

Fast accuracy estimation of deep learning based multi-class musical source separation

Alexandru Mocanu^{*1, 2}, Benjamin Ricaud^{1, 3}, and Milos Cernak²

¹LTS2, École Polytechnique Fédérale de Lausanne, Station 11, CH-1015 Lausanne, Switzerland

²Logitech Europe S.A., CH-1015, Lausanne, Switzerland

³Department of Physics and Technology, UiT The Arctic University of Norway, Tromsø NO-9037, Norway

Abstract

Music source separation represents the task of extracting all the instruments from a given song. Recent breakthroughs on this challenge have gravitated around a single dataset, MUSDB, only limited to four instrument classes. Larger datasets and more instruments are costly and time-consuming in collecting data and training deep neural networks (DNNs). In this work, we propose a fast method to evaluate the separability of instruments in any dataset without training and tuning a DNN. This separability measure helps to select appropriate samples for the efficient training of neural networks. Based on the oracle principle with an ideal ratio mask, our approach is an excellent proxy to estimate the separation performances of state-of-the-art deep learning approaches such as TasNet or Open-Unmix. Our results contribute to revealing two essential points for audio source separation: 1) the ideal ratio mask, although light and straightforward, provides an accurate measure of the audio separability performance of recent neural nets, and 2) new end-to-end learning methods such as Tasnet, that operate directly on waveforms, are, in fact, internally building a Time-Frequency (TF) representation, so that they encounter the same limitations as the TF based-methods when separating audio pattern overlapping in the TF plane.

1 Introduction

Music, along with speech, is a main point of interest for the field of audio processing. The extraction of individual instrument tracks from a song has been of great interest for a couple of decades. Most of the initial efforts were based on signal processing techniques and used unsupervised methods to exploit signal characteristics such as sparsity [11], harmonicity and inharmonicity [14], harmonic and percussive spectra [8, 2], beat [7] or repetitiveness [12, 3]. However, these methods do not separate the instruments but rather some sound patterns produced by them. With the rise of deep learning, the latest models [10, 1] provide a better separation quality and manage to extract the instruments with all their sound specificities. Nevertheless, a substantial inconvenience for the deep learning models comes from their supervised learning training procedure that requests large amounts of labeled data, namely songs and their separate instrument tracks.

The most widespread dataset used for training deep learning models for music source separation, MUSDB18 [13], only provides four instrument classes (vocals, drums, bass, and others). This is a significant limitation for the community. Creating a new dataset is time-consuming and may have a high cost. Moreover, after spending much time training deep networks on it, it may be inappropriate to train particular classes of instruments.

The proposed method is based on the Ideal Ratio Mask (IRM) oracle [17] and does not require any training. It outputs a separation score for each in-

*Work completed during an internship at Logitech.

strument class of the input tracks from a song and its separated track.

Our contribution is twofold. Firstly, we show that the oracle is a good indicator for source separation performance. We do this by comparing the separation score of the oracle to that of a trained proof-of-concept TasNet-based model. The results of the two models are strongly correlated. We tried the separation on electric guitar, which overlaps in the time-frequency domain with bass and vocals. It confirms that the separation of new instrument classes (beyond the MUSDB18 4 classes) can be obtained successfully with deep learning models. Secondly, by using a TasNet-based model and the IRM oracle, we show a connection between waveform-based and spectrogram-based models. The IRM looks at separability in terms of spectro-temporal overlap, and the Tasnet-based models operate directly on the waveform. In this context, a correlation between the two seemed much less trivial. There seems to be a more interesting relationship that indicates that even if TasNet is end-to-end learning, internally, it is a TF-based approach. Thus it cannot separate well sounds that are overlapping in the TF plane.

The structure of the paper is as follows: Section 2 introduces related work. Section 3 describes our proposed method, and Section 4 presents the experimental setup and obtained results. Section 5 concludes the paper.

2 Related work

Deep learning methods have been prevalent in music source separation in the last few years. Most of the models separate only the four classes provided in the MUSDB [13] dataset. Takahashi et al. [19] combine convolutional dense blocks and Long short-term memory blocks and apply them on multiple frequency bands to capture features specific to each band. This model achieved the best performance in the SiSEC 2018 [17] competition but used additional proprietary training data. Stöter et al. [18] proposed a recurrent model that trains a separate network for each instrument. They made the code open source and trained their model only on MUSDB to serve as a benchmark for future research.

Many deep learning models operate on the mag-

nitude spectrogram, as this provides a representation that is both intuitive and has proven helpful in several audio processing tasks. End-to-end systems, which also learn how to encode and decode the raw signal, have also attracted more attention. While Stoller et al. [16], who adapted a U-Net architecture, did not manage to outperform the best spectrogram-based models, Samuel et al. [15] managed to reach state-of-the-art performance. Their model reuses the temporal convolutional network from Conv-TasNet [10], for which it learns to generate parameters based on instrument embeddings. It also includes a complex encoder that leverages convolutions and spectrograms, and it operates in a multi-stage fashion, receiving as input mixture signals resampled at three different sampling rates. Deffosez et al. [1] also adapt the Conv-TasNet architecture and obtain even better performance than Meta-TasNet, a model with more parameters. These recent increases in performances are at the cost of more complex architectures that require more computing power.

There have also been a few attempts at separating music into more classes. Hennequin et al. [4] use a U-Net [6] architecture and add the piano class as well. However, they do not release the results for piano separation. Hung and Lerch [5] propose a multitask learning solution. They also detect instrument activity, extend the number of instruments to six, and add besides the piano, electric, and acoustic guitars. They use open-source data to train their models. Nevertheless, they use three large datasets of songs together with data augmentation, which requires more computing power.

3 Method

The IRM oracle [17] is typically used as a baseline when evaluating the performance of a source separation model. It acts as an upper bound for spectrogram-based models that apply masks on the spectrogram to perform the separation. The IRM is defined as follows: given the spectrograms $y_j(f, t, i)$ for source signals s_j , $1 \leq j \leq J$, where J is the number of sources, and (f, t) are the TF bin coordinates for channel i , the IRM separation mask for

source j is computed as:

$$M_j(f, t, i) = \frac{v_j(f, t, i)}{\sum_{j'=1}^J v_{j'}(f, t, i)} \quad (1)$$

where $v_j(f, t, i) = |y_j(f, t, i)|^\alpha$, $\alpha \geq 0$. In our case, we choose $\alpha = 2$ to use power spectrograms, which usually give the best performing oracle. Figure 1 depicts the structure of the oracle source separator. For a given song, we compute the spectrograms for both the mixture signal and the source signals. Then, we take the magnitude spectrograms for the sources and feed them to the IRM that outputs masks for each source. After that, the masks are applied to the mixture spectrogram, and finally, the estimated source spectrograms are converted into source signal estimates through an inverse short-time Fourier transform.

To test a configuration of instrument classes, we use the IRM oracle for performing the separation. Then we evaluate the following metrics:

$$\text{SDR} = 10 \log_{10} \frac{\|s\|^2}{\|e_{spat} + e_{interf} + e_{artif}\|^2} \quad (2)$$

$$\text{ISR} = 10 \log_{10} \frac{\|s\|^2}{\|e_{spat}\|^2} \quad (3)$$

$$\text{SIR} = 10 \log_{10} \frac{\|s + e_{spat}\|^2}{\|e_{interf}\|^2} \quad (4)$$

$$\text{SAR} = 10 \log_{10} \frac{\|s + e_{spat} + e_{interf}\|^2}{\|e_{artif}\|^2} \quad (5)$$

$$\text{SI-SDR} = 10 \log_{10} \frac{\|\alpha s\|^2}{\|\alpha s - \hat{s}\|^2}, \alpha = \frac{\hat{s}^T s}{\|s\|^2} \quad (6)$$

where s is the clean source signal, \hat{s} is the estimated source signal given by $\hat{s} = s + e_{spat} + e_{interf} + e_{artif}$ with e_{spat} the error due to spatial distortions, e_{interf} the error due to interference with other sources and e_{artif} the error due to artifacts. The Source to Distortion Ratio (SDR), Image to Spatial Distortion Ratio (ISR), Source to Interference Ratio (SIR), and Source to Artifacts Ratio (SAR) are computed using the *museval* Python package. In contrast, the Scale-Invariant SDR (SI-SDR) [9] is computed in the same way as the other metrics from *museval*, by first computing for each song the median score over one-second windows and then taking the median over the songs. The primary metric of interest is usually the SDR. However, as the relative scale of the real and estimated source signals influences

		SI-SDR	SDR	SIR	ISR	SAR
IRM	vocals	8.43	9.02	17.42	16.38	9.26
	drums	5.11	6.24	14.40	10.31	5.74
	bass	9.55	9.99	14.09	15.38	11.44
	eguitar	2.83	4.62	10.61	9.08	4.86
	wind	-2.76	1.06	-9.48	7.81	0.11
	other	3.45	5.04	10.43	9.24	5.24
Meta-TasNet	vocals	1.36	3.52	8.27	6.81	3.74
	drums	1.36	3.62	15.39	5.30	3.88
	bass	9.00	9.42	14.17	13.56	10.35
	eguitar	-6.82	0.46	-2.83	2.73	2.11
	wind	-34.55	-11.22	-18.90	0.38	0.66
	other	-3.77	0.92	-1.12	4.75	3.49

Table 1: IRM and Meta-TasNet scores on Karaoke Version dataset.

its value, we examine the SI-SDR when deciding what instrument classes to choose.

4 Experiments

4.1 Experimental Setup

We use our custom dataset for training and testing, obtained from downloading songs from the Karaoke Version website¹. We use 241 train, 46 validation, and 50 test songs sampled at 44.1 kHz, totaling 23 hours of music. We also experiment with training on smaller subsets of 24 songs. We included the 4 MUSDB classes for the experiments and added two new ones: electric guitar and wind instruments (like saxophone or flute). We generated equally loud source tracks for vocals, drums, bass, electric guitar, wind, and other instrument classes by mixing instrument signals appropriately for each of the songs.

To validate the IRM oracle’s role as an indicator for the performance of a deep learning model, we train a variation of Meta-TasNet² described in [15]. The original model consists of three stages: separate signals resampled at 8, 16, and 32 kHz and reuse information from the previous stage. We only keep the first stage, training a single network that separates signals resampled at 8 kHz. The rest of the network is left unchanged. At evaluation time, the separated 8 kHz signals are upsampled to 44.1

¹<https://www.karaoke-version.com>

²<https://github.com/pfnet-research/meta-tasnet>

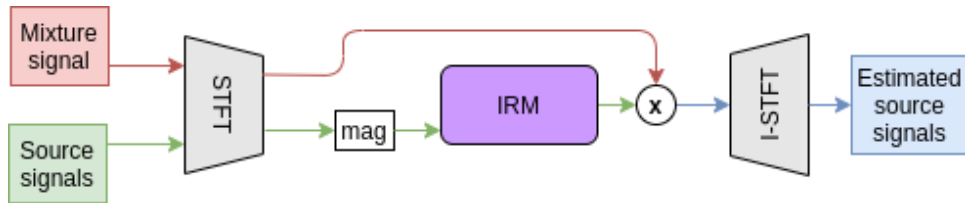


Figure 1: Oracle source separator based on Ideal Ratio Mask. The TF representations of both the mixture and source signals are computed using the Short Time Fourier Transform (STFT). Then, masks are applied from the magnitude (mag) of the source signals. The masked signals are represented back in the time domain with an Inverse STFT (I-STFT).

kHz, the original sampling rate of the data. We train our 8kHz-Meta-TasNet model for 80 epochs. The separation performances are lower than the full Meta-TasNet at 44.1 kHz (due to smaller network and downsampling to 8kHz, which lose the high-frequency components) but are still decent. It permits to obtain a good compromise between the network accuracy and computing time. Our goal is to estimate the ability to separate instruments and not to obtain the best separation possible.

4.2 Results

Table 1 shows the results of the IRM oracle applied to the Karaoke Version test set. We notice that the initial classes, vocals, drums, and bass, yield the best scores. However, compared to the results obtained on MUSDB [13], drums obtain an SI-SDR score of more than 3 dB lower. We still expect drums to be quite well separated on this dataset, so we assume we can separate instruments with scores that may be as low as about 2.5 dB. In this case, we predict that we can have decent separation for all classes except wind. In the wind instruments class, the problem originates from the severe interference with other instruments, as shown by the SIR score.

After running the Meta-TasNet model on the test set, we obtained the bottom of Table 1. Given the small size of the model, the scores are lower than for the full-size state-of-the-art model but still provide intuition for what would happen when scaling up the model size. For example, we notice that vocals, drums, and bass do indeed maintain excellent scores even in these conditions. In contrast, electric guitar and other classes still do better than having noise as loud as the clean signal, i.e., having an SDR score of 0 dB. On the other hand, the wind in-

		SI-SDR	SDR	SIR	ISR	SAR
Pearson	vocals	0.96	0.81	0.89	0.86	0.84
	drums	0.46	0.56	0.88	0.91	0.81
	bass	0.87	0.68	0.88	0.82	0.72
	eguitar	0.88	0.67	0.90	0.74	0.60
	wind	0.80	0.76	0.87	0.81	0.40
	other	0.71	0.47	0.66	0.57	0.61
Spearman	vocals	0.84	0.83	0.53	0.54	0.87
	drums	0.78	0.76	0.67	0.73	0.80
	bass	0.82	0.82	0.80	0.42	0.75
	eguitar	0.76	0.78	0.49	0.41	0.56
	wind	0.75	0.89	0.84	0.82	0.36
	other	0.69	0.57	0.37	0.51	0.49

Table 2: Correlations between IRM and Meta-TasNet scores.

struments suffer from interference that drags down the SDR score, just as predicted by IRM. Therefore, IRM indicates which instrument classes would have good or bad separation scores when using deep learning approaches.

We computed the Pearson and Spearman correlations of the two models, for each metric, on the scores obtained *on every individual song* from the test set. Table 2 shows these correlations between the results of IRM and Meta-TasNet on the 50 random test tracks for each of the instrument classes. In both cases, the correlations that we obtain for most of the metrics are strong (above 0.6) or very strong (above 0.8). This confirms that IRM serves as a good indicator for source separation performance on the dataset.

To further confirm the relationship between the scores provided by the IRM oracle and those obtained by a Meta-TasNet model, we trained on

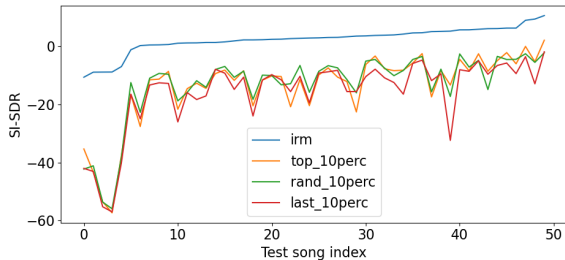


Figure 2: SI-SDR scores for electric guitar on the test set sorted by IRM performance for Meta-TasNet models trained on subsets of the training set.

samples consisting of 10% of the songs of the initial training set. We selected samples by three different criteria: the “best” songs (top_10perc), random songs (rand_10perc), the “worst” songs (last_10perc) according to the SI-SDR score of the IRM model. Figure 2 presents the SI-SDR scores for electric guitar on the test set for the IRM model and the Meta-TasNet model trained with the three different samples. The songs are sorted by the SI-SDR score of the IRM model. We see that no matter what data we train on, the results of the models still follow the trend predicted by the IRM model. Table 3 shows the correlations between the Meta-TasNet results and the IRM results, solidifying our statement. We can also note from Fig. 2 the close separability scores obtained from the different training subsets. It highlights the relative stability of the separation process when training on different training sets, with different time-frequency separabilities. We notice however, on average, a slightly smaller SI-SDR score for the neural network trained on the worst IRM separable samples (red curve).

An interesting aspect to note is that IRM is an oracle that uses a TF representation to compute its metrics, while the Meta-TasNet model learns directly from waveforms. This implies that IRM acts as a universal predictor, being useful for both TF-based and waveform-based models.

Hung and Lerch [5] choose the instrument categories based on their frequency of appearance in the dataset. A natural question is whether the IRM oracle would be a better indicator of the separation performance than the simple frequency of an instrument in the dataset. To analyze this situ-

	Pearson	Spearman
top_10perc	0.88	0.76
rand_10perc	0.88	0.70
last_10perc	0.85	0.69

Table 3: Correlations between IRM and Meta-TasNet SI-SDR scores on electric guitar.

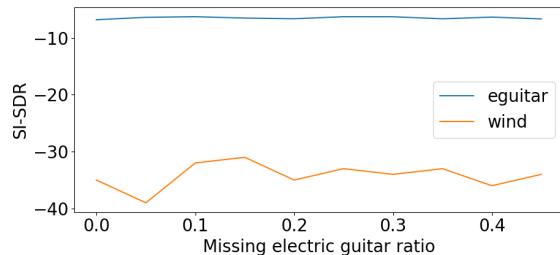


Figure 3: Meta-TasNet results for different ratios of muted electric guitar and wind instruments.

ation, we took the training set and muted some electric guitar tracks. The test set was kept intact. We then trained our Meta-TasNet model on this data. We repeated this process for ratios of muted tracks ranging from 0.00 to 0.45 with a step of 0.05. Fig. 3 displays the SI-SDR score of the Meta-TasNet model on the electric guitar and wind instrument classes. We see that even when muting the electric guitar considerably, the Meta-TasNet performance stays stable. The electric guitar scores maintain a large margin of about 28 dB compared to the wind instruments. For the wind instruments, no drop in performance was expected, as no tracks were muted. Since the IRM oracle is not a learning process, its performance is independent of the number of samples used in the training set. We thus confirmed that the IRM oracle is a much better indicator of separation performance than the frequency of an instrument in the dataset. This conclusion is valid for the Meta-TasNet architecture and within a reasonable minimal number of samples for each instrument category.

5 Conclusion

Increasing the size of the dataset or the network often leads to better performance for deep learning models. This is the actual trend, which implies

increasing computing resources not everyone can afford. We propose a new way for evaluating the separability of audio sources in recordings, prior to the training of any neural network. To assess the quality of samples, we take advantage of the TF masking approach and the possibility to make an oracle in this particular configuration.

We reveal the close relationship between mask-based learning approaches such as Meta-TasNet and the TF representation. Our results indicate that even if TasNet has the freedom to learn a latent space where instruments would be separated efficiently by a masking process, this space is no better than the TF representation. Hence, in order to improve the separation of sources overlapping in the TF plane, new architectures are needed.

6 Acknowledgements

We would like to thank Lucien Barret for helping us run the experiments on the muted electric guitar tracks.

References

- [1] A. Défossez, N. Usunier, L. Bottou, and F. Bach. Music source separation in the waveform domain. *arXiv preprint arXiv:1911.13254*, 2019. URL <https://arxiv.org/pdf/1911.13254>.
- [2] D. Fitzgerald. Harmonic/percussive separation using median filtering. In *Proceedings of the International Conference on Digital Audio Effects (DAFx)*, volume 13, 2010. URL https://www.dafx.de/paper-archive/2010/DAFx10/DerryFitzGerald_DAFx10_P15.pdf.
- [3] D. FitzGerald. Vocal separation using nearest neighbours and median filtering. 2012. doi:10.1049/ic.2012.0225.
- [4] R. Hennequin, A. Khlif, F. Voituret, and M. Moussallam. Spleeter: a fast and efficient music source separation tool with pre-trained models. *Journal of Open Source Software*, 5(50):2154, 2020. doi:10.21105/joss.02154.
- [5] Y.-N. Hung and A. Lerch. Multitask learning for instrument activation aware music source separation. *arXiv preprint arXiv:2008.00616*, 2020. URL https://program.ismir2020.net/static/final_papers/334.pdf.
- [6] A. Jansson, E. Humphrey, N. Montecchio, R. Bittner, A. Kumar, and T. Weyde. Singing voice separation with deep u-net convolutional networks. 2017. URL <https://openaccess.city.ac.uk/id/eprint/19289/1/>.
- [7] H. Kameoka, M. Nakano, K. Ochiai, Y. Imoto, K. Kashino, and S. Sagayama. Constrained and regularized variants of non-negative matrix factorization incorporating music-specific constraints. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5365–5368. IEEE, 2012. doi:10.1109/ICASSP.2012.6289133.
- [8] C. Laroche, M. Kowalski, H. Papadopoulos, and G. Richard. A structured nonnegative matrix factorization for source separation. In *2015 23rd European Signal Processing Conference (EUSIPCO)*, pages 2033–2037. IEEE, 2015. doi:10.1109/EUSIPCO.2015.7362741.
- [9] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey. Sdr-half-baked or well done? In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 626–630. IEEE, 2019. URL <https://ieeexplore.ieee.org/abstract/document/8683855/>.
- [10] Y. Luo and N. Mesgarani. Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation. *IEEE/ACM transactions on audio, speech, and language processing*, 27(8):1256–1266, 2019. doi:10.1109/TASLP.2019.2915167.
- [11] M. D. Plumbley, T. Blumensath, L. Daudet, R. Gribonval, and M. E. Davies. Sparse representations in audio and music: from coding to source separation. *Proceedings of the IEEE*, 98(6):995–1005, 2009. doi:10.1109/JPROC.2009.2030345.
- [12] Z. Rafii and B. Pardo. A simple music/voice separation method based on the

- extraction of the repeating musical structure. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 221–224. IEEE, 2011. doi:10.1109/ICASSP.2011.5946380.
- [13] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimi-lakis, and R. Bittner. The MUSDB18 corpus for music separation, Dec. 2017. URL <https://doi.org/10.5281/zenodo.1117372>.
- [14] F. Rigaud, A. Falaize, B. David, and L. Daudet. Does inharmonicity improve an nmf-based piano transcription model? In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 11–15. IEEE, 2013. doi:10.1109/ICASSP.2013.6637599.
- [15] D. Samuel, A. Ganeshan, and J. Naradowsky. Meta-learning extractors for music source separation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 816–820. IEEE, 2020. doi:10.1109/ICASSP40776.2020.9053513.
- [16] D. Stoller, S. Ewert, and S. Dixon. Wave-u-net: A multi-scale neural network for end-to-end audio source separation. *arXiv preprint arXiv:1806.03185*, 2018. URL http://ismir2018.ircam.fr/doc/pdfs/205_Paper.pdf.
- [17] F.-R. Stöter, A. Liutkus, and N. Ito. The 2018 signal separation evaluation campaign. In *International Conference on Latent Variable Analysis and Signal Separation*, pages 293–305. Springer, 2018. doi:10.1007/978-3-319-93764-9_28.
- [18] F.-R. Stöter, S. Uhlich, A. Liutkus, and Y. Mitsufuji. Open-unmix - a reference implementation for music source separation. *Journal of Open Source Software*, 4(41):1667, 2019. doi:10.21105/joss.01667. URL <https://doi.org/10.21105/joss.01667>.
- [19] N. Takahashi, N. Goswami, and Y. Mitsufuji. Mmdenselstm: An efficient combination of convolutional and recurrent neural networks for audio source separation. In *2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC)*, pages 106–110. IEEE, 2018. doi:10.1109/IWAENC.2018.8521383.