

# Musical Source Separation Bake-Off: Comparing Objective Metrics with Human Perception

Noah Jaffe, John Ashley Burgoyne

Institute for Logic, Language, and Computation, University of Amsterdam, The Netherlands

**Abstract**—Music source separation aims to extract individual sound sources (e.g., vocals, drums, guitar) from a mixed music recording. However, evaluating the quality of separated audio remains challenging, as commonly used metrics like the source-to-distortion ratio (SDR) do not always align with human perception. In this study, we conducted a large-scale listener evaluation on the MUSDB18 test set, collecting approximately 30 ratings per track from seven distinct listener groups. We compared several objective energy-ratio metrics, including legacy measures (BSSEval v4, SI-SDR variants), and embedding-based alternatives (Fréchet Audio Distance using CLAP-LAION-music, EnCodec, VGGish, Wave2Vec2, and HuBERT). While SDR remains the best-performing metric for vocal estimates, our results show that the scale-invariant signal-to-artifacts ratio (SI-SAR) better predicts listener ratings for drums and bass stems. Fréchet Audio Distance (FAD) computed with the CLAP-LAION-music embedding also performs competitively—achieving Kendall’s  $\tau$  values of 0.25 for drums and 0.19 for bass—matching or surpassing energy-based metrics for those stems. However, none of the embedding-based metrics, including CLAP, correlate positively with human perception for vocal estimates. These findings highlight the need for stem-specific evaluation strategies and suggest that no single metric reliably reflects perceptual quality across all source types. We release our raw listener ratings to support reproducibility and further research.

## 1. INTRODUCTION

Music source separation involves extracting individual sound sources—called stems—such as vocals, drums, and guitar, from a mixed music recording. Practical applications include karaoke or a reprocessing step for many music information retrieval (MIR) tasks. Music source separation systems are typically evaluated using objective metrics such as the source-to-distortion ratio (SDR), source-to-artifacts ratio (SAR), and source-to-interferences ratio (SIR) [1]. Despite their widespread use, these energy-ratio metrics often correlate poorly with human perception, leading to discrepancies between objective evaluations and listener judgments [2]–[4]. Recent work attempted to address shortcomings by proposing scale-invariant alternatives, such as the scale-invariant signal-to-distortion ratio (SI-SDR) and scale-invariant signal-to-artifacts ratio (SI-SAR) [5]. However, even with these improvements, a significant gap remains between numerical scores and listener perception.

In evaluating the performance of source separation systems, a set of original, un-mixed sources – known as stems or ground truths – are needed for comparison with the estimated stems. A state-of-the-art musical source separation system for pop music typically produces four estimates: *vocals*, *bass*, *drums*, and *other*—encompassing remaining instruments such as guitar or piano.

In this work, we present a comprehensive large-scale listening study that re-evaluates a range of objective metrics on the MUSDB18 test dataset. Our study includes several separation systems covering both legacy methods (e.g., REP1) and state-of-the-art deep-learning models (e.g., Open-Unmix, HTDemucs-ft, SCNet-large), as well as an oracle method (IRM1). Using a webMUSHRA test [6], we collected approximately 30 ratings per track from participants with a minimum of two years of musical experience.

The key contributions of this paper are as follows.

- We conduct a robust, large-scale listener evaluation.

- We compare legacy metrics (e.g., SDR, SI-SDR) with Fréchet Audio Distance (FAD) computed using different embeddings.
- We release the complete raw data from the listening study, promoting transparency and enabling future meta-analyses.

## 2. OVERVIEW OF METRICS

### 2.1. Energy-ratio metrics

Musical source separation systems are evaluated using a dataset of *stems*, which are then mixed together to create the final recording. Objective energy-ratio metrics such as SDR and its variants quantify the extent to which an estimate differs from its ground truth stem. The estimate  $\hat{s}$  is decomposed into its ground-truth content  $s_{\text{target}}$  and residual error  $e_{\text{residual}}$ .

$$\hat{s} = s_{\text{target}} + e_{\text{residual}} \quad (1)$$

The error term is decomposed further based on its origin and can be used to calculate metrics such as the *sources to artifacts ratio* (SAR) and *signal to interference ratio* (SIR).

For traditional SDR, SIR, and SAR in BSSEval v4,<sup>1</sup>

$$e_{\text{residual}} = e_{\text{spatial}} + e_{\text{interference}} + e_{\text{artifact}} \quad (2)$$

These components form the basis for the log ratio for source-to-distortion ratio (SDR) [1]:

$$\text{SDR} = 10 \log_{10} \frac{\|s_{\text{target}}\|^2}{\|e_{\text{spatial}} + e_{\text{interference}} + e_{\text{artifact}}\|^2} \quad (3)$$

As expressed in (3), SDR compares the energy originating from the original, un-mixed stem with noise sources present in the extracted estimate [1].

In Le Roux et al.’s SI-SDR paper, the authors demonstrate that standard SDR is highly sensitive to gain changes: rescaling an estimate by any constant factor (without altering its perceptual quality) still yields different SDR scores [5]. To avoid this, they introduced scale-invariant variants (SI-SDR, SI-SIR, SI-SAR) that normalize out signal energy differences so the metric reflects only perceptual fidelity rather than arbitrary amplitude scaling. SI-SDR also addresses shortcomings in BSSEval’s time-invariant 512-tap decomposition filter, which in traditional SDR can mask errors by placing spectral nulls and thus forgive distortions in those bands. These metrics, computed on mono tracks, decompose the error term exclusively into interference and artifact components, as illustrated in (4) relative to (2). In particular, SI-SDR introduces an explicit scaling step: the optimal scaling factor,  $\alpha$ , defined as  $\alpha = \hat{s}^T s / \|s\|^2$  from which the scaled reference is defined as  $e_{\text{target}} = \alpha s$ . The estimate is then split as  $\hat{s} = e_{\text{target}} + e_{\text{res}}$ , where  $e_{\text{res}}$  is the residual error, yielding the expanded SI-SDR (and by extension SI-SAR, SI-SIR) formulas as seen in (5), (6), and (7).

$$e_{\text{residual}} = e_{\text{interference}} + e_{\text{artifact}} \quad (4)$$

<sup>1</sup><https://github.com/sigsep/sigsep-mus-eval>

$$\text{SI-SDR} = 10 \log_{10} \frac{\|e_{\text{target}}\|^2}{\|e_{\text{interference}} + e_{\text{artifact}}\|^2} \quad (5)$$

$$\text{SI-SIR} = 10 \log_{10} \frac{\|e_{\text{target}}\|^2}{\|e_{\text{interference}}\|^2} \quad (6)$$

$$\text{SI-SAR} = 10 \log_{10} \frac{\|e_{\text{target}}\|^2}{\|e_{\text{artifact}}\|^2} \quad (7)$$

In addition, SI-SIR (6) and SI-SAR (7) are simplified and made orthogonal—original formulas for SAR contained noise error terms in the numerator, which made the metric problematic.

Although SI-SDR has seen some adoption, SDR remains the de facto metric used in music source separation evaluation.

## 2.2. Perception-informed metrics

A first effort to create a listener-informed objective metric was the development of the Perceptual Evaluation methods for Audio Source Separation (PEASS) toolkit in 2010 [7]. Implemented in MATLAB, PEASS uses nonlinear neural networks to generate multiple perceptual scores aimed at predicting human judgments of separated audio signals. Although innovative in its attempt to capture complex, perceptually relevant aspects of audio quality, PEASS has not achieved widespread adoption.

The Fréchet Audio Distance (FAD) was originally introduced by Kilgour et al. [8] as a novel deep-learning metric that compares the statistical distributions of audio embeddings, rather than merely relying on an objective energy-based ratio. This approach can theoretically offer a better assessment of perceptual audio quality because it captures subtle, high-level characteristics that energy metrics might overlook. In 2024, Microsoft Research augmented FAD with `fadt_k` [9],<sup>2</sup> which extends its application beyond the original VGGish [10] embedding to include alternatives like CLAP-LAION-music [11], EnCodec [12], Wave2Vec2 [13], and HuBERT [14]. Using an embedding model pretrained on audio more similar to the target material—such as music (CLAP, EnCodec) versus speech (HuBERT, Wave2Vec2)—can lead to more accurate representations of relevant features, potentially improving correlation with human perception and the reliability of quality assessments.

## 3. METHODOLOGY

### 3.1. Listener Study

To obtain perceptual ground truth for evaluating objective metrics, we conducted a listener study. Participants evaluated every stem from every track in the MUSDB18 test dataset (50 tracks). Our study builds on the work conducted in the 2018 Signal Separation Evaluation Campaign (SiSEC2018), which established standard benchmarks and evaluation procedures for music source separation [15]. For each track, a 10-second fragment was extracted beginning at the start time specified in the SiSEC2018 cutlist<sup>3</sup>.

We use estimates spanning a range of expected quality levels produced by several separation systems. Estimates were generated using state-of-the-art models such as Hybrid Transformer Demucs (HT-Demucs) [16] and Sparse Compression Network (SCNet) [17]. In addition, estimates were produced using Open-Unmix [18], while the SiSEC2018 submissions<sup>4</sup> provided estimates for the Ideal Ratio

Mask (IRM1)—an oracle method [15]—and the Repeating Pattern Extraction Technique (REPET)<sup>5</sup>—a legacy BSS method [19].

To reduce listener fatigue and ensure high-quality responses, we conducted the evaluation online using the webMUSHRA platform [20] in seven separate batches. In contrast to the official MUSHRA paradigm, we employed non-expert listeners using their own audio equipment—with clear instructions to use headphones in a quiet environment—to achieve a cost-effective yet robust evaluation. Previous studies have shown that online testing using webMUSHRA is a reliable alternative to traditional in-person evaluations, even in informal or non-laboratory settings [4], [21]–[23]. Each participant rated the estimates during a session lasting approximately 30 minutes and received compensation of €5. Participants were recruited via Prolific and were required to have at least one year of musical experience, such as playing an instrument. For each participant, the presentation of individual estimates was randomly shuffled to mitigate ordering effects and reduce potential bias.

We implemented four quality checks on each listener’s data: (1) the difference between the hidden reference and the hidden low-pass filtered anchor had to exceed 10 points (on a 0–100 scale), (2) the reference score had to be at least 90, (3) the standard deviation of all ratings given by a single user had to be at least 20, and (4) the time spent rating all estimates of a single stem had to fall within a reasonable time range (e.g., 20–213 seconds). In our study, the hidden reference is the original, unmodified audio sample provided to listeners, while the hidden anchor is a degraded version of the reference, created by applying a low-pass filter. Specifically, for vocals, drums, and other stems, the anchor was produced using a 3500 Hz low-pass filter; for bass, a 175 Hz low-pass filter was used to ensure that the difference would be clearly perceptible to listeners. These criteria were selected to ensure participants were attending to key quality differences and to screen out inattentive responses.

A total of 5,889 stem ratings were collected with the following quality check violation distribution: 2,371 ratings (40%) had zero violations, 1,778 ratings (30%) had one violation, 1,284 ratings (22%) had two violations, 396 ratings (7%) had three violations, and 60 ratings (1%) had four violations. We chose to include data for which two or fewer quality checks failed; furthermore, we repeated our experiments with strict quality control (no failed quality checks) and found that the overall findings remained unchanged. After applying our quality control criteria, approximately 8% of the data was dropped.

### 3.2. Correlation Studies

In line with [24], [25], we computed Kendall’s  $\tau$  for each user’s 0–100 ranking of estimates for a given stem and compared these with the corresponding scores produced by a specific objective metric (e.g., SDR). Kendall’s  $\tau$  is simply the proportion of concordant pairs between subjective and objective rankings, thereby providing a robust and averageable measure of rank correlation. By calculating Kendall’s  $\tau$  on a per-user, per-stem basis, we effectively mitigate confounding factors such as mix bias and individual listener idiosyncrasies. The Kendall’s  $\tau$  values were then averaged across the stem-type classes (Vocals, Drums, Bass, and Other) to obtain an overall measure for each metric.


To further explore the influence of error weighting on the correlation between objective metrics and listener expectations, we performed a grid search over the error components of SI-SDR. In our approach, we varied the weight assigned to the interference error (SI-SIR) while assigning its complement to the artifact error (SI-SAR), ensuring

<sup>2</sup>[https://github.com/microsoft/fadt\\_k](https://github.com/microsoft/fadt_k)

<sup>3</sup><https://github.com/sigsep/sigsep-mus-cutlist-generator>

<sup>4</sup><https://zenodo.org/records/1256003>

<sup>5</sup>Only vocals estimates were used from REPET.



**Fig. 1:** Rating distribution of state-of-the-art systems, by stem type. State-of-the-art models often struggle to produce bass estimates that match the reference, in part because bass instruments, while containing some high-frequency content, are easily masked by other sounds. This makes them especially vulnerable to interference in dense musical textures, leading to noticeably degraded estimates compared to the reference.

that the sum of weights was always 1. For each weight pair, we computed the energy of the interference and artifact errors, combined these energies geometrically using the assigned weights, and then calculated a re-weighted ratio in decibels as described in (??). Finally, we computed Kendall’s  $\tau$  between these re-weighted ratios and the listener rankings, thereby quantifying how variations in the weighting of error components correlate with human perception.

$$\text{Reweighted SI-SDR} = 10 \log_{10} \left( \frac{\|e_{\text{target}}\|^2}{e_{\text{interference}}^w \cdot e_{\text{artifact}}^{1-w}} \right)$$


For Fréchet Audio Distance, we computed the distance between each stem estimate and its ground truth. Because smaller distances indicate greater similarity, we inverted the distances before calculating Kendall’s  $\tau$  against listener scores.

## 4. RESULTS AND DISCUSSION


### 4.1. Perceptual Rankings of Separation Systems

Figure 1 shows that listeners rate bass estimates produced by state-of-the-art systems (i.e., SCNet-large and HTDemucs-ft) as lower in quality than those for the other stems. We acknowledge that the bass estimates exhibit limited variability compared to other stems—such as vocals or drums—where the best estimates achieve higher scores and the worst are much lower. This suggests that current separation systems are inherently less effective at producing high-quality bass outputs. One contributing factor is that bass instruments, while containing some high-frequency content, are easily masked by other sources in the mix. Additionally, because the bass often functions as part of the rhythm section, its note onsets frequently coincide with broad-spectrum transients from instruments like drums and rhythm guitar, making accurate isolation particularly challenging.

Figure 2 shows that listeners have a clear preference for estimates produced by state-of-the-art systems, compared to legacy ones. The results confirm that listeners reliably distinguish between separation quality tiers: hidden references are rated highest, followed by state-of-the-art systems (SCNet-large, HTDemucs-ft), then legacy methods (Open-Unmix, REP1). This suggests that even in a web-based setting, participants can meaningfully discriminate separation quality.



**Fig. 2:** Rating distribution of all stems, by separation system. Listeners clearly distinguish between references, legacy methods such as Open-Unmix and REP1, and state-of-the-art systems such as SCNet-large and HTDemucs-ft; surprisingly, the oracle method IRM1 does not outperform the state-of-the-art systems.



**Fig. 3:** Reweightings of the scale-invariant SDR error terms (SI-SIR vs. SI-SAR). Kendall’s  $\tau$  decreases almost monotonically as weight shifts from artifact toward interference for all stems except bass, indicating interference errors better predict perception; bass remains flat, likely due to the narrow quality range of bass estimates.

Surprisingly, the IRM1 oracle method does not outperform state-of-the-art models. While IRM1 is designed to produce ideal ratio masks based on ground truth data, its outputs may sound overly filtered or musically unnatural—limiting perceived quality. This challenges the assumption that oracle methods reflect an upper bound of perceptual performance, and reinforces the importance of perceptually grounded evaluation methods in music source separation research. IRM1 may benefit from additional post-processing to improve perceptual plausibility, which wasn’t applied here.

### 4.2. Error-Terms Correlation

Figure 3 shows how Kendall’s  $\tau$  changes as we sweep the weight from an interference-only error term (SI-SIR) to an artifact-only error term (SI-SAR) in our reweighted SI-SDR formula (Eq.(??)). For

**Table 1:** Kendall’s  $\tau$  for BSSEval v4

Stem Type	SDR	ISR	SAR	SIR
Vocals	<b>0.316</b>	0.262	0.258	0.236
Drums	0.165	<b>0.169</b>	0.124	0.067
Bass	0.086	0.058	<b>0.181</b>	0.093
Other	<b>0.273</b>	0.176	0.199	0.213
<b>Average</b>	<b>0.211</b>	0.167	0.190	0.152

**Table 2:** Kendall’s  $\tau$  for SD-SDR, SI-SDR, SI-SAR, and SI-ISR

Stem Type	SD-SDR	SI-SDR	SI-SAR	SI-ISR
Vocals	0.150	0.197	<b>0.246</b>	0.007
Drums	0.133	0.203	<b>0.240</b>	0.017
Bass	0.080	0.084	<b>0.116</b>	0.084
Other	0.261	<b>0.277</b>	0.271	0.174
<b>Average</b>	0.156	0.190	<b>0.218</b>	0.071

Vocals, Drums, and Other stems,  $\tau$  decreases almost monotonically. In contrast, the Bass curve is nearly flat, reflecting the lack of quality variation in our bass estimates. These findings indicate that artifact errors drive listener judgments more strongly than interference errors.

### 4.3. Legacy Metrics

Table 1 shows the correlation between BSSEval v4 metrics and human ratings from our listening study. Notably, traditional SDR is the most predictive metric for vocal estimates in all that we tested in this study. We attribute this to the exceptionally high quality of vocal separation in state-of-the-art systems: spatial artifacts and gain mismatches—captured by BSSEval v4’s formulation—become more perceptually salient as the estimates improve. In contrast, scale-invariant (SI-SDR) and scale-dependent SDR (SD-SDR) do not account for spatial error due to their mono-only implementation, likely reducing their effectiveness for evaluating vocals.

While SDR also performs best on average across all stems, some alternative metrics show better alignment for specific sources. For bass, SAR correlates more strongly with listener ratings than any other BSSEval v4 metric, suggesting that perceived artifacts are a primary driver of quality judgments in this stem.

In contrast, Table 2 presents correlations for scale-invariant metrics as implemented in the Northwestern University Source Separation Library (nussl) [26]. While SI-SAR shows improved performance overall—particularly for instruments—it falls short of traditional SDR for vocals, reinforcing the importance of scale and spatial fidelity in that stem. Interestingly, for drums and bass, SI-SAR outperforms both SDR and SAR from BSSEval, indicating its utility for sources where distortion and interference dominate perceptual quality. This is likely because SI-SAR excludes interference and noise from the numerator—unlike traditional SAR—making it a more focused and predictive measure of artifact-related degradation.

Overall, these comparisons underscore that no single metric uniformly captures human perception across all stem types. While traditional SDR remains the most faithful for vocals, modern metrics like SI-SAR offer better alignment for instrument stems—particularly when masking and artifacts dominate perceptual judgments.

### 4.4. Fréchet Audio Distance

None of the evaluated embeddings produced a positive Kendall’s  $\tau$  for vocal estimates. We initially hypothesized that embeddings trained on speech, such as HuBERT and Wave2Vec2, might better capture degradations in vocal stems—but this was not supported by our results. However, CLAP-LAION-music consistently outperformed others on instrument stems—substantially improving over VGGish, the default

**Table 3:** Kendall’s  $\tau$  for Various Embeddings by Stem Type

CLAP-LAION-music outperforms all other embeddings across every stem type. However, none of the embeddings yield positive correlation with listener ratings for vocal estimates.

Stem Type	CLAP-LAION-music	EnCodec	VGGish	Wave2Vec2	HuBERT
Vocals	0.000	-0.092	-0.129	-0.096	-0.118
Drums	<b>0.253</b>	0.130	0.046	-0.010	0.114
Bass	<b>0.188</b>	0.151	0.098	0.056	0.064
Other	<b>0.198</b>	0.017	0.099	0.022	0.073
<b>Average</b>	<b>0.160</b>	0.052	0.029	-0.008	0.034

embedding in FAD. This aligns with recent work showing that CLAP-based FAD correlates better with perceptual audio quality than its VGGish-based counterpart [27], [28]. For drums, a Kendall’s  $\tau$  of 0.25 matches the best-performing energy metric (SI-SAR); for bass,  $\tau = 0.19$  is on par with SAR and well above standard SDR. For the other estimates, however, energy-based metrics (SDR, SI-SDR, SI-SAR) yield Kendall’s  $\tau$  values around 0.27—slightly outperforming the 0.20 achieved by the CLAP-LAION-music embedding.

While VGGish captures aggregated mel-spectral textures, CLAP adds high-level semantic content, HuBERT focuses on phonetic patterns, and Wave2Vec2 is tuned to phonetic articulations. However, none of these embeddings necessarily capture errors related to spatial distortion, interference, or artifacts. Systematic ablations—for example, masking embedding dimensions associated with pitch or noise—could help identify the perceptual drivers within the embedding space and guide improvements to FAD for source separation evaluation.

## 5. CONCLUSION

Our work builds on existing literature by conducting a large-scale listener study on the MUSDB18 dataset, comparing a diverse set of separation systems—including state-of-the-art deep-learning models (e.g., HTDemucs-ft, SCNet-large), an oracle method (IRM1), and legacy methods (REP1, Open-Unmix)—against a range of objective metrics. Our comprehensive evaluation, using a webMUSHRA-based online test with carefully implemented quality checks and rigorous statistical analysis, confirms that traditional metrics such as SDR often fail to capture perceptual quality. In particular, we found that the scale-invariant signal-to-artifacts ratio (SI-SAR) correlates best with human ratings across instrument stems, especially for drums and bass, where it either matches or surpasses legacy metrics.

Furthermore, our grid search over the weighting parameters of SI-SDR revealed that artifact errors are more highly correlated with listener perception than interference errors. Our analysis of Fréchet Audio Distance embeddings shows that, while the CLAP-LAION-music embedding yields promising results for instrument stems—with drum estimates achieving Kendall’s  $\tau = 0.25$  and bass estimates  $\tau = 0.19$ , on par with the best bass metric (SAR)—none of the embeddings produced positive correlations for vocal estimates. This highlights that the embeddings used by FAD do not capture the perceptually relevant details of vocal quality, and further refinement of perceptual metrics is necessary.

Overall, our findings underscore the critical importance of using robust, listener-informed evaluations to benchmark objective metrics in music source separation. By releasing our complete raw listener data, we aim to promote transparency and support future meta-analyses that can further improve both evaluation techniques and separation technologies. All data used in this study are publicly available at <https://doi.org/10.5281/zenodo.15843081>

## REFERENCES

- [1] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [2] E. Cano, D. FitzGerald, and K. Brandenburg, "Evaluation of quality of sound source separation algorithms: Human perception vs quantitative metrics," in *24th European Signal Processing Conference (EUSIPCO)*, 2016, pp. 1758–1762.
- [3] V. Emiya, E. Vincent, N. Harlander, and V. Hohmann, "Subjective and objective quality assessment of audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2046–2057, 2011.
- [4] E. Rumbold, G. Tzanetakis, and B. Pardo, "Correlations between objective and subjective evaluations of music source separation," 2024. [Online]. Available: [https://smcnetwork.org/smc2024/papers/SMC2024\\_paper\\_id153.pdf](https://smcnetwork.org/smc2024/papers/SMC2024_paper_id153.pdf)
- [5] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR – half-baked or well done?" in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 626–630.
- [6] International Telecommunication Union, "Method for the subjective assessment of intermediate quality level of coded audio," ITU-R, Tech. Rep., 2015.
- [7] V. Emiya, E. Vincent, N. Harlander, and V. Hohmann, "Subjective and objective quality assessment of audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2046–2057, 2011, iNRIA Research Report: inria-00567152. [Online]. Available: <https://inria.hal.science/inria-00567152/document>
- [8] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, "Fréchet audio distance: A metric for evaluating music enhancement algorithms," in *Proceedings of Interspeech 2019*, Graz, Austria, 2019. [Online]. Available: [https://www.isca-archive.org/interspeech\\_2019/kilgour19\\_interspeech.pdf](https://www.isca-archive.org/interspeech_2019/kilgour19_interspeech.pdf)
- [9] A. Gui, H. Gamper, S. Braun, and D. Emmanouilidou, "Adapting Fréchet audio distance for generative music evaluation," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 1331–1335.
- [10] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson, "Cnn architectures for large-scale audio classification," 2017. [Online]. Available: <https://arxiv.org/abs/1609.09430>
- [11] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov, "Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [12] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," 2022, arXiv preprint arXiv:2210.13438. [Online]. Available: <https://arxiv.org/abs/2210.13438>
- [13] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," 2020. [Online]. Available: <https://arxiv.org/abs/2006.11477>
- [14] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," 2021. [Online]. Available: <https://arxiv.org/abs/2106.07447>
- [15] F.-R. Stöter, A. Liutkus, and N. Ito, "The 2018 signal separation evaluation campaign," 2018. [Online]. Available: <https://arxiv.org/abs/1804.06267>
- [16] S. Rouard, F. Massa, and A. Défossez, "Hybrid transformers for music source separation," 2022. [Online]. Available: <https://arxiv.org/abs/2211.08553>
- [17] W. Tong, J. Zhu, J. Chen, S. Kang, T. Jiang, Y. Li, Z. Wu, and H. Meng, "SCNet: Sparse compression network for music source separation," 2024. [Online]. Available: <https://arxiv.org/abs/2401.13276>
- [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*, vol. 4, no. 41, p. 1667, 2019. [Online]. Available: <https://doi.org/10.21105/joss.01667>
- [19] Z. Rafii and B. Pardo, "Repeating pattern extraction technique (REPET): A simple method for music/voice separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 1, p. 71, 2013.
- [20] M. Schoeffler, F.-R. Stöter, B. Edler, and J. Herre, "Towards the next generation of web-based experiments: A case study assessing basic audio quality following the ITU-R recommendation BS.1534 (MUSHRA)," in *Proceedings of the 1st Web Audio Conference (WAC)*. Paris, France: IRCAM, 2015. [Online]. Available: <http://webaudioconf.com/>
- [21] E. Gusó, J. Pons, S. Pascual, and J. Serrà, "On loss functions and evaluation metrics for music source separation," 2022. [Online]. Available: <https://arxiv.org/abs/2202.07968>
- [22] M. Cartwright, B. Pardo, G. J. Mysore, and M. Hoffman, "Fast and easy crowdsourced perceptual audio evaluation," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 619–623.
- [23] D. Ward, R. D. Mason, C. Kim, F.-R. Stöter, A. Liutkus, and M. D. Plumbley, "SISEC 2018: State of the Art in Musical Audio Source Separation – Subjective Selection of the Best Algorithm," in *Proceedings of the 4th Workshop on Intelligent Music Production (WIMP)*. Huddersfield, UK: University of Huddersfield, 2018.
- [24] M. Cartwright, B. Pardo, and G. J. Mysore, "Crowdsourced pairwise-comparison for source separation evaluation," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 606–610.
- [25] M. Torcoli, T. Kastner, and J. Herre, "Objective measures of perceptual audio quality reviewed: An evaluation of their application domain dependence," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1530–1541, 2021. [Online]. Available: <https://doi.org/10.1109/TASLP.2021.3069302>
- [26] E. Manilow, P. Seetharaman, and B. Pardo, "The Northwestern University Source Separation Library," in *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR)*. Paris, France: ISMIR, 2018, pp. 297–305, northwestern University.
- [27] S. Deshmukh, D. Alharthi, B. Elizalde, H. Gamper, M. A. Ismail, R. Singh, B. Raj, and H. Wang, "PAM: Prompting audio-language models for audio quality assessment," in *Proceedings of Interspeech 2024*, Kos, Greece, 2024, pp. 3320–3323.
- [28] S. Hwang, S. Kang, K. Kim, S. Ahn, and K. Lee, "Dose: Drum one-shot extraction from music mixture," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.