



Universal Score-based Speech Enhancement with High Content Preservation

Robin Scheibler, Yusuke Fujita, Yuma Shirahata, Tatsuya Komatsu

LY Corporation, Japan

robin.scheibler@ieee.org

Abstract

We propose UNIVERSE++, a universal speech enhancement method based on score-based diffusion and adversarial training. Specifically, we improve the existing UNIVERSE model that decouples clean speech feature extraction and diffusion. Our contributions are three-fold. First, we make several modifications to the network architecture, improving training stability and final performance. Second, we introduce an adversarial loss to promote learning high quality speech features. Third, we propose a low-rank adaptation scheme with a phoneme fidelity loss to improve content preservation in the enhanced speech. In the experiments, we train a universal enhancement model on a large scale dataset of speech degraded by noise, reverberation, and various distortions. The results on multiple public benchmark datasets demonstrate that UNIVERSE++ compares favorably to both discriminative and generative baselines for a wide range of qualitative and intelligibility metrics.

Index Terms: universal, speech enhancement, score-based generative model, adversarial training

1. Introduction

Speech enhancement (SE) is the task of restoring clean speech from a degraded signal [1]. Many SE methods operate on the magnitude spectrogram [2] with a mask estimated by statistical methods [3] or, recently, deep neural networks (DNN) [4, 5, 6]. See [7] for an overview of current methods. DNN-based SE can be divided into time-domain [8] and time-frequency methods [9]. Generally, the model is discriminatively trained to regress to the clean speech from the degraded signal. However, this type of training is known to be prone to residual noise and artifacts [10]. An alternative approach is to use generative models, derived from text-to-speech technology, in the hope of generating high quality speech without residual noise. Generative adversarial networks (GAN) [11] are suitable for SE. The denoising network, playing the role of generator, is jointly trained with a discriminator that learns to distinguish between clean and enhanced speech [12, 13]. Score-based diffusion [14, 15] is another generative paradigm that has recently gained traction for SE [16, 17, 18, 19]. Diffusion-based SE models operate by progressively transforming normally distributed noise into clean speech conditioned on the input degraded speech.

This work focuses on universal speech enhancement (USE, also known as speech restoration), the extension of SE to all types of signal degradation, including reverberation, low-pass filtering, clipping, etc [19, 20, 21]. It is believed that generative models are more suitable for USE as the model output is not uniquely determined by the input. For example, there are many plausible reconstructions for a low-pass signal. UNIVERSE [19] is a diffusion-based method for USE that attracted

attention due to a high profile demonstration featuring very high quality enhanced speech. However, during preliminary experiments, we found it difficult to train, and prone to hallucinations such as mumbled speech. We address these issues in the UNIVERSE++ model with a number of improvements over the original. First, we introduce several architectural upgrades such as normalization and anti-aliasing filters in the down/up-sampling layers of the network. Second, we combine the score matching training with the adversarial loss of HiFi-GAN [22] that promotes the extraction of high-quality features of the clean speech to condition the diffusion process. Finally, we propose a lightweight low-rank adaptation (LoRA) [23] procedure to improve the retention of linguistic content in the enhanced speech. Extensive experiments on multiple benchmarks demonstrate the versatility and effectiveness of the proposed method. In particular, we show that it produces highly natural speech without compromising on the intelligibility of the content. Our implementation of UNIVERSE++ is available as open-source¹.

2. Background

The task of *universal* speech enhancement (USE) consists in recovering a clean speech signal \mathbf{x} that is only available through a degraded signal \mathbf{y} . Unlike conventional SE, the degradation is not limited to additive noise, but may include reverberation, clipping, filtering, coding artifacts, etc. Throughout the paper, bold upper and lower case letters are for matrices and vectors, respectively. The Euclidean norm of $\mathbf{v} \in \mathbb{R}^n$ is $\|\mathbf{v}\| = (\mathbf{v}^\top \mathbf{v})^{1/2}$. UNIVERSE has been proposed to tackle this task [19]. It is composed of two networks, as shown in Fig. 1. A conditioning network $C(\cdot)$ extracts clean speech features from the degraded speech. Then, these features are used to condition a diffusion process synthesizing the clean speech directly. The second network is the score denoising model $S(\cdot)$ used as part of the diffusion process. This division of labor in the network makes it easy to add losses to the conditioning network to control the extracted features.

UNIVERSE follows the stochastic differential equation (SDE) formulation of the diffusion models [15]. This framework defines a continuous time process transforming the clean speech \mathbf{x}_0 into normally distributed noise until the original signal is imperceptible. According to theory of SDE, under mild conditions, the reverse process exists and depends only on the so-called *score function*, the derivative of the logarithm of the marginal distribution at time t of the process. Although unknown in practice, a neural network can be trained to approximate the score function by minimizing the following objective,

$$\mathcal{L}_{\text{score}} = \mathbb{E}_{t, \mathbf{z}, \mathbf{x}} [\|\sigma_t S(\mathbf{x} + \sigma_t \mathbf{z}, \sigma_t, \mathbf{c}) + \mathbf{z}\|^2], \quad (1)$$

¹<https://github.com/line/open-universe>




Figure 1: Overview of the UNIVERSE network architecture with proposed (purple box) and original (yellow box) loss functions.

where σ_t^2 is the noise variance at time $t \in [0, 1]$, and $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Furthermore, $\mathbf{c} = C(\mathbf{y})$ is the feature vector produced by the conditioning network. The noise variance follows an exponential schedule $\sigma_t^2 = \sigma_{\min}^2 (\sigma_{\max}^2 / \sigma_{\min}^2)^t$, where σ_{\min} and σ_{\max} are hyperparameters. At inference time, the estimated clean speech is produced by solving the reverse diffusion process with the noise-consistent Langevin dynamics [24],

$$\mathbf{x}_{t-\Delta} = \mathbf{x}_t + \eta \sigma_t^2 S(\mathbf{x}_t, \mathbf{c}, \sigma_t) + \beta \sigma_{t-\Delta} \mathbf{z}. \quad (2)$$

For N steps, $\Delta = 1/N$. The parameters $\eta = 1 - \gamma^\epsilon$ and $\beta = \sqrt{1 - \gamma^{2(\epsilon-1)}}$ with $\gamma = \sigma_{\max} / \sigma_{\min}$, and ϵ is a hyperparameter.

Model Architecture The two networks share a similar encoder/decoder structure, each with four stages, that gradually reduce/increase the sampling rate and increase/reduce the number of channels. Each stage is composed of three convolutional blocks, one of which changes the sampling rate. The convolutional blocks are interleaved with PReLU activation functions. Gated recurrent units are used to model the temporal dependencies in the bottleneck of the network with two and one layers in the conditioning and score networks, respectively. The intermediate features produced at each stage of the decoder in the conditioning network are injected in the corresponding stage of the score network via linear projection layers. In the score network, residual connections are used between corresponding stages of the encoder and decoder. In the conditioning networks, residual connections connect each stage of the encoder to the bottleneck features via a convolutional adaptation layer. The noise variance of the diffusion process is injected at all stages of the score network via a random Fourier embedding and FiLM layers. See Fig. 1 for an overview, and [19] for the full details.

Loss Functions and Training In addition to the score matching loss (1), UNIVERSE uses a number of mixture density network (MDN) losses [25] to promote the learning of good features. One MDN matches the bottleneck features of C to the clean speech mel-spectrogram. Another, connected to the last layer of the decoder, matches the output features to the clean speech waveform. In the original UNIVERSE, a number of other losses with respect to speech features such as pitch and voice activity and other speech features are also used. However, we did not consider them as they were not shown to improve the metrics in [19]. The model was trained on a large private dataset of speech and noise coupled to dynamic distortions [19].

3. UNIVERSE++

Our modifications to the UNIVERSE methodology are three-fold. First, we introduce improvements in the network to follow best practices for diffusion and synthesis. Second, we intro-




Figure 2: Proposed anti-aliasing filters at down/up-sampling stages of the score network UNet.

duce adversarial training to improve the quality of the extracted features. Finally, we propose a fine-tuning procedure directly targeting the linguistic content.

3.1. Proposed Network Improvements

Normalization Due to the exponential noise schedule in the diffusion process, the score network has to deal with inputs and outputs with widely different scales spanning up to four orders of magnitude. This issue is tackled by Karras et al. [26] who propose to re-parameterize the score network as

$$S(\mathbf{x}, \mathbf{c}, \sigma_t) = c_t^{\text{skip}} \mathbf{x} + c_t^{\text{out}} S'(c_t^{\text{in}} \mathbf{x}, \mathbf{c}, \sigma_t), \quad (3)$$

where the weights depending on the variance at time t ,

$$c_t^{\text{skip}} = \frac{\sigma_{\text{data}}^2}{\sigma_{\text{data}}^2 + \sigma_t^2}, \quad c_t^{\text{out}} = \sigma_t \sqrt{c_t^{\text{skip}}}, \quad c_t^{\text{in}} = 1 / \sqrt{\sigma_{\text{data}}^2 + \sigma_t^2},$$

are such that the network inputs and training targets have unit variance. The σ_{data} is the variance of the clean speech data.

Anti-aliasing In the score network UNet, down- and up-sampling stages introduce aliasing artifacts which have been shown to be detrimental to translation invariance in image generation networks [27]. Following the practice of image generation [27], we introduce anti-aliasing filters before any change of sampling rate. As shown in Fig. 2, the high-frequency content is propagated only via the residual connection and its processing is clearly restricted to the upper stages of the network. We did not use anti-aliasing in the conditioning network, that lacks residual connections, as we observed a performance degradation.

Miscellaneous We use the weight normalization parameterization [28] for the convolutional and linear layers. The original embeddings for σ_t use random Fourier features with a three layer multi-layer perceptron (MLP). Since the embeddings are followed by yet another projection layer, the MLP seemed redundant and we use non-random Fourier embedding instead

$$\mathbf{e}(\sigma_t) = [\dots \cos(2\pi f(\sigma_t)m) \dots \sin(2\pi f(\sigma_t)m) \dots]^\top.$$

where $f(\sigma_t) = \alpha \log(\sigma_t) + \beta$, α, β are trainable.

3.2. HiFi-GAN Adversarial Loss

Despite its probabilistic formulation, the original MDN loss is still applied sample-wise to the target speech in a discriminative fashion. Instead, we replace it by the adversarial loss of HiFi-GAN [22] (purple box in Fig. 1). It uses two discriminators, multi-period and multi-resolution, trained with the least-squares GAN loss [29]. A mel-spectrogram loss between the output of the adapter and the clean speech improves the training efficiency and the fidelity of the features produced. Finally, feature matching losses between the discriminators' feature maps of a clean sample and the output of the conditioning network allow to capture high-level features of the clean speech content.




Figure 3: Results on the Voicebank-DEMAND dataset at 16 kHz. The values for HiFi++ and SGMSE+M are those reported in [13] and [35], respectively.

3.3. Linguistic Content Enhancing Low-rank Adaptation

As reported in [10, 30], we noticed that while the generatively enhanced speech sounds very natural, hallucinations tend to appear in low signal-to-noise ratio (SNR) segments. We address this issue with a fine-tuning procedure to improve the retention of linguistic content. We select LoRA for the fine-tuning [23]. Due to its low-memory footprint, we can perform the fine-tuning on the inference stage of the diffusion process. We back-propagate through the last two steps, similarly to [31], and the conditioning network. We use a phoneme prediction model [32], which is language agnostic, and avoids requiring transcripts of the training data, to predict the phonemes of the clean and enhanced speech. A connectionist temporal classification (CTC) loss [33] is used between the predicted phonemes of the clean speech and those of the enhanced speech. We note that phoneme losses have been used with text labels for speech synthesis [34]. The speech integrity is preserved by the frozen HiFi-GAN losses, and a multi-resolution spectrogram loss.

4. Experiments

First, we evaluate UNIVERSE++ for additive noise reduction on the Voicebank+DEMAND (VB) dataset [36], enabling comparison to a large number of SE methods. We also perform an ablation study to evaluate the gains due to architecture and loss function. Our second experiment studies the performance of UNIVERSE++ for USE on several public benchmark datasets.

4.1. Evaluation Metrics

Generative methods in SE require the use of a variety of metrics [10]. The waveform reconstruction is measured by PESQ [37] and log-spectral distance (LSD), and intelligibility by extended STOI [38]. To evaluate the linguistic content preservation of the algorithms, we use the Levenshtein phoneme similarity (LPS), i.e., the phoneme accuracy of the enhanced versus clean speech according to a phoneme predictor [10], and the word error rate (WER) of the Whisper large-v3 speech recognition model [39]. Naturalness is measured with the non-intrusive neural mean opinion score DNSMOS [40].

4.2. Noise Reduction on Voicebank+DEMAND

Dataset The Voicebank+DEMAND dataset is originally split into train and test [36]. Following standard practice, we form a validation set with speakers p226 and p287 of the training set. This results in 8.75 h for training, 38 min for validation, and 34 min for testing. The sampling frequency is 16 kHz.

Training We train the model with the AdamW optimizer [41]. The learning rate starts at 10^{-6} and linearly increases to 10^{-4}

over the first 10000 steps. A cosine schedule reduces the learning rate to 10^{-6} from steps 200000 to 300000 when the training stops. The batch size is 40. An exponential moving average of the weights with forgetting factor 0.999 is used for inference.

Baselines HiFi++ is a generative method enhancing the mel-spectrogram and using a HiFi-GAN neural vocoder to reconstruct the clean speech from it [13]. SGMSE+M is a generative speech enhancement method using diffusion based on the Ornstein-Uhlenbeck SDE [35]. StoRM is a hybrid method that uses a jointly, but discriminatively, trained network to initialize the diffusion process of SGMSE+M [35].

Results Fig. 3 shows that omitting from UNIVERSE++ the adversarial loss (-HiFi-GAN) and further removing the network improvements (-Network) reduce PESQ by 0.05 and 0.07 points, respectively. UNIVERSE++ outperforms the baseline methods for PESQ and DNSMOS, while being less than 0.01 point STOI behind StoRM.

4.3. Universal Speech Enhancement

4.3.1. Datasets

Training We train our universal enhancement model at 24 kHz on a large scale dataset of speech degraded by noise, reverberation, and various distortions applied dynamically during training. The clean speech dataset is composed of the following sources. The training set of the URGENT challenge [42], which was itself created from the English read speech part of the deep noise suppression (DNS) challenge [43], the Common Voice English v11 [44], LibriTTS [45], and VCTK [46]. The dataset was filtered to remove samples originally collected at less than 24 kHz or of low quality, as evaluated using the DNSMOS metric [40], resulting in 300 h of high quality speech. See [42] for the full details. In addition, we use an internal dataset of 237 h of studio quality Japanese speech from 17 speakers. This totals 537 h of clean speech. The noise dataset consists of 177 h from Audio Set [47] (from the DNS challenge training set), 78 h from the Wham! dataset [48], and 346 h of an internal background music dataset. The noise is mixed with an SNR uniformly distributed between -5 dB to 30 dB, except for the background music for which it is between -10 dB to 5 dB. We dynamically apply one of the 132037 real room impulse responses of the Arni [49] dataset, or of 100000 synthesized using Pyroomacoustics [50]. We also apply band limitation, equalization distortion, clipping (clamp, tanh, sigmoid), random attenuation, packet loss, or codec distortion (MP3). One to five types of distortions are applied at random following approximately the methodology of UNIVERSE [19]. The two internal datasets were necessary for our target downstream task.

We use the following test sets, all resampled at 24 kHz.

Voicebank+DEMAND (VB) This is the test split of the VB dataset described in Section 4.2 but sampled at 24 kHz.

VB + Bandwidth Extension (VB-BWE) Same as above, but we apply in addition a low-pass filter with cut-off at 4 kHz to test the bandwidth extension capability of the model.

Packet Loss Concealment (PLC) Validation set of the ICASSP 2024 Audio Deep Packet Loss Concealment Challenge [51]. The dataset contains 800 samples of speech with burst of packet loss of length up to 3 s realistically distributed.

Signal Improvement Challenge (SIG) ICASSP 2023 Speech Signal Improvement Challenge test set [52]. The dataset contains 500 samples of real-world recorded distorted speech in five different languages. It includes various distortions such as noise (background, circuit, or coding), reverberation, clipping, etc.

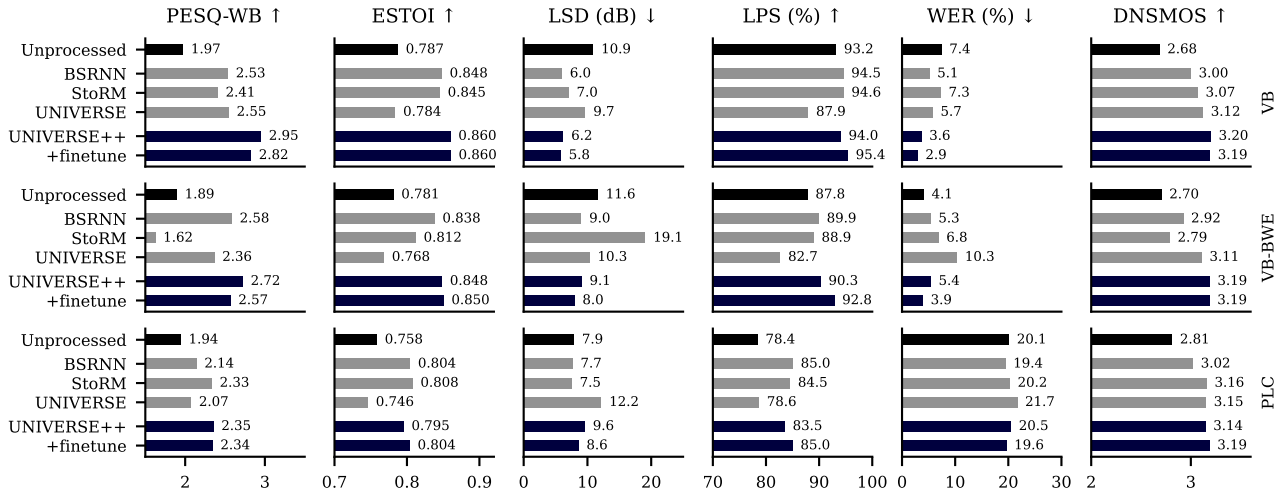


Figure 4: Results on the Voicebank+DEMAND (VB) dataset, VB low-pass filtered at 4 kHz (VB-BWE), and packet loss concealment challenge validation set (PLC). Arrows indicate if the metric is better when increasing (\uparrow) or decreasing (\downarrow).




Figure 5: Results on the Signal Improvement Challenge non-blind test set, for which reference clean speech is not available.

4.4. Model Details and Training

The rate change factors in the four stages of the encoder/decoder are set to 2, 3, 5, 8 so that the rate at the bottleneck layer is 10 Hz as in [19]. The initial number of channels is 48. The model has 107.5 M parameters and is trained for 1500000 steps with the AdamW optimizer [41]. The HiFi-GAN losses add 41.4 M parameters during training. The learning rate linearly goes from 10^{-6} to 10^{-4} over the first 50000 steps, then back to 10^{-6} with a cosine schedule over the last 500000. The batch size is 40. The training samples are truncated to 2s. During fine-tuning, the rank of the adaptation matrix is 16. Matrices with either dimension smaller are not updated. The learning rate is 10^{-5} and the samples are truncated to 4s. The batch size is 16. The number of diffusion steps for inference is 8 and $\epsilon = 1.3$.

4.5. Baselines

Band-split recurrent neural network (BSRNN) is a strong baseline for discriminative SE [9]. We train a model with 128 channels and 6 layers. The STFT uses window size and shift of 20 ms and 5 ms, respectively. We split the spectrum into 10, 12, 8, and 4 bands in the range of 0 kHz to 1 kHz, 1 kHz to 4 kHz, 4 kHz to 8 kHz, and 8 kHz to 12 kHz, respectively. The number of parameters is 11.5 M. We trained for 750000 steps with batch size 48 and the loss of [9]. The checkpoint with highest validation PESQ was selected.

StoRM is trained using the code of [35] modified to work with the dataset of Section 4.3.1. The model has 39.2 M trainable parameters. The model is trained for 1500000 steps with a batch size of 32. Inference is done with 50 diffusion steps.

UNIVERSE is our re-implementation of [19] without the proposed modifications of Section 3 (107.7 M parameters and another 7.6 M for the MDN losses during training).

4.6. Results

The results of evaluation on the tests sets are shown in Fig. 4 and 5. The purely discriminatively trained BSRNN has high content retention demonstrated by low LSD, low WER, and high LPS on all datasets. However, it is markedly worse in terms of naturalness, demonstrated by the low DNSMOS (OVRL) values. StoRM exhibits attributes of both generative and discriminative methods with high naturalness, as well as good overall content preservation. However, its performance is not consistent over all datasets, e.g., it does not do well on the bandwidth extension task. The UNIVERSE baseline has higher DNSMOS than BSRNN, but is overall worse than all other methods. High LSD and WER, and low STOI and LSD suggest it may change content or drop segments, which is confirmed by informal listening. The improved UNIVERSE++ is the most versatile in that it performs well with respect to all metrics and on all datasets. It does not damage the linguistic content, as shown by low WER at the same level, or below, that of the unprocessed speech. Fine-tuning further improves the WER and LPS, at the expense of some PESQ. Anecdotally, UNIVERSE(++) is much faster than StoRM, as it only requires 8 diffusion steps compared to 50 for StoRM.

5. Conclusion

We proposed UNIVERSE++, a universal speech enhancement method using score-based diffusion and adversarial training. Experiments demonstrate that the proposed model improves over the original UNIVERSE and also outperforms conventional methods on several test sets covering a wide range of speech distortions. The adversarial loss significantly improves both the naturalness of the enhanced speech and the linguistic content preservation. Together with the fine-tuning procedure, the method achieves the same performance as discriminative methods in terms of content preservation, but produces significantly more natural enhanced speech. In future work, we will explore the use of phoneme loss during the main training stage.

6. References

- [1] P. C. Loizou, *Speech Enhancement: Theory and Practice, Second Edition*, 2nd ed. CRC Press, 2017.
- [2] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 27, no. 2, pp. 113–120, 1979.
- [3] Y. Ephraim and D. Malah, "Speech enhancement using a minimum-mean square error short-time spectral amplitude estimator," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 32, no. 6, pp. 1109–1121, 1984.
- [4] S. Tamura and A. Waibel, "Noise reduction using connectionist models," in *ICASSP*, 1988, pp. 553–556.
- [5] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, "Speech enhancement based on deep denoising autoencoder," in *INTERSPEECH*, 2013, pp. 436–440.
- [6] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 23, no. 1, pp. 7–19, Jan. 2015.
- [7] Y.-J. Lu *et al.*, "ESPnet-SE++: Speech enhancement for robust speech recognition, translation, and understanding," in *INTERSPEECH*, Sep. 2022, pp. 5458–5462.
- [8] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, Aug. 2019.
- [9] J. Yu *et al.*, "High fidelity speech enhancement with band-split RNN," in *INTERSPEECH*, 2023, pp. 2483–2487.
- [10] J. Pirklbauer, M. Sach, and K. Fluyt, "Evaluation metrics for generative speech enhancement methods: Issues and perspectives," in *Proc. of Speech Commun. - 15th ITG Conf.*, Sep. 2023.
- [11] I. Goodfellow *et al.*, "Generative adversarial nets," in *NIPS*, 2014.
- [12] S. Pascual, A. Bonafonte, and J. Serrà, "SEGAN - speech enhancement generative adversarial network," in *INTERSPEECH*, Jan. 2017, pp. 3642–3646.
- [13] P. Andreev, A. Alanov, O. Ivanov, and D. Vetrov, "HiFi++: A Unified Framework for Bandwidth Extension and Speech Enhancement," in *ICASSP*, 2023, pp. 1–5.
- [14] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *NeurIPS*, vol. 33, 2020, pp. 6840–6851.
- [15] Y. Song *et al.*, "Score-based generative modeling through stochastic differential equations," in *ICLR*, May 2021.
- [16] Y.-J. Lu *et al.*, "Conditional diffusion probabilistic model for speech enhancement," in *ICASSP*, May 2022, pp. 7402–7406.
- [17] J. Richter *et al.*, "Speech Enhancement and Dereverberation with Diffusion-based Generative Models," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 31, pp. 2351–2364, 2023.
- [18] H. Yen, F. G. Germain, G. Wichern, and J. L. Roux, "Cold diffusion for speech enhancement," in *ICASSP*, Jun. 2023, pp. 1–5.
- [19] J. Serrà *et al.*, "Universal speech enhancement with score-based diffusion," *arXiv:2206.03065*, Jun. 2022.
- [20] J. Zhang, S. Jayasuriya, and V. Berisha, "Restoring degraded speech via a modified diffusion model," in *INTERSPEECH*, Apr. 2021, pp. 221–225.
- [21] J. Byun *et al.*, "An empirical study on speech restoration guided by self-supervised speech representation," in *ICASSP*, Jun. 2023, pp. 1–5.
- [22] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS*, 2020, pp. 17 022–17 033.
- [23] E. J. Hu *et al.*, "LoRA: Low-rank adaptation of large language models," in *ICLR*, Jan. 2022.
- [24] A. Jolicœur-Martineau, R. Piché-Taillefer, I. Mitliagkas, and R. T. des Combes, "Adversarial score matching and improved sampling for image generation," in *ICLR*, May 2021.
- [25] C. M. Bishop, "Mixture Density Networks," Aston University, UK, Tech. Rep., 1994.
- [26] T. Karras, M. Aittala, T. Aila, and S. Laine, "Elucidating the Design Space of Diffusion-Based Generative Models," in *NeurIPS*, 2022.
- [27] T. Karras *et al.*, "Alias-Free Generative Adversarial Networks," in *NeurIPS*, 2021.
- [28] T. Salimans and D. P. Kingma, "Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks," in *NIPS*, 2016.
- [29] X. Mao *et al.*, "Least Squares Generative Adversarial Networks," in *ICCV*, 2017, pp. 2794–2802.
- [30] R. Scheibler *et al.*, "Diffusion-based generative speech source separation," in *ICASSP*, Jun. 2023, pp. 1–5.
- [31] B. Lay, S. Welker, J. Richter, and T. Gerkmann, "Reducing the Prior Mismatch of Stochastic Differential Equations for Diffusion-based Speech Enhancement," in *INTERSPEECH*, 2023, pp. 3809–3813.
- [32] Q. Xu, A. Baevski, and M. Auli, "Simple and Effective Zero-shot Cross-lingual Phoneme Recognition," in *INTERSPEECH*, 2022, pp. 2113–2117.
- [33] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *ICML*, Jun. 2006, pp. 369–376.
- [34] Y. Zhang *et al.*, "VISinger: Variational inference with adversarial learning for end-to-end singing voice synthesis," in *ICASSP*, Singapore, Singapore, May 2022, pp. 7237–7241.
- [35] J.-M. Lemerrier, J. Richter, S. Welker, and T. Gerkmann, "StoRM: A diffusion-based stochastic regeneration model for speech enhancement and dereverberation," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 31, pp. 2724–2737, 2023.
- [36] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks," in *INTERSPEECH*, Sep. 2016, pp. 352–356.
- [37] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ)," in *ICASSP*, Jan. 2001, pp. 749–752.
- [38] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 24, no. 11, pp. 2009–2022, Nov. 2016.
- [39] A. Radford *et al.*, "Robust speech recognition via large-scale weak supervision," in *ICML*, Jul. 2023, pp. 28 492–28 518.
- [40] C. K. A. 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.
- [41] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, Sep. 2018.
- [42] W. Zhang *et al.*, "URGENT challenge: Universality, robustness, and generalizability for speech enhancement," in *INTERSPEECH*, Sep. 2024.
- [43] H. Dubey *et al.*, "ICASSP 2022 deep noise suppression challenge," *arXiv:2202.13288*, Feb. 2022.
- [44] R. Ardila *et al.*, "Common Voice: A massively-multilingual speech corpus," in *Proc. of the 12th Lang. Resour. Evaluation Conf.*, May 2020, pp. 4218–4222.
- [45] H. Zen *et al.*, "LibriTTS: A corpus derived from LibriSpeech for text-to-speech," in *INTERSPEECH*, 2019, pp. 1526–1530.
- [46] 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*, Nov. 2013, pp. 1–4.
- [47] J. F. Gemmeke *et al.*, "Audio Set: An ontology and human-labeled dataset for audio events," in *ICASSP*, Mar. 2017, pp. 776–780.
- [48] G. Wichern *et al.*, "WHAM!: Extending speech separation to noisy environments," in *INTERSPEECH*, Sep. 2019.
- [49] P. Karolina, S. J. Schlecht, and V. Välimäki, "Dataset of impulse responses from variable acoustics room Arni at Aalto Acoustic Labs," Aug. 2022, <https://doi.org/10.5281/zenodo.6985104>.
- [50] R. Scheibler, E. Bezzam, and I. Dokmanić, "Pyroomacoustics: A Python package for audio room simulations and array processing algorithms," in *ICASSP*, 2018.
- [51] L. Diener, S. Branets, A. Saabas, and R. Cutler, "The ICASSP 2024 audio deep packet loss concealment grand challenge," *arXiv:2402.16927*, 2024.
- [52] R. Cutler, A. Saabas, B. Naderi, and S. Branets, "ICASSP 2023 speech signal improvement challenge," *arXiv:2303.06566*, 2023.