Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

EffiFusion-GAN: Efficient Fusion Generative Adversarial Network for Speech Enhancement

BIN WEN
Universiti Sains Malaysia
Tien-Ping Tan
tienping@usm.my
Universiti Sains Malaysia

Article

Keywords: speech enhancement, generative adversarial network, depthwise-separable convolution, model pruning, Conformer

Posted Date: August 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10363711/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

EffiFusion-GAN: Efficient Fusion Generative Adversarial Network for Speech Enhancement

Bin Wen, Tien-Ping Tan*
Universiti Sains Malaysia
Penang, Malaysia
* Corresponding author: tienping@usm.my

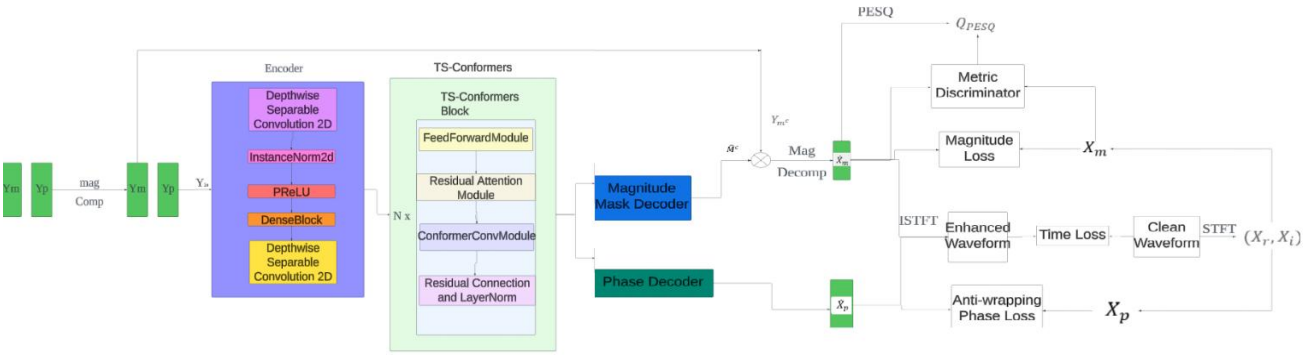
Abstract—Speech enhancement models increasingly estimate magnitude and phase jointly, but the resulting architectures can be difficult to deploy under tight model-size constraints. We present EffiFusion-GAN, a time–frequency generative adversarial network designed to preserve enhancement quality with fewer trainable parameters. The generator combines a depthwise-separable dense encoder, residual time–frequency Conformer blocks, and parallel magnitude-mask and phase decoders; L1 unstructured pruning removes 30% of the selected convolutional weights. On VoiceBank+DEMAND, EffiFusion-GAN achieves PESQ 3.45, CSIG 4.71, CBAK 3.91, COVL 4.18, and STOI 0.96 with a reported parameter count of 1.08 million. Its PESQ is 0.05 below MP-SENet while its reported parameter count is 47% lower. Ablation results show that replacing depthwise-separable convolution with standard convolution nearly doubles the parameter count for a 0.02 PESQ increase, whereas removing residual connections reduces PESQ by 0.10. These findings support a favorable parameter–quality trade-off. Future work should confirm latency, memory use, and robustness on deployment hardware and unseen acoustic conditions.

Keywords—speech enhancement; generative adversarial network; depthwise-separable convolution; model pruning; Conformer

I. INTRODUCTION

Speech enhancement aims to recover intelligible, natural speech from recordings degraded by environmental noise. Deep models approach this problem either in the waveform domain [1–3] or in a time–frequency representation obtained by the short-time Fourier transform (STFT) [4]. Waveform models avoid explicit spectral reconstruction, whereas time–frequency models can separately model spectral structure that varies over time. Many early time–frequency systems emphasized magnitude estimation and reused the noisy phase because phase is wrapped and difficult to model. However, phase errors can materially affect perceptual quality, particularly at low signal-to-noise ratios [5]. Recent magnitude–phase systems therefore improve reconstruction quality, but their additional branches and attention blocks can increase model size.

This work addresses the resulting design tension: a useful model should recover both magnitude and phase without requiring an unnecessarily large parameter budget. Standard convolutions provide strong channel mixing but can dominate the parameter count of dense time–frequency encoders. Depthwise-separable convolution reduces this burden by separating spatial filtering from channel mixing [6–8], although aggressive factorization may weaken feature interaction. EffiFusion-GAN combines a depthwise-separable dense encoder with residual time–frequency Conformer blocks, parallel magnitude and phase decoders, and L1 unstructured pruning. The central question is therefore not whether the model obtains the single highest score, but whether it retains competitive enhancement quality at a smaller reported parameter count.



The main contributions are as follows:

- A compact magnitude–phase generator combines depthwise-separable dense encoding with parallel decoders, reducing the reported parameter count while retaining multi-scale time–frequency feature extraction.
- Residual time–frequency Conformer blocks integrate local convolutional structure with long-range attention; the ablation study isolates the contribution of the residual paths.
- Evaluation on VoiceBank+DEMAND quantifies the parameter–quality trade-off and separately examines depthwise-separable convolution, residual connections, and 30% L1 unstructured pruning.

II. Related Work

Adversarial speech enhancement has progressed from waveform generation in SEGAN [9] to metric-oriented optimization in MetricGAN and MetricGAN+ [10,11]. Time–frequency models have increasingly incorporated structured spectral modeling: PHASEN explicitly models phase [12], DPT-FSNet fuses full-band and sub-band information [13], CMGAN combines Conformer blocks with metric-guided adversarial learning [14], and MP-SENet processes magnitude and phase in parallel [15]. These systems demonstrate the value of phase-aware and attention-based enhancement, but they also motivate a closer examination of model footprint. EffiFusion-GAN builds on this line of work by making the parameter–quality trade-off the primary design objective rather than treating compactness as a secondary property.

III. Methods

Figure 1 summarizes the proposed time–frequency generator. A noisy waveform is transformed by STFT into magnitude Ym and phase Yp. Power-law-compressed magnitude and phase features are stacked as Yin and encoded by depthwise-separable convolutional blocks. Four residual time–frequency Conformer blocks then model dependencies along both axes. Parallel decoders estimate a magnitude mask and the clean phase; inverse STFT reconstructs the enhanced waveform. During training, 30% L1 unstructured pruning is applied to selected generator convolutional weights. This design concentrates model capacity on joint magnitude–phase reconstruction while controlling the reported parameter count.

A. *Model Architecture*

1) *Encoder*

As shown in Supplementary Fig. S2, the encoder maps Yin to a compact time–frequency representation. It contains an initial depthwise-separable convolutional block, a four-layer dilated

DenseNet module [16], and a final depthwise-separable block. Each block uses two-dimensional depthwise-separable convolution, instance normalization, and parametric ReLU (PReLU) [17].

The first block expands the channel dimension. The DenseNet layers use dilation rates of 1, 2, 4, and 8 along the temporal axis, enlarging the receptive field without a proportional increase in depth. Dense connections concatenate intermediate representations so that features at different temporal scales remain available to later layers.

The final block downsamples the encoded representation by adjusting the convolution stride. The resulting compact feature map is passed to the residual attention blocks and the two decoders.

2) ***Decoder***

The decoder has parallel magnitude-mask and phase branches (Supplementary Fig. S1) [15]. The magnitude branch predicts a mask that is applied to the noisy magnitude spectrum. Because the direct ratio M = Xm/Ym is unbounded, both clean and noisy magnitudes are power-law compressed with factor c = 0.3. A learnable sigmoid (LSigmoid) [11] then predicts the compressed mask without imposing a fixed conventional sigmoid mapping.

The phase branch predicts pseudo-real Xp(r) and pseudo-imaginary Xp(i) components using two parallel two-dimensional convolutional layers. The clean phase Xp is recovered with the two-argument arctangent. Predicting these components avoids direct regression across the phase wrap boundary and supports joint magnitude–phase waveform reconstruction [18].

B. ***Discriminator***

The discriminator evaluates the enhanced magnitude spectrum relative to the clean reference and returns a scalar quality score (Figure 2). Clean and enhanced magnitudes are concatenated along the channel dimension to form an input of shape [B, 2, T, F].

A stack of two-dimensional convolutions and downsampling layers extracts local spectral and temporal patterns. Spectral normalization constrains convolutional weights, while instance normalization and PReLU activations support stable feature learning across samples.

Adaptive max pooling converts variable-duration feature maps to a fixed-size representation. Fully connected layers with dropout then map this representation to a scalar, and a learnable sigmoid constrains the output to [0,1].

During adversarial training, the discriminator encourages the generator to produce spectra that receive a quality score close to that of the clean reference. The generator objective also contains waveform, magnitude, complex-spectrum, and phase reconstruction terms.

The combined objective balances signal fidelity in complementary domains rather than relying on a single reconstruction criterion.

The individual loss terms are defined as follows.

The time-domain L1 loss penalizes waveform reconstruction error:

$$LTime = \|x - \hat{x}\|_1 \quad (1)$$

The magnitude-domain mean-squared error is

$$LMag = \|Xm - \hat{X}m\|_2{}^2 \quad (2)$$

The complex-spectrum loss applies mean-squared error to the real and imaginary components, where Xr and Xi denote the clean components and the hatted variables denote predictions:

$$LCom = \|Xr - \hat{X}r\|^2 + \|Xi - \hat{X}i\|^2 \quad (3)$$

The phase objective combines instantaneous phase (LIP), group delay (LGD), and instantaneous angular frequency (LIAF) losses to reduce errors associated with phase wrapping:

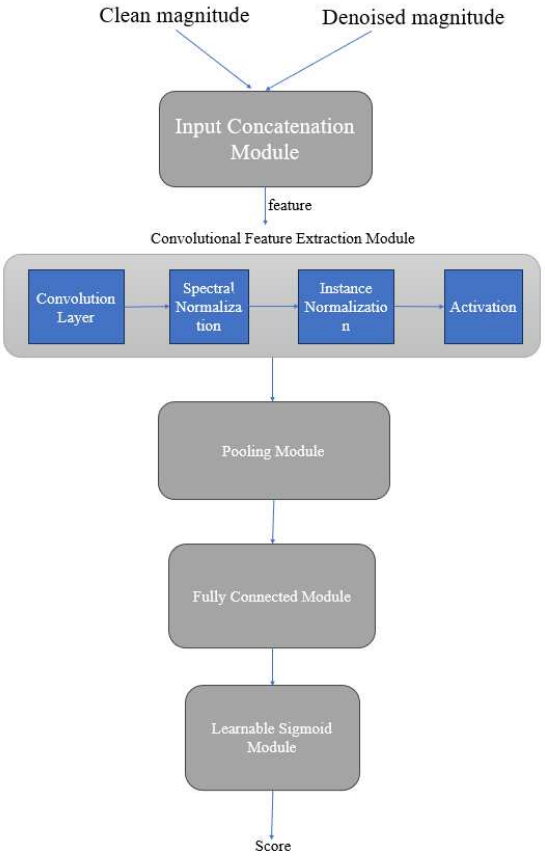
$$LPha = LIP + LGD + LIAF \quad (4)$$

The perceptual metric loss uses the discriminator score to encourage outputs with a target quality label:

$$LMetric = F \times mse_loss(metricg \times flatten(\cdot), one_labels) \quad (5)$$

The complete generator objective is a weighted sum of the five terms:

$$LGen = 0.05 \cdot LMetric + 0.9 \cdot LMag + 0.3 \cdot LPha + 0.1 \cdot LCom + 0.2 \cdot LTime \quad (6)$$



Use of generative AI

OpenAI Codex was used to assist with manuscript restructuring and English-language editing. It was not used to generate or alter data, analyses, or figures. The authors reviewed and verified all scientific content, analyses, interpretations, and references and take full responsibility for the final manuscript.

IV. EXPERIMENTS AND RESULTS

A. Dataset and Experimental Setup

Experiments used the publicly available VoiceBank+DEMAND benchmark [19,20]. The original 48-kHz recordings were resampled to 16 kHz.

The official split was retained. The training set contains 11,572 utterances from 28 speakers; the test set contains 872 utterances from two unseen speakers mixed with five unseen noise types at 2.5, 7.5, 12.5, and 17.5 dB. The test set was used for final reporting. PESQ was also monitored every 1,000 steps during training, and the reported checkpoint was selected from the late training stage using average performance rather than a single peak. Because a separate development split was not defined, this procedure may yield optimistic benchmark estimates; future evaluation should reserve an independent validation set.

EffiFusion-GAN contains a generator and a metric discriminator. The generator comprises a depthwise-separable dense encoder with dilation rates 1, 2, 4, and 8; four residual time–frequency Conformer blocks with four attention heads; and parallel magnitude-mask and phase decoders. The discriminator uses spectrally normalized convolutions and instance normalization to produce the adversarial quality score.

Audio was processed at 16 kHz using an STFT with a 512-point FFT, a 512-sample window, and a 128-sample hop. The magnitude compression factor was 0.3, and training segments contained 48,000 samples (3 s). Both networks were trained for 200 epochs with AdamW and an ExponentialLR schedule. At initialization, L1 unstructured pruning set 30% of the selected generator convolutional weights to zero. The generator loss weights were 0.9 for magnitude, 0.3 for phase, 0.1 for complex spectrum, 0.2 for waveform L1, and 0.05 for the adversarial metric term. Checkpoints were saved every 5,000 steps and PESQ was evaluated every 1,000 steps.

B. ***Comparison with Published Methods***

Table 1 compares EffiFusion-GAN with representative waveform and time–frequency systems on VoiceBank+DEMAND. We report PESQ, segmental signal-to-noise ratio (SSNR), short-time objective intelligibility (STOI), and the composite CSIG, CBAK, and COVL measures. Higher values indicate better performance. Values for comparison methods are taken from the cited publications; missing values are marked with a dash.

The analysis focuses on enhancement quality relative to reported parameter count. Parameter count is a model-footprint indicator [21], but it does not by itself establish latency, memory traffic, or energy consumption.

- Parameter–quality trade-off: EffiFusion-GAN reaches PESQ 3.45 with 1.08M reported parameters, compared with PESQ 3.50 and 2.05M parameters for MP-SENet. Thus, the proposed model gives up 0.05 PESQ while reducing the reported parameter count by approximately 47%. This result supports a smaller model footprint, although deployment efficiency requires direct hardware measurements.
- Comparison at a similar scale: DPT-FSNet is the closest listed model in size (0.88M parameters). EffiFusion-GAN reports higher PESQ (3.45 versus 3.33), CSIG (4.71 versus 4.58), CBAK (3.91 versus 3.72), and COVL (4.18 versus 4.00). Because these results come from separate published systems, the comparison should be interpreted as benchmark-level evidence rather than a controlled head-to-head experiment.
- Performance across metrics: EffiFusion-GAN obtains CSIG 4.71, CBAK 3.91, COVL 4.18, SSNR 10.12, and STOI 0.96. The STOI value matches the strongest values reported in Table 1, while the other quality measures remain close to MP-SENet. Together, the metrics indicate that the reduction in

reported parameter count is not accompanied by a large decline in the evaluated speech-quality measures.

Overall, the benchmark and ablation results support the intended design objective: competitive speech-enhancement scores at a smaller reported parameter count. The evidence supports a parameter-efficiency claim, but not yet a general claim of faster or lower-energy inference.

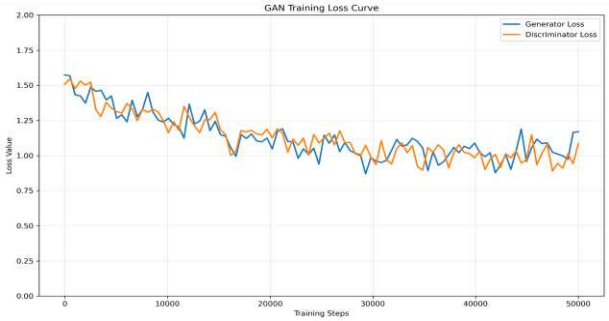
C. *Training Dynamics*

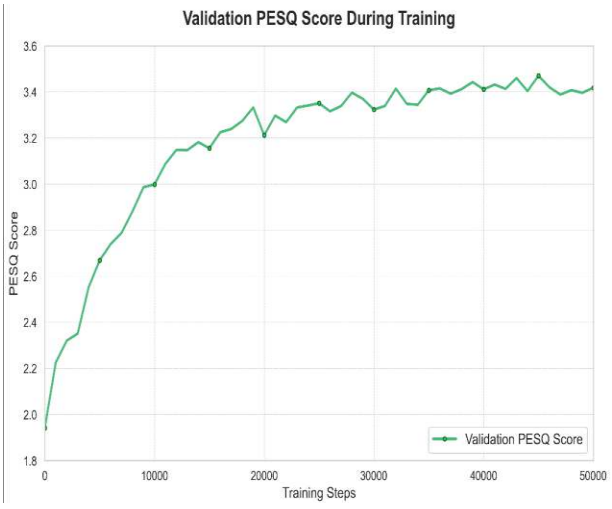
Generator and discriminator losses were tracked for 50,000 training steps. Figure 3 shows the observed adversarial training trajectories, and Figure 4 shows the corresponding PESQ trend.

Both losses begin near 1.55 and decrease during approximately the first 10,000 steps. This pattern is consistent with early optimization of the reconstruction and discrimination objectives, although absolute generator and discriminator loss values are not directly comparable.

After the initial decline, the curves oscillate and cross as the two networks continue to adapt. Such oscillation is expected in adversarial optimization; curve crossings alone do not prove an equilibrium, so stability is assessed primarily by the absence of sustained divergence and by the accompanying PESQ trajectory.

After roughly 20,000 steps, both losses remain within a bounded range and the PESQ curve approaches a plateau. These observations indicate convergence for this run. Residual paths, attention, normalization, and pruning may contribute to the observed behavior, but the loss curves alone do not isolate their causal effects; the ablation results provide the more direct evidence for individual components.





D. *Descriptive Training-Stability Comparison*

Supplementary Figs. S4 and S5 provide a descriptive comparison of the reported loss trajectories for EffiFusion-GAN, SEGAN, MetricGAN, CMGAN, and MP-SENet. Because adversarial losses can differ in definition, scale, and implementation across models, these plots are used to describe curve behavior rather than to establish a formal statistical ranking of training stability.

Two qualitative properties are examined:

- Discriminator-loss trajectory: boundedness and the amplitude of oscillation.
- Generator-loss trajectory: downward trend, boundedness, and late-stage variability.

In Supplementary Fig. S4, EffiFusion-GAN exhibits a comparatively narrow oscillation range over the displayed training period. Within the limits of this descriptive comparison, the curve is consistent with bounded discriminator optimization.

Supplementary Fig. S5 likewise shows a gradual decline and relatively small late-stage fluctuations for the EffiFusion-GAN generator loss. This observation is consistent with stable optimization in the reported run, but repeated runs are needed to quantify variability across random seeds.

Taken together, the curves show no sustained divergence for EffiFusion-GAN. They complement, but do not replace, the component-level ablation study. A stronger stability analysis would report repeated runs, common loss definitions, and summary measures such as late-stage variance with confidence intervals.

E. *Ablation Analysis*

Table 2 isolates the contributions of depthwise-separable convolution, residual connections, and L1 unstructured pruning. Each variant changes one or more components and reports both model size and enhancement quality.

- Depthwise-separable convolution: Replacing it with standard convolution increases the reported parameter count from 1.08M to 2.04M, while PESQ rises from 3.45 to 3.47. The 0.02 gain is small relative to the 0.96M increase, supporting the use of depthwise-separable layers when model footprint is the priority.
- Residual connections: Removing residual paths leaves the parameter count unchanged but reduces PESQ from 3.45 to 3.35 and lowers CSIG, CBAK, and COVL. The consistent decline indicates that the residual paths contribute to enhancement quality in this architecture.
- L1 pruning: Without pruning, the reported parameter count rises from 1.08M to 1.75M and PESQ changes from 3.45 to 3.41. This pattern is consistent with a regularizing effect, but a single ablation cannot determine whether regularization, optimization variability, or another factor caused the difference.
- Combined components: The standard-convolution, unpruned variant reaches the highest PESQ in the table (3.48) with 2.04M parameters. EffiFusion-GAN is 0.03 lower in PESQ with approximately 47% fewer reported parameters, which summarizes the intended parameter–quality trade-off.

The ablation results therefore assign distinct roles to the three design choices: depthwise-separable convolution and pruning reduce the reported model footprint, while residual connections preserve enhancement quality. Results across multiple random seeds would be useful to quantify uncertainty in these differences.

V. CONCLUSION

EffiFusion-GAN combines a depthwise-separable dense encoder, residual time–frequency Conformer blocks, parallel magnitude and phase decoders, and L1 unstructured pruning. On VoiceBank+DEMAND, the model achieves PESQ 3.45 with a reported 1.08M parameters. Relative to MP-SENet, this is a 0.05 reduction in PESQ and an approximately 47% reduction in reported parameter count. The ablation study shows that standard convolution yields only a small PESQ gain at nearly twice the parameter count, whereas removing residual connections produces the clearest quality decline. These findings support EffiFusion-GAN as a compact magnitude–phase enhancement model. The present evidence is limited to one benchmark and uses parameter count as the main efficiency proxy. Future work should evaluate independent validation protocols, repeated runs, unseen acoustic conditions, and direct latency, memory, and energy measurements on target hardware.

FUNDING

The authors received no funding for this work.

AUTHOR CONTRIBUTIONS

B.W. developed the method, conducted the experiments, analysed the results, and drafted the manuscript. T.P.T. supervised the work, contributed to the study design and interpretation, and revised the manuscript. Both authors reviewed and approved the final manuscript.

DATA AVAILABILITY

The VoiceBank+DEMAND corpus used in this study is publicly available through the University of Edinburgh DataShare repository, and the DEMAND noise recordings are available from Zenodo. The trained model checkpoints, enhanced test outputs, per-utterance evaluation scores, and implementation code are available from the corresponding author upon reasonable request.

COMPETING INTERESTS

The authors declare no competing interests.

REFERENCES

1. Pandey, A. & Wang, D. TCNN: Temporal convolutional neural network for real-time speech enhancement in the time domain. In Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 6875–6879 (2019).
2. Defossez, A., Synnaeve, G. & Adi, Y. Real time speech enhancement in the waveform domain. In Proc. Interspeech 3291–3295 (2020).
3. Kim, E. & Seo, H. SE-Conformer: Time-domain speech enhancement using conformer. In Proc. Interspeech 2736–2740 (2021).
4. Strauss, M. & Edler, B. A flow-based neural network for time domain speech enhancement. In Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 5754–5758 (2021).
5. Paliwal, K., Wojcicki, K. & Shannon, B. The importance of phase in speech enhancement. *Speech Commun.* **53**, 465–494 (2011).
6. Chollet, F. Xception: Deep learning with depthwise separable convolutions. In Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 1251–1258 (2017).
7. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. & Chen, L. C. MobileNetV2: Inverted residuals and linear bottlenecks. In Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 4510–4520 (2018).
8. Zhang, X., Zhou, X., Lin, M. & Sun, J. ShuffleNet: An extremely efficient convolutional neural network for mobile devices. In Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 6848–6856 (2018).
9. Pascual, S., Bonafonte, A. & Serra, J. SEGAN: Speech enhancement generative adversarial network. In Proc. Interspeech 3642–3646 (2017).

10. Fu, S.-W., Liao, C.-F., Tsao, Y. & Lin, S.-D. MetricGAN: Generative adversarial networks based black-box metric scores optimization for speech enhancement. In Proc. International Conference on Machine Learning (ICML) 2031–2041 (2019).
11. Fu, S.-W. et al. MetricGAN+: An improved version of MetricGAN for speech enhancement. In Proc. Interspeech 201–205 (2021).
12. Yin, D., Luo, C., Xiong, Z. & Zeng, W. PHASEN: A phase-and-harmonics-aware speech enhancement network. *Proc. AAAI Conf. Artif. Intell.* **34**, 9458–9465 (2020).
13. Dang, F., Chen, H. & Zhang, P. DPT-FSNet: Dual-path transformer based full-band and sub-band fusion network for speech enhancement. In Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 6857–6861 (2022).
14. Cao, R., Abdulatif, S. & Yang, B. CMGAN: Conformer-based Metric GAN for speech enhancement. In Proc. Interspeech 936–940 (2022).
15. Lu, Y., Ai, Y. & Ling, Z.-H. MP-SENet: A speech enhancement model with parallel denoising of magnitude and phase spectra. Preprint at https://arxiv.org/abs/2305.13686 (2023).
16. Pandey, A. & Wang, D. Densely connected neural network with dilated convolutions for real-time speech enhancement in the time domain. In Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 6629–6633 (2020).
17. He, K., Zhang, X., Ren, S. & Sun, J. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In Proc. IEEE International Conference on Computer Vision (ICCV) 1026–1034 (2015).
18. Ai, Y. & Ling, Z.-H. Neural speech phase prediction based on parallel estimation architecture and anti-wrapping losses. In Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (2023).
19. Valentini-Botinhao, C., Wang, X., Takaki, S. & Yamagishi, J. Investigating RNN-based speech enhancement methods for noise-robust text-to-speech. In Proc. Speech Synthesis Workshop (SSW) 146–152 (2016).
20. Thiemann, J., Ito, N. & Vincent, E. The diverse environments multi-channel acoustic noise database (DEMAND): A database of multichannel environmental noise recordings. *Proc. Meet. Acoust.* **19**, 035081 (2013).
21. Drgas, S. A survey on low-latency DNN-based speech enhancement. *Sensors* **23**, 1380 (2023).

FIGURE LEGENDS

Figure 1. Overview of the proposed EffiFusion-GAN architecture. A noisy waveform is transformed into magnitude and phase features, encoded by depthwise-separable convolutional blocks, processed by residual time–frequency Conformer modules, and reconstructed through parallel magnitude-mask and phase decoders. The discriminator supplies the perceptual metric objective.

Figure 2. Discriminator architecture. Clean and enhanced magnitudes are concatenated, processed by normalized convolutional feature-extraction layers, pooled, and mapped to a scalar quality score.

Figure 3. Generator and discriminator losses during training. Loss trajectories for the reported EffiFusion-GAN run over 50,000 optimization steps.

Figure 4. PESQ trajectory during training. PESQ values evaluated at 1,000-step intervals over 50,000 steps.

TABLES

Table 1. Comparison on VoiceBank+DEMAND. Values for prior methods are drawn from the cited papers; "–" indicates an unreported result.

Method	Year	Input	Param	PESQ	CSIG	CBAK	COVL	SSNR	STOI
Noisy	–	–	–	1.97	3.35	2.44	2.63	1.68	0.91
SEGAN [9]	2017	Waveform	43.18M	2.16	3.48	2.94	2.80	7.73	0.92
MetricGAN [10]	2019	Magnitude	–	2.86	3.99	3.18	3.42	–	–
MetricGAN+ [11]	2021	Magnitude	–	3.15	4.14	3.16	3.64	–	0.96
DPT-FSNet [13]	2021	Complex	0.88M	3.33	4.58	3.72	4.00	–	0.96
CMGAN [14]	2022	Magnitude+Complex	1.83M	3.41	4.63	3.94	4.12	11.10	0.96
PHASEN [12]	2020	Magnitude+Phase	–	2.99	4.21	3.55	3.62	10.08	–
MP-SENet [15]	2023	Magnitude+Phase	2.05M	3.50	4.73	3.95	4.22	10.64	0.96
EffiFusion-GAN	2024	Magnitude+Phase	1.08M	3.45	4.71	3.91	4.18	10.12	0.96

Table 2. Ablation results for the principal architectural components

Method	Param	PESQ	CSIG	CBAK	COVL
EffiFusion-GAN	1.08M	3.45	4.71	3.91	4.18
w/o Depthwise	2.04M	3.47	4.76	3.90	4.17
w/o Res	1.08M	3.35	4.64	3.86	4.08
w/o pruning	1.75M	3.41	4.69	3.87	4.12
w/o Depthwise & w/o Res	2.04M	3.38	4.67	3.85	4.10
w/o Depthwise & w/o pruning	2.04M	3.48	4.78	3.92	4.19
w/o Res & w/o pruning	1.75M	3.32	4.61	3.83	4.05

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- EffiFusionGANSupplementaryInformation.docx