

SepIt: Approaching a Single Channel Speech Separation Bound

Shahar Lutati¹, Eliya Nachmani^{1,2}, Lior Wolf¹

¹Tel-Aviv University ²Meta AI Research

shahar761@gmail.com, enk100@gmail.com, wolf@cs.tau.ac.il

Abstract

We present an upper bound for the Single Channel Speech Separation task, which is based on an assumption regarding the nature of short segments of speech. Using the bound, we are able to show that while the recent methods have made great progress for a few speakers, there is room for improvement for five and ten speakers. We then introduce a Deep neural network, SepIt, that iteratively improves the different speakers' estimation. At test time, SepIt has a varying number of iterations per test sample, based on a mutual information criterion that arises from our analysis. In an extensive set of experiments, SepIt outperforms the state of the art neural networks for 2, 3, 5, and 10 speakers.

Index Terms: speech separation, single channel, deep learning

1. Introduction

Single Channel Speech Separation (SCSS) problem is a specific setting of the general Blind Source Separation (BSS) problem. By employing deep neural networks, great improvement has been achieved in the last few years in separating two and three speakers [1, 2], leading to the very recent state of the art [3]. Somewhat disappointingly, in terms of model sizes versus improvement obtained, the model size has more than tripled (26M [3] vs 7.5M [2]), while the improvement as measured when training on same dataset without dynamic mixing is only 0.1dB. One may, therefore, wonder whether additional improvement in performance is possible. This drives the need for a theoretical upper bound for SCSS.

SCSS has unique characteristics in comparison to other BSS problems. The speech signal is not stationary, unless short segments are considered. Jensen et al. [4] have shown empirically that when the length of the speech segment is longer than 20[ms], the distribution is closer to the Laplace distribution than to the normal distribution. As a result, the known bounds for BSS [5] do not hold. In this work, we derive such a bound by employing the assumption of a Laplacian distribution. Treating the speech mixture as a random process, we derive a bound that expresses the maximum achievable Signal to Distortion Ratio (SDR) by any neural network. We then present a deep learning method called SepIt, which uses the bound during its training.

In an extensive set of experiments, we validate the assumption made in our analysis, as well as the bound itself. We then show that the SepIt model outperforms the state of the art in separating two and three speakers for the WSJ [6] benchmark and five and ten speakers for the LibriMix [7] benchmark.

Related Work SCSS using deep learning techniques has been explored intensively in the last years. Erdogan et al. [8] introduced a phase-sensitive loss function trained using a LSTM neural network. Hershey et al. [9] developed the WSJ-2mix benchmark and proposed a neural separation network with a clustering-based embedding. The TasNet architecture [10] employs a time domain encoder-decoder. Subsequently, Luo et al. [11] introduced the ConvTasNet architecture based on a convo-

lutional neural network with masking. ConvTasNet was further improved with the dual-path recurrent neural network (DPRNN) architecture by [12]. Zeghidour et al. [13] introduced Wavesplit, which uses clustering on speaker representation to separate the mixture. Nachmani et al. [2] proposed the VSUNS model, which is based on the DPRNN model, but removes the masking sub-network. The SepFormer architecture is a transformer-based architecture that captures both short- and long-term dependencies [3]. Yao et al. [14] introduced a coarse-to-fine framework called the Stepwise-Refining Speech Separation Network (SRSSN). Hu et al. [15] presented the Fully Recurrent Convolutional Neural Network (FRCNN) architecture, which uses lateral connections to fuse information processed at various timescales.

Independently, upper bounds have been developed for the BSS problem. Sahlin et al. [16] derived an asymptotic Cramer Rao bound for the MIMO case, which assumes a full-rank channel matrix. Doron et al. [5] presented a Cramer Rao bound for Gaussian sources. Kautsky et al. [17] perform an analysis for both the Gaussian and non-Gaussian sources, presenting a Cramer Rao induced bound. However, the analysis only addresses linear operations over the mixture signal. As far as we can ascertain, no upper bound for non-Gaussian sources exists for non-linear methods such as Deep Neural Networks.

2. SCSS Upper bound

We start by formulating the required lemmas and stating the assumption we employ.

Lemma 2.1 (Fisher Information Upper Bound). *From [18] the Fisher Information $J(\theta)$ is upper bounded by the Mutual Information $I(X, \theta)$:*

$$J(\theta) \leq I(X, \theta) \quad (1)$$

Lemma 2.2 (Joint Estimation Lemma). *For N i.i.d random variables with parameter θ , the joint Fisher Information is*

$$J(\theta_{joint}) = N \cdot J(\theta) \quad (2)$$

Jensen et al. [4] demonstrated empirically that speech, divided into short segments, neglecting quiet segments, is well modeled as a stationary process. Furthermore, short segments of a single speaker follow a Laplace random variable.

Assumption 2.3. Short speech segments can be captured by a Laplace random variable with zero mean and scale parameter b .

$$f(\text{voice} = x) = 1/2b^{-1}e^{-|x|/b} \quad (3)$$

Mixture Model Let the different speakers' signals in the mixture compose a matrix $V = [v_0, v_1, \dots, v_{C-1}]^T$, where $v_i \in \mathbb{R}^L$ is a single speaker signal in the time domain, with length L . To maintain the stationary requirements we use windowing of size w as in Assumption 2.3. The mixture $m = \sum_i a_i \cdot v_i$ is modeled as a linear combination of the columns of V , where $a_i \in \mathbb{R}$ are the mixture coefficients.

Definition 2.4 (Signal to Distortion Ratio (SDR)). Let the error ϵ be defined as $\epsilon = v_0 - \bar{v}_0$. The Signal to Distortion Ratio (SDR) is,

$$SDR = 10 \log_{10} \frac{Var(v_0)}{Var(\epsilon)} \quad (4)$$

Denote a single segment of size w from m as m_r , and a single segment of v_i as $v_{i,r}$, where $0 \leq r \leq \frac{L}{w}$. Without loss of generality, assume that the goal is to separate v_0 from the mixture m . Denote by \bar{v}_0 the estimation of v_0 , using deep neural network, D .

Lemma 2.5. Given a single segment $v_{0,r}, m_r$ and deep neural network, D , the following inequality holds,

$$\frac{L}{w} \cdot I(m_r, v_{0,r}) \geq J(\bar{v}_0) \quad (5)$$

Proof. Considering the Markov chain $v_{0,r} \rightarrow m_r \xrightarrow{D} \bar{v}_{0,r}$ and using the Data Processing Theorem,

$$I(m_r, \bar{v}_{0,r}) < I(m_r, v_{0,r}) \quad (6)$$

For neural network D that estimates jointly all of the segments we have from Lemma 2.2,

$$J(\bar{v}_0) = \frac{L}{w} \cdot J(\bar{v}_{0,r}) \quad (7)$$

Recall that \bar{v}_0 is the estimator of v_0 from m . Referring v_0 as θ in Lemma 2.1 gives,

$$I(m_r, v_{0,r}) \geq J(\bar{v}_{0,r}) \quad (8)$$

Plugging Eq.8, and Eq.7 gives,

$$\frac{L}{w} \cdot I(m_r, v_{0,r}) \geq J(\bar{v}_0) \quad (9)$$

□

Based on this lemma, we obtain the following bound.

Theorem 2.6 (SDR upper bound). Let $m_r, v_{0,r}, L, w$ be defined as previously stated. For any neural network, the SDR upper bound that separates a mixture of C speakers is given by:

$$SDR(v_0, \bar{v}_0) \leq 10 \log_{10} \left(\frac{L}{w} \cdot Var(v_0) \cdot I(m_r, v_{0,r}) \right) \quad (10)$$

Proof. The Cramer Rao Lower Bound [19] states that for any unbiased estimator there is a lower bound on the estimator square error i.e. for estimator of $\bar{v}_{0,r}$

$$Var(\bar{v}_{0,r}) \geq \frac{1}{J(\bar{v}_{0,r})} \quad (11)$$

Applying Eq. 11 to Definition. 2.4 gives an upper bound for estimating the single segment:

$$SDR(v_{0,r}, \bar{v}_{0,r}) \leq 10 \log_{10} (Var(v_{0,r}) J(\bar{v}_{0,r})) \quad (12)$$

Using the upper bound for the Fisher information Lemma.2.1 yields:

$$SDR(v_{0,r}, \bar{v}_{0,r}) \leq 10 \log_{10} (Var(v_{0,r}) I(m_r, \bar{v}_{0,r})) \quad (13)$$

The data processing inequality states the the mutual information is reduced after each processing, thus we obtain the following:

$$SDR(v_{0,r}, \bar{v}_{0,r}) \leq 10 \log_{10} (Var(v_{0,r}) I(m_r, v_{0,r})) \quad (14)$$

Using Lemma.2.2 gives the upper bound:

$$SDR(v_0, \bar{v}_0) \leq 10 \log_{10} \left(\frac{L}{w} \cdot Var(v_0) I(m_r, v_{0,r}) \right) \quad (15)$$

□

Theorem 2.6 states that the upper bound for separation algorithm, is a function of: (i) the number of speakers C , (ii) the number of segments that jointly estimated $\frac{L}{w}$, (iii) the mutual information $I(m_r, v_{0,r})$.

Specifically, the bound is linear with L , the length of the sequence. Furthermore, it is inversely proportional to the segment length w . There is a trade-off between w and the lowest frequencies in each segment, thus, decreasing w also decreases $Var(v_{0,r})$.

A clear limitation of the bound is that it holds only for a network D that jointly process i.i.d stationary segments. This is not the case if a neural network processes the entire signal without segmentation. However, current architectures, including the attention architecture, tend to be myopic.

3. Method

Our method, denoted as SepIt, consists of sequential processing of the estimated signals, where each iteration contains a replica of the basic model. Consider again the mixture signal, m , consisting of C speakers and of length L . First, a backbone network, B , is used for initial separation estimation. The backbone can change between different data sets, and it does not learn during the training phase of the SepIt model. Expanding the notation in Sec. 2, denote the estimation of the i -th speaker in the j -th iteration as \bar{v}_i^j . Collectively, we denote the list of obtained speakers as $\bar{v}^j = [\bar{v}_0^j, \bar{v}_1^j, \dots, \bar{v}_{C-1}^j]$. The output of the backbone, B , is v^0 .

$$\bar{v}^0 = B(m) \quad (16)$$

Where $\bar{v} \in \mathbb{R}^{C \times L}$ is all speaker estimation.

Both the current estimation, \bar{v}^{j-1} , and the mixture, m , pass through an encoder, E , of a 1-D convolution with N channels, kernel size K and stride $\frac{K}{2}$, followed by a ReLU activation function:

$$\hat{m} = E(m) \quad (17)$$

$$\hat{v}^{j-1} = E(\bar{v}^{j-1}) \quad (18)$$

Where $\hat{m} \in \mathbb{R}^{N \times L'}$, $\hat{v}^{j-1} \in \mathbb{R}^{C \times N \times L'}$ are the latent space for the mixture and current speech estimations, respectively, and

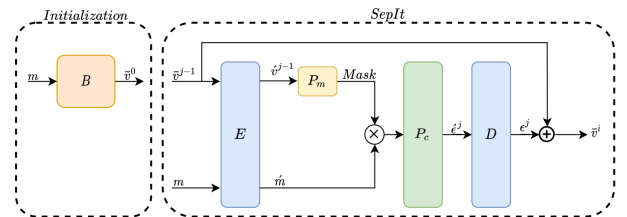


Figure 1: The SepIt block architecture.

$L' = \frac{L}{K}$ is the strided latent dimension. The encoder of m shares its weights with the encoder of \bar{v}^{j-1} .

The latent estimation \hat{v}^{j-1} is treated as a prior on the latent masking. A deep neural network, P_m , is used to evaluate a latent space mask. P_m consists of 3 ResBlocks, each a 1-D version of the Residual block from ResNet, each with N channels, kernel size of 3, and ReLU activation function. The latent mask is multiplied element-wise by the mixture latent representation. To subtract residuals from other sources, a 1×1 Convolution layer P_c with $N \times C$ channels is then introduced. The procedure is:

$$\epsilon^j = P_c(P_m(\hat{v}^{j-1}) \circ \hat{m}), \quad (19)$$

where $\epsilon^j \in \mathbb{R}^{C \times N \times L'}$ is the latent space residual, and \circ is the element-wise product operator. The latent residual ϵ is passed to a decoder layer, D . Later, a skip-connection is introduced, such that the j -th iteration estimation is the summation of the decoded ϵ and the previous estimation,

$$\bar{y}^j = \bar{y}^{j-1} + D(\epsilon^j) \quad (20)$$

A single SepIt block is depicted in Fig 1. The training algorithm is depicted in Alg. 1. The algorithm starts with the initialization of the threshold u to 1. The zero-th iteration is set using the backbone model as in line 3. Next, the algorithm evaluates the SepIt neural network over the previous iteration, and takes a gradient descent step over the loss function. The threshold u is updated by the difference between $I(m, \bar{v}^j)$ and $I(m, \bar{v}^{j-1})$.

The algorithm proceeds to apply the SepIt neural network until the maximum number of iterations is reached or the threshold is lower or equal to zero. This stopping criterion follows the derivation of Thm.2.6, which indicates that the mutual information between the speaker estimation \bar{v}^j and the mixture m can indicate the maximal achievable SDR. Thus, a per-sample stopping criterion is derived when $I(m, \bar{v}^j)$ stops increasing relative to previous iteration $I(m, \bar{v}^{j-1})$. This stopping criterion can be computed also on the test-set samples since it does not require the ground truth.

Loss Function Following [20], a Scale Invariant Signal to Distortion Ratio (SI-SDR) is used as the loss function, which is a similar loss to SDR, but demonstrated better convergence rate. The overall loss function is given by:

$$\mathcal{L}(v, \bar{v}) = 10 \log_{10} \left(\frac{\|\tilde{v}\|^2}{\|\tilde{e}\|^2} \right) \quad (21)$$

where \bar{v} is the estimated source, $\tilde{v} = \frac{\langle v, \bar{v} \rangle v}{\|v\|^2}$, and $\tilde{e} = \bar{v} - \tilde{v}$.

4. Experiments

Validating Assumption 2.3 This assumption is extensively validated by [4], yet we re-evaluated it again over the LibriMix training set. The LibriMix dataset containing over 100 hours of speech is investigated. Each speaker signal is split into different non-overlapping segments, each 20[ms] long. Fig. 2 compares the empirical PDF with the best fitted Laplace distribution and normal distribution. Evidently, the Laplacian distribution provides a much better fit. The Kullback–Leibler that presented in the figure supports this conclusion too.

Separation results For all experiments, the Adam [21] optimizer is used with a learning rate of $5 \cdot 10^{-4}$ and a decay factor of 0.95 every epoch. The window size $w = 20$ [ms] and speech length $L = 4$ [s]. Other hyperparameters are summarized in Tab. 1. The experiments are divided into two categories based on their Backbone architecture: (i) Transformer -based and (ii)

Algorithm 1 The SepIt algorithm.

Input: m - mixture signal, C - number of speakers, B - Backbone network, MAXITER - maximum number of iterations
Output: \bar{v}^j - estimated separated speakers.

```

1:  $\bar{v}^0 \leftarrow B(m)$ 
2:  $j \leftarrow 0$ 
3:  $u \leftarrow 1$ 
4:  $l \leftarrow 0$ 
5: while  $u \geq 0$  and  $j \leq \text{MAXITER}$  do
6:    $\bar{v}^j \leftarrow \text{SepIt}_j(\bar{v}^{j-1}, m)$ 
7:    $l \leftarrow \text{SDR}(\bar{v}^j, v)$ 
8:   Take gradient descent step on  $-l$ 
9:    $j \leftarrow j + 1$ 
10:   $u \leftarrow I(m, \bar{v}^j) - I(m, \bar{v}^{j-1})$ 
11: end while
12: return  $\bar{v}^j$ 

```

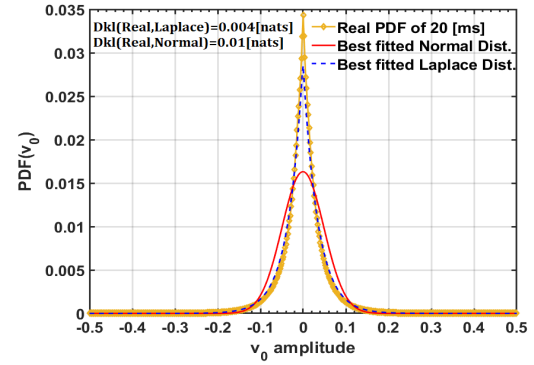


Figure 2: Distribution comparison of empirical PDF, best fitted normal and Laplace distribution. X axis - the signal sample amplitude, Y-axis the PDF of the sample amplitude. The Laplace distribution is the best fit to the empirical PDF

LSTM-based. Each backbone is used in the range for which it is currently the state of the art. We use dynamic mixing augmentation, introduced in [13], which consist of creating new mixtures in run time, noted as **DM**.

Table 1: Hyperparameters for the different number of speakers.

No. speakers	N	K	Params[M]
2	256	4	4.6
3	256	4	7.2
5	128	4	1.95
10	128	4	2.7

For all experiments, 5 iterations are conducted, where **SC** implies that the stopping criterion per sample is activated. Fig. 4 depicts the behavior of the criterion on a typical sample from the test set.

Transformer Backbone As seen from the upper bound, the current state-of-the-art results are close to the bound. Thus, only a small improvement is expected when experimenting above it with SepIt. The results using a current state-of-the-art Transformer network [3] as B over the Wall Street Journal Mix 2/3 (WSJ0-2/3 Mix) are summarized in Tab. 2.

First, we observe that the current state of the art is approach-




Figure 3: The upper bound for single channel source separation (blue curve). The upper bound decreases when the number of speakers C increases (Theorem 2.6). Furthermore, the score of SDR improvement for different separation algorithms is presented. SepIt, denoted by a pink triangle, produces state-of-the-art results.

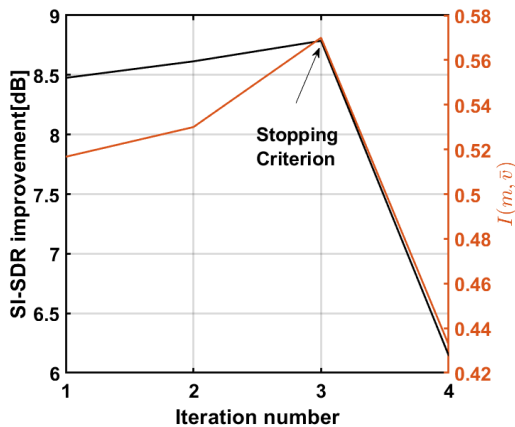


Figure 4: SDR improvement for 5 speakers and the Mutual information between the mixture and the estimated speaker. training stops when mutual information is no longer increases, evidently the SI-SDR improvement degrades too.

Table 2: The performance obtained by our method, as well as the Transformer backbone network. See text for details.

Method	SI-SDR _i [dB] (↑)	
	WSJ0-2Mix	WSJ0-3Mix
Upper Bound	23.1	21.2
(i) SepFormer	20.4	17.6
(ii) SepFormer + DM	22.3	19.8
(iii) SepIt + SepFormer + DM	22.4	20.1
(iv) SepIt + SepFormer + DM + SC	22.4	20.1

ing the upper bound. For 2 and 3 speakers, the gap between the upper bound and current results is 0.8dB and 1.4dB, respectively. Second, as expected, our SepIt model is able to improve only slightly the result with 2 speakers, with an improvement of 0.1dB over the current state of the art, while for 3 speakers, SepIt improves the result by 0.3dB. We note that for 2/3 speakers

only few examples had early stopping with the stopping criterion (SC), which resulted in similar results.

LSTM Backbone In LSTM-based architectures the state-of-the-art network for a large number of speakers was provided by [22]. We apply the SepIt network over the Libri5Mix and Libri10Mix datasets. The results are summarized in Tab. 3. As expected, here the improvement over the current state-of-the-art architecture is more prominent than in the Transformer case, where for 5 speakers, SepIt improves results by 1.0dB, and for 10 speakers by 0.5dB.

Table 3: The performance obtained by our method, as well as the LSTM backbone network. See text for details.

Method	SI-SDR _i [dB] (↑)	
	Libri5Mix	Libri10Mix
Upper Bound	14.5	12.0
Gated LSTM [22]	12.7	7.7
SepIt+ Gated LSTM	13.2	8.0
SepIt+ Gated LSTM + DM	13.6	8.1
SepIt+ Gated LSTM + DM + SC	13.7	8.2

5. Summary

In this work, a general upper bound for the Single Channel Speech Separation (SCSS) problem is derived. The upper bound is obtained using the Cramer-Rao bound, with fundamental modeling of the speech signal. We show that the gap between the current state-of-the-art network for 2 and 3 speakers and the obtained bound is relatively small, 0.8[dB] and 1.4[dB] respectively. This may indicate that future research should focus more on reducing the model size or the required training set and less on improving overall accuracy. The gap for 5 and 10 speakers is larger, and there is still room for substantial improvement in separation accuracy. Using the upper bound, a new neural network named SepIt is introduced. SepIt takes a backbone network estimation, and iteratively improves the estimation, until a stopping criterion based on the upper bound is met. SepIt is shown to outperform current state-of-the-art networks for 2, 3, 5 and 10 speakers.

6. References

- [1] Y. Luo, Z. Chen, and T. Yoshioka, "Dual-path rnn: efficient long sequence modeling for time-domain single-channel speech separation," 2020.
- [2] E. Nachmani, Y. Adi, and L. Wolf, "Voice separation with an unknown number of multiple speakers," 2020.
- [3] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, "Attention is all you need in speech separation," 2021.
- [4] J. Jensen, I. Batina, R. Hendriks, and H. Richard, "A study of the distribution of time-domain speech samples and discrete fourier coefficients," 01 2005.
- [5] E. Doron, A. Yeredor, and P. Tichavsky, "Cramér-rao-induced bound for blind separation of stationary parametric gaussian sources," *IEEE Signal Processing Letters*, vol. 14, no. 6, pp. 417–420, 2007.
- [6] J. Garofolo, D. Graff, D. Paul, and D. Pallett, "Csr-i (wsj0) complete ldc93s6a," *Web Download. Philadelphia: Linguistic Data Consortium*, vol. 83, 1993.
- [7] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "Librimix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [8] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 708–712.
- [9] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [10] Y. Luo and N. Mesgarani, "Tasnet: time-domain audio separation network for real-time, single-channel speech separation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 696–700.
- [11] —, "Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [12] Y. Luo, Z. Chen, and T. Yoshioka, "Dual-path rnn: efficient long sequence modeling for time-domain single-channel speech separation," *arXiv preprint arXiv:1910.06379*, 2019.
- [13] N. Zeghidour and D. Grangier, "Wavesplit: End-to-end speech separation by speaker clustering," *arXiv preprint arXiv:2002.08933*, 2020.
- [14] Z. Yao, W. Pei, F. Chen, G. Lu, and D. Zhang, "Stepwise-refining speech separation network via fine-grained encoding in high-order latent domain," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2022.
- [15] X. Hu, K. Li, W. Zhang, Y. Luo, J.-M. Lemercier, and T. Gerkmann, "Speech separation using an asynchronous fully recurrent convolutional neural network," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [16] H. Sahlin and U. Lindgren, "The asymptotic cramer-rao lower bound for blind signal separation," in *Proceedings of 8th Workshop on Statistical Signal and Array Processing*, 1996, pp. 328–331.
- [17] V. Kautsky, Z. Koldovsky, P. Tichavsky, and V. Zarzoso, "Cramér-rao bounds for complex-valued independent component extraction: Determined and piecewise determined mixing models," *IEEE Transactions on Signal Processing*, vol. 68, p. 5230–5243, 2020. [Online]. Available: <http://dx.doi.org/10.1109/TSP.2020.3022827>
- [18] N. Brunel and J.-P. Nadal, "Mutual information, fisher information, and population coding," *Neural computation*, vol. 10, pp. 1731–57, 11 1998.
- [19] C. R. Rao, *Information and the Accuracy Attainable in the Estimation of Statistical Parameters*. New York, NY: Springer New York, 1992, pp. 235–247. [Online]. Available: https://doi.org/10.1007/978-1-4612-0919-5_16
- [20] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr - half-baked or well done?" 2018.
- [21] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017.
- [22] S. Dovrat, E. Nachmani, and L. Wolf, "Many-speakers single channel speech separation with optimal permutation training," 2021.