



Distance-Based Sound Separation

Katharine Patterson, Kevin Wilson, Scott Wisdom, John R. Hershey

Google Research, Cambridge MA

{kappi, kwilson, scottwisdom, johnhershey}@google.com

Abstract

We propose the novel task of *distance-based sound separation*, where sounds are separated based only on their distance from a single microphone. In the context of assisted listening devices, proximity provides a simple criterion for sound selection in noisy environments that would allow the user to focus on sounds relevant to a local conversation. We demonstrate the feasibility of this approach by training a neural network to separate near sounds from far sounds in single channel synthetic reverberant mixtures, relative to a threshold distance defining the boundary between near and far. With a single nearby speaker and four distant speakers, the model improves scale-invariant signal to noise ratio by 4.4 dB for near sounds and 6.8 dB for far sounds.

Index Terms: distance-based sound separation

1. Introduction

Extracting estimates of clean speech in the presence of interference is a long-standing research problem in signal processing. This task is referred to as *speech enhancement* when the interference is non-speech, and *speech separation* when the interference is speech. More generally, *sound separation* refers to the extraction of a subset of sounds from a mixture of sounds.

Hearing aids and other assisted listening devices are an important application for these methods, e.g. [1, 2, 3], with a typical task being to aid in conversations held in a noisy space. But implementing sound separation for assisted listening is challenging, requiring low algorithmic latency and limited computation and memory. Further, it is not always clear which sounds the user would like to hear.

The problem of selecting which sounds to enhance has been approached by focusing only on speech [2], or otherwise classifying the sounds [4, 5], as well as using visual input [6, 7, 8] to select the sounds of interest. However, these methods aren't always appropriate. Selective listening methods may be cumbersome and require user effort to control which sounds are enhanced. They also typically exclude sounds other than speech that the user may want to hear, such as nearby music, the clinking of wine glasses at dinner, or the tell-tale sound of a dropped set of keys. Class-based methods that try to include non-speech sounds can fail whenever there are interfering sounds of the same classes as the desired sounds.

We propose an alternative approach called *distance-based sound separation*, illustrated in Figure 1, in which we assume the user would like to hear any sounds that occur within a local region around them, and block sounds coming from farther away. A system that accomplishes this would allow the user to engage in normal conversations without the interference of a crowded environment, and without becoming deaf to non-speech sounds in their immediate area. One hypothesis is that there are cues for distance that are related to superficial characteristics of the sound, rather than to its fine spectro-temporal




Figure 1: Illustration of the proposed task of distance-based sound separation.

structure. If so, a system that relies on such cues might have both a computational advantage, and an advantage in terms of generalization, compared to a system that has to perform deep pattern recognition in order to separate sounds. Although humans have the perceptual ability to estimate the relative distance of sounds [9], there has been no prior work showing that distance-based separation is feasible.

Possible cues for distance perception can be traced to physical effects. The intensity I of the *direct path* component of the sound varies with distance d according to inverse-square law $I \propto 1/d^2$. In contrast, in an enclosed area, the intensity of the sound's reverberation, is roughly independent of its distance to the microphone. Thus the *direct-to-reverberation* (DRR) ratio decreases with distance, and has been shown to be a cue for human distance perception [9]. There is also a *proximity effect* with directional microphones in which low-frequencies are more accentuated for closer sources, but the effect is only strong for sources less than one meter away [10, 11]. Absorption by the air is a frequency-dependent effect of distance, and plays a strong role over distances beyond tens of meters [12]. Other effects of distance come from spatial effects that could be detected using an array of microphones. In the present work, however, we focus on what can be inferred using a single microphone.

To validate the concept of distance-based sound separation, we train neural network separation models using mixtures of near and far sounds, where the acoustic properties of the sounds have been emulated using an acoustic room simulator. This allows us to have ground-truth targets for the near and far signals relative to a distance threshold. In this work, we focus on the case where all sounds, both near and far, are speech. We leave the case where there are non-speech sounds for future work. Our results show that it is possible to perform separation solely based on distance. In particular, in some scenarios, such as with a distance threshold of 1.5 m, a single nearby speaker, and as many as four distant speakers, distance-based sound separation can achieve improvements of up to 4.4 dB in scale-invariant signal to noise ratio.

2. Related Work

Beamforming [13] uses multiple microphones and generally uses the direction of arrival (DOA) of sources as a cue to separate them. Near-field beamforming methods can additionally estimate source distance using the spherical nature of acoustic wavefronts [14, 15], though they can only distinguish distances of nearby sounds within a range limited by the size of the microphone array. In contrast, our method can learn to operate over a wider range of distances. Distance estimation has also been investigated using neural networks with two microphones [16, 17, 18, 19]. However, our method only requires a single microphone, and it performs separation.

Prior work has demonstrated the feasibility of estimating reverberation parameters such as T60 and DRR [9]. Neural networks can also be trained to estimate reverberation parameters from a single microphone [20, 21], as well as room volume [22]. The success of these methods suggests that neural networks are able to perceive cues from raw audio to accurately predict properties of acoustic transfer functions, and it is likely that neural networks could also be trained to estimate the distance of a source. Rather than predict reverb parameters, our proposed model implicitly leverages acoustic cues to identify near and far sources in order to separate them.

Single-channel dereverberation methods [23, 24, 25, 26] solve a somewhat related task to distance-based sound separation, in that they try to separate the direct path of a reverberant signal (the shortest path of sound propagation) from the reverberant components. In contrast, the goal of distance-based sound separation is to preserve the reverberance of sources, and to group sources together based only on whether they are near or far based on a distance threshold. Also, dereverberation methods often make the assumption of a single reverberant source, while our method can handle multiple sources.

To our knowledge, the work introduced here is the first to propose and demonstrate the separation of speech from a single microphone based solely on distance cues.

3. Methods

3.1. Model

To separate based on distance, we experimented with an architecture [27], that uses short-time Fourier transform (STFT) masking. For the STFT, we use a 32 ms square-root Hann window with 16 ms hop. The 0.3-power-compressed magnitude of the STFT, Y , of the time-domain mixture, y , is fed as input to L layers of uni-directional LSTMs [28] with N units each. Then, a fully-connected layer with a sigmoid activation is applied to create two masks for the input STFT Y , M_{near} and M_{far} . To ensure STFT consistency [27], the masked STFTs for near and far are each passed through inverse and forward STFT operations: $\hat{X}_{\text{near|far}} = \text{STFT}\{\text{iSTFT}(M_{\text{near|far}} \odot Y)\}$.

The training loss is mean-squared error between 0.3-power-compressed magnitude of target X and estimate \hat{X} [27]:

$$\mathcal{L}(X, \hat{X}) = \sum_{f,t} (|X_{f,t}|^{0.3} - |\hat{X}_{f,t}|^{0.3})^2 \quad (1)$$

We use weights of 0.8 on the near loss and 0.2 on the far loss, $0.8\mathcal{L}(X_{\text{near}}, \hat{X}_{\text{near}}) + 0.2\mathcal{L}(X_{\text{far}}, \hat{X}_{\text{far}})$, which encourages the model to focus on the performance of near targets, which is more likely to be desired as the output in a practical application. For all experiments, we used the Adam optimizer [29] with a learning rate of 3×10^{-5} , batch size 128, and trained for one million steps on 16 Google Cloud TPU v3 cores.




Figure 2: Source distribution by distance from microphone, and spatial distribution in 2D.

3.2. Acoustic Simulation

To create a large amount of training data, we use an image-method [30] acoustic room simulator with frequency-dependent wall filters [31] to generate reverber impulse responses (RIRs) for rooms with varied acoustic properties, with randomized microphone and source locations. Basic distance-related phenomena such as the amplitude and DRR effects are well simulated. However, our simulation does not model the proximity effect and the air absorption effect, which might provide additional cues for distance-based separation in the real world. Clean speech recordings are randomly assigned to the source locations within each room and convolved with the corresponding RIRs. We created training examples for each room by combining all sources within a threshold distance into a near target x_{near} , and all sources beyond a threshold distance into a far target x_{far} . The near and far targets are added to create mixture $y = x_{\text{near}} + x_{\text{far}}$, and fed into the model to be separated.

Within each of the randomly generated rooms, we randomly generate 5 speaker locations and 1 microphone location, such that the distance between the microphone and each source was uniformly distributed, subject to the rejection of samples falling outside each room. The resulting distribution of distances from the microphones is illustrated in Figure 2. Note that the spatial distribution differs from what might be expected in, for example, a restaurant setting, where the average number of sources at a given distance radius will increase as the distance increases.

4. Experiments

4.1. Data Preparation

For the experiments, we use speech recordings from the Libri-light dataset [32] for training, with validation, and test partitions coming from LibriSpeech [33], so that speaker IDs are unique for each partition. For our synthetic room data, rooms are generated with dimensions varying from $3.0 \times 4.0 \times 2.13$ meters to $7.0 \times 8.0 \times 3.05$ meters.

During training, randomly chosen Libri-light clips are reverberated and mixed according to a randomly chosen RIR to create clips that are 10 seconds in duration. Source utterances shorter than 10 seconds are offset at random intervals within the 10 second clip, and source utterances longer than 10 seconds are clipped to a random 10 second interior segment.

The source and microphone locations are used to determine which sources are near or far in relation to the microphone in the given room, for a given threshold distance. To vary the number of sources, we apply a *source presence probability* (SPP) to each source, so that the total number of sources in a room varies from 0 to 5, with a distribution dependent on the chosen SPP.

Because the rooms are generated independently of a specific distance threshold, the choice of distance threshold affects the number of sources considered near versus far for any given room. In particular, it affects the fraction of examples where all the sources are considered near *or* all sources are considered far, leading to silent far targets *or* silent near targets. Distances thresholds of 0.8/1.5/3.0 meters result in training data with silent far target fractions of 0.0/0.001/0.04, respectively and silent near target fractions of 0.59/0.30/0.03 respectively. It is notable that some thresholds result in a majority of training examples containing a silent target, even for an SPP of 1.0.

4.2. Metrics

For training examples where both near and far targets have sources present (i.e. neither target is silent), we use scale-invariant SDR improvement [34] (SI-SDRi), $\text{SI-SDRi}(x, \hat{x}) = \text{SI-SDR}(x, \hat{x}) - \text{SI-SDR}(x, y)$. SI-SDR measures signal fidelity with respect to a reference signal while allowing for a gain mismatch:

$$\text{SI-SDR}(x, \hat{x}) = 10 \log_{10} \left(\frac{\|\alpha x\|^2}{\|\alpha x - \hat{x}^s\|^2} \right), \quad (2)$$

where $\alpha = \text{argmin}_\alpha \| \alpha x - \hat{x} \|^2 = x^T \hat{x} / \|x\|^2$.

SI-SDRi diverges when one of the targets is silent, because the non-silent target will be exactly equal to the input mixture and achieve ∞ dB SI-SDR, and the silent target will have $-\infty$ dB SI-SDR. This makes any improvement calculation meaningless. For examples with a silent near target, we use a noise reduction metric, to measure how much of the far sound leaks into the silent near output:

$$\text{NoiseReduction}(y, \hat{x}_{\text{near}}) = 10 \log_{10} \left(\frac{\|y\|^2}{\|\hat{x}_{\text{near}}\|^2} \right), \quad (3)$$

which measures the power reduction of the separated output \hat{x}_{near} relative to the input audio mixture power y .

5. Results

Results are shown in Table 1 for an evaluation set of 1000 examples with a distance threshold of 1.5 meters, model size of $L = 4$ and $N = 400$ (6.3M parameters), SPP of 0.5, and with all 5 speakers present. They are bucketed according to number of near speakers (1 near speaker implies 4 far speakers, etc).

We compare our distance-based separation network to a conventional speech separation baseline in which we train the same network architecture (but with more masking outputs and a permutation-invariant loss) to output 5 separate source estimates, one for each possible speaker. To upper-bound the performance of this baseline, we use oracle knowledge of the true near and far targets to find the best grouping of the 5 separate source estimates into near and far estimates. In all but one condition, direct distance-based separation outperforms this upper bound on our “separate-all” baseline.

Figure 3 shows scatter plots of near input SI-SDR versus SI-SDR improvement for individual examples. The mean noise reduction (3) for this model across the 271 evaluation examples with silent near targets (n/a in Table 1) is 48.9 dB.

Demos are online at <https://google-research.github.io/sound-separation/papers/distance-based-separation-interspeech2022>.

5.1. Effect of distance threshold

The results of training and evaluating with different distance thresholds is shown in Table 2.

Table 1: Mean SI-SDRi for distance threshold of 1.5 m, LSTM depth $L = 4$, width $N = 400$, and training SPP of 0.5, on an eval set of 1000 5-speaker examples with SPP of 1.0. “Distance-based” is our method of directly training to separate near sounds from far. “Separate-all” is a baseline in which same network architecture is trained to output 5 individual sources, which are then grouped into near and far estimates.

| # near sources | Distance-based | | Separate-all | | # |
|----------------|----------------|------------|--------------|------------|-------|
| | Near | Far | Near | Far | |
| 1 | 4.4 | 6.8 | 2.5 | 3.9 | (372) |
| 2 | 2.2 | 7.1 | 0.4 | 4.7 | (248) |
| 3 | -0.3 | 6.7 | -2.1 | 4.9 | (92) |
| 4 | -2.6 | 5.5 | -6.6 | 5.7 | (15) |

Table 2: Performance as a function of distance threshold, for models identical to that in Table 1, but trained with varied distance thresholds. SI-SDRi is given for all examples with 1 non-silent near source. Noise reduction is calculated only for the 271 examples with silent near targets.

| Distance (m) | SI-SDRi, 1 near source | | Noise reduction |
|--------------|------------------------|-----|-----------------|
| | Near | Far | |
| 0.8 | 4.0 | 8.3 | 63.9 |
| 1.5 | 4.4 | 6.8 | 48.9 |
| 3.0 | 1.9 | 2.7 | 22.8 |

Performance generally degrades with increased distance threshold values, despite the fact that higher thresholds resulted in more non-silent training targets. This can be seen in Figure 2 and Section 4.1 where a threshold at 3 meters is closer to the center of the distribution than a threshold of 0.8 or 1.0, and thus will change the training examples to have fewer silent targets.

The correlation between increased performance and decreased silent targets seen in Table 4 suggests that the degradation seen in Table 2 with increased distance threshold is a true artifact of the distance, as opposed to an effect of the change in training data between models.

5.2. Effect of model size

All combinations of LSTM depth $L \in \{2, 4, 6\}$ and width $N \in \{200, 400, 600\}$ are evaluated, except $L = 6$ and $N = 600$, which was too large for our resource constraints. Overall, larger models performed better, with LSTM width N contributing more to performance increases. The SI-SDR difference between the smallest model and the largest model for distance thresholds of 0.8 m, 1.5 m, and 3.0 m was 0.9 dB near / 1.2 dB far, 1.5 dB near / 1.6 dB far, and 1.8 dB near / 1.4 dB far, respectively. Note that for increasing model size, SI-SDRi improves more for higher distance thresholds.

5.3. Effect of varied speaker count in evaluation set

To explore whether the model uses the relationship between the number of near sources and the number of far sources, we vary the SPP. For instance, an SPP of 1.0 indicates that all 5 sources are present in the mixture, and in general a SPP of p induces a binomial distribution $\mathcal{B}(5, p)$ over number of sources. This also modifies the distribution of near and far speaker counts. Our baseline model evaluated in Table 1 was trained with a SPP of 0.5, but our evaluation set used in other sections uses a SPP of 1.0 for ease of slicing data by number of near speakers. This means that training data and the evaluation data differ in terms

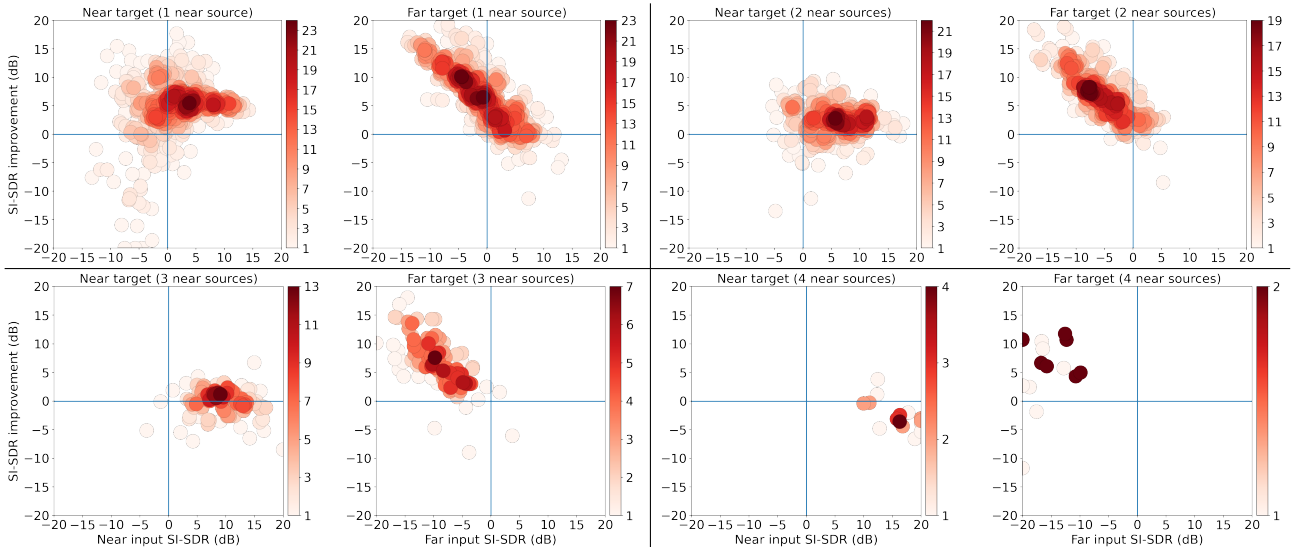


Figure 3: $SI-SDR_i$ in dB as a function of input $SI-SDR$ for near and far targets, for the model evaluated in Table 1 on test examples with a distance threshold of 1.5 m. Results are divided by the number of sources in the near target.

of near and far speaker counts.

In order to study this mismatch, we also created a version of the evaluation set with a SPP of 0.5. This provides a set with 527 silent near target examples, 411 examples with sound in near and far (i.e. 1-4 speakers in each target, $\text{sum} \leq 5$), and 62 examples with silent far targets. We also trained our baseline model with SPP of 1.0 instead of 0.5. Results for the baseline model and eval data at SPPs of 0.5 and 1.0 are in Table 3.

Table 3: Performance as a function of SPP, for same model as Table 1, except for SPP. For Eval SPP=1, $SI-SDR_i$ is given for all 727 examples with non-silent near and far targets. Noise reduction is calculated only for the 271 examples with silent near targets. For Eval SPP=0.5, $SI-SDR_i$ is for 394 non-silent near and far target examples, noise reduction for 538 silent near targets examples.

| Model SPP | Eval SPP | SI-SDR _i all non-silent | | Noise reduction |
|-----------|----------|------------------------------------|-----|-----------------|
| | | Near | Far | |
| 0.5 | 1 | 2.9 | 6.8 | 48.9 |
| 0.5 | 0.5 | 3.2 | 9.6 | 56.3 |
| 1 | 1 | 3.4 | 7.2 | 40.6 |
| 1 | 0.5 | 2.8 | 2.6 | 46.8 |

Models do best when their training SPP matches the eval data’s SPP, as expected. However, the effect of mismatch for near $SI-SDR_i$ is relatively minor. For far $SI-SDR_i$, we observe more degradation with mismatched SPPs, indicating that perhaps source count is an important cue for far source estimation.

5.4. Effect of silent targets

To study the effect of silent targets, we pre-filtered our RIRs to achieve certain percentages of silent targets (before applying SPP). Results are shown in Table 4 for SPPs of 1.0 and 0.5. Note that SPP and percentage of rooms with silent targets are compounding effects. As expected, showing the model more examples with non-silent targets increases $SI-SDR_i$ on non-silent examples, but degrades noise reduction on silent near examples.

Table 4: Performance as a function of percentage silent training targets. “% rooms silent target” is how many rooms in the training set had silent near or far targets, before the SPP was applied. Noise reduction is calculated only for the 271 eval examples with silent near targets.

| SPP | % rooms silent target | SI-SDR _i | | Noise reduction |
|-----|-----------------------|---------------------|-----|-----------------|
| | | Near | Far | |
| 1.0 | 0 | 6.7 | 7.8 | 23.7 |
| 1.0 | 30 | 5.0 | 6.9 | 40.6 |
| 0.5 | 0 | 4.6 | 6.5 | 35.4 |
| 0.5 | 30 | 4.4 | 6.8 | 48.9 |

6. Conclusions

We proposed the novel task of separating mixtures of signals based on their distance from a single microphone. To address this task we applied an existing high-performing sound separation model architecture, trained on synthetic data prepared to exemplify the primary qualities of sound that vary with distance. Experiments showed promising initial performance on this task, with our model benchmark 1.5 meter distance threshold model producing an average improvement of 4.4 dB in $SI-SDR$ for nearby sounds.

In future work, we plan to handle more realistic scenarios with both speech and non-speech and to investigate multi-channel approaches. We also aim to elucidate the relative importance of acoustic cues, such as relative loudness and DRR, in distance-based separation networks, and evaluate their performance in real acoustic conditions. Training with more realistic acoustic effects, including air absorption and proximity effects, may also be useful for generalization to real world conditions, and we plan to investigate this using improved simulations.

7. Acknowledgements

We would like to acknowledge Dick Lyon and Malcolm Slaney for helpful discussions during the preparation of this work, and Hakan Erdogan for useful comments on the manuscript.

8. References

- [1] D. Wang, "Time-frequency masking for speech separation and its potential for hearing aid design," *Trends in amplification*, vol. 12, no. 4, pp. 332–353, 2008.
- [2] S. Wang, G. Naithani, and T. Virtanen, "Low-latency deep clustering for speech separation," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2019, pp. 76–80.
- [3] I. Fedorov, M. Stamenovic, C. Jensen, L.-C. Yang, A. Mandell, Y. Gan, M. Mattina, and P. N. Whatmough, "Tynylstms: Efficient neural speech enhancement for hearing aids," *Proc. Interspeech 2020*, pp. 4054–4058, 2020.
- [4] Q. Kong, Y. Wang, X. Song, Y. Cao, W. Wang, and M. D. Plumbley, "Source separation with weakly labelled data: An approach to computational auditory scene analysis," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2020, pp. 101–105.
- [5] S. Wisdom, A. Jansen, R. J. Weiss, H. Erdogan, and J. R. Hershey, "Sparse, efficient, and semantic mixture invariant training: Taming in-the-wild unsupervised sound separation," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2021, pp. 51–55.
- [6] A. Hussain, J. Barker, R. Marxer, A. Adeel, W. Whitmer, R. Watt, and P. Derleth, "Towards multi-modal hearing aid design and evaluation in realistic audio-visual settings: Challenges and opportunities," in *Proc. International Workshop on Challenges in Hearing Assistive Technology*, 2017.
- [7] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–11, 2018.
- [8] M. Gogate, K. Dashtipour, A. Adeel, and A. Hussain, "Cochleanet: A robust language-independent audio-visual model for real-time speech enhancement," *Information Fusion*, vol. 63, pp. 273–285, 2020.
- [9] A. J. Kolarik, B. C. Moore, P. Zahorik, S. Cirstea, and S. Pardhan, "Auditory distance perception in humans: a review of cues, development, neuronal bases, and effects of sensory loss," *Attention, Perception, & Psychophysics*, vol. 78, no. 2, pp. 373–395, 2016.
- [10] E. Milanova and E. Milanov, "Proximity effect of microphone," in *Audio Engineering Society Convention 110*. Audio Engineering Society, 2001.
- [11] E. N. Milanov and E. B. Milanova, "Proximity effect frequency characteristics of directional microphones," in *Audio Engineering Society Convention 108*. Audio Engineering Society, 2000.
- [12] C. M. Harris, "Absorption of sound in air versus humidity and temperature," *The Journal of the Acoustical Society of America*, vol. 40, no. 1, pp. 148–159, 1966.
- [13] H. L. Van Trees, *Optimum array processing: Part IV of detection, estimation, and modulation theory*. John Wiley & Sons, 2004.
- [14] J. G. Ryan and R. A. Goubran, "Near-field beamforming for microphone arrays," in *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, 1997, pp. 363–366.
- [15] Y. R. Zheng, R. A. Goubran, and M. El-Tanany, "Robust near-field adaptive beamforming with distance discrimination," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 12, no. 5, pp. 478–488, 2004.
- [16] M. Yiwere and E. J. Rhee, "Distance estimation and localization of sound sources in reverberant conditions using deep neural networks," *Int. J. Appl. Eng. Res.*, vol. 12, no. 22, pp. 12 384–12 389, 2017.
- [17] E. L. Ferguson, S. B. Williams, and C. T. Jin, "Sound source localization in a multipath environment using convolutional neural networks," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2386–2390.
- [18] M. Yiwere and E. J. Rhee, "Sound source distance estimation using deep learning: an image classification approach," *Sensors*, vol. 20, no. 1, p. 172, 2019.
- [19] D. A. Krause, A. Politis, and A. Mesaros, "Joint direction and proximity classification of overlapping sound events from binaural audio," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2021, pp. 331–335.
- [20] H. Gamper and I. J. Tashev, "Blind reverberation time estimation using a convolutional neural network," in *Proc. IEEE International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2018, pp. 136–140.
- [21] N. J. Bryan, "Impulse response data augmentation and deep neural networks for blind room acoustic parameter estimation," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2020, pp. 1–5.
- [22] A. F. Genovese, H. Gamper, V. Pulkki, N. Raghuvanshi, and I. J. Tashev, "Blind room volume estimation from single-channel noisy speech," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2019, pp. 231–235.
- [23] T. Nakatani, K. Kinoshita, and M. Miyoshi, "Harmonicity-based blind dereverberation for single-channel speech signals," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 1, pp. 80–95, 2006.
- [24] E. A. Habets, "Speech dereverberation using statistical reverberation models," in *Speech Dereverberation*. Springer, 2010, pp. 57–93.
- [25] Y. Luo and N. Mesgarani, "Real-time single-channel dereverberation and separation with time-domain audio separation network," in *Proc. Interspeech*, 2018, pp. 342–346.
- [26] O. Ernst, S. E. Chazan, S. Gannot, and J. Goldberger, "Speech dereverberation using fully convolutional networks," in *Proc. European Signal Processing Conference (EUSIPCO)*, 2018, pp. 390–394.
- [27] S. Wisdom, J. R. Hershey, K. Wilson, J. Thorpe, M. Chinen, B. Patton, and R. A. Saurous, "Differentiable consistency constraints for improved deep speech enhancement," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2019.
- [28] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.
- [30] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," in *The Journal of the Acoustical Society of America*, 1979, p. 65(4):943–950.
- [31] Z.-Q. Wang, H. Erdogan, S. Wisdom, K. Wilson, and J. R. Hershey, "Sequential Multi-Frame Neural Beamforming for Speech Separation and Enhancement," in *Proc. Spoken Language Technology Workshop (SLT)*, 2021.
- [32] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, "Libri-light: A benchmark for asr with limited or no supervision," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2020, pp. 7669–7673, <https://github.com/facebookresearch/libri-light>.
- [33] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [34] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR-half-baked or well done?" in *Proc. ICASSP*, 2019, pp. 626–630.