

CleanUMamba**A Compact Mamba Network for Speech Denoising using Channel Pruning**

Groot, Sjoerd; Chen, Qinyu; Van Gemert, Jan C.; Gao, Chang

DOI[10.1109/ISCAS56072.2025.11043389](https://doi.org/10.1109/ISCAS56072.2025.11043389)**Publication date**

2025

Document Version

Final published version

Published in

Proceedings of the 2025 IEEE International Symposium on Circuits and Systems (ISCAS)

Citation (APA)

Groot, S., Chen, Q., Van Gemert, J. C., & Gao, C. (2025). CleanUMamba: A Compact Mamba Network for Speech Denoising using Channel Pruning. In *Proceedings of the 2025 IEEE International Symposium on Circuits and Systems (ISCAS)* (Proceedings - IEEE International Symposium on Circuits and Systems). IEEE. <https://doi.org/10.1109/ISCAS56072.2025.11043389>

Important note

To cite this publication, please use the final published version (if applicable).

Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

CleanUMamba: A Compact Mamba Network for Speech Denoising using Channel Pruning

Sjoerd Groot, Qinyu Chen, Jan C. van Gemert, Chang Gao

Abstract—This paper presents **CleanUMamba**, a time-domain neural network architecture designed for real-time causal audio denoising directly applied to raw waveforms. **CleanUMamba** leverages a U-Net encoder-decoder structure, incorporating the Mamba state-space model in the bottleneck layer. By replacing conventional self-attention and LSTM mechanisms with Mamba, our architecture offers superior denoising performance while maintaining a constant memory footprint, enabling streaming operation. To enhance efficiency, we applied structured channel pruning, achieving an 8X reduction in model size without compromising audio quality. Our model demonstrates strong results in the Interspeech 2020 Deep Noise Suppression challenge. Specifically, **CleanUMamba** achieves a PESQ score of 2.42 and STOI of 95.1% with only 442K parameters and 468M MACs, matching or outperforming larger models in real-time performance. Code will be available at: <https://github.com/lab-emi/CleanUMamba>

Index Terms—deep learning, speech enhancement, audio denoising, state-space model, convolutional neural network

I. INTRODUCTION

Audio denoising, or speech enhancement, removes background noise from speech recordings while preserving quality and intelligibility. This technology is important for applications such as hearing aids, audio calls, and speech recognition systems. Traditional methods like spectral subtraction [1] and Wiener filtering [2] struggle in dynamic environments, especially when noise overlaps with speech.

Deep neural networks (DNNs) have significantly advanced this field in recent years. Various architectures, including convolutional neural networks (CNNs) [3]–[6], recurrent neural networks (RNNs) [7]–[9], and transformers [10], [11], have been used to enhance speech. Although transformers achieve high-quality results, they are computationally expensive at longer input sequences.

The Mamba state space model [12] offers a promising solution for sequence modeling and time series prediction. Mamba enables parallel computation during training and recurrent processing during inference without sequence length constraints, making it a strong candidate for audio denoising.

In this work, we introduce **CleanUMamba**, a neural network designed for real-time audio denoising that processes raw waveforms. **CleanUMamba** adopts the U-Net encoder-decoder architecture from [10], replacing self-attention with Mamba state-space blocks, which enables a more compact

Corresponding Author: Chang Gao (chang.gao@tudelft.nl)

S. Groot, J. C. van Gemert and C. Gao are with the Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, The Netherlands.

Q. Chen is with the Leiden Institute of Advanced Computer Science (LIACS), Leiden University, The Netherlands.

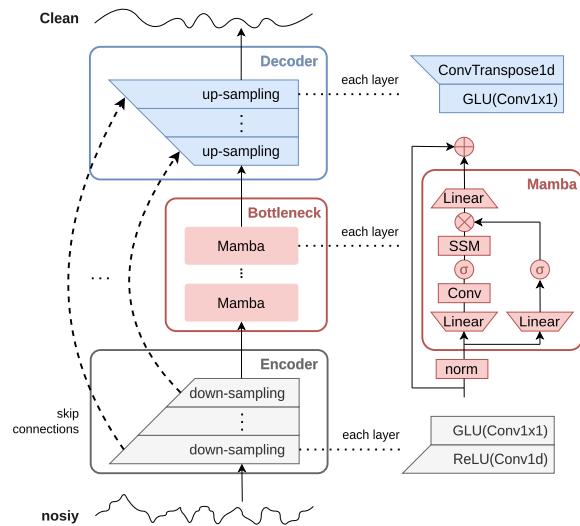


Fig. 1: CleanUMamba architecture consisting of convolutional encoder-decoder with 3 sequential Mamba blocks in the bottleneck.

model with reduced algorithmic latency. Additionally, we implement structured pruning to reduce the model size while maintaining high performance.

Our main contributions are:

- 1) A novel Mamba-based architecture for time-domain speech enhancement with a 12 ms real-time algorithmic latency.
- 2) A comparative analysis of Mamba, self-attention, LSTM, and Mamba-S4 for audio denoising.
- 3) An efficient, structured pruning strategy using periodic calibration of GroupTaylor importance [13].

II. RELATED WORKS

Recent research has explored the application of Mamba to audio denoising, with several concurrent works emerging. Many of these studies focus on non-causal speech enhancement for offline settings, where the entire audio track is available for processing. Notable works include SPMamba [14], DPMamba [15], Mamba-TasNet [16], and SEMamba [17], which build upon TF-Gridnet [11], SlowFast [18], Dual-path RNN [19], Conv-TasNet [4], and MP-SENet [20] respectively. These approaches incorporate bidirectional Mamba implementations for efficient global audio processing. TRAMBA [21] enhances a time domain U-Net Temporal

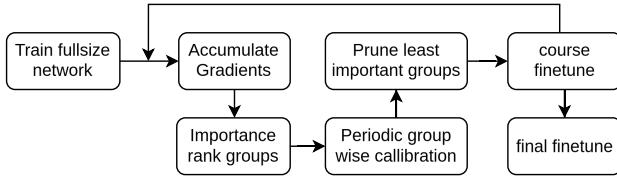


Fig. 2: Pruning Pipeline

TABLE I: Multi-Head Attention (MHA), LSTM, Mamba, and Mamba S4 compared in the same size unit model and training conditions on the DNS no-reverb test set.

Model	Params	MACs	PESQ (WB)	PESQ (NB)	STOI (%)	pred. CSIG	pred. CBAK	pred. COVRL
Mamba	442K	468M	2.42	2.95	95.1	3.98	3.25	3.21
Mamba S4	451K	468M	2.36	2.90	94.9	3.93	<u>3.20</u>	3.15
MHA	443K	470M	<u>2.37</u>	<u>2.92</u>	94.9	<u>3.94</u>	<u>3.20</u>	<u>3.16</u>
LSTM	443K	463M	2.32	2.88	94.7	3.90	3.19	3.12

FiLM [22] by integrating Mamba in the bottleneck and using self-attention for feature-wise linear modulation. This network is applied to audio super-resolution and reconstruction from bone-conducting microphone and accelerometer data. For long-term streaming applications, oSpatialNet-Mamba [23] extends SpatialNet [24] by replacing self-attention with masked self-attention, Retention, or Mamba. Operating in the time-frequency domain with 2-4 second window sizes, Mamba outperforms other variations. Zhang et al. [25] conduct an ablation study comparing different backbones, including Mamba and two bidirectional Mamba implementations, in the time-frequency domain.

In contrast to these works, our research focuses on real-time causal speech enhancement in the time domain. By applying Mamba to the U-Net architecture, we reduce computational load compared to DPMamba and Mamba-Tasnet through processing in a lower-resolution latent space.

III. PROPOSED METHOD

A. Problem Definition

Audio denoising aims to recover the clean speech signal $x \in R^T$ from the noisy signal $y = x + v$, where v represents zero-mean noise uncorrelated with x . A causal model for denoising reconstructs the clean speech $\hat{x}_t = f(y_{1:t}) \approx x_t$ using noisy samples up to time t . In real-world applications, a slight look-ahead, such as a delay of 5-6 ms for hearing aids [26], [27], or up to 200 ms for video calls [28], is acceptable. This work uses 16kHz audio with CleanUMamba, achieving an algorithmic delay of 48 ms and 12 ms for 8 and 6 encoder layers, respectively.

B. Network Architecture

Figure 1 shows the architecture of CleanUMamba, consisting of an encoder-decoder structure with E layers. Each encoder layer employs a 1D convolution with a kernel size of 4 and stride of 2, followed by a ReLU activation and a 1x1 convolution with a GLU activation. These layers halve the

temporal resolution and include bypass connections to their respective decoder layers. Decoder layers reverse this process using 1D transposed convolutions, doubling the temporal resolution. The channels per layer start at $H = 64$ and increase up to a maximum of 768.

In the bottleneck, three Mamba blocks perform sequence modeling on the latent representation, which has a model dimension of $D = 512$. Each Mamba block is preceded by layer normalization and a residual connection, while the inner dimension is set to $I = 2048$, corresponding to CleanUNet’s fully connected layer output [10]. The state-space model uses $S = 64$ channels.

C. Pruning Pipeline

Pruning aims to minimize the loss increase per pruned parameter. Let $\mathcal{L}(\theta)$ represent the loss function of the network with parameters θ , and $\Delta\mathcal{L}$ denote the change in loss due to pruning. The objective is:

$$\min_{\mathcal{S}} \frac{\Delta\mathcal{L}(\mathcal{S})}{|\mathcal{S}|} \quad (1)$$

where \mathcal{S} is the set of parameters to prune, and $|\mathcal{S}|$ is its size. Since exact loss sampling for every group is computationally infeasible, the Group Taylor importance metric [13] is used. It estimates the importance of parameter sets with both absolute and squared gradients:

$$I_S = \sum_{s \in \mathcal{S}} |g_s w_s| \quad (2)$$

$$I_S = \sum_{s \in \mathcal{S}} (g_s w_s)^2 \quad (3)$$

$$I_S = \sum_{s \in \mathcal{S}} |w_s| \quad (4)$$

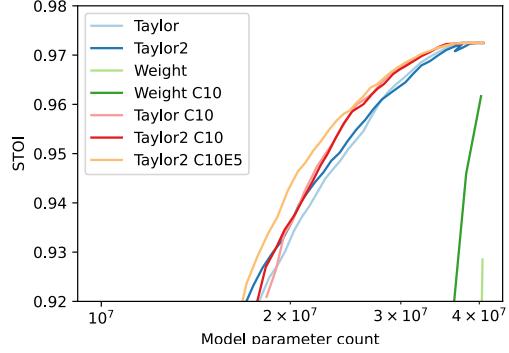
Here, g_s and w_s represent the gradient and weight for group s . Gradients are accumulated via backpropagation. The pruning pipeline (Fig. 2) follows a standard train-prune-finetune approach [29], accumulating micro-batch gradients [30]. Pruning targets a percentage of groups, selecting those with the lowest Group Taylor importance, while ensuring compatibility with Mamba’s causal convolutions by pruning in multiples of 8 channels.

Though effective, the Taylor importance metric may over- or under-penalize groups of varying sizes or depths. To address this, periodic calibration is applied by pruning 20% of groups in a layer, measuring the loss increase, and adjusting the global metric accordingly. An exponential moving average filter smooths out any noise, preventing a single outlier from disproportionately affecting the pruning process.

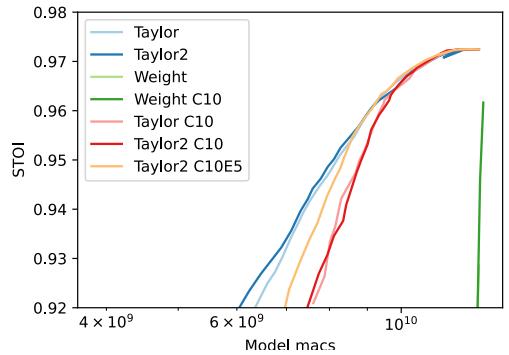
IV. EXPERIMENTAL SETUP

A. Evaluation Metrics

We evaluated the model using several objective metrics: Perceptual Evaluation of Speech Quality (PESQ) [31], Short-Time Objective Intelligibility (STOI) [32], and the Mean Option Score (MOS) for signal distortion (CSIG), background noise intrusiveness (CBAK), and overall quality (COVL) [33].



(a) Model parameter count vs STOI



(b) Model MACs vs STOI

Fig. 3: Comparison of importance metrics for pruning CleanUMamba without fine-tuning.

B. Dataset

The experiments used the Interspeech 2020 Deep Noise Suppression (DNS) dataset [34]. This dataset includes speech from 2150 speakers and 65,000 noise clips, yielding 500 hours of training data, with signal-to-noise ratios (SNR) between -5 and 25 dB across 31 levels. During training, random 10-second crops were selected.

C. Training Setup

We employed a loss function combining L1 loss and multi-resolution short-time Fourier transform (STFT) loss [35]:

$$L(x, \hat{x}) = \|x - \hat{x}\| + STFT(x, \hat{x}) \quad (5)$$

$$STFT(x, \hat{x}) = \sum_{i=1}^m \left(\frac{\|s(x; \theta_i) - s(\hat{x}; \theta_i)\|_F}{\|s(x; \theta_i)\|_F} + \frac{1}{T} \|\log \frac{s(x; \theta_i)}{s(\hat{x}; \theta_i)}\| \right) \quad (6)$$

STFT loss was calculated for FFT bins {512, 1024, 2048}, hop sizes {50, 120, 240}, and window lengths {240, 600, 1200}. Both full-band and high-frequency (4kHz-8kHz) STFT losses were tested [10].

The network was trained with the ADAM optimizer (learning rate 0.0002, $\beta_1 = 0.9$, $\beta_2 = 0.999$) using a linear warm-up for the first 5% of training followed by cosine decay. PyTorch AMP was applied for mixed-precision training.

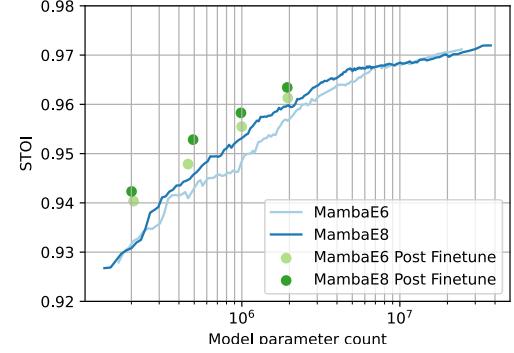


Fig. 4: Pruning CleanUMamba with 6 and 8 encoder layers with fine-tuning: Model parameter count vs STOI

A Wiener filter baseline [36], implemented with the Pyroomacoustics library [37], used an LPC order of 10, window size of 256, 2 iterations, smoothing alpha of 0.5, and threshold set for 16% noise classification. The output was mixed with 50% unfiltered audio to optimize quality.

V. EXPERIMENTAL RESULTS

A. Ablation Study of Mamba

In Table I, we compare the performance of Mamba in the bottleneck layer against Multi-head Attention, Mamba S4, and LSTM. All models use the same encoder-decoder architecture with $E = 8$ layers, starting with $H = 32$ channels in the first layer and $H = 64$ channels in subsequent layers. Each model employs a bottleneck with a model dimension of $D = 64$.

For the multi-head attention variant, 4 heads were used, each with a dimension of 16, followed by an MLP expanding to $d_{inner} = 128$, in line with the architecture in [10], but with fewer parameters. The Mamba model uses an inner dimension of $d_{inner} = 128$ and a state size $S = 16$, while Mamba S4 uses the same configuration but replaces the selective state space with the S4 state space [39]. The LSTM model consists of 3 layers with a hidden size of 64.

All models were trained on the DNS dataset for 1M iterations with a batch size of 16. The ADAMW optimizer was employed, with a weight decay of 0.1, using the full STFT loss. Autocast was disabled for Mamba S4 due to training errors.

B. CleanUMamba

In Table II, we present the results of training CleanUMamba with the same procedure as CleanUNet, using $N = 3$ bottleneck layers and encoder depths of $E = 6$ and $E = 8$. When trained with high-band loss, CleanUMamba outperforms both the same-size and larger CleanUNet models. However, under full-band loss, CleanUMamba slightly underperforms compared to CleanUNet.

The model with $E = 6$ performs marginally worse than the $E = 8$ version but has fewer than 75% of the parameters.

TABLE II: Evaluation results for denoising on the DNS no-reverb testset.

Model	Params	MACs	look ahead	PESQ (WB)	PESQ (NB)	STOI (%)	pred. CSIG	pred. CBAK	pred. COVRL
raw noisy waveform	-	-	-	1.585	2.164	91.6	3.091	2.539	2.304
wiener filter	-	-	32ms	1.645	2.237	91.6	3.099	2.558	2.336
DEMUCS [7]	33.53M	-	40ms	2.659	3.229	96.6	4.145	3.627	3.419
FullSubNet [38]	5.6 M	-	32ms	2.777	3.305	96.1	-	-	-
CleanUNet N5 full [10]	46.07M	13.66G	48ms	3.146	3.551	97.7	4.619	3.892	3.932
CleanUNet N5 high [10]	46.07M	13.66G	48ms	3.011	3.460	97.3	4.513	3.812	3.800
CleanUNet N3 full [10]	39.77M	13.18G	48ms	3.128	3.539	97.6	-	-	-
CleanUNet N3 high [10]	39.77M	13.18G	48ms	3.006	3.453	97.3	-	-	-
CleanUMamba E8 full (ours)	41.37M	13.38G	48ms	3.067	3.507	97.4	4.502	3.870	3.829
CleanUMamba E8 high (ours)	41.37M	13.38G	48ms	3.017	3.471	97.2	4.456	3.840	3.775
Pruned CleanUMamba E8 high (ours)	14.90M	6.08G	48ms	2.910	3.397	97.0	4.393	3.742	3.682
	6.00M	3.87G		2.888	3.359	96.9	4.380	3.711	3.661
	3.22M	2.09G		2.746	3.253	96.4	4.272	3.599	3.530
	1.94M	1.97G		2.707	3.222	96.3	4.244	3.570	3.495
	0.99M	1.29G		2.558	3.102	95.8	4.125	3.466	3.356
	492K	807M		2.426	2.980	95.3	3.996	3.351	3.219
CleanUMamba E6 high (ours)	201K	403M	12ms	2.189	2.745	94.2	3.784	3.163	2.986
	27.21M	13.48G		2.935	3.400	97.1	4.415	3.785	3.710
Pruned CleanUMamba E6 high (ours)	13.50M	7.49G	12ms	2.855	3.346	96.9	4.366	3.710	3.637
	7.31M	4.97G		2.799	3.291	96.8	4.331	3.659	3.590
	1.95M	2.14G		2.602	3.128	96.1	4.171	3.499	3.402
	1.00M	1.36G		2.431	2.967	95.5	4.033	3.365	3.241
	457K	858M		2.237	2.796	94.8	3.855	3.218	3.048
	207K	483M		2.096	2.660	94.0	3.715	3.104	2.902

Despite this, the number of multiply accumulates per second (MAC/s) remains comparable due to the longer input to the bottleneck layer.

With $E = 8$, most of the computations occur in the encoder and decoder, leading to a doubling of MAC/s after 18 minutes, making Mamba’s fixed state size less significant. Reducing the encoder depth to $E = 6$ quadruples the bottleneck sequence length and reduces MAC/s in the encoder-decoder. As a result, after just 80 seconds, MAC/s with multi-head attention doubles compared to Mamba.

C. Different Importance Metrics without Fine-Tuning

To assess the effectiveness of different importance metrics, the CleanUMamba network trained with high loss was pruned without fine-tuning, as shown in Figure 3. At each pruning step, 24 of the least important groups were removed, with 128 audio samples used for gradient accumulation in the Taylor importance metric. The calibrated runs, denoted $C10$, recalibrate the importance metric every 10 pruning steps, and the exponential moving average, $E5$, applies a smoothing factor of 0.5.

Results indicate that the Group Taylor importance metric effectively reflects global importance. The squared Taylor loss slightly outperforms the absolute Taylor loss at higher pruning levels. In contrast, weight magnitude performs poorly as a global importance metric, with significant performance drops after pruning the first few layers.

With calibration, Taylor loss initially shows better results, but stability decreases over time. Filtering the calibration improves quality per parameter, while uncalibrated Taylor loss consistently delivers superior quality per MAC/s.

D. Pruning with Fine-Tuning

For pruning with fine-tuning, CleanUMamba models with $E = 6$ and $E = 8$ encoder layers pre-trained using high loss were pruned using squared Taylor importance. The metric was recalibrated every 20 steps with a smoothing factor of 0.5. At each step, gradients were accumulated over 128 samples, and the model was fine-tuned on 40,960 training samples every 5 steps. The importance pruned was limited to $3e-13$, with 4 channels pruned per step once this limit was reached.

However, both E6 and E8 models became under-trained below 6M and 4M parameters, respectively. To address this, models at 200K, 500K, 1M, and 2M parameters were fine-tuned for an additional 100K iterations with a batch size of 16.

Figure 4 shows the STOI scores during pruning, and Table II presents a subset of the full evaluation. At 3.22M parameters, CleanUMamba matches the performance of DEMUCS with $8 \times$ fewer parameters and nearly matches FullSubNet with $2 \times$ fewer parameters.

VI. CONCLUSION

In this study, we introduce CleanUMamba, a real-time speech denoising model that operates in the waveform domain and investigate its size reduction through pruning. We evaluated the model on the Interspeech 2020 Deep Noise Suppression challenge and compared it to the self-attention-based CleanUNet. While both models perform similarly with an encoder depth of 8, Mamba’s linear time complexity becomes more advantageous with a reduced depth of 6, achieving a 12 ms latency. Our pruning pipeline further reduces model size, achieving performance comparable to LSTM-based DEMUCs with $8 \times$ fewer parameters.

REFERENCES