

DISENTANGLING SPEECH FROM SURROUNDINGS WITH NEURAL EMBEDDINGS

Ahmed Omran, Neil Zeghidour, Zalán Borsos,
Félix de Chaumont Quitry, Malcolm Slaney, Marco Tagliasacchi

Google Research

ABSTRACT

We present a method to separate speech signals from noisy environments in the embedding space of a neural audio codec. We introduce a new training procedure that allows our model to produce structured encodings of audio waveforms given by embedding vectors, where one part of the embedding vector represents the speech signal, and the rest represent the environment. We achieve this by partitioning the embeddings of different input waveforms and training the model to faithfully reconstruct audio from mixed partitions, thereby ensuring each partition encodes a separate audio attribute. As use cases, we demonstrate the separation of speech from background noise or from reverberation characteristics. Our method also allows for targeted adjustments of the audio output characteristics.

Index Terms— Disentangled representations, audio compression, speech enhancement

1. INTRODUCTION

Recent progress in generative models of audio have allowed for the development of neural codecs [1, 2, 3, 4] that are competitive with traditional codecs. However, while traditional codecs such as Opus [5] cascade blocks with specific roles (e.g., voice activity detection, noise shaping), neural codecs typically stack learned representations that are not easily interpretable. This can be detrimental to downstream tasks, such as audio event detection, and selective transmission of parts of the signal — such tasks can potentially benefit from disentangled representations in the compressed domain, where different attributes of the signal are represented separately.

In this work, we train a neural audio codec based on *SoundStream* [4] to represent audio signals in a structured way in the compressed domain. The neural network maps input audio into an embedding space, which we split into partitions that are each to capture a different attribute of the input audio. We then use targeted augmentations of the data as well as a custom loss function to introduce a strong inductive bias in the model to specifically allocate each partition to a given attribute. The compressed representations can then be tailored to make best use of the available bit rate, by prioritizing certain types of content over others, or only transmitting the desired part of the signal, such as the clean speech content.

We demonstrate two different types of disentangling tasks: The first is separating speech from background noise. The noise signal is an additive component that varies on similar time scales as speech changes. But in contrast to conventional source separation and speech enhancement [6, 7, 8], the separation happens in the embedding space. The second is disentangling speech from the reverberation characteristics of the recording environment. Here, the reverberation signal cannot be simply subtracted from the original signal in the time domain, and is instead modeled as a convolution with a time-invariant room impulse response that spans an extended period of time. An

important property we exploit in our method is that the reverberation characteristics typically change more slowly than the speech signal. For both cases, we demonstrate the separation of speech from the environment in the embedding space, and how this disentanglement enables fine-grained control over the synthesized output audio.

A standard approach for disentangling attributes relies on learning a factorized latent space of a generative model [9], e.g., in the bottleneck of an autoencoder. Several works have used an autoencoder with a classifier that tries to predict an attribute from the embeddings, while the encoder is trained to fool the classifier by generating embeddings that are invariant under a change in this attribute. This has been applied to disentanglement of image features [10], sensor anonymization [11], voice conversion [12, 13] and music translation [14]. This idea has also been proposed for disentangling multiple speech attributes by chaining several autoencoders trained in this fashion [15]. However, this adversarial training adds complexity and potential instability, and the necessary labels are difficult to obtain in the audio domain. Here, we instead rely on a suitable choice of input and target data and training objectives to get sufficient inductive biases, which are necessary for disentanglement [16].

To avoid issues of adversarial training, some works rely on constraining the information capacity by a careful choice of embedding dimensions that ensures each channel contains only one factor of variation, such as AutoVC [17] and SpeechSplit [18]. However this comes with a trade-off in the generated output quality, and does not provide flexibility in choosing the embedding dimension for a signal attribute of interest, whereas our method provides more freedom in choosing the embedding dimensions.

Polyak et al. [19] used self-supervised learning to train multiple encoders to obtain discrete representations of speech inputs, each one capturing a different speech attribute. These representations are resynthesized to high-quality speech, enabling very low-bitrate speech transmission. This is only limited to speech characteristics and relies on domain-specific encoders, whereas our codec can process a broader audio domain. Yang et al. [20] demonstrated source separation and neural coding in the latent space, an approach that is only limited to signals that can be added together in the time domain.

Disentangled representations can be learned by partitioning feature vectors in the latent space and mixing the partitions while training the model to produce realistic decoded outputs, thereby capturing one factor of variation in each partition [21] or to synthesize new samples that share attributes between different input samples [22]. A related approach is to supply input data that differs in some attribute and constrain different subspaces of the latent vectors to capture those changes, as was shown in the vision domain with dataset batching [23] or targeted augmentations [24], or for the SPICE audio pitch estimator [25]. Inspired by these mixing approaches, we present how to separate speech signals from their environment, and demonstrate disentanglement of speech from background noise or reverberation in a way that enables fine-grained control over the output.

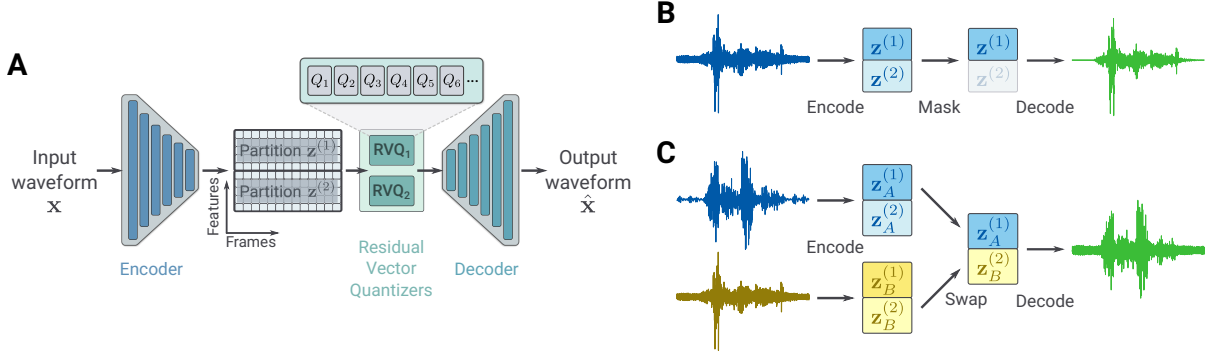


Fig. 1. The architecture of our model. **A)** An input waveform is processed by the encoder to produce one embedding vector per frame. Each partition is processed using a separate residual vector quantizer. The quantized embeddings are converted back into an audio waveform. **B)** The embedding partitions can be individually set to zero before decoding, which removes the corresponding attribute from the output audio, e.g., the background noise. **C)** An audio waveform can be encoded and an embedding partition swapped out with that of a different waveform, thereby replacing an audio attribute of the first waveform with that of the second.

2. METHODS

Our model is based on *SoundStream* [4], a streamable neural audio codec with a vector-quantized autoencoder architecture. However, unlike the baseline *SoundStream* model, we have the following additional requirements: We need to set the dimension allocated to each embedding partition, and for compression we need a way of setting the effective bitrate for each channel. The codec should allow for streaming audio while producing structured embeddings in real time, such that certain types of content can have their bitrate reduced or be removed entirely on demand.

2.1. Model

Our model architecture is shown in Fig. 1A. Single-channel audio waveforms $\mathbf{x} \in \mathbb{R}^T$ are processed by the encoder, which is a convolutional network with blocks consisting of residual units and a strided convolutional layer. All convolutions have causal padding to allow for real-time streaming of the input waveforms. The encoder produces embeddings $\mathbf{z} = \text{enc}(\mathbf{x}) \in \mathbb{R}^{F \times D}$, where F is the number of frames, and D is the number of features per frame, set by the number of convolutional filters at the bottleneck. The number of frames F is given by the number of waveform samples T divided by the product of all convolution striding factors.

We apply residual vector quantization to each embedding partition separately. A residual vector quantizer [4] is a stack of vector quantization layers, each of which replaces its input by a vector from a learned discrete codebook and passes the residual error to the next layer. The effective bitrate is given by $R = R_f N n_q$, where R_f is the frame rate, N is the codebook size in bits, and n_q is the number of quantizer layers. Each embedding partition, a subset of the D total dimensions, is processed by such a quantization module, allowing us to set a separate bitrate for each partition, particularly when different attributes change over different time scales. For a finite bitrate budget, we can also prioritize speech content over other attributes like the room characteristics.

The quantized embeddings are concatenated together along the feature axis and fed into the decoder, which converts them into a reconstructed audio waveform $\hat{\mathbf{x}} = \text{dec}(\mathbf{z}) \in \mathbb{R}^T$. The decoder is very similar in structure to the encoder, where convolutions are replaced by transposed convolutions [4].

2.2. Disentanglement scheme

Disentangling speech from its environment requires careful design of the training objectives, in order to both achieve high-quality audio transmission and create structured embeddings with no content leakage between different partitions.

We apply an augmentation $f(\cdot, \theta)$ to a clean speech waveform \mathbf{x}^c to produce $\mathbf{x} = f(\mathbf{x}^c, \theta)$. Here we consider parameters θ that are used as additive signals to the clean speech $\mathbf{x} = \mathbf{x}^c + \theta$, or as convolution kernels $\mathbf{x}(t) = \int \theta(\tau) \mathbf{x}^c(t - \tau) d\tau$. The aim is to produce embeddings where \mathbf{x}^c is encoded in a dedicated partition $\mathbf{z}^{(1)}$, and θ is encoded in a separate partition $\mathbf{z}^{(2)}$, such that $\mathbf{z} = \text{concat}(\mathbf{z}^{(1)}, \mathbf{z}^{(2)})$. Additionally, we require a good reconstruction quality of the audio and speech signal of interest. Therefore, during each training iteration we create three different types of augmented signals, which are used sequentially as input to the training objective in Section 2.3:

- (i) Reconstruct the augmented input audio \mathbf{x} using the full embedding vectors \mathbf{z} .
- (ii) Reconstruct the clean speech component \mathbf{x}^c from the augmented input \mathbf{x} . Here, the embeddings are multiplied by a binary mask to set the second partition $\mathbf{z}^{(2)}$ that encodes the augmentation to zero (Fig. 1B).
- (iii) Encode two different augmented input waveforms $\mathbf{x}_A = f(\mathbf{x}_A^c, \theta_A)$ and $\mathbf{x}_B = f(\mathbf{x}_B^c, \theta_B)$, where θ_A and θ_B represent different augmentations. This produces embeddings \mathbf{z}_A and \mathbf{z}_B , which are partitioned as $[\mathbf{z}_A^{(1)}, \mathbf{z}_A^{(2)}]$ and $[\mathbf{z}_B^{(1)}, \mathbf{z}_B^{(2)}]$, respectively. Then, the second partition of the embedding vectors is swapped, and the new embedding vectors $[\mathbf{z}_A^{(1)}, \mathbf{z}_B^{(2)}]$ are decoded to reconstruct a new version $\hat{\mathbf{x}}_A = f(\mathbf{x}_A^c, \theta_B)$ that inherits the attribute of interest from \mathbf{x}_B (Fig. 1C).

2.3. Training objectives

To train the model, we minimize a multi-scale spectral reconstruction loss for each step [26, 27]:

$$\mathcal{L}_{\text{rec}}(\mathbf{x}, \mathbf{x}') = \sum_{\log_2 s=6}^{11} \left(\sum_t \|\mathcal{S}_t^s(\mathbf{x}) - \mathcal{S}_t^s(\mathbf{x}')\|_1 + \sqrt{\frac{s}{2}} \sum_t \|\log \mathcal{S}_t^s(\mathbf{x}) - \log \mathcal{S}_t^s(\mathbf{x}')\|_2 \right), \quad (1)$$

ViSQOL	Multichannel quantization (6.3 kbps/channel)	Single channel quantization (12 kbps)	No quantization
Reconstruction	3.88 ± 0.03	3.84 ± 0.04	4.32 ± 0.03
Speech partition	3.16 ± 0.08	3.08 ± 0.08	3.06 ± 0.08
Noise partition	2.54 ± 0.05	2.55 ± 0.06	2.57 ± 0.05

Table 1. Audio quality after decoding the full embeddings and individual partitions as measured with ViSQOL, where the expected target audio is used as a reference, that is, noisy speech, clean speech and noise, respectively. All three cases have identical training objectives and only differ in the quantization method. The quantization imposes an information bottleneck in the compressed domain, which suppresses noise in the speech partition and vice versa. Values are averaged over 250 noisy waveforms of 3 s duration. Uncertainties denote 95% confidence intervals.

where $S_t^s(\cdot)$ denotes the t -th frame of a 64-bin mel-spectrogram with a window length of s and hop length of $s/4$. The L_1 objective in the first term encourages matching the highest-energy features of the signal, while the L_2 logarithmic term provides attention to the quieter details. For each of the three training steps, we compute this loss function between the reconstructed waveform $\hat{\mathbf{x}}$ and the corresponding target waveform, i.e. $\mathcal{L}_{\text{rec}}(\mathbf{x}, \hat{\mathbf{x}})$ for training step (i), $\mathcal{L}_{\text{rec}}(\mathbf{x}^c, \hat{\mathbf{x}})$ for step (ii), and $\mathcal{L}_{\text{rec}}(\tilde{\mathbf{x}}, \hat{\mathbf{x}})$ for step (iii).

When relying only on the reconstruction losses (1), we find that the output audio suffers from robotic artifacts. To alleviate these, we use a set of discriminators operating on the audio waveforms at multiple scales [4], and train them in an adversarial fashion to encourage the *SoundStream* model to produce more realistically sounding audio. In contrast to previous work, we do not use discriminators operating in the STFT domain, as we found they add no performance benefit, and instead our discriminators operate on multiple scales on the time-domain audio. In addition to minimizing the adversarial loss, we also rely on a feature loss from the absolute differences of the discriminator layer outputs between the original and reconstructed audio [4]. However, we observed that simultaneous training of all components often yields poor output quality and find better results by first training the model on all three tasks using reconstruction losses only, then freezing the parameters of the encoder and quantizer, and finally training the decoder alongside the discriminators to improve the output quality. The optimizers, learning rate and loss weights are identical to those used for training *SoundStream* [4].

3. RESULTS

3.1. Separating speech from noise

We use speech waveforms from the LibriVox dataset [28], and synthesize noisy speech waveforms by mixing these with noise from the Freesound dataset [29], which was screened to remove speech content and only keep clips with a CC0 license. Both the speech and noise waveforms are normalized, then the noise is multiplied by a random gain drawn from a normal distribution with mean of -5 dB and standard deviation of 10 dB before being mixed with the speech waveform. In the input pipeline, we supply pairs of waveforms $(\mathbf{x}_A, \mathbf{x}_A^c)$ and $(\mathbf{x}_B, \mathbf{x}_B^c)$, each with the noisy and clean version, and where \mathbf{x}_A and \mathbf{x}_B differ in speech and background noise signals.

We encode each audio frame of 320 samples into a 256-dimensional embedding vector, which is partitioned into two equal halves, one to carry information about the speech and the other about background noise. Each embedding channel is processed with a dedicated residual vector quantizer, with $n_q = 14$ layers of $\log_2 N = 9$ bits of depth each. Embedding vectors are produced with a frame rate of $R_f = 50$ Hz, leading to an effective bitrate per

channel of $R = 6.3$ kbps. We then train the model on the three objectives detailed in Section 2.2 using waveforms \mathbf{x}_A and \mathbf{x}_B .

Table 1 shows the reconstruction quality of the model as measured with ViSQOL [30, 31], an objective quality metric that correlates well with subjective evaluations. We show the scores for the reconstruction of the full embeddings, each individual partition, and the reconstructed waveform from mixed embeddings of two separate inputs. There, our trained model with separate quantization of each channel is compared with two other models: One with a global residual vector quantizer operating on all embeddings with a bitrate of $R = 12$ kbps, and one where no vector quantization is performed. The model with individually quantized channels performs better on average than the one with a global quantization of the embeddings. The model without quantization yields better quality when reconstructing the noisy speech signal owing to the absence of quantization noise, but it offers no signal compression, and the clean speech reconstruction is similar to the other models.

When reconstructing clean speech from the first embedding partition, the background noise is efficiently removed without substantially hurting speech intelligibility,¹ but broadband stationary noise in the waveform is not fully suppressed, limiting the values of quantitative metrics like ViSQOL. We observe that decoding the noise embedding partition while setting the speech embeddings to zero leads to a reconstruction of the noise component, albeit at lower quality as this reconstruction was not an explicit training task. We also find that multiplying a partition with a weight factor between 0 and 1 leads to a reduction in volume of the corresponding audio content, a feature that can be used for adjusting the target level of denoising.¹

3.2. Separating speech from reverberation

Synthetic reverberated speech differs from speech with background noise in that reverberation is not an additive signal, but rather the room impulse response is convolved with the speech. Another difference is that the information about the room impulse response typically changes much more slowly than the underlying speech content, and for the purpose of this work is assumed to be time-invariant.

We can use the same general framework to train a new codec that separates speech from reverberation. Here we allocate 54 embedding features to speech, which proved sufficient for this purpose. To capture the difference in temporal variations speech and reverberation, we modify the encoder to include an additional convolutional layer striding over multiple frames that outputs a second set of 10-dimensional embedding vectors at a 10-times lower frame rate than the speech embeddings, to capture the typical time scales of the room impulse response. We quantize the reverberation partition with only

¹Audio examples at https://google-research.github.io/seanet/disentangling_soundstream/examples/

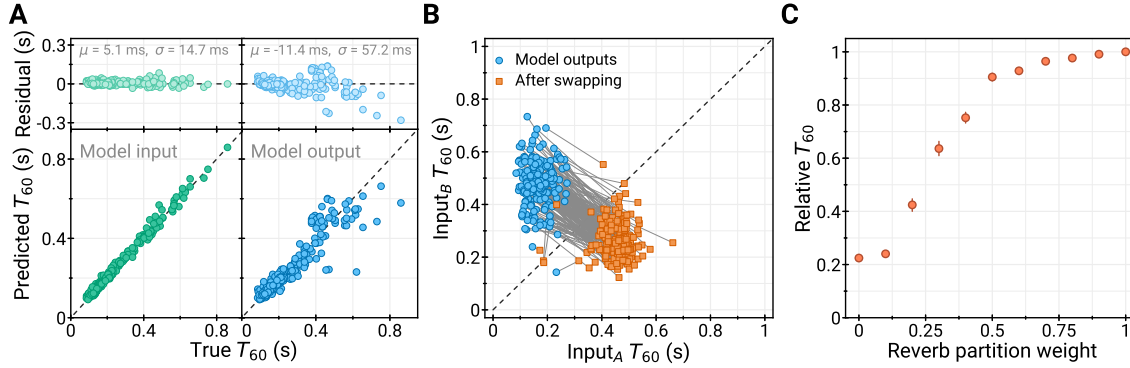


Fig. 2. **A)** Benchmark of the T_{60} estimator and our model’s reconstruction fidelity. In the lower plots, the classified values are compared with the true T_{60} values known from the augmentation, both for the model inputs and reconstructed outputs, where data points with a perfect reproduction of the T_{60} value should lie on the diagonal dashed line. The upper plots show the residuals, the means of the error and their standard deviations. **B)** T_{60} values of pairs of inputs after reconstruction by the model (blue circles) and with a swapped reverb embedding partition (orange squares). Solid gray lines connect points that belong to the same pair of inputs. **C)** Demonstration of reverb tuning. 200 strongly reverberant waveforms ($T_{60} > 400$ ms) are encoded, and the reverb partition is multiplied by a global weight factor before decoding. The measured T_{60} values are normalized by the value at weight factor of 1 and then averaged. Error bars denote 95% confidence intervals.

$n_q = 4$ layers of $\log_2 N = 9$ bit depth, yielding an effective bitrate of $R = 180$ bps. These “slow” embeddings allocated for reverberation are then upsampled by a factor of 10 and concatenated with the “fast” speech embeddings along the feature axis.

We train on clean speech waveforms from the LibriVox dataset that are each augmented with a time-invariant, synthetic room impulse response, which represents a stationary speaker and receiver. The room impulse responses have characteristic T_{60} times between 70 ms and 1.2s with an average of 214 ms and median of 272 ms. In the input pipeline, we provide pairs of waveforms $(\mathbf{x}_A, \mathbf{x}_A^c)$ and $(\mathbf{x}_B, \tilde{\mathbf{x}}_B)$, where \mathbf{x}_A^c is the non-reverberant version of \mathbf{x}_A , and $\tilde{\mathbf{x}}_B = f(\mathbf{x}_B^c, \theta_A)$ is the version of \mathbf{x}_B that has the same room impulse response as \mathbf{x}_A . To make the task more explicit, \mathbf{x}_A is augmented with strong reverberation levels ($T_{60}(\mathbf{x}_A) \geq 400$ ms) and \mathbf{x}_B with weaker reverberation ($T_{60}(\mathbf{x}_B) \leq 250$ ms). We then train on the three objectives detailed in Section 2.2.

We measure the reverberation characteristics of the waveforms using a pretrained estimator of T_{60} , the time reverberated sound takes to decay by 60 dB. Its accuracy is shown in Fig. 2A. Input waveforms \mathbf{x} are fed to the estimator, which yields estimated T_{60} values that do not differ significantly from the ground-truth values used in the augmentation. We also evaluate the reconstructed waveforms $\hat{\mathbf{x}}$ generated by our model, and find a faithful reconstruction of the reverberation characteristics except for very long reverberation times.

To verify the separation of reverberation from speech in the embedding space, we input pairs of audio waveforms $(\mathbf{x}_A, \mathbf{x}_B)$ to our model to obtain reconstructed waveforms $(\hat{\mathbf{x}}_A, \hat{\mathbf{x}}_B)$ and evaluate their T_{60} values. We then take the encoded embeddings of both waveforms and swap their reverberation embedding partitions before decoding and estimate the T_{60} values of the new waveforms $(\tilde{\mathbf{x}}_A, \tilde{\mathbf{x}}_B)$. Fig. 2B shows the T_{60} values before and after this swapping operation, where most pairs have their T_{60} values successfully swapped, indicating the second partition includes all the information about the reverberation.

We further demonstrate the tunability of the decoded outputs by encoding 200 strongly reverberated examples ($T_{60} \geq 400$ ms), multiplying the partition encoding reverberation by a weight factor between 0 and 1 and estimating the decoded output’s T_{60} time (Fig. 2C). We observe that the weight factor smoothly reduces T_{60} in

a predictable way, indicating that we can fine-tune the reverberation of speech on demand.

4. CONCLUSION

In this work, we introduce a training scheme to separate speech from background noise or the room reverberation in the embedding space of a neural audio codec. By alternating between reconstructing input audio, reconstructing only the speech component, and reconstructing new synthetic audio by decoding different combinations of embedding partitions, we achieve effective disentanglement of these contents. We further show how the separation in embedding space leads to a tunable output, for example, by removing the noise component altogether, tuning out the room reverberation, or inheriting a different room impulse response from a separate reference waveform. In future studies, we will scale up this scheme to disentangle multiple attributes at the same time and extend our approach to other factors of variation like pitch and speaker identity.

5. ACKNOWLEDGMENTS

The authors thank Olivier Bachem, Hannah Muckenhirn, John Hershey, Ben Laurie, Dominik Roblek, Beat Gfeller, Yunpeng Li for helpful discussions, and Dan Ellis and Dick Lyon for technical discussions and helpful feedback on the manuscript.

6. REFERENCES

- [1] S. Morishima, H. Harashima, and Y. Katayama, “Speech coding based on a multi-layer neural network,” in *IEEE International Conference on Communications*, 1990, vol. 2, pp. 429–433.
- [2] Srihari Kankanahalli, “End-To-end optimized speech coding with deep neural networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 2521–2525.
- [3] Cristina Gârbasea, Aäron van den Oord, Yazhe Li, Felicia S C Lim, Alejandro Luebs, Oriol Vinyals, and Thomas C Walters,

- “Low bit-rate speech coding with VQ-VAE and a WaveNet decoder,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 735–739.
- [4] Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi, “SoundStream: An end-to-end neural audio codec,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2021.
- [5] Jean-Marc Valin, Koen Vos, and Tim Terriberry, “Definition of the Opus audio codec,” Request for Comments RFC 6716, Internet Engineering Task Force, 2012.
- [6] Philipos C. Loizou, *Speech Enhancement: Theory and Practice*, CRC Press, Boca Raton, 2007.
- [7] Ilya Kavalero, Scott Wisdom, Hakan Erdogan, Brian Patton, Kevin Wilson, Jonathan Le Roux, and John R. Hershey, “Universal sound separation,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2019, pp. 175–179.
- [8] Scott Wisdom, Efthymios Tzinis, Hakan Erdogan, Ron J. Weiss, Kevin Wilson, and John R. Hershey, “Unsupervised sound separation using mixture invariant training,” in *Advances in Neural Information Processing Systems*, 2020, vol. 33, pp. 3846–3857.
- [9] Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner, “Towards a definition of disentangled representations,” *arXiv:1812.02230*, 2018.
- [10] Guillaume Lampe, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic Denoyer, and Marc’Aurelio Ranzato, “Fader networks: Manipulating images by sliding attributes,” in *Advances in Neural Information Processing Systems*, 2017, vol. 30.
- [11] Omid Hajihassnai, Omid Ardakanian, and Hamzeh Khazaei, “ObscureNet: Learning attribute-invariant latent representation for anonymizing sensor data,” in *International Conference on Internet-of-Things Design and Implementation*, New York, NY, USA, 2021, pp. 40–52.
- [12] Ju-chieh Chou, Cheng-chieh Yeh, Hung-yi Lee, and Lin-shan Lee, “Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations,” *arXiv:1804.02812*, 2018.
- [13] Bac Nguyen and Fabien Cardinaux, “NVC-Net: End-to-end adversarial voice conversion,” *arXiv:2106.00992*, 2021.
- [14] Noam Mor, Lior Wolf, Adam Polyak, and Yaniv Taigman, “A universal music translation network,” *arXiv:1805.07848*, 2018.
- [15] Yuan Gong and Christian Poellabauer, “Towards learning fine-grained disentangled representations from speech,” *arXiv:1808.02939*, 2018.
- [16] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem, “A sober look at the unsupervised learning of disentangled representations and their evaluation,” *arXiv:2010.14766*, 2020.
- [17] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson, “AutoVC: Zero-shot voice style transfer with only autoencoder loss,” in *International Conference on Machine Learning*, 2019, vol. 36, pp. 5210–5219.
- [18] Kaizhi Qian, Yang Zhang, Shiyu Chang, David Cox, and Mark Hasegawa-Johnson, “Unsupervised speech decomposition via triple information bottleneck,” in *International Conference on Machine Learning*, 2020, vol. 37, pp. 7836–7846.
- [19] Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharonov, Kushal Lakhotia, Wei-Ning Hsu, Abdelrahman Mohamed, and Emmanuel Dupoux, “Speech resynthesis from discrete disentangled self-supervised representations,” *arXiv:2104.00355*, 2021.
- [20] Haici Yang, Kai Zhen, Seungkwon Beack, and Minje Kim, “Source-aware neural speech coding for noisy speech compression,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 706–710.
- [21] Qiyang Hu, Attila Szabó, Tiziano Portenier, Matthias Zwicker, and Paolo Favaro, “Disentangling factors of variation by mixing them,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 3399–3407.
- [22] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A. Efros, and Richard Zhang, “Swapping autoencoder for deep image manipulation,” in *Advances in Neural Information Processing Systems*, 2020, vol. 33, pp. 7198–7211.
- [23] Tejas D. Kulkarni, Will Whitney, Pushmeet Kohli, and Joshua B. Tenenbaum, “Deep convolutional inverse graphics network,” in *Advances in Neural Information Processing Systems*, 2015, vol. 28.
- [24] Sarthak Bhagat, Vishaal Udandarao, Shagun Uppal, and Saket Anand, “DisCont: Self-supervised visual attribute disentanglement using context vectors,” in *European Conference on Computer Vision Workshops*, Adrien Bartoli and Andrea Fusiello, Eds., Cham, 2020, pp. 549–553.
- [25] Beat Gfeller, Christian Frank, Dominik Roblek, Matt Sharifi, Marco Tagliasacchi, and Mihajlo Velimirović, “SPICE: Self-supervised pitch estimation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1118–1128, 2020.
- [26] Jesse Engel, Lamtharn Hantrakul, Chenjie Gu, and Adam Roberts, “DDSP: Differentiable digital signal processing,” *arXiv:2001.04643*, 2020.
- [27] Alexey A. Gritsenko, Tim Salimans, Rianne van den Berg, Jasper Snoek, and Nal Kalchbrenner, “A spectral energy distance for parallel speech synthesis,” in *Advances in Neural Information Processing Systems*, 2020, vol. 33, pp. 13062–13072.
- [28] Jodi Kearns, “LibriVox: Free public domain audiobooks,” *Reference Reviews*, vol. 28, pp. 7–8, 2014.
- [29] Eduardo Fonseca, Jordi Pons Puig, Xavier Favory, Frederic Font Corbera, Dmitry Bogdanov, Andrés Ferraro, Sergio Oramas, Alastair Porter, and Xavier Serra, “Freesound datasets: A platform for the creation of open audio datasets,” in *International Society for Music Information Retrieval Conference*, 2017, pp. 486–493.
- [30] Andrew Hines, Jan Skoglund, Anil C. Kokaram, and Naomi Harte, “ViSQOL: An objective speech quality model,” *EURASIP Journal on Audio, Speech, and Music Processing*, p. 13, 2015.
- [31] Michael Chinen, Felicia S. C. Lim, Jan Skoglund, Nikita Gureev, Feargus O’Gorman, and Andrew Hines, “ViSQOL v3: An open source production ready objective speech and audio metric,” in *International Conference on Quality of Multimedia Experience (QoMEX)*, 2020, pp. 1–6.