reset gates and candidate hidden states. Since the computation for the update gate \mathbf{z} cannot be saved, the total computation of the D-GRU is $(1 + 2\mathcal{P}/100)/3$ of that in the conventional GRU.

4. Experimental Setup

The D-GRU is tested on a set of speech enhancement models described in Section 4.1 while the dataset and tested SNR ranges are described in Section 4.2.

4.1. D-GRU Based Speech Enhancement

We applied D-GRU on three strong causal speech enhancement models that here we call 1) *GRU* [20]; 2) *MPT* [15]; and 3) the time-frequency domain *DPRNN* [11].

- The *GRU* model comprises a Fully Connected (FC) layer, two GRU layers, and another FC layer. The output dimensions of the first FC layer and the two GRU layers are 320. The final FC layer outputs the ideal ratio mask (IRM) [20].
- The *MPT* model [15] consists of 2 MPT blocks, each containing four GRUs, with input and output dimensions of 24 and 48, respectively. Following the MPT paper, this model outputs both the IRM and magnitude deep-filtering coefficients [21].
- The time-frequency domain *DPRNN* [11] directly outputs the estimated clean speech’s complex spectrogram. It includes

four dual-path blocks, each composed of three GRUs, where both input and output dimensions of each GRU are 64.

For the *GRU* and *MPT* models which enhance only the magnitude, the input is the noisy magnitude spectrum, and the loss functions are the mean square error (MSE) between the enhanced and target magnitude spectrograms. For the *DPRNN*, with a complex spectrogram as input, the training loss function is an average of the magnitude MSE and the complex spectrogram MSE. For each model, the neuron update percentage \mathcal{P} is set to [100, 75, 50, 25] and the effective computation cost for the D-GRU at each \mathcal{P} are [100%, 83.33%, 66.66%, 50%], respectively (see Section 3.3). The parameter size and the number of MAC operations, with a breakdown between the non-GRU and GRU layers costs, for all three types of models and different \mathcal{P} values are detailed in Table 1.

We highlight that each of the three model types is characterized by their lightweight design, marked by either parameter efficiency, low computation cost, or a combination of both. These characteristics pose a challenge in demonstrating the benefits of the D-GRU, given that these models are already designed for minimal computation cost and size compactness.

4.2. Dataset

The speech enhancement models are trained on the DNS challenge dataset [17], a benchmark dataset used for assessment of speech enhancement algorithms. The dataset comprises 500 hours of clean speech from 2150 speakers and 180 hours of noise signals. Noisy speech is dynamically mixed during the training phase, with the SNR values ranging from -5 dB to 15 dB. The trained models are first evaluated using the no-reverb test set, which consists of 150 test samples. Each file in this set contains a 10-second speech clip. The SNR of the test files in DNS test set ranges from 0 dB to 19 dB with 1 dB interval and the mean SNR of the test set is 9 dB. Additionally, to investigate the influence of SNRs on the D-GRU-based models, the network is further tested on a dataset with noisy speech SNRs at $\{-5, 0, 5\}$ dB, with 150 test utterances at each SNR level.

All training and test signals are sampled at 16 kHz. The complex spectrum for each signal is derived by applying a 320-point Fast Fourier Transform (FFT) on each frame with a size of 20 ms and an overlap of 10 ms.




Figure 2: Mean PESQ improvement (Δ_{PESQ}) of different update percentages \mathcal{P} on different models under $\{-5, 0, 5\}$ dB SNRs.

Table 2: Results of hypothesis tests (see Section 5) in terms of p -values for different models on each $\mathcal{P}_A - \mathcal{P}_B$ pair. Only **bold** values denote the PESQ difference of two tested models is statistically significant.

| Model | \mathcal{P}_A (%) | \mathcal{P}_B (%) | | |
|-------|---------------------|---------------------|-------|--------------|
| | | 75 | 50 | 25 |
| GRU | 100 | 0.258 | 0.508 | 0.035 |
| | 75 | - | 0.265 | 0.009 |
| | 50 | - | - | 0.032 |
| MPT | 100 | 0.401 | 0.377 | 0.121 |
| | 75 | - | 0.469 | 0.172 |
| | 50 | - | - | 0.205 |
| DPRNN | 100 | 0.622 | 0.358 | 0.042 |
| | 75 | - | 0.241 | 0.021 |
| | 50 | - | - | 0.084 |

4.3. Evaluation Metrics

The output of the speech enhancement algorithm is evaluated in terms of: 1) wide band PESQ [22]; 2) ESTOI [23]; 3) Scale-invariant signal-to-noise ratio (SISNR) and 4) DNS OVAL [24]. For all evaluation metrics, higher values denote better enhancement performance.

5. Results

Table 1 shows the results of speech enhancement for all tested models on the DNS test set. The results for the conventional GRU-based models (where $\mathcal{P} = 100\%$) form baselines. Overall, there are only slight differences in speech enhancement across different \mathcal{P} values for different models. Only when \mathcal{P} is reduced to 25% we observe a drop in PESQ scores. We conducted a series of hypothesis tests¹, particularly focusing on PESQ where the most noticeable variations were observed, to determine if these variations are statistically significant. These tests compare the PESQ results at each \mathcal{P} -pair ($\mathcal{P}_A - \mathcal{P}_B$) within the same model. Test results (p -value) below 0.05 signifies that the differences between two tested \mathcal{P} setups are statistically significant.

Table 2 shows the detailed results of these hypothesis tests. For all models, when $\mathcal{P} \geq 50\%$, the differences between the PESQ results are not significant. At these settings, 33% of the computation costs from the GRUs can be saved, and the overall computation costs of the *GRU*, *MPT*, and *DPRNN* models can be reduced to 69%, 75%, and 71%, respectively. For the *GRU* and *DPRNN* models, there is a significant decrease in PESQ when comparing $\mathcal{P} = 25$ to $\mathcal{P} = 100$. This suggests that the

¹Specifically employing the Mann-Whitney U test [25], due to the non-normality of the test data.

observed degradation is more likely due to the inherent limitations of the D-GRUs, where updating only a limited number of neurons may lead to failures in capturing input changes. Additionally, infrequent updates of some neurons might result in inaccurate representations of the current state, introducing internal noise and worsening the speech enhancement. In contrast, for the *MPT* model, no significant differences are observed between different \mathcal{P} setups. This could be attributed to the presence of more non-GRU layers in the model, which might compensate for inaccuracies in the information modeled by the D-GRUs, albeit with a sacrifice of less computation savings.

Fig 2 shows the mean PESQ improvements of each model under $\{-5, 0, 5\}$ dB SNRs. It reveals a consistent trend observed in results on the DNS test set and further demonstrates that even under lower SNR, the dynamic GRU-based models maintain similar PESQ scores compared to the conventional GRU when $\mathcal{P} \geq 50\%$.

6. Related Works and Discussions

Previous proposed models such as the Delta RNNs [18, 26, 27] also reduce computation by allowing only inputs whose change across two steps have exceeded a threshold, to update the neurons. By contrast, our DG-RNN method achieves computation savings by reducing the output dimension through the proposed binary select gate that updates only the \mathcal{P} percent of neurons with the largest update gate, z , values. It would be interesting to integrate both models in a unified framework in future work. Another model, the Phased LSTM [28], also allows only a subset of neurons to be updated on each step. In this model, the neuron is updated only on a particular phase of an oscillation cycle. Additionally, the number of updated neurons varies across steps in Phased LSTM, resulting in dynamically changing computation loads per step. Similarly the Skip RNN model [29, 30, 31] is designed to skip the update of all neurons in some steps, therefore the computational load also changes across steps. In addition, these models have been tested primarily on classification tasks, and their performance on regression tasks, such as speech enhancement, remains to be studied.

7. Conclusion

We propose the DG-RNN architecture for compute-efficient networks running on embedded platforms with limited resource. The DG-RNN reduces computes by updating only a fixed percentage of neurons per step. The test results show that, even with highly efficient RNN-based models and in SNR conditions down to -5 dB, updating only 50% of the GRU neurons at each step can still achieve the same speech enhancement results. In the future, our aim is to investigate the applicability of these models to other modalities and tasks.

8. References

- [1] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [2] K. Tan and D. Wang, "A convolutional recurrent neural network for real-time speech enhancement," in *Proc. INTERSPEECH 2018 – 19th Annual Conference of the International Speech Communication Association*, 2018.
- [3] —, "Complex spectral mapping with a convolutional recurrent network for monaural speech enhancement," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6865–6869.
- [4] Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, and L. Xie, "DCCRN: Deep complex convolution recurrent network for phase-aware speech enhancement," in *Proc. INTERSPEECH 2020 – 21th Annual Conference of the International Speech Communication Association*, 2020.
- [5] A. Li, W. Liu, C. Zheng, C. Fan, and X. Li, "Two heads are better than one: A two-stage complex spectral mapping approach for monaural speech enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1829–1843, 2021.
- [6] X. Le, H. Chen, K. Chen, and J. Lu, "DPCRN: Dual-path convolution recurrent network for single channel speech enhancement," in *Proc. INTERSPEECH 2021 – 22th Annual Conference of the International Speech Communication Association*, 2021.
- [7] A. Defossez, G. Synnaeve, and Y. Adi, "Real time speech enhancement in the waveform domain," in *Proc. INTERSPEECH 2020 – 21th Annual Conference of the International Speech Communication Association*, 2020.
- [8] X. Hao, X. Su, R. Horaud, and X. Li, "Fullsubnet: A full-band and sub-band fusion model for real-time single-channel speech enhancement," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6633–6637.
- [9] J. Chen, Z. Wang, D. Tuo, Z. Wu, S. Kang, and H. Meng, "Fullsubnet+: Channel attention fullsubnet with complex spectrograms for speech enhancement," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 7857–7861.
- [10] A. Li, G. Yu, C. Zheng, W. Liu, and X. Li, "A general unfolding speech enhancement method motivated by Taylor's theorem," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [11] Y. Luo, Z. Chen, and T. Yoshioka, "Dual-path rnn: efficient long sequence modeling for time-domain single-channel speech separation," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 46–50.
- [12] F. Dang, H. Chen, and P. Zhang, "DPT-FSNet: Dual-path transformer based full-band and sub-band fusion network for speech enhancement," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6857–6861.
- [13] A. Pandey, B. Xu, A. Kumar, J. Donley, P. Calamia, and D. Wang, "TPARN: Triple-path attentive recurrent network for time-domain multichannel speech enhancement," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6497–6501.
- [14] J. Yu, Y. Luo, H. Chen, R. Gu, and C. Weng, "High fidelity speech enhancement with band-split rnn," in *Proc. INTERSPEECH 2023 – 24th Annual Conference of the International Speech Communication Association*, 2023.
- [15] H. Chen, J. Yu, and C. Weng, "Complexity scaling for speech denoising," *arXiv preprint arXiv:2309.07757*, 2023.
- [16] H. Chen, J. Yu, Y. Luo, R. Gu, W. Li, Z. Lu, and C. Weng, "Ultra dual-path compression for joint echo cancellation and noise suppression," in *Proc. INTERSPEECH 2023 – 24th Annual Conference of the International Speech Communication Association*, 2023.
- [17] C. K. Reddy, V. Gopal, R. Cutler, E. Beyrami, R. Cheng, H. Dubey, S. Matuskevych, R. Aichner, A. Aazami, S. Braun *et al.*, "The interspeech 2020 deep noise suppression challenge: Datasets, subjective testing framework, and challenge results," in *Proc. INTERSPEECH 2020 – 21th Annual Conference of the International Speech Communication Association*, 2020.
- [18] D. Neil, J. H. Lee, T. Delbruck, and S.-C. Liu, "Delta networks for optimized recurrent network computation," in *International Conference on Machine Learning (ICML)*, 2017, pp. 2584–2593.
- [19] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*. MIT press, 2022.
- [20] A. Narayanan and D. Wang, "Ideal ratio mask estimation using deep neural networks for robust speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7092–7096.
- [21] H. Schroter, A. N. Escalante-B, T. Rosenkranz, and A. Maier, "DeepFilterNet: A low complexity speech enhancement framework for full-band audio based on deep filtering," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 7407–7411.
- [22] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2001, pp. 749–752.
- [23] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 11, pp. 2009–2022, 2016.
- [24] C. K. Reddy, V. Gopal, and R. Cutler, "DNSMOS: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [25] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The Annals of Mathematical Statistics*, pp. 50–60, 1947.
- [26] Z. Jelčićová, R. Jones, D. T. Blix, M. Verhelst, and J. Sparsø, "PeakRNN and StatsRNN: Dynamic pruning in recurrent neural networks," in *European Signal Processing Conference (EU-SIPCO)*, 2021, pp. 416–420.
- [27] C. Gao, T. Delbruck, and S.-C. Liu, "Spartus: A 9.4 TOP/s FPGA-based LSTM accelerator exploiting spatio-temporal sparsity," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 1, pp. 1098–1112, 2024.
- [28] D. Neil, M. Pfeiffer, and S.-C. Liu, "Phased LSTM: Accelerating recurrent network training for long or event-based sequences," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [29] V. Campos, B. Jou, X. Giró-i Nieto, J. Torres, and S.-F. Chang, "Skip RNN: Learning to skip state updates in recurrent neural networks," *ICLR*, 2018.
- [30] I. Fedorov, M. Stamenovic, C. Jensen, L.-C. Yang, A. Mandell, Y. Gan, M. Mattina, and P. N. Whatmough, "TinyLSTMs: Efficient neural speech enhancement for hearing aids," in *Proc. INTERSPEECH 2020 – 21th Annual Conference of the International Speech Communication Association*, 2020.
- [31] X. Le, T. Lei, K. Chen, and J. Lu, "Inference skipping for more efficient real-time speech enhancement with parallel rnns," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2411–2421, 2022.