



URGENT Challenge: Universality, Robustness, and Generalizability For Speech Enhancement

Wangyou Zhang^{1,2}, Robin Scheibler³, Kohei Saijo⁴, Samuele Cornell¹, Chenda Li^{1,2}, Zhaoheng Ni⁵,
Jan Pirklbauer⁶, Marvin Sach⁶, Shinji Watanabe¹, Tim Fingscheidt⁶, Yanmin Qian²

¹Carnegie Mellon University, USA ²Shanghai Jiao Tong University, China ³LY Corp., Japan
⁴Waseda University, Japan ⁵Meta, USA ⁶Technische Universität Braunschweig, Germany

wyz-97@sjtu.edu.cn, robin.scheibler@linecorp.com, saijo@pcl.cs.waseda.ac.jp

Abstract

The last decade has witnessed significant advancements in deep learning-based speech enhancement (SE). However, most existing SE research has limitations on the coverage of SE sub-tasks, data diversity and amount, and evaluation metrics. To fill this gap and promote research toward universal SE, we establish a new SE challenge, named URGENT, to focus on the universality, robustness, and generalizability of SE. We aim to extend the SE definition to cover different sub-tasks to explore the limits of SE models, starting from denoising, dereverberation, bandwidth extension, and declipping. A novel framework is proposed to unify all these sub-tasks in a single model, allowing the use of all existing SE approaches. We collected public speech and noise data from different domains to construct diverse evaluation data. Finally, we discuss the insights gained from our preliminary baseline experiments based on both generative and discriminative SE methods with 12 curated metrics.

Index Terms: speech enhancement, universality, robustness, generalizability

1. Introduction

Speech enhancement (SE) is the task of improving a speech signal that has been subject to distortions such as additive noise, acoustic interference, reverberation, or bandwidth limitation. In recent years, we have witnessed the rapid development of deep learning-based SE techniques, with impressive performance under matched conditions [1]. However, most conventional SE approaches focus only on denoising or dereverberation in a *limited range of conditions*, such as single-channel, multi-channel, anechoic, etc. Usually, they tend to only train and evaluate SE models on one or two common datasets, such as the Voice-Bank+DEMAND [2] and Deep Noise Suppression (DNS) Challenge datasets [3–7]. The evaluation is often restricted to simulated conditions similar to those of training. This greatly impedes a comprehensive understanding of the generalizability and robustness of SE methods. In addition, such practice can impact the model design process as it can favor models that are only suitable for limited conditions or have limited capacity to handle more complicated scenarios.

Apart from conventional discriminative methods, generative approaches have also attracted a lot of attention. They are good at handling different distortions with a single model [8,9] and tend to generalize better than discriminative methods [10]. However, their capability and universality have not yet been fully understood through a comprehensive benchmark. Meanwhile, recent efforts [11] have shown the possibility of building a single system to handle various input formats, such as different sampling frequencies and numbers of microphones. However, there *lacks a well-established benchmark* covering a wide range of conditions, and, crucially, *no systematic comparison*

has been made yet between state-of-the-art (SOTA) discriminative and generative methods regarding their generalizability.

We believe that the community should focus on this problem urgently. And we propose a new challenge, which is called URGENT, to boost the research on **Universality, Robustness, and Generalizability** for speech **Enhancement**. The key contributions and innovations of this challenge are listed below:

1) *Broader definition of the SE task:* In most real-world scenarios, speech is likely to be degraded by several of the distortions mentioned previously, and the recording devices may also vary in the sampling frequency. So, it is important to build a universal SE model that can handle different distortions and input formats. Although it is possible to build a separate SE model for each distortion and each input format, having a single universal SE model is more efficient and simpler to deploy. Crucially, it can avoid the error propagation that occurs when cascading several specialized models. It may also improve the overall performance by sharing knowledge among different sub-tasks, a direction to be explored during this challenge. To facilitate this exploration, we further propose a technically novel framework (see Section 3.2) that is general for different SE approaches.

2) *Larger scale and more diverse data with training data mandated and limited:* As mentioned above, SE models are often evaluated on fixed, small (e.g., ~10 h) or medium datasets (e.g., ~100 h). It is very likely that recent SOTA models heavily overfit these datasets. Furthermore, the test sets associated with these datasets are often matched in terms of speech quality, linguistic content, noise family, and other characteristics. The matter is complicated by the scarcity of truly high-quality, anechoic speech recordings. Large-scale speech datasets are typically recorded with diverse types of equipment under diverse “sub-optimal” conditions and are not equalized. It is unclear how current SE models can scale with a larger amount of “sub-optimal” data and whether they can generalize well to unseen conditions. In this challenge, we aim to explore this aspect with public data. But unlike most earlier challenges, we mandate and limit the (still large amount of) training material, giving us better insights into the actual capabilities of the various network architectures.

3) *Extensive evaluation metrics:* Existing challenges often only adopt one or two objective metrics for evaluation, which cannot provide a comprehensive understanding of the SE models. For example, some models are trained with a particular evaluation metric (e.g., scale-invariant signal-to-noise ratio [12]) discriminatively leading to biased evaluation. Beyond a variety of task-dependent intrusive and non-intrusive metrics, as a challenge novelty, we also adopt metrics that are downstream task-independent (e.g., phoneme similarity). This not




Figure 1: URGENT speech enhancement task definition.

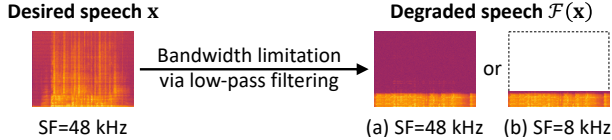


Figure 2: Example of bandwidth limitation (48 kHz \rightarrow 8 kHz).

only fits perfectly to our multi-task challenge, but also promises to allow a better comparison of generative and discriminative SE methods.

2. Related Works

Existing SE challenges have fostered the development of SE models for specific scenarios, such as denoising and dereverberation [3–7], speech restoration [13, 14], packet loss concealment [15], acoustic echo cancellation [16–19], hearing aids [20, 21], 3D SE [22–24], far-field multi-channel SE for video conferencing [25], and unsupervised domain adaptation for denoising [26]. These challenges have greatly advanced SE studies. The URGENT challenge uniquely focuses on universality, generalizability, and robustness in a wide range of scenarios and evaluation metrics, complementing existing challenges.

3. Challenge Description

3.1. Task definition

As shown in Figure 1, we define an SE model $\text{SE}(\cdot)$ in the URGENT challenge as the following general form:

$$\hat{x} = \text{SE}(\mathcal{F}(\mathbf{x})), \quad (1)$$

where \mathbf{x} and \hat{x} are the desired and enhanced speech signals, respectively. $\mathcal{F}(\cdot)$ is the distortion model that degrades the desired signal. The resultant degraded speech $\mathcal{F}(\mathbf{x})$ serves as the input to SE models. Our definition differs from the commonly-adopted definition in the literature in two aspects. First, the model input $\mathcal{F}(\mathbf{x})$ can have *various sampling frequencies (SF)*, while conventional SE systems often only consider a fixed SF. Second, the distortion model \mathcal{F} covers *diverse distortions* (i.e., additive noise, reverberation, clipping, and bandwidth limitation), while conventional SE systems mostly apply noise or reverberation. Note that multichannel signals are not considered in this challenge to simplify the problem. We leave them for our future challenges.

3.2. Baseline systems unifying various SFs and sub-tasks

The URGENT challenge prepares several baseline systems, which have been carefully designed to unify different SE sub-tasks in a simple manner. As shown in Figure 1, the model input $\mathcal{F}(\mathbf{x})$ is simulated by applying different distortions to the original desired speech \mathbf{x} . In this procedure, there are two conditions that involve SF variations and require special treatment.

First, when bandwidth limitation is applied to the desired speech \mathbf{x} , corresponding to the bandwidth extension (BWE) sub-task, its high-frequency components are removed via low-

Table 1: Detailed information of the corpora used in our baseline experiments². \dagger denotes the data is not used in this paper.

Type	Training Set	Validation Set	Non-blind Test Set
Speech	LibriVox data from DNS5 challenge [7]	Same as left	Added one unseen corpus
	LibriTTS reading speech [32]		
	CommonVoice 11.0 English portion [33] [†]		
	VCTK reading speech [34]		
Noise	WSJ reading speech [35, 36] [†]	Same as left	Added two unseen corpora
	Audioset+FreeSound noise in DNS5 challenge		
RIR	WHAM! noise [37]	Same as left	Real recorded RIRs
	Simulated RIRs from DNS5 challenge		
	Other simulated RIRs [†]		

pass filtering. This process typically corresponds to downsampling of the signal, as shown in Figure 2 (b). However, it can also appear in a high SF with its upper frequencies missing due to poor microphone devices, as shown in Figure 2 (a). To unify these two scenarios and to be consistent with other distortions, we always keep the SF unchanged in this procedure as illustrated in Figure 2 (a). We further design the enhanced speech \hat{x} to have the same SF as input \mathbf{x} in Figure 1 so that we can easily unify the data format in BWE and other SE sub-tasks.

Second, the model input $\mathcal{F}(\mathbf{x})$ can have various SFs as mentioned in Section 3.1, which cannot be handled by most conventional SE models directly. One solution is to adopt the so-called sampling-frequency-independent (SFI) SE approaches [11, 27–29], as shown in the upper right-hand corner of Figure 1. The SFI SE models feature a strong generalizability to different SFs, even though only trained on a fixed SF or limited SFs. Here, we adopt the SFI short-time Fourier transform (STFT) based design due to its zero-shot capability [11, 28], which uses *fixed-duration* window and hop sizes in STFT and iSTFT. Specifically, we apply this design to two recently proposed time-frequency dual-path SE models, e.g., BSRNN¹ [29] and TF-GridNet [1] to achieve SFI processing.

On the other hand, we also adopt a simple yet effective solution for most existing SE approaches that only support a single SF [30]. As shown in the lower right-hand corner of Figure 1, we always upsample the model input $\mathcal{F}(\mathbf{x})$ to the highest SF (48 kHz) as pre-processing, so the model only takes 48 kHz data as input. The generated output is also 48 kHz, which will be downsampled to the original input SF for loss calculation as well as for generating the final enhanced speech \hat{x} . We apply this design to Conv-TasNet [31] as an additional baseline.

With this framework (i.e., data and model design), we can easily build an SE system to handle different sub-tasks and SFs.

3.3. Data

We collect diverse speech, noise, and room impulse response (RIR) samples from public corpora to construct the datasets for this challenge. As shown in Table 1, we combine 4 public speech corpora, 2 noise corpora, and 1 RIR corpus for preparing the training and validation sets. For the non-blind test set, we additionally add 1 unseen speech corpus, 2 unseen noise corpora, and real RIR samples to evaluate the generalizability. To generate simulation datasets for both training and evaluation, we first preprocess the data as introduced in Section 3.3.1 and then simulate the data according to Section 3.3.2. The data preparation scripts will be made publicly available.

¹This differs from BSRNN’s original SFI design, where the input signal is always upsampled to 48 kHz. We verified that our design could achieve comparable performance with better generalizability.

²Although both LibriTTS and DNS5 speech data in Table 1 come from LibriVox, they only occupy $\sim 40\%$ of the total speech data.

3.3.1. Preprocessing

Since the speech and noise samples are collected from different sources with diverse devices, the effective bandwidth may not be equal to their default SF due to resampling and device discrepancies. Meanwhile, our baselines in Section 3.2 rely on the accurate bandwidth information (ground-truth SF) to perform BWE and other sub-tasks. And it also allows more accurate metric computation with the actual bandwidth information. Therefore, it is critical to detect the true bandwidth of each sample and resample it accordingly. In addition, we observe that speech samples from diverse corpora may be actually non-speech, or can contain noise, or have low quality. It is thus important to filter out such samples, especially for generative approaches. To tackle the above issues, we adopt the following procedure as data preprocessing³:

- 1) We first follow the algorithm proposed in [38] to estimate the effective bandwidth of each speech and noise sample, and then resample it to the best matching SF⁴.
- 2) We use a voice activity detection (VAD) algorithm⁵ to filter out speech samples that are detected to be non-speech or dominated by silence.
- 3) We calculate the DNSMOS scores (OVRL, SIG, BAK) [39] for each speech sample and set a threshold for each score to filter out noisy and low-quality speech samples.

Some interesting observations in this stage are:

- While the original LibriVox English data from DNS5 challenge [7] should be in 48 kHz, after the above preprocessing, we found out that about 50% of them have an SF of 32 kHz, and about 20% of them have SFs between 8 kHz and 24 kHz. Similar phenomena are also observed in LibriTTS [32] and CommonVoice [33].
- 19 speech samples in the LibriVox portion of the DNS5 Challenge data are detected to be actually non-speech.
- Some speech samples in the LibriVox data from the DNS5 Challenge are found to contain multiple speakers sequentially⁶. However, we decide not to filter out such samples to allow the SE models to learn to cope with them.

Through this process, we finally obtain a curated list of speech (~1300 hours) and noise (~250 hours) samples that will be used for data simulation of training and test data in Section 3.3.2.

3.3.2. Simulation

We design the data simulation process by considering both speed and reproducibility. For fixed data simulation, a manifest is firstly generated from the given list of speech, noise, and RIR samples. It specifies how each sample will be simulated, including the type of distortion to be applied, the speech/noise/RIR sample to be used, the signal-to-noise ratio (SNR), the random seed, and so on. Then, the simulation can be done in parallel for different samples according to the manifest while ensuring reproducibility. This procedure can be used to generate training, validation, and non-blind test datasets. For the training set, we also recommend dynamically generating degraded speech samples during training to increase the data diversity. It should be noted that only the listed corpora in Table 1 shall be used

³The detailed procedure can be found at https://github.com/urgent-challenge/urgent2024_challenge.

⁴The best matching SF is defined as the lowest SF that can fully cover the effective frequency range.

⁵<https://github.com/wiseman/py-webrtcvad>

⁶We only detected such samples manually.

to generate the training and validation data. This is to ensure a fair comparison and proper understanding of various SE approaches. This rule particularly deviates from DNS Challenges which allowed the use of arbitrary training data.

3.4. Evaluation metrics

To comprehensively evaluate the baseline models, we adopt a wide range of evaluation metrics, including⁷

- *intrusive SE metrics*: POLQA [40], PESQ [41], extended short-time objective intelligibility (ESTOI) [42], signal-to-distortion ratio (SDR) [43], mel cepstral distortion (MCD) [44], log-spectral distance (LSD) [45];
- *non-intrusive SE metrics*: DNSMOS [39], NISQA [46];
- *downstream-task-independent metrics*: phoneme similarity (PhnSim, equal to “1-LPD” in [47]), SpeechBERTScore [48];
- *downstream-task-dependent metrics*: speaker similarity (SpkSim), word accuracy (WAcc)⁸.

Among them, a lower value in MCD and LSD indicates better performance, while in all other metrics a higher value corresponds to better performance. The intrusive SE metrics require well-aligned reference speech for calculation, and can reflect the objective quality of enhanced speech. The non-intrusive SE metrics are calculated by pre-trained neural networks which do not require reference speech. They are useful to evaluate the generative approaches or when no aligned reference speech is available. The downstream task independent metrics compare the enhanced speech and reference speech based on some task-agnostic representation (e.g., phoneme prediction and discrete tokens). Note that although they require reference speech as an additional input, no strict alignment is needed. The PhnSim metric captures frame-wise phone information in the enhanced speech, which is useful for comparing generative and discriminative approaches in the correctness of their generated contents. The SpeechBERTScore metric measures the similarity between semantic embeddings of reference and enhanced speech. The downstream task related metrics use a pre-trained model to evaluate the downstream task performance such as speaker similarity and WAcc. These allow us to further exploit real-recorded data for evaluation. We use the RawNet3 [49] model pre-trained on VoxCeleb datasets for cosine-based speaker similarity calculation and the OWSM v3.1 [50] model for WAcc calculation.

The rank of different SE systems will be obtained by considering all these metrics. More specific ranking rules will be updated on the challenge website <https://urgent-challenge.github.io/urgent2024/>, which are facilitated by our investigation in this paper.

4. Experiments

We report hereafter a preliminary investigation on different baselines as mentioned in Section 3.2 for the challenge. This includes comparing generative and discriminative methods using a wide range of metrics introduced in Section 3.4.

4.1. Data configuration

As a preliminary investigation before the challenge, we primarily conducted experiments on a fixed simulation dataset to compare different baseline approaches. The fixed simulation dataset is generated following the procedure described in Section 3.3.2, resulting in ~400 hours of training samples, ~30 hours of validation samples, and ~15 hours of test samples. Note that this

⁷The final evaluation metrics in the challenge may differ slightly.

⁸WAcc is equal to 1 - word error rate (WER).

Table 2: Evaluation on non-blind test data. Results with * are not fully comparable due to different data and training setups.

Model	#Param	#MACs (48 kHz)	Non-intrusive SE metrics		Intrusive SE metrics					Downstream-task-independent		Downstream-task-dependent		
			DNSMOS ↑	NISQA ↑	POLQA ↑	PESQ ↑	ESTOI (×100) ↑	SDR (dB) ↑	MCD ↓	LSD ↓	SpeechBERTScore ↑	PhnSim ↑	SpkSim ↑	WAcc (%) ↑
Noisy input	-	-	1.64	1.76	2.50	1.63	70.40	6.11	6.76	3.99	0.87	0.68	0.72	82.18
OM-LSA [51]	-	-	2.19	2.09	2.37	1.81	70.24	10.88	5.26	3.64	0.85	0.71	0.65	78.61
VoiceFixer [9]*	116.8 M	-	2.93	3.65	1.97	1.50	52.71	-9.59	9.16	7.54	0.81	0.59	0.54	66.19
Conv-TasNet	40.0 M	38 G/s	2.31	2.71	3.12	2.42	79.91	14.42	3.23	2.73	0.85	0.73	0.70	76.82
BSRNN	37.8 M	78 G/s	2.41	3.05	3.49	2.66	83.29	14.89	2.75	2.66	0.87	0.80	0.77	82.53
TF-GridNet	8.5 M	401 G/s	2.43	3.06	3.54	2.76	84.05	15.42	2.70	2.39	0.87	0.81	0.78	82.87




Figure 3: Radar plot of different baseline models.

data is only used for our preliminary exploration, while the final challenge data may be updated according to our findings. The SNR ranges from -5 dB to 20 dB, and reverberation is added to each sample with a probability of 0.5 . Note that the generated dataset covers a wide range of sampling frequencies, i.e., $\{8, 16, 22.05, 24, 32, 44.1, 48\}$ kHz. During training, we always segment each sample into 4 s chunks for better efficiency. All models are trained using the L1-based time-domain plus frequency-domain multi-resolution loss [52], where we adopt four STFT window sizes $\{256, 512, 768, 1024\}$ to obtain different time-frequency resolutions. All baseline experiments have been done using the ESPnet [53] toolkit.

4.2. Model configuration

As mentioned in Section 3.2, we evaluate the performance of four different baseline models, including BSRNN, TF-GridNet, and Conv-TasNet. We follow the best model configuration in the original papers for TF-GridNet and Conv-TasNet, except that the encoder/decoder kernel sizes are scaled to match a 48 kHz input. In BSRNN, the STFT window and hop sizes are set to 20 ms and 10 ms, respectively. We stack 6 BSRNN blocks with a relatively large embedding dimension (196) to enhance the model capacity. Due to the space limitation, we omit the details for the model configuration, which can be found in the official repository⁹ for reproducibility.

4.3. Experimental results and discussion

As shown in Table 2, we compare the performance of different SE models with a wide range of evaluation metrics on the simulated non-blind test set. Our baselines trained on the challenge data show consistent improvement in most metrics. Among them, the masking-based SE approach (Conv-TasNet) has the worst performance, which is clearly illustrated in Figure 3. The breakdown results imply that it cannot work well on

⁹https://github.com/urgent-challenge/urgent2024_challenge

band-limited samples, which is attributed to the inherent limitation of masking-based approaches¹⁰. In contrast, mapping-based methods (i.e., BSRNN and TF-GridNet) show better performance in all SE sub-tasks, demonstrating their potential for unifying multiple SE sub-tasks. It is also interesting that all metrics share a similar tendency among the discriminative approaches. This verifies the feasibility of building a universal SE system with a high-capacity SE model.

In addition to the baseline models described in Section 3.2, we also present the evaluation results of OM-LSA [51] and VoiceFixer [9]. The former is a representative denoising method based on signal processing, which serves as a weak baseline since it can only handle the denoising sub-task. VoiceFixer is a vocoder-based generative SE approach¹¹, which is trained to process the same set of distortions as mentioned in Section 3.1. Note that VoiceFixer is trained to only process data in 44.1 kHz, so we always resample the input to 44.1 kHz and the output back to the original SF for evaluation. Since VoiceFixer is trained on a different dataset, no fair comparison can be made. Thus, this model only serves as a reference to check the effectiveness of the other baselines. We leave a comparable generative baseline for future work. As expected, it achieves unmatched DNSMOS and NISQA scores, confirming the strength of generative SE approaches to generate natural speech. Meanwhile, our diverse metrics also demonstrate their capability of detecting the ‘‘hallucination’’ of generative SE approaches. For example, the low PhnSim and SpkSim scores can indicate inconsistent contents and speaker traits in the generated speech. This also verifies the necessity of using a wide range of evaluation metrics to capture various properties of discriminative and generative SE methods.

5. Conclusion

In this paper, we have introduced a new SE challenge, URGENT, which aims to promote research towards universal SE with strong generalizability and robustness. This new challenge features a broad definition of the SE task, large scale and diverse data based on public corpora, and extensive evaluation metrics. A novel framework has been proposed to facilitate this exploration, allowing easy extension of existing SE models to handle multiple SE sub-tasks and different sampling frequencies. We open source all scripts for data preparation, baseline training, and extensive evaluation. As a preliminary investigation, we conducted experiments with several baselines on the simulated data. The results verified the potential of both generative and discriminative SE approaches, each dominating a different set of metrics. Our goal is to attract more research towards building universal SE models with strong robustness and good generalizability. In future work, we aim to extend the challenge to cover more scenarios, such as additional distortions, more microphone channels, multiple speakers, etc.

¹⁰Not presented here due to space limitation. We provide the breakdown results at https://drive.google.com/file/d/1lu9dbHPxmexG_AHDrmxU75kNzqk12UQy/view.

¹¹Available at <https://github.com/haoheliu/voicefixer>. We adopted ‘‘mode 0’’ as it performs best.

6. Acknowledgment

The experiments were done using the PSC Bridges2 system via ACCESS allocation CIS210014, supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296. Samuele Cornell was supported by IC Postdoctoral Research Fellowship Program at Carnegie Mellon University via ORISE. Wangyou Zhang, Chenda Li, and Yanmin Qian were supported in part by China STI 2030-Major Projects under Grant No. 2021ZD0201500, in part by China NSFC projects under Grants 62122050 and 62071288, and in part by Shanghai Municipal Science and Technology Commission Project under Grant 2021SHZDZX0102.

7. References

- [1] Z.-Q. Wang *et al.*, “TF-GridNet: Integrating full-and sub-band modeling for speech separation,” *IEEE/ACM Trans. ASLP*, vol. 31, pp. 3221–3236, 2023.
- [2] C. Valentini-Botinhao *et al.*, “Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks,” in *Interspeech*, 2016, pp. 352–356.
- [3] C. K. Reddy *et al.*, “The INTERSPEECH 2020 deep noise suppression challenge: Datasets, subjective testing framework, and challenge results,” in *Interspeech*, 2020, pp. 2492–2496.
- [4] C. K. Reddy *et al.*, “ICASSP 2021 deep noise suppression challenge,” in *ICASSP*, 2021, pp. 6623–6627.
- [5] C. K. Reddy *et al.*, “INTERSPPEECH 2021 deep noise suppression challenge,” in *Interspeech*, 2021, pp. 2796–2800.
- [6] H. Dubey *et al.*, “ICASSP 2022 deep noise suppression challenge,” in *ICASSP*, 2022, pp. 9271–9275.
- [7] H. Dubey *et al.*, “ICASSP 2023 deep noise suppression challenge,” *IEEE Open Journal of Signal Processing*, pp. 1–13, 2024.
- [8] J. Serrà *et al.*, “Universal speech enhancement with score-based diffusion,” *arXiv preprint arXiv:2206.03065*, 2022.
- [9] H. Liu *et al.*, “VoiceFixer: A unified framework for high-fidelity speech restoration,” in *Interspeech*, 2022, pp. 4232–4236.
- [10] Y.-J. Lu *et al.*, “Conditional diffusion probabilistic model for speech enhancement,” in *ICASSP*, 2022, pp. 7402–7406.
- [11] W. Zhang *et al.*, “Toward universal speech enhancement for diverse input conditions,” in *Proc. IEEE ASRU*, 2023.
- [12] J. Le Roux *et al.*, “SDR—half-baked or well done?” in *ICASSP*, 2019, pp. 626–630.
- [13] R. Cutler *et al.*, “ICASSP 2023 speech signal improvement challenge,” *IEEE Open Journal of Signal Processing*, pp. 1–12, 2024.
- [14] N. C. Ristea *et al.*, “ICASSP 2024 speech signal improvement challenge,” *arXiv preprint arXiv:2401.14444*, 2024.
- [15] L. Diener *et al.*, “INTERSPPEECH 2022 audio deep packet loss concealment challenge,” in *Interspeech*, 2022, pp. 580–584.
- [16] R. Cutler *et al.*, “INTERSPPEECH 2021 acoustic echo cancellation challenge,” in *Interspeech*, 2021, pp. 4748–4752.
- [17] K. Sridhar *et al.*, “ICASSP 2021 acoustic echo cancellation challenge: Datasets, testing framework, and results,” in *ICASSP*, 2021, pp. 151–155.
- [18] R. Cutler *et al.*, “ICASSP 2022 acoustic echo cancellation challenge,” in *ICASSP*, 2022, pp. 9107–9111.
- [19] R. Cutler *et al.*, “ICASSP 2023 acoustic echo cancellation challenge,” *IEEE Open Journal of Signal Processing*, 2024.
- [20] S. Graetzer *et al.*, “Clarity-2021 challenges: Machine learning challenges for advancing hearing aid processing,” in *Interspeech*, 2021, pp. 686–690.
- [21] M. A. Akeroyd *et al.*, “The 2nd clarity enhancement challenge for hearing aid speech intelligibility enhancement: Overview and outcomes,” in *ICASSP*, 2023.
- [22] E. Guizzo *et al.*, “L3DAS21 challenge: Machine learning for 3D audio signal processing,” in *Proc. MLSP*. IEEE, 2021, pp. 1–6.
- [23] E. Guizzo *et al.*, “L3DAS22 challenge: Learning 3D audio sources in a real office environment,” in *ICASSP*, 2022, pp. 9186–9190.
- [24] C. Marinoni *et al.*, “Overview of the L3DAS23 challenge on audio-visual extended reality,” in *ICASSP*, 2023.
- [25] W. Rao *et al.*, “ConferencingSpeech challenge: Towards far-field multi-channel speech enhancement for video conferencing,” in *Proc. IEEE ASRU*, 2021, pp. 679–686.
- [26] S. Leglaive *et al.*, “The CHiME-7 UDASE task: Unsupervised domain adaptation for conversational speech enhancement,” in *Proc. CHiME*, 2023.
- [27] K. Saito *et al.*, “Sampling-frequency-independent audio source separation using convolution layer based on impulse invariant method,” in *Proc. EUSIPCO*, 2021, pp. 321–325.
- [28] J. Paulus and M. Torcoli, “Sampling frequency independent dialogue separation,” in *Proc. EUSIPCO*, 2022, pp. 160–164.
- [29] J. Yu and Y. Luo, “Efficient monaural speech enhancement with universal sample rate band-split RNN,” in *ICASSP*, 2023.
- [30] J.-M. Valin *et al.*, “A perceptually-motivated approach for low-complexity, real-time enhancement of fullband speech,” in *Interspeech*, 2020, pp. 2482–2486.
- [31] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Trans. ASLP*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [32] H. Zen *et al.*, “LibriTTS: A corpus derived from LibriSpeech for text-to-speech,” in *Interspeech*, 2019, pp. 1526–1530.
- [33] R. Ardila *et al.*, “Common voice: A massively-multilingual speech corpus,” in *Proc. LREC*, 2020, pp. 4218–4222.
- [34] C. Veaux, J. Yamagishi, and S. King, “The Voice Bank corpus: Design, collection and data analysis of a large regional accent speech database,” in *Proc. O-COCOSDA/CASLRE*, 2013, pp. 1–4.
- [35] LDC, *LDC Catalog: CSR-I (WSJ0) Complete*, University of Pennsylvania, 1993.
- [36] Philadelphia: Linguistic Data Consortium, *LDC Catalog: CSR-II (WSJ1) Complete LDC94S13A*, 1994.
- [37] G. Wichern *et al.*, “WHAM!: Extending speech separation to noisy environments,” in *Interspeech*, 2019, pp. 1368–1372.
- [38] E. Bakhturina *et al.*, “Hi-Fi multi-speaker English TTS dataset,” in *Interspeech*, 2021, pp. 2776–2780.
- [39] C. K. Reddy, V. Gopal, and R. Cutler, “DNSMOS P.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors,” in *ICASSP*, 2022, pp. 886–890.
- [40] J. G. Beerends *et al.*, “Perceptual objective listening quality assessment (POLQA), the third generation ITU-T standard for end-to-end speech quality measurement part I—temporal alignment,” *AES Journal*, vol. 61, no. 6, pp. 366–384, 2013.
- [41] A. W. Rix *et al.*, “Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs,” in *ICASSP*, 2001, pp. 749–752.
- [42] J. Jensen and C. H. Taal, “An algorithm for predicting the intelligibility of speech masked by modulated noise maskers,” *IEEE/ACM Trans. ASLP*, vol. 24, no. 11, pp. 2009–2022, 2016.
- [43] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE Trans. ASLP*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [44] R. Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proc. PACRIM*, 1993, pp. 125–128.
- [45] A. Gray and J. Markel, “Distance measures for speech processing,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 5, pp. 380–391, 1976.
- [46] G. Mittag *et al.*, “NISQA: A deep CNN-self-attention model for multidimensional speech quality prediction with crowdsourced datasets,” in *Interspeech*, 2021, pp. 2127–2131.
- [47] J. Pirklbauer *et al.*, “Evaluation metrics for generative speech enhancement methods: Issues and perspectives,” in *Speech Communication; 15th ITG Conference*, 2023, pp. 265–269.
- [48] T. Saeki *et al.*, “SpeechBERTScore: Reference-aware automatic evaluation of speech generation leveraging NLP evaluation metrics,” *arXiv preprint arXiv:2401.16812*, 2024.
- [49] J.-w. Jung *et al.*, “Pushing the limits of raw waveform speaker recognition,” in *Interspeech*, 2022, pp. 2228–2232.
- [50] Y. Peng *et al.*, “OWSM v3. 1: Better and faster open Whisper-style speech models based on E-Branchformer,” *arXiv preprint arXiv:2401.16658*, 2024.
- [51] I. Cohen and B. Berdugo, “Speech enhancement for non-stationary noise environments,” *Signal Processing*, vol. 81, no. 11, pp. 2403–2418, 2001.
- [52] Y.-J. Lu *et al.*, “Towards low-distortion multi-channel speech enhancement: The ESPnet-SE submission to the L3DAS22 challenge,” in *ICASSP*, 2022, pp. 9201–9205.
- [53] C. Li *et al.*, “ESPnet-SE: End-to-end speech enhancement and separation toolkit designed for ASR integration,” in *Proc. IEEE SLT*, 2021, pp. 785–792.