



Music Demixing Challenge 2021

Yuki Mitsufuji^{1*}, Giorgio Fabbro², Stefan Uhlich², Fabian-Robert Stöter²,
Alexandre Défossez³, Minseok Kim⁴, Woosung Choi⁴, Chin-Yun Yu⁵ and Kin-Wai Cheuk⁶

¹R&D Center, Sony Group Corporation, Tokyo, Japan, ²R&D Center, Sony Europe B.V., Stuttgart, Germany, ³Facebook AI Research, Paris, France, ⁴Korea University, Seoul, South Korea, ⁵Independent Researcher, Taipei, Taiwan, ⁶Information Systems and Technology Design, Singapore University of Technology and Design, Singapore, Singapore

Music source separation has been intensively studied in the last decade and tremendous progress with the advent of deep learning could be observed. Evaluation campaigns such as MIREX or SiSEC connected state-of-the-art models and corresponding papers, which can help researchers integrate the best practices into their models. In recent years, the widely used MUSDB18 dataset played an important role in measuring the performance of music source separation. While the dataset made a considerable contribution to the advancement of the field, it is also subject to several biases resulting from a focus on Western pop music and a limited number of mixing engineers being involved. To address these issues, we designed the Music Demixing Challenge on a crowd-based machine learning competition platform where the task is to separate stereo songs into four instrument stems (Vocals, Drums, Bass, Other). The main differences compared with the past challenges are 1) the competition is designed to more easily allow machine learning practitioners from other disciplines to participate, 2) evaluation is done on a hidden test set created by music professionals dedicated exclusively to the challenge to assure the transparency of the challenge, i.e., the test set is not accessible from anyone except the challenge organizers, and 3) the dataset provides a wider range of music genres and involved a greater number of mixing engineers. In this paper, we provide the details of the datasets, baselines, evaluation metrics, evaluation results, and technical challenges for future competitions.

OPEN ACCESS

Edited by:

Sebastian Ewert,
Spotify GmbH, Germany

Reviewed by:

Marius Miron,
Pompeu Fabra University, Spain
Hongqing Liu,
Chongqing University of Posts and
Telecommunications, China

*Correspondence:

Yuki Mitsufuji
Yuhki.Mitsufuji@sony.com

Specialty section:

This article was submitted to
Signal Processing Theory,
a section of the journal
Frontiers in Signal Processing

Received: 03 November 2021

Accepted: 23 December 2021

Published: 28 January 2022

Citation:

Mitsufuji Y, Fabbro G, Uhlich S,
Stöter F-R, Défossez A, Kim M,
Choi W, Yu C-Y and Cheuk K-W (2022)
Music Demixing Challenge 2021.
Front. Sig. Proc. 1:808395.
doi: 10.3389/frsip.2021.808395

Keywords: music source separation, music demixing, machine learning challenge, hybrid models, remixing, karaoke

1 INTRODUCTION

Audio source separation has been studied extensively for decades as it brings benefits in our daily life, driven by many practical applications, e.g., hearing aids, denoising in video conferences, etc. Additionally, music source separation (MSS) attracts professional creators because it enables the remixing or reviving of songs to a level, never achieved with conventional approaches such as equalizers. Further, suppressing vocals in songs can improve the experience of a karaoke application, where people can enjoy singing together on top of the original song (where the vocals were suppressed), instead of relying on content developed specifically for karaoke applications. Despite the potential benefits, the research community struggled to achieve a separation quality required by commercial applications. These demanding requirements were also aggravated by the under-determined settings the problem was formulated in since the number of provided channels in the audio recording is less than the number of sound objects that need to be separated.

In the last decade, the separation quality of MSS has mainly been improved owing to the advent of deep learning. A significant improvement in an MSS task was observed at the Signal Separation Evaluation Campaign (SiSEC) 2015 (Ono et al., 2015), where a simple feed-forward network (Uhlich et al., 2015) able to perform four-instruments separation achieved the best signal-to-distortion ratio (SDR) scores, surpassing all other methods that did not use deep learning. The use of deep learning in MSS was accelerated ever since and led to improved SDR results year after year in the successive SiSEC editions, held in 2016 (Liutkus et al., 2017) and 2018 (Stöter et al., 2018). An important component of this success story was the release of publicly available datasets such as Rafii et al. (2017) which, compared to previous datasets such as Bittner et al. (2014), was created specifically for MSS tasks. MUSDB18 consists of 150 music tracks in four stems and is up until now widely used due to a lack of alternatives¹. The dataset also has a number of limitations such as its limited number of genres (mostly pop/rock) and its biases concerning mixing characteristics (most stems were produced by the same engineers). Since the last evaluation campaign took place, many new papers were published claiming state-of-the-art, based on MUSDB18 test data, however, it is unclear if generalization performance did improve at the same pace or if some models overfit on MUSDB18. To keep scientific MSS research relevant and sustainable, we want to address some of the limitations of current evaluation frameworks by using:

- a fully automatic evaluation system enabling straightforward participation for machine learning practitioners from other disciplines.
- a new professionally produced dataset containing unseen data dedicated exclusively to the challenge to ensure transparency in the competition (i.e., the test set is not accessible from anyone except the challenge organizers).

With these contributions, we designed a new competition called Music Demixing (MDX) Challenge², where a call for participants was conducted on a crowd-based machine learning competition platform. A hidden dataset crafted exclusively for this challenge was employed in a system that automatically evaluated all MSS systems submitted to the competition. The MDX Challenge is regarded as a follow-up event of the professionally-produced music (MUS) task of the past SiSEC editions; to continue the tradition of the past MUS, participants were asked to separate stereo songs into stems of four instruments (Vocals, Drums, Bass, Other). Two leaderboards are used to rank the submissions: A) methods trained on MUSDB18(-HQ) and B) methods trained with extra data. Leaderboard A gives the possibility to any participant, independently on the data they possess, to train a MSS system

(since MUSDB18 is open) and includes systems that can, in a later stage, be compared with the existing literature, as they share the same training data commonly used in research; leaderboard B permits models to be used to their full potential and therefore shows the highest achievable scores as of today.

In the following, the paper provides the details about the test dataset in **Section 2**, the leaderboards in **Section 3**, the evaluation metrics in **Section 4**, the baselines in **Section 5**, the evaluation results in **Section 6**, and the technical challenges for future competitions in **Section 7**.

2 MDXDB21

For the specific purpose of this challenge, we introduced a new test set, called MDXDB21. This test set is made of 30 songs, created by Sony Music Entertainment (Japan) Inc. (SMEJ) with the specific intent to use it for the evaluation of the MDX Challenge. The dataset was hidden from the participants, only the organizers of the challenge could access it. This allowed a fair comparison of the submissions. Here we provide details on the creation of the dataset:

- More than 20 songwriters were involved in the making of the 30 songs in the dataset, so that there is no overlap with existing songs in terms of composition and lyrics;
- The copyright of all 30 songs is managed by SMEJ so that MDXDB21 can be integrated easily with other datasets in the future, without any issue arising from copyright management;
- More than 10 mixing engineers were involved in the dataset creation with the aim of diversifying the mixing styles of the included songs;
- The loudness and tone across different songs were not normalized to any reference level, since these songs are not meant to be distributed on commercial platforms;
- To follow the tradition of past competitions like SiSEC, the mixture signal (i.e., the input to the models at evaluation time) is obtained as the simple summation of the individual target stems.

Table 1 shows a list of the songs included in MDXDB21. To give more diversity in genre and language, the dataset features also non-western and non-English songs. The table also provides a list of the instruments present in each song; this can help researchers understand under which conditions their models fail to perform the separation. More information about the loudness as well as stereo information for each song and its stems are given in the Appendix.

3 LEADERBOARDS AND CHALLENGE ROUNDS

For a fair comparison between systems trained with different data, we designed two different leaderboards for the MDX Challenge:

¹An overview of available dataset for music source separation is given in <https://source-separation.github.io/tutorial/data/datasets.html>.

²<https://www.aicrowd.com/challenges/music-demixing-challenge-ismir-2021>.

TABLE 1 | List of songs in MDXDB21. Abbreviations of the instrument names comply with the International Music Score Library Project (IMSLP)⁹.

Song ID	Genre	Language	Title	Instruments in <i>Other</i>
SS_001	Jpop	Japanese	Give Me One More Day	ch, db, egtr, gtr, hn, perc, pf, sfx, va, vc, vn
SS_002	Acid Jazz	English	Do It Right	egtr, epf, fl, perc, pf, sfx, str
SS_003	Ballad	English	I Reach For You	ch, pf, pk, va, vc, vn
SS_004	Funk	English	Diamond In A Rough	asax, gl, gtr, hpd, key, perc, syn, tbn, tpt, tsax, va, vc, vn
SS_005	Ballad (Duet)	English	You And I	cl, db, fl, gl, hn, hp, ob, pf, va, vc, vn
SS_006	Jazz	English	City Swing	asax, gtr, key, pf, tbn, tpt, tsax
SS_007	Bigband	English	Cut Loose, Then Ya Vamoose	asax, clap, gtr, key, tpt, tbn, tsax
SS_008	Rock	English	Take You Out	gtr, perc, pf
SS_009	Country	English	For Your Smile	egtr, gtr, mand, org, pad, perc, pf, sfx
SS_010	Hiphop	English	Monsta	bell, sfx, syn
SS_011	EDM	English	Tormented Soul	bell, egtr, hpd, perc, pf, sfx
SS_012	Metal	English	Time To Let Go	gtr, key
SS_013	Jazz	English	This Is Day	bell, gtr, pad, pf, sfx, syn
SS_014	AOR	English	Moments	gtr, mar, perc, syn, xyl
SS_015	Indian Pop	Hindi	Dil Ki Baatein	syn, egtr
—	—	—	(Conversations Of The Heart)	—
SS_016	Healing	Japanese	Waraeya Utaeya	epf, pad, perc, pf, str, shakuhachi, sfx, syn
SS_017	Reggae	English	Jump Up	ch, gtr, org, perc, pf
SS_018	Pop	English	I Wanna See	gl, gtr, pad, pf, syn, vib, va, vc, vn
SS_019	Pop	English	Tonight	bell, egtr, epf, pad, sfx, str, syn
SS_020	Rock	English	Six Feet Between	org, pad, pf, syn
SS_021	R&B	Portuguese	Minha Sina	epf, pad, pf
SS_022	Jpop	Japanese	Spring (Three Hundred Years Later)	bell, egtr, gl, koto, perc, pf, sfx, shakuhachi, syn, va, vc, vn
SS_023	Jpop	Japanese	Scissorhands	egtr, gtr, perc
SS_024	Digital Rock	English	I Can See You Calling Me	pad, sfx, syn
SS_025	Kpop	Korean	We Can Fly Away	pf, sfx, syn
SS_026	Electro	English	A.I. Robot	epf, gtr, pad, perc, pf, sfx, str, syn
SS_027	Trance	English	By My Side	gtr, pad, sfx, syn
SS_028	House	English	Break Free (Dance Together)	perc, pf, syn
SS_029	Pop	English	By Your Side	gtr, pad, pf, syn
SS_030	Jpop	Japanese	Kikai Jikake No Koenaki Katsubou	pf, sfx, str, syn

- Leaderboard A accepted MSS systems that are trained exclusively on MUSDB18-HQ (Rafii et al., 2019)³. Our main purpose was to give everyone the opportunity to start training a MSS model and take part in the competition, independently on the data they have. On top of that, since MUSDB18-HQ is the standard training dataset for MSS in literature, models trained with it can also be compared with the current state-of-the-art in publications, by evaluating their performance on the test set of MUSDB18-HQ and using the metrics included in the BSS Eval v4 package, as done for example by Défossez (2021) and Kim et al. (2021)⁹.
- Leaderboard B did not pose any constraints on the used training dataset. This allowed participants to train bigger models, exploiting the power of all the data at their disposal.

To avoid some participants overfitting to the MDXDB21 dataset, we split the dataset into three equal-sized parts and designed two challenge rounds: in the first round, participants could access the scores of their models computed only on the first portion of MDXDB21. In the second round, the second portion of MDXDB21 was added to the evaluation and participants could see how well their models generalized on new data. After the challenge ended, the overall score was computed on all songs⁴. These overall scores were also used for the final ranking of the submissions.

4 EVALUATION METRICS

In the following section, we will introduce the metric that was used for the MDX Challenge as the ranking criterion and compare it to other metrics that have been used in past competitions.

³Participants were encouraged to use MUSDB18-HQ as opposed to MUSDB18, as the former is not limited to 16kHz, a limit imposed by the audio codec of MUSDB18. Nevertheless, systems trained on MUSDB18 were also eligible for leaderboard A as MUSDB18 can be seen as a “derived” version of MUSDB18-HQ.
⁹https://imslp.org/wiki/IMSLP:abbreviations_for_Instruments

⁴We excluded three songs, which we used for demo purposes, e.g., to provide feedback to the participants about their submissions on AICrowd. These three songs are SS_008, SS_015, and SS_018.

4.1 Signal-To-Distortion Ratio (SDR)

As an evaluation metric, we chose the multichannel signal-to-distortion ratio (SDR) (Vincent et al., 2007), also called signal-to-noise ratio (SNR), which is defined as

$$SDR = 10 \log_{10} \frac{\sum_n \|s(n)\|^2 + \epsilon}{\sum_n \|s(n) - \hat{s}(n)\|^2 + \epsilon}, \quad (1)$$

where $s(n) \in \mathbb{R}^2$ denotes the waveform of the ground truth and $\hat{s}(n)$ the waveform of the estimate for one of the sources with n being the (discrete) time index. We use a small constant $\epsilon = 10^{-7}$ in (1) to avoid divisions by zero. The higher the SDR score is, the better the output of the system is.

For each song, this allows computing the average SDR, SDR_{Song} , given by

$$SDR_{\text{Song}} = \frac{1}{4} (SDR_{\text{Bass}} + SDR_{\text{Drums}} + SDR_{\text{Other}} + SDR_{\text{Vocals}}). \quad (2)$$

Finally, the systems are ranked by averaging SDR_{Song} over all songs in the hidden test set.

As Eq. 1 considers the full audio waveform at once, we will denote it as a “global” metric. In contrast, we will denote a metric as “frame-wise” if the waveform is split into shorter frames before analyzing it. Using the global metric Eq. 1 has two advantages. First, it is not expensive to compute as opposed to more complex measures like BSS Eval v4 which also outputs the image-to-spatial distortion ratio (ISR), signal-to-interference ratio (SIR), and signal-to-artifacts ratio (SAR). Second, there is also no problem with frames where at least one source or estimate is all-zero. Such frames are discarded in the computation of BSS Eval v4 as otherwise, e.g., SIR can not be computed. This, however, yields the unwanted side-effect that the SDR values of the different sources of BSS Eval v4 are not independent of each other which is not desired⁵. The global SDR Eq. 1 does not suffer from this cross-dependency between source estimates.

4.2 Comparison With Other Metrics

Before deciding to choose the global SDR Eq. 1, we did a comparison with other metrics for audio source separation for the best system from SiSEC 2018 (“TAU1”). Other common metrics are.

- (a) Global/frame-wise SDR,
- (b) Global/frame-wise SI-SDR (Le Roux et al., 2019),
- (c) Global/frame-wise mean absolute error (MAE),
- (d) Global/frame-wise mean squared error (MSE),
- (e) SDR of multi-channel BSS Eval v3 (evaluation metric of SiSEC 2015) (Vincent et al., 2007; Vincent et al., 2012; Ono et al., 2015),⁶
- (f) Mean/median of frame-wise multi-channel BSS Eval v3 (evaluation metric of SiSEC 2016) (Liutkus et al., 2017),
- (g) Mean/median of frame-wise multi-channel BSS Eval v4

(evaluation metric of SiSEC 2018, available as museval Python package) (Stöter et al., 2018).

“Global” refers to computing the metric on the full song whereas “frame-wise” denotes a computation of the metric on shorter frames which are then averaged to obtain a value for the song. For the frame-wise metrics, we used in our experiment a frame size as well as a hop size of one second, which is the default for museval (Stöter et al., 2018), except for the frame-wise SDR of SiSEC 2016 where we used a frame size of 30 s and a hop size of 15 s as in Liutkus et al. (2017).

Figure 1 shows the correlations of the global SDR Eq. 1 (used as reference) with the different metrics on MUSDB18 Test for “TAU1”, the best system from SiSEC 2018 Stöter et al. (2018). For each metric, we compute the correlation coefficient for each source to the reference metric and show the minimum, average, and maximum correlation over all four sources. Please note that some metrics are similarity metrics (“higher is better”) whereas others measure the distance (“smaller is better”). As we use the global SDR as reference, the correlation coefficient becomes negative if the correlation with a distance metric is computed. We can see that there is a strong correlation between the used global SDR Eq. 1 and the median-averaged SDR from BSS Eval v3 and v4 as the Pearson and Spearman correlations are on average larger than 0.9. This analysis confirms that the global SDR Eq. 1 is a good choice as it has a high correlation to the evaluation metrics of SiSEC 2016, i.e., metric (f), and SiSEC 2018, i.e., metric (g), while being at the same time simple to compute and yielding per-source values which are independent of the estimates for other sources. In the following, we will refer to the global SDR Eq. 1 as “SDR”.

5 BASELINE SYSTEMS

The MDX Challenge featured two baselines: Open-Unmix (UMX) and CrossNet-UMX (X-UMX). A description of UMX can be found in Stöter et al. (2019), where the network is based on a BiLSTM architecture that was studied in Uhlich et al. (2017). X-UMX (Sawata et al., 2021) is an enhanced version of UMX.

Figure 2 shows the architectures of the two models. The main difference between them is that UMX can be trained independently for any instrument while X-UMX requires all the networks together during training to allow the exchange of gradients between them, at the cost of more memory. During inference, there is almost no difference between UMX and X-UMX regarding the model size or the computation time as the averaging operations in X-UMX do not introduce additional learnable parameters.

6 MDX CHALLENGE 2021 RESULTS AND KEY TAKEAWAYS

In this section, we will first give results for various systems known from the literature on MDXDB21 before summarizing the outcome of the MDX Challenge 2021.

⁵See for example <https://github.com/sigsep/bsseval/issues/4>.

⁶This metric is equivalent to (1) except for the small constant ϵ .

		Pearson Correlation Coefficient			Spearman Correlation Coefficient		
		min	mean	max	min	mean	max
	Global SDR [Reference]	1.000	1.000	1.000	1.000	1.000	1.000
(a)	Mean of framewise SDR	0.018	0.380	0.657	0.150	0.433	0.697
	Median of framewise SDR	-0.039	0.593	0.950	0.615	0.794	0.927
	Global SI-SDR	0.694	0.888	0.996	0.994	0.996	0.997
(b)	Mean of framewise SI-SDR	0.193	0.495	0.760	0.300	0.521	0.762
	Median of framewise SI-SDR	0.163	0.659	0.929	0.626	0.807	0.926
	Global MSE	-0.530	-0.380	-0.257	-0.493	-0.339	-0.202
(c)	Mean of framewise MSE	-0.530	-0.380	-0.257	-0.493	-0.338	-0.202
	Median of framewise MSE	-0.534	-0.298	-0.068	-0.561	-0.279	-0.054
	Global MAE	-0.610	-0.308	-0.050	-0.522	-0.262	-0.086
(d)	Mean of framewise MAE	-0.610	-0.308	-0.049	-0.523	-0.263	-0.089
	Median of framewise MAE	-0.647	-0.277	0.079	-0.569	-0.246	-0.004
(e)	SDR of BSS Eval v3	1.000	1.000	1.000	1.000	1.000	1.000
(f)	Mean of framewise SDR of BSS Eval v3	0.908	0.933	0.966	0.954	0.963	0.975
	Median of framewise SDR of BSS Eval v3	0.964	0.980	0.992	0.966	0.977	0.987
(g)	Mean of SDR of BSS Eval v4	0.601	0.711	0.820	0.753	0.797	0.853
	Median of SDR of BSS Eval v4	0.871	0.914	0.942	0.898	0.915	0.934

FIGURE 1 | Comparison of MSS metrics using Pearson and Spearman correlation. Metrics are compared to “median of framewise SDR”, i.e., to the metric that was used for SISEC 2018.

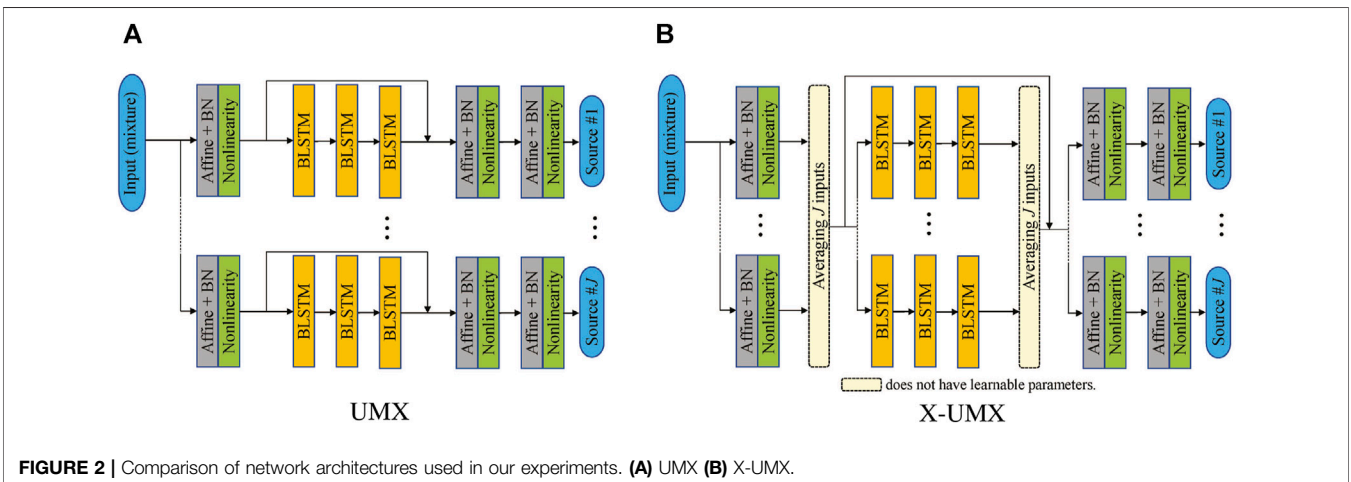


FIGURE 2 | Comparison of network architectures used in our experiments. (A) UMX (B) X-UMX.

6.1 Preliminary Experiments on MDXDB21

The two baselines described in Section 5, as well as state-of-the-art MSS methods, were evaluated on MDXDB21. Table 2 shows SDR results averaged over all 27 songs on leaderboard A, where all the listed models were trained only either on MUSDB18 or on MUSDB-HQ. For leaderboard B, UMX and X-UMX were trained on both, training and test set of MUSDB18-HQ, using the same 14 songs for validation as if only the train part of MUSDB18 would have been used. Table 3 shows the SDR results on leaderboard B. Since the extra dataset used in each model is different, we cannot directly compare their scores, but we nonetheless listed the results to see how well the SOTA

models can perform on new data created by music professionals. It can already be seen that the difference in SDR between these models is smaller than what is reported on the well-established MUSDB18 test set. This indicates that generalization is indeed an issue for many of these systems and it will be interesting to see if other contributions can outperform the SOTA systems.

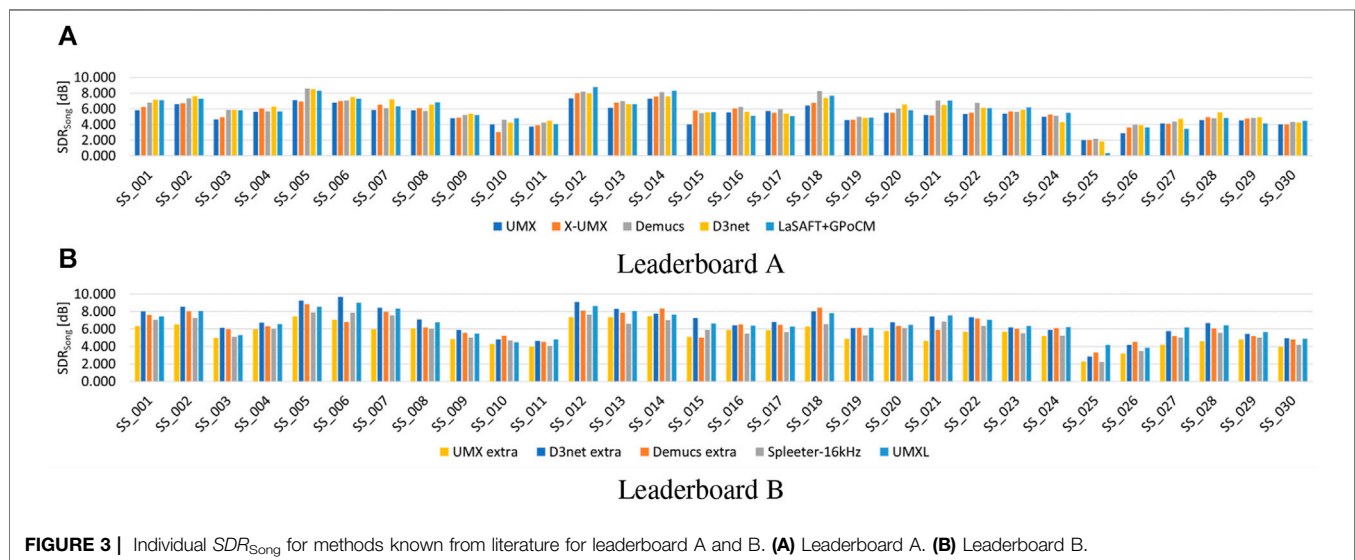
Figures 3A,B show SDR_{song} for all 30 songs for the baselines as well as the currently best methods known from literature. There is one exception in the computation of SDR_{song} for SS_015: For this song, it was computed by excluding *bass* and averaging over three instruments only as this song has an all-silent *bass* track and

TABLE 2 | SDR results on MDXDB21 for oracle baselines as well as systems known from literature which are eligible for leaderboard A, i.e., systems trained only on MUSDB18/MUSDB18-HQ.

System	Training set	SDR_{Song}	SDR_{Bass}	SDR_{Drums}	SDR_{Other}	SDR_{Vocals}
Ideal SWF	Oracle	9.63	9.01	9.54	8.63	11.33
Ideal MWF	Oracle	9.78	9.39	9.59	8.84	11.30
UMX Stöter et al. (2019)	MUSDB18-HQ	5.18	5.40	5.71	3.56	6.07
X-UMX Sawata et al. (2021)	MUSDB18-HQ	5.37	5.62	5.81	3.72	6.34
LaSAFT + GPoCM Choi et al. (2021)	MUSDB18	5.64	5.78	5.23	4.37	7.16
D3Net Takahashi and Mitsufuji (2021)	MUSDB18	5.80	5.74	6.18	4.30	6.97
Demucs Défossez et al. (2021)	MUSDB18	5.81	6.48	6.44	3.96	6.37

TABLE 3 | SDR results on MDXDB21 for methods known from literature which are eligible for leaderboard B, i.e., systems allowed to train with extra data.

System	Training set	SDR_{Song}	SDR_{Bass}	SDR_{Drums}	SDR_{Other}	SDR_{Vocals}
UMX extra. Stöter et al. (2019)	MUSDB18-HQ + 50 songs	5.42	5.60	5.94	3.75	6.36
X-UMX extra. Sawata et al. (2021)	MUSDB18-HQ + 50 songs	5.59	5.79	6.26	3.90	6.41
Demucs extra. Défossez et al. (2021)	MUSDB18 + 150 songs	6.33	6.96	6.92	4.51	6.94
D3Net extra. Takahashi and Mitsufuji (2021)	MUSDB18 + 1.5k songs	6.67	6.58	6.82	5.24	8.06
Spleeter (11 kHz). Hennequin et al. (2020)	25k songs	5.72	5.77	5.85	4.32	6.94
Spleeter (16 kHz) Hennequin et al. (2020)	25k songs	5.77	5.77	5.97	4.33	6.99
UMXL. Stöter et al. (2019)	200 h MP3 songs	6.52	6.62	6.84	4.89	7.73

**FIGURE 3** | Individual SDR_{Song} for methods known from literature for leaderboard A and B. (A) Leaderboard A. (B) Leaderboard B.

SDR_{Bass} would aggravate the average SDR over four instruments. For this reason this song was made part of the set of three songs (SS_008, SS_015, SS_018) that were left out of the evaluation so that they could be provided to the participants as demo (see note 5).

The results for SS_025–026 are considerably worse than for the other songs; we assume that this is because these songs contain more electronic sounds than the others.

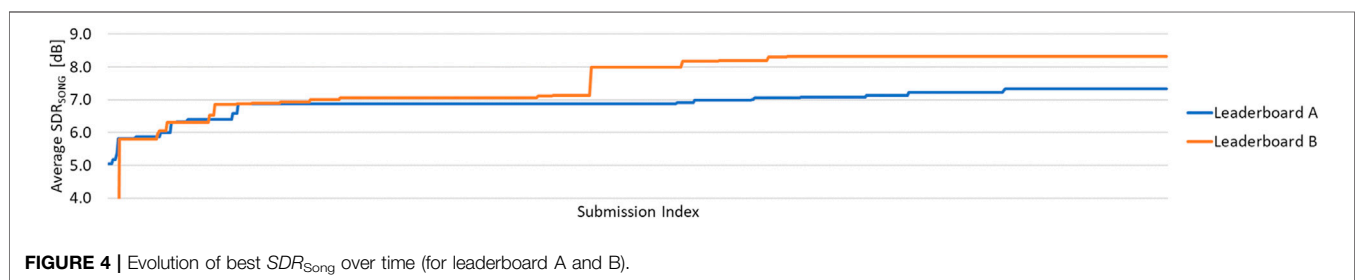
6.2 MDX Challenge 2021 Results

The challenge was well received by the community and, in total, we received 1,541 submissions from 61 teams around the world. **Table 4** shows the final leaderboards of the MDX Challenge. It

gives the mean SDR from (1) for all 27 songs as well as the standard-deviation for the mean SDR over the three splits (each with nine songs) that were used in the different challenge rounds as described in **Section 3**. Comparing these numbers with the results in **Table 2, 3**, we can observe a considerable SDR improvement of approximately 1.5dB throughout the contest. The evolution of the best SDR over time is shown in **Figure 4**. As several baselines for leaderboard A were provided at the start of the challenge, progress could be first observed for leaderboard A and the submissions for leaderboard B were not significantly better. With the start of round 2, the performance gap between leaderboard A and leaderboard B increased as participants managed to train bigger models with additional data.

TABLE 4 | Final leaderboards with top-5 submissions evaluated on MDXDB21.

Rank	Team/User name	Model name	SDR_{Song}	SDR_{Bass}	SDR_{Drums}	SDR_{Other}	SDR_{Vocals}
(a) Leaderboard A							
1	defossez	Hybrid Demucs	7.328 ± 0.46	8.115 ± 0.54	8.037 ± 0.36	5.193 ± 1.02	7.968 ± 1.03
2	kuielab	KUIELab-MDX-Net	7.236 ± 0.59	7.232 ± 0.42	7.173 ± 0.46	5.636 ± 0.89	8.901 ± 1.41
3	Music_AI	—	6.882 ± 0.23	7.273 ± 0.37	7.371 ± 0.62	5.091 ± 0.73	7.792 ± 0.86
4	Kazane_Ryo_no_Danna	Danna-Sep	6.649 ± 0.41	6.993 ± 0.56	7.018 ± 0.58	4.901 ± 0.91	7.686 ± 1.15
5	ByteMSS	—	6.514 ± 0.48	6.602 ± 0.42	6.545 ± 0.48	4.830 ± 0.86	8.079 ± 1.24
(b) Leaderboard B							
1	Audioshake	—	8.326 ± 0.44	8.342 ± 0.45	8.664 ± 0.37	6.505 ± 0.85	9.793 ± 1.13
2	defossez	Hybrid Demucs	8.110 ± 0.27	8.856 ± 0.65	8.850 ± 0.22	5.978 ± 0.72	8.756 ± 0.95
3	kuielab	KUIELab-MDX-Net	7.370 ± 0.47	7.495 ± 0.18	7.554 ± 0.57	5.533 ± 0.82	8.896 ± 1.30
4	Music_AI	—	6.882 ± 0.23	7.273 ± 0.37	7.371 ± 0.62	5.091 ± 0.73	7.792 ± 0.86
5	PyMDX	—	6.680 ± 0.32	6.619 ± 0.56	6.838 ± 0.46	4.891 ± 0.75	8.374 ± 1.07



This improvement was not only achieved by blending several existing models but also by architectural changes. The following sections give more details for the winning models and they are written by each team, respectively.

6.3 Hybrid Demucs (defossez)

Hybrid Demucs extends the original Demucs architecture (Défossez et al., 2021) with multi-domain analysis and prediction capabilities. The original model consisted of an encoder/decoder in the time domain, with U-Net skip connection (Ronneberger et al., 2015). In the hybrid version, a spectrogram branch is added, which is fed with the input spectrogram, either represented by its amplitude, or real and imaginary part, a.k.a Complex-As-Channels (CAC) (Choi et al., 2020). Unlike the temporal branch, the spectrogram branch applies convolutions along the frequency axis, reducing the number of frequency bins by a factor of 4 with every encoder layer. Starting from 2048 frequency bins, obtained with a 4,096 steps STFT with a hop-length of 1,024 and excluding the last bin for simplicity, the input to the fifth layer of the spectral branch has only eight frequency bins remaining, which are collapsed to a single one with a single convolution. On the other hand, the input to the fifth layer of the temporal branch has an overall stride of $4^4 = 256$, which is aligned with the stride of 1,024 of the spectral branch with a single convolution. At this point, the two representations have the same shape and are summed before going through a common layer that further reduces the number of time steps by 2. Symmetrically, the first layer in the decoder is shared before being fed both to the temporal and spectral

decoder, each with its own set of U-Net skip connections. The overall structure is represented in **Figure 5**.

The output of the spectrogram branch is inverted to a waveform, either directly with the ISTFT when CAC is used, or thanks to Wiener filtering (Nugraha et al., 2016) for the amplitude representation, using Open-Unmix differentiable implementation (Stöter et al., 2019). The final loss is applied directly in the time domain. This allows for end-to-end hybrid domain training, with the model being free to combine both domains. To account for the fact that musical signals are not equivariant with respect to the frequency axis, we either inject a frequency embedding after the first spectral layer, following Isik et al. (2020), or we allow for different weights depending on the frequency band, as done by Takahashi and Mitsufuji (2017).

Further improvements come from inserting residual branches in each of the encoder layers. The branches operate with a reduced number of dimensions (scaled down by a factor of 4), using dilated convolutions and group normalization (Wu and He, 2018), and for the two innermost layers, BiLSTM and local attention. Local attention is based on regular attention (Vaswani et al., 2017), but replacing positional embedding with a controllable penalty limits its scope to nearby time steps. All ReLUs in the network were replaced by GELUs (Hendrycks and Gimpel, 2016). Finally, we achieve better generalization and stability by penalizing the largest singular values of each layer (Yoshida and Miyato, 2017). We achieved further gains (between 0.1 and 0.2 dB) by fine-tuning the models on a specifically crafted dataset, and with longer training samples (30 s instead of 10). This dataset was built by combining stems

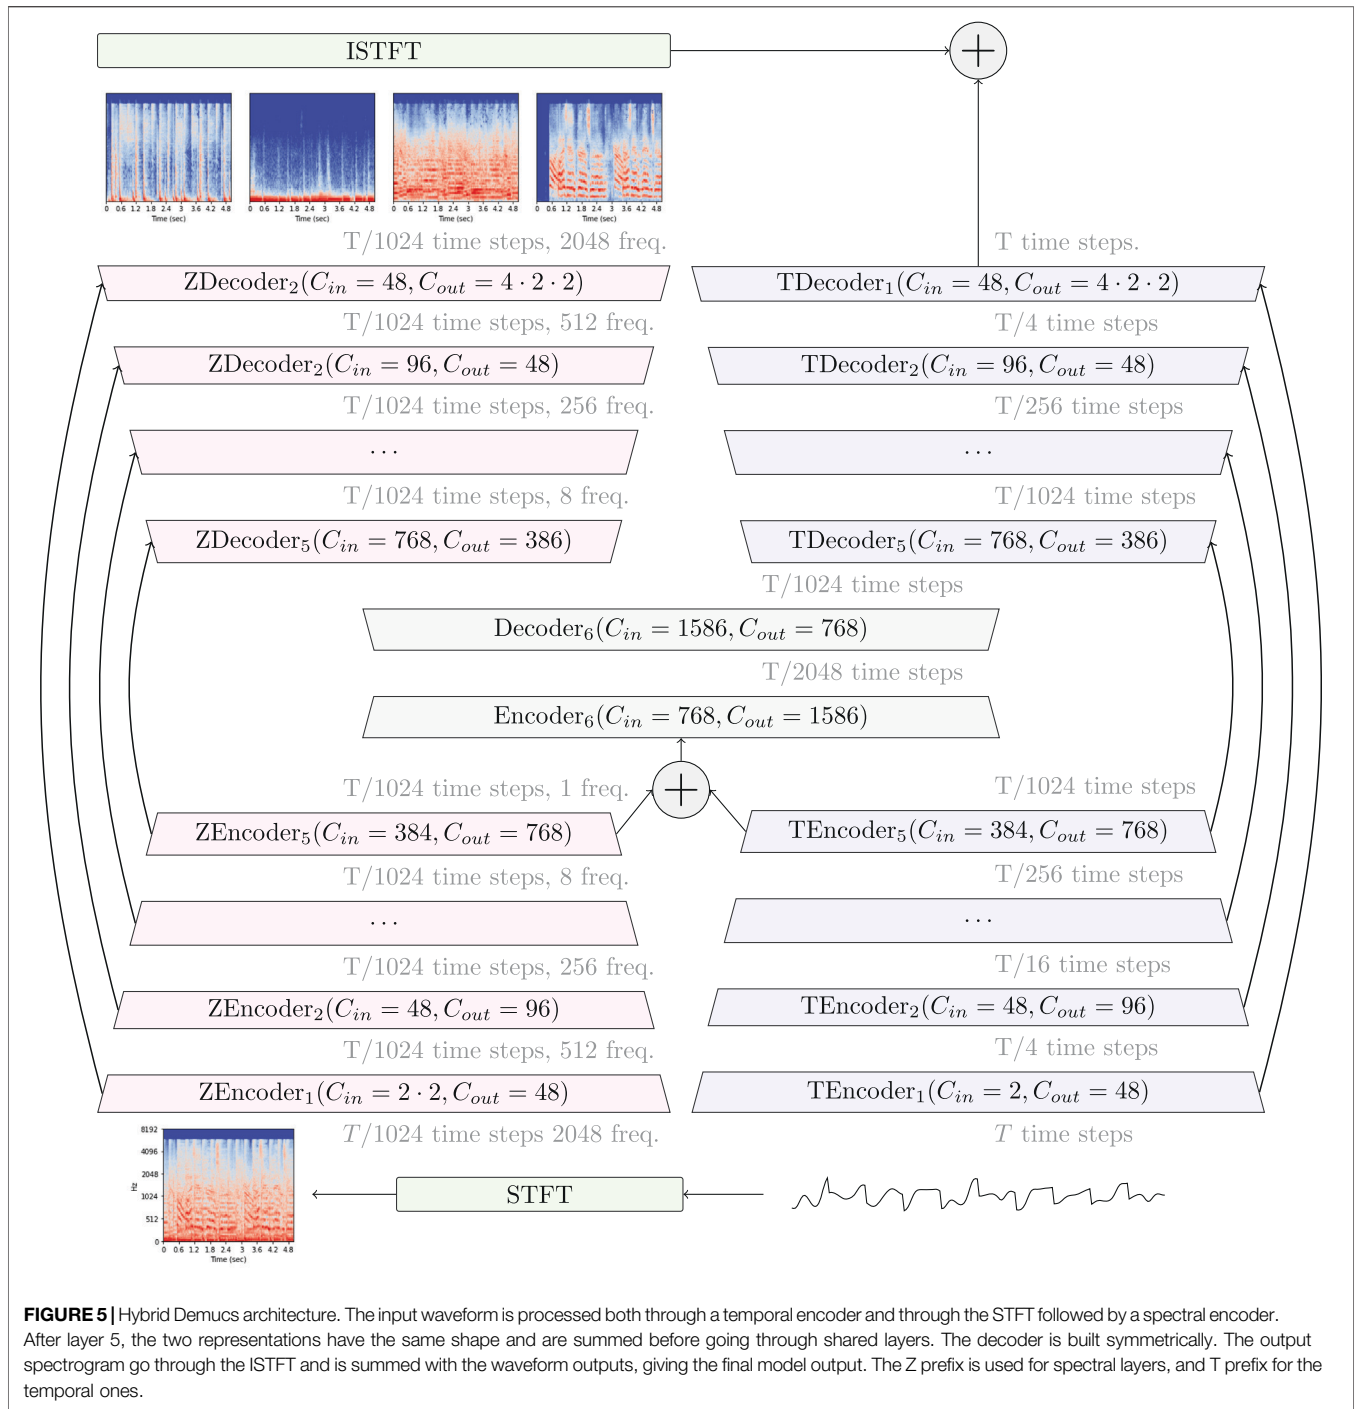


FIGURE 5 | Hybrid Demucs architecture. The input waveform is processed both through a temporal encoder and through the STFT followed by a spectral encoder. After layer 5, the two representations have the same shape and are summed before going through shared layers. The decoder is built symmetrically. The output spectrogram go through the ISTFT and is summed with the waveform outputs, giving the final model output. The Z prefix is used for spectral layers, and T prefix for the temporal ones.

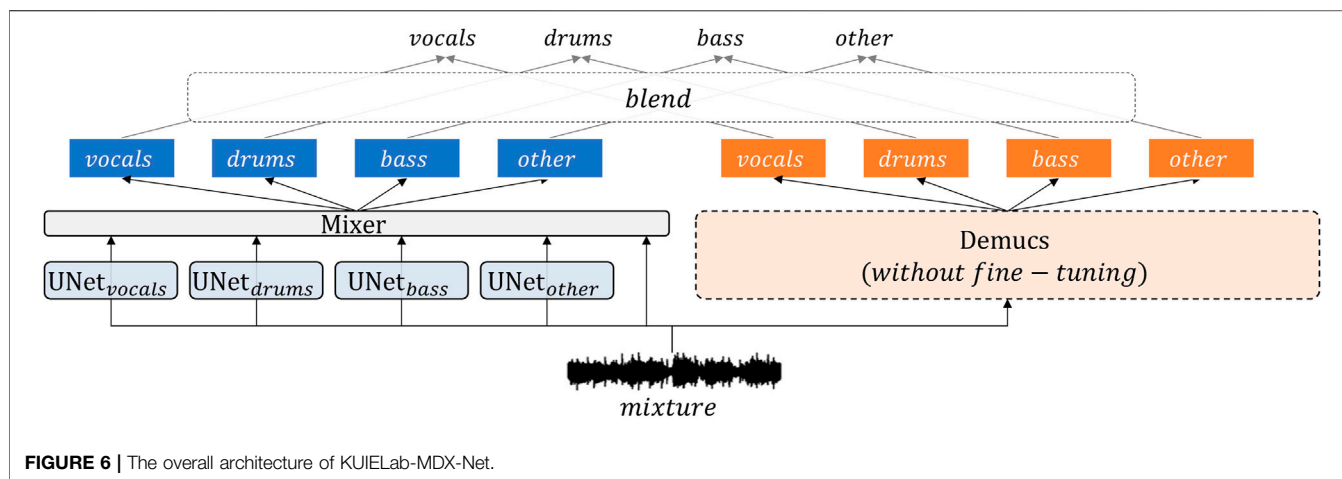
from separate tracks, while respecting a number of conditions, in particular beat matching and pitch compatibility, allowing only for small pitch or tempo corrections.

The final models submitted to the competition are bags of four models. For leaderboard A, it is a combination of temporal only and hybrid Demucs models, given that with only MUSDB18-HQ as a train set, we observed a regression on the bass source. For leaderboard B, all models are hybrid, as the extra training data made the hybrid version better for all sources than its time-only

version. We refer the reader to our Github repository facebookresearch/demucs for the exact hyper-parameters used. More details and experimental results, including subjective evaluations, are provided in the Hybrid Demucs paper (Défossez, 2021).

6.4 KUIELab-MDX-Net (kuielab)

Similar to Hybrid Demucs, KUIELab-MDX-Net (Kim et al., 2021) uses a two-branched approach. As shown in Figure 6, it



has a time-frequency branch (left-side) and a time-domain branch (right-side). Each branch estimates four stems (i.e., vocals, drums, bass, and other). The *blend* module (Uhlich et al., 2017) outputs the average of estimated stems from two branches for each source. While branches of Hybrid Demucs were jointly trained end-to-end, each branch of KUIELab-MDX-Net was trained independently.

For the time-domain branch, it uses the original Demucs (Défossez et al., 2021), which was pre-trained on MUSDB18. It was not fine-tuned on MUSDB18-HQ, preserving the original parameters.

For the time-frequency-domain branch, it uses five sub-networks. Four sub-networks were independently trained to separate four stems, respectively. For each stem separation, an enhanced version of *TFC-TDF-U-Net* (Choi et al., 2020) was used. We call the enhanced one *TFC-TDF-U-Net v2* for the rest of the paper. Another sub-network called *Mixer* was trained to output enhanced sources by taking and refining the estimated stems.

TFC-TDF-U-Net is a variant of U-Net (Ronneberger et al., 2015) architecture. It improves source separation by employing TFC-TDF (Choi et al., 2020) as building blocks instead of fully convolutional layers. Architectural/training changes were made for TFC-TDF-U-Net v2 to the original as follows:

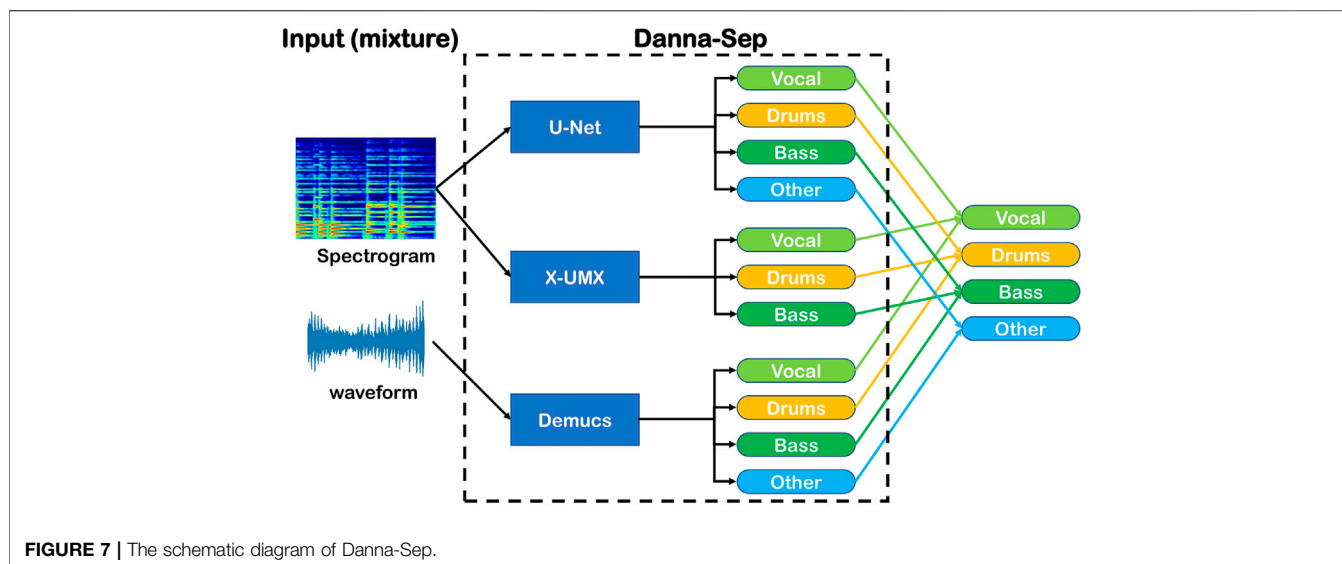
- For skip connections between encoder and decoder, multiplication was used instead of concatenation.
- The other skip connections (e.g., dense skip connections in a dense block (Takahashi and Mitsufuji, 2017)) were removed.
- While the number of channels is not changed after down/up-sampling in the original, channels are increased/decreased when downsampling/up-sampling in v2.
- While the original was trained to minimize time-frequency domain loss, v2 was trained to minimize time-domain loss (i.e., l_1 loss between the estimated waveform and the ground-truth)

Since dense skip connections based on concatenation usually require a large amount of GPU memory, as discussed in Chen et al. (2017), TFC-TDF-U-Net v2 was designed to use simpler

modules. Kim et al. (2021) found that replacing concatenation with multiplication for each skip connection does not severely degrade the performance of the TFC-TDF-U-Net structure. Also, they observed that removing dense skip connections in each block does not significantly degrade the performance if we use TFC-TDF blocks. A single Time Distributed Fully connected (TDF) block, contained in a TFC-TDF, has an entire receptive field in the frequency dimension. Thus, U-Nets with TFC-TDFs can show promising results even with shallow or simple structures, as discussed in (Choi et al., 2020). To compensate for the lack of parameters by this design shift, it enlarges the number of channels for every downsampling, which is more general in the conventional U-Net (Ronneberger et al., 2015). It was trained to minimize time-domain loss for a direct end to end optimization.

Also, KUIELab-MDX-Net applies a frequency cut-off trick, introduced in Stöter et al. (2019) to increase the window size of STFT (or *fft size in short*) as a source-specific preprocessing. It cuts off high frequencies above the target source's expected frequency range from the mixture spectrogram. In this way, *fft size* could be increased while using the same input spectrogram size (which we needed to constrain for the separation time limit) for the model. Since using a larger *fft size* usually leads to better SDR, this approach can improve quality effectively with a proper cut-off frequency. It is also why we did not use a multi-target model (i.e., a single model to separate all the sources once), where we could not apply source-specific frequency cutting.

Training one separation model for each source can have the benefit of source-specific preprocessing and model configurations. However, these sub-networks lack the knowledge that they are separating using the same mixture because they cannot communicate with each other to share information. An additional sub-network called *Mixer* could further enhance the “independently” estimated sources. For example, estimated “vocals” often have drum snare noises left. The Mixer can learn to remove sounds from ‘vocals’ that are also present in the estimated “drums” or vice versa. Very shallow models (such as a single convolution layer) have been tried for the Mixer due to the time limit. One can try more complex models in the future, since even a single 1×1 convolution layer was enough



to make some improvement on total SDR. Mixer used in KUIELab-MDX-Net is a point-wise convolution that is applied to the waveform domain. It takes a multi-channel waveform input containing four estimated stereo stem channels and the original stereo mixture channel. It outputs four different stereo stems. It can be viewed as a U-Net blending module with learnable parameters. An ablation study is provided in Kim et al. (2021) for interested readers.

Finally, KUIELab-MDX-Net takes the weighted average of estimated stems from two branches for each source. In other words, it blends (Uhlich et al., 2017) results from two branches. For Leaderboard A, we trained KUIELab-MDX-Net which adopted TFC-TDF-U-Net v2, Mixer, and blending with the original Demucs, after training on MUSDB18-HQ (Rafii et al., 2019) training dataset with pitch/tempo shift (Défossez et al., 2021) and random track mixing augmentation (Uhlich et al., 2017). For Leaderboard B, we used KUIELab-MDX-Net without Mixer but with validation and test dataset of MUSDB18-HQ. The source code for training KUIELab-MDX-Net is available at the Github repository [kuielab/mdx-net](https://github.com/kuielab/mdx-net).

6.5 Danna-Sep (Kazane_Ryo_no_Danna)

Danna-Sep is, not surprisingly, also a hybrid model blending the outputs from three source separation models across different feature domains. Two of them receive magnitude spectrogram as input, while the other one use waveforms, as shown in Figure 7. The design principle is to combine the complementary strengths of both waveform-based and spectrogram-based approaches.

The first one of spectrogram-based models is a X-UMX trained with complex value frequency domain loss to better address the distance between ground truth and estimation spectrograms. In addition, we also incorporated differentiable Wiener filtering into training with our own PyTorch implementation⁷, similar to how Hybrid Demucs did. We

initialized this model with the official pre-trained weights⁸ before we start training. The second one is a U-Net with six layers consisting of D3 Blocks from D3Net (Takahashi and Mitsufuji, 2021) and two layers of 2D local attention (Parmar et al., 2018) at the bottleneck. We also experimented with using biaxial BiLSTM along the time and frequency axes as the bottleneck layers, but it took slightly longer to train yet offered a negligible improvement. We used the same loss function as our X-UMX during training but with Wiener filtering being disabled.

The waveform-based model is a variant of 48 channels Demucs, where the decoder is replaced by four independent decoders responsible for four respective sources. We believe that the modification can help each decoder focusing on their target source without interfering with others. Each decoder has the same architecture as the original decoder, except for the size of the hidden channel which was reduced to 24. This makes the total number of parameters compared with the original Demucs. The training loss aggregates the L1-norm between estimated and ground-truth waveforms of the four sources. We didn't apply the shift trick (Défossez et al., 2021) to this variant because of limited computation resources set by the competition, but in our experiments, we found it still slightly outperformed the 48 channels Demucs.

The aforementioned models were all trained separately on MUSDB18-HQ with pitch/tempo shift and random track mixes augmentation.

Finally, we calculated the weighted average of individual outputs from all models. Experiments were conducted to search for optimal weighting. One iteration of Wiener filtering was used for our X-UMX and U-Net before averaging. To see the exact configurations we used, reader can refer to our repository [yoyololicon/music-demixing-challenge-ismir-2021-entry](https://github.com/yoyololicon/music-demixing-challenge-ismir-2021-entry) on Github. Experimental results of each individual model

⁷<https://github.com/yoyololicon/norbent>.

⁸https://zenodo.org/record/4740378/files/pretrained_xumx_musdb18HQ.pth.

compare to the unmodified baselines are provided in our workshop paper (Yu and Cheuk, 2021).

7 ORGANIZING THE CHALLENGE AND FUTURE EDITIONS

7.1 General Considerations

Organizing a research challenge is a multi-faceted task: the intention is to bring people from different research areas together, enabling them to tackle a task new to them in the easiest way possible; at the same time, it is an opportunity for more experienced researchers to measure themselves once more against the state-of-the-art, providing them with new challenges and space to further improve their previous work. All this, while making sure that the competition remains fair for all participants, both novice and experienced ones.

To satisfy the first point above, we decided to host the challenge on an open crowd-based machine learning competition website. This allowed the competition to be visible by researchers outside of the source separation community. The participants were encouraged to communicate, exchange ideas and even create teams, all through AICrowd's forum. To make the task accessible to everyone, we prepared introductory material and collected useful resources available on the Internet so that beginners could easily start developing their systems and actively take part in the challenge.

We also increased the interest of experienced researchers because the new evaluation set represented an opportunity for them to evaluate their existing models on new data. This music was created for the specific purpose of being used as the test set in this challenge: this stimulated the interest of the existing MSS community because old models could be tested again and, possibly, improved.

The hidden test set also allowed us to organize a competition where fairness had a very high priority. Nobody could adapt their model to the test set as nobody except the organizers could access it. The whole challenge was divided into two rounds, where different subsets of the test set were used for the evaluation: this prevented users from adapting their best model to the test data by running multiple submissions. Only at the very end of the competition, the final scores on the complete test set were computed and the winners were selected: at that moment, no new model could be submitted anymore. Cash prizes are offered for the winners in return for open-sourcing their submission.

7.2 Bleeding Sounds Among Microphones

The hidden test set was created with the task of MSS in mind, which meant that we had to ensure as much diversity in the audio data as possible while maintaining a professional level of production and also deal with issues that are usually not present when producing music in a studio.

Depending on the genre of the music being produced, bleeding among different microphones (i.e., among the recorded tracks of different instruments) is more or less tolerated. For example, when recording an orchestra, to maintain a natural performance, the different sections play in the same room: even if each section

may be captured with a different set of microphones, each microphone will still capture the sound of all the other sections. This phenomenon is not desirable for the task of MSS: if a model learns to separate each instrument very well but the test data contains bleeding, the model will be wrongly penalized during the evaluation. For this reason, when producing the songs contained in the dataset, we had to ensure that specific recording conditions were respected. All our efforts were successful, except for one song (SS_008): even taking appropriate measures during the recording process (e.g., placing absorbing walls between instruments), the tracks contained some bleeding between drums and bass. For this reason, we removed this track from the evaluation and used it as a demo song instead.

Bleeding can be an issue also for training data: we cannot expect models to have better performance than the data they are trained upon unless some specific measures are taken. We did not explicitly address this aspect in the MDX Challenge; nevertheless, designing systems that are robust to bleeding in the training data is a desirable feature, not just for MSS. We could envision that future editions of the MDX Challenge will have a track reserved for systems trained exclusively on data that suffer from bleeding issues, to focus the research community on this aspect as well.

7.3 Silent Sources

Recording conditions are not the only factor that can make a source separation evaluation pipeline fail. Not all instruments are necessarily present in a track: depending on the choices of the composer, arranger, or producer, some instruments may be missing. This is an important aspect, as some evaluation metrics, like the one we chose, may not be robust to the case of a silent target. We decided to exclude from MDXDB21 one song (SS_015) as it has a silent target for the bass track. Please note that this issue was not present in previous competitions like SiSEC, as the test set of MUSDB18 does not feature songs with silent sources. Nevertheless, we think it would be an important improvement if the challenge evaluation could handle songs where one of the instruments is missing (e.g., instrumental songs without vocals, acapella interpretations, etc.).

Such an issue arises from the definition of clear identities for the targets in the source separation task: the evaluation pipeline suffers from such strict definition and causes some songs to be unnecessarily detrimental to the overall scores. This is a motivation to move towards source separation tasks that do not feature clear identities for the targets, such as universal sound source separation: in that case, the models need to be able to separate sounds, independently on their identity. An appropriately designed pipeline for such a task would not suffer from the issue above. For this reason, for future editions of the challenge, we may consider including tasks similar to universal sound source separation.

7.4 Future Editions

We believe that hosting the MDX Challenge strengthened and expanded the source separation community, by providing a new playground to do research and attracting researchers from other communities and areas, allowing them to share knowledge and skills: all this focused on solving one of the most interesting

TABLE 5 | Statistics about songs and their stems in MDXDB21. “Max-Abs” and “Loudness” values are given for the left and right channel, respectively.

Song ID	Mixture			Bass			Drums			Other			Vocals		
	Max-Abs	Loudness	Corr	Max-Abs	Loudness	Corr	Max-Abs	Loudness	Corr	Max-Abs	Loudness	Corr	Max-Abs	Loudness	Corr
SS_001	0.77,	-20.3,	0.74	0.17,	-30.7,	1.00	0.53,	-27.5,	0.99	0.43,	-24.4,	0.41	0.57,	-21.2,	0.93
	0.80	-20.2		0.17	-30.7		0.53	-27.6		0.43	-24.2		0.60	-21.2	
SS_002	0.86,	-19.3,	0.89	0.31,	-25.2,	1.00	0.45,	-25.1,	0.98	0.45,	-25.1,	0.59	0.44,	-22.5,	0.90
	0.87	-19.3		0.31	-25.2		0.45	-24.5		0.42	-25.7		0.43	-22.6	
SS_003	0.85,	-20.3,	0.64	0.15,	-28.5,	1.00	0.23,	-35.7,	0.82	0.44,	-26.8,	0.12	0.68,	-21.1,	0.82
	0.69	-21.4		0.15	-28.5		0.22	-35.6		0.42	-27.5		0.58	-22.4	
SS_004	0.88,	-18.2,	0.77	0.35,	-29.6,	1.00	0.70,	-25.0,	0.81	0.52,	-23.9,	0.47	0.67,	-21.1,	0.85
	0.94	-18.6		0.35	-29.6		0.67	-25.6		0.65	-24.5		0.63	-21.2	
SS_005	0.84,	-19.8,	0.74	0.23,	-28.5,	1.00	0.52,	-29.7,	0.94	0.40,	-26.0,	0.10	0.71,	-21.0,	0.93
	0.91	-20.1		0.23	-28.5		0.56	-29.3		0.46	-26.3		0.63	-21.3	
SS_006	0.99,	-18.5,	0.78	0.30,	-25.2,	0.98	0.45,	-28.1,	0.93	0.52,	-23.4,	0.32	0.49,	-21.1,	0.94
	0.95	-18.7		0.28	-25.5		0.42	-28.0		0.57	-23.7		0.49	-21.0	
SS_007	0.95,	-18.2,	0.88	0.32,	-23.0,	1.00	0.61,	-24.7,	0.91	0.55,	-22.1,	0.53	0.42,	-23.0,	0.96
	0.87	-18.4		0.32	-23.0		0.54	-24.7		0.62	-22.8		0.44	-22.8	
SS_008	0.86,	-18.3,	0.78	0.31,	-25.9,	0.96	0.46,	-25.4,	0.82	0.43,	-24.6,	0.49	0.43,	-21.1,	0.87
	0.84	-18.2		0.27	-26.3		0.43	-26.2		0.50	-23.5		0.43	-21.3	
SS_009	0.93,	-18.7,	0.62	0.19,	-26.4,	1.00	0.75,	-27.2,	0.84	0.50,	-22.8,	0.35	0.54,	-21.9,	0.58
	0.96	-18.6		0.19	-26.4		0.70	-27.6		0.50	-22.9		0.49	-21.6	
SS_010	0.87,	-18.9,	0.64	0.45,	-25.2,	0.55	0.50,	-22.7,	0.95	0.45,	-24.8,	0.05	0.29,	-25.0,	0.88
	0.81	-18.8		0.43	-25.7		0.51	-22.5		0.44	-24.5		0.29	-25.0	
SS_011	0.83,	-17.3,	0.72	0.26,	-24.3,	1.00	0.45,	-22.7,	0.77	0.64,	-21.5,	0.35	0.39,	-23.1,	0.88
	0.97	-17.0		0.26	-24.5		0.50	-21.4		0.53	-21.9		0.40	-23.1	
SS_012	0.91,	-18.5,	0.75	0.29,	-28.8,	0.98	0.49,	-30.0,	0.80	0.54,	-24.3,	0.28	0.63,	-19.1,	0.96
	0.94	-18.5		0.27	-28.9		0.48	-30.0		0.57	-24.3		0.63	-19.1	
SS_013	0.68,	-20.0,	0.66	0.20,	-29.6,	1.00	0.46,	-27.5,	0.95	0.52,	-23.1,	0.36	0.41,	-24.6,	0.78
	0.72	-20.1		0.20	-29.6		0.40	-27.6		0.49	-23.4		0.40	-24.5	
SS_014	0.90,	-19.8,	0.77	0.27,	-24.9,	0.95	0.72,	-25.4,	0.91	0.37,	-25.5,	0.02	0.50,	-23.1,	0.96
	0.89	-19.8		0.25	-25.0		0.83	-24.6		0.36	-26.7		0.54	-23.1	
SS_015	0.87,	-20.4,	0.68	0.00,	-inf, -inf	0.00	0.63,	-24.9,	0.62	0.51,	-24.9,	0.62	0.40,	-23.9,	0.82
	0.99	-19.9		0.00			0.79	-23.6		0.51	-24.6		0.42	-23.7	
SS_016	0.65,	-20.6,	0.64	0.26,	-26.2,	0.99	0.40,	-27.8,	0.95	0.29,	-25.6,	0.08	0.46,	-26.0,	0.43
	0.66	-20.2		0.26	-26.1		0.38	-28.2		0.31	-25.1		0.42	-25.2	
SS_017	0.72,	-22.7,	0.98	0.15,	-30.8,	1.00	0.57,	-28.5,	0.99	0.38,	-28.1,	0.97	0.48,	-27.0,	0.93
	0.72	-22.6		0.17	-29.9		0.60	-28.2		0.39	-28.1		0.41	-27.3	
SS_018	0.90,	-18.9,	0.84	0.44,	-23.2,	1.00	0.48,	-23.7,	0.95	0.43,	-24.5,	0.37	0.35,	-26.3,	0.91
	0.92	-18.9		0.44	-23.2		0.48	-23.6		0.47	-24.6		0.34	-26.4	
SS_019	0.92,	-18.1,	0.81	0.16,	-26.9,	1.00	0.49,	-21.4,	0.97	0.37,	-26.1,	0.19	0.46,	-22.3,	0.77
	0.99	-18.1		0.16	-27.0		0.48	-21.9		0.47	-25.3		0.47	-22.1	
SS_020	0.91,	-20.7,	0.74	0.26,	-25.5,	1.00	0.72,	-26.3,	0.95	0.38,	-23.9,	0.24	0.22,	-27.7,	0.95
	0.98	-21.1		0.26	-25.7		0.76	-26.2		0.39	-24.8		0.21	-27.6	
SS_021	0.93,	-18.8,	0.80	0.39,	-23.9,	1.00	0.59,	-24.2,	0.91	0.52,	-25.1,	-0.07	0.40,	-22.6,	0.87
	0.97	-18.7		0.39	-23.8		0.55	-24.0		0.60	-25.2		0.38	-22.6	
SS_022	0.92,	-17.2,	0.79	0.28,	-24.6,	1.00	0.48,	-22.6,	0.94	0.41,	-22.2,	0.28	0.75,	-21.4,	0.92
	0.94	-17.3		0.30	-24.6		0.45	-22.7		0.43	-22.8		0.74	-21.4	
SS_023	0.83,	-18.4,	0.77	0.28,	-26.6,	0.97	0.61,	-25.1,	0.83	0.34,	-24.4,	0.46	0.74,	-21.1,	0.81
	0.85	-18.2		0.27	-26.4		0.58	-25.3		0.40	-24.5		0.74	-21.0	
SS_024	0.91,	-17.2,	0.69	0.42,	-22.7,	0.70	0.81,	-21.6,	0.92	0.45,	-24.5,	0.02	0.56,	-21.8,	0.84
	0.95	-17.2		0.45	-22.5		0.82	-21.9		0.46	-24.2		0.57	-21.8	
SS_025	0.91,	-21.9,	0.82	0.25,	-28.0,	1.00	0.41,	-28.6,	0.86	0.57,	-26.3,	0.68	0.37,	-24.9,	0.85
	0.91	-21.9		0.25	-28.0		0.43	-28.3		0.56	-26.2		0.37	-25.2	
SS_026	0.98,	-20.8,	0.70	0.49,	-26.2,	0.84	0.63,	-26.1,	0.88	0.80,	-24.0,	0.38	0.13,	-28.6,	0.74
	0.79	-21.2		0.48	-26.2		0.64	-26.0		0.50	-25.0		0.14	-28.6	
SS_027	0.64,	-21.0,	0.57	0.19,	-29.9,	1.00	0.43,	-25.6,	0.94	0.50,	-23.6,	0.27	0.24,	-26.1,	0.63
	0.70	-21.0		0.19	-29.9		0.36	-25.6		0.51	-23.7		0.25	-26.1	
SS_028	0.59,	-23.2,	0.85	0.14,	-32.6,	1.00	0.43,	-27.0,	0.96	0.18,	-32.3,	0.39	0.30,	-27.0,	0.69
	0.61	-23.0		0.14	-32.6		0.47	-26.5		0.24	-31.8		0.25	-27.1	
SS_029	0.66,	-21.0,	0.66	0.25,	-28.7,	0.99	0.36,	-25.7,	0.96	0.37,	-25.8,	0.09	0.33,	-26.4,	0.73
	0.63	-20.9		0.26	-28.9		0.40	-25.6		0.38	-25.6		0.30	-26.6	
SS_030	0.73,	-18.3,	0.71	0.23,	-26.0,	1.00	0.38,	-24.3,	0.85	0.44,	-22.7,	0.33	0.40,	-24.0,	0.71
	0.76	-18.7		0.23	-26.1		0.40	-24.7		0.40	-23.3		0.37	-24.3	

research problems in the area of sound processing. Source separation can still bring benefits to many application and research areas: this motivates the need for future editions of this competition.

In our view, this year's edition of the MDX Challenge allows us to start a hopefully long tradition of source separation competitions. The focus of this year was MSS on four instruments: given the role this task played in past competitions, this was a convenient starting point, that provided us with enough feedback and experience on how to make the competition grow and improve.

We encountered difficulties when compromising between how the source separation task expects data to be created and the professional techniques for music production: for instance, to keep the competition fair, we had to make sure that no crosstalk between target recordings was present. The same argument about crosstalk also highlighted the need for source separation systems that can be trained on data that suffer from this issue: this potentially opens access to training material not available before and can be another source of improvement for existing models. We realized how brittle the design of a simple evaluation system is when dealing with the vast amount of choices that artists can make when producing music: even the simple absence of an instrument in a track can have dramatic consequences in the competition results.

This knowledge will eventually influence the decisions we will take when designing the next editions of the MDX Challenge. In particular, we will:

- design an evaluation system around a metric that is robust to bleeding sounds between targets in the test ground truth data
- direct the attention of researchers to the robustness of models concerning bleeding among targets in the training data, possibly reserving a separate track to systems trained exclusively on such data
- partially moves towards source separation tasks where there is no predefined identity for the targets, such as universal sound source separation.

Furthermore, motivated by the pervasiveness of audio source separation, we will consider reserving special tracks to other types of audio signals, such as speech and ambient noise. The majority of techniques developed nowadays provide useful insights independently on whether they are applied to music, speech, or other kinds of sound. In the interest of scientific advancement, we will try to make the challenge as diverse as possible, to have the highest number of researchers cooperate, interact and ultimately compete for the winning system.

8 CONCLUSION

With the MDX Challenge 2021, we continued the successful series of SiSEC MUS challenges. By using a crowd-sourced platform to host the competition, we tried to make it easy for ML practitioners from other disciplines to enter this field. Furthermore, we introduced a newly created dataset, called

MDXDB21, which served as the hidden evaluation set for this challenge. Using it allows to fairly compare all recently published models as it shows their generalization capabilities towards unseen songs.

We hope that this MDX Challenge will be the first one in a long series of competitions.

ADDITIONAL INFORMATION FOR MDXDB21

Table 5 provides more information about the songs in MDXDB21, which allows a better interpretation of the SDR results that the participants obtained. In particular, **Table 5** contains for each song as well as its stems the following statistics:

- Maximum absolute peak value (left/right channel),
- Loudness according to BS.1770 (left/right channel),
- Correlation coefficient between left/right channel (time-domain).

DATA AVAILABILITY STATEMENT

The datasets presented in this article are not readily available because Generated data has to remain hidden to allow a fair comparison of challenge submissions in the future editions of Music Demixing Challenge. Requests to access the datasets should be directed to YM, Yuhki.Mitsufuji@sony.com.

ETHICS STATEMENT

Bioethical review and/or approval is not required for this research according to the following reasons: this research is not a medical or health research involving human and/or animal subjects. This research does not utilize specimen and/or information obtained and om human and/or animal subjects.

AUTHOR CONTRIBUTIONS

The first author, YM, created the dataset used in the challenge together with Sony Music and established Music Demixing (MDX) Challenge on a crowd-based competition system. He is responsible for the entire sections in the paper, especially for Section 1, 2, 3, and 5. The second author, GF, is the main contributor of running MDX Challenge, and he is responsible for Sec. 7 where he summarize what the future editions of MDX Challenge need to tackle with. The third and fourth authors, SU and F-RS, evaluated source separation metrics that should be run fast in a crowd-based system. They also successfully implemented the evaluation system as well as baseline systems with which participants could easily join the challenge. They are responsible for Sec. 4, 6. The winners of MDX Challenge, AD, MK, WC, C-YY, and K-WC, described each system in Sec. 6.3, 6.4, and 6.5, respectively.

REFERENCES

- Bittner, R. M., Salamon, J., Tierney, M., Mauch, M., Cannam, C., and Bello, J. P. (2014). “MedleyDB: A Multitrack Dataset for Annotation-Intensive MIR Research,” in Proceedings: International Society for Music Information Retrieval Conference (ISMIR), Taipei, Taiwan, October 27–31, 2014, 155–160.
- Chen, Y., Li, J., Xiao, H., Jin, X., Yan, S., and Feng, J. (2017). “Dual Path Networks,” in Proceedings: Annual Conference on Neural Information Processing Systems (NeurIPS), Long Beach, CA, December 4–9, 2017, 4467–4475.
- Choi, W., Kim, M., Chung, J., Lee, D., and Jung, S. (2020). “Investigating U-Nets with Various Intermediate Blocks for Spectrogram-Based Singing Voice Separation,” in Proceedings: International Society for Music Information Retrieval Conference (ISMIR), Montreal, QC, October 11–16, 2020, 192–198.
- Choi, W., Kim, M., Chung, J., and Jung, S. (2021). “LaSAFT: Latent Source Attentive Frequency Transformation for Conditioned Source Separation,” in Proceedings: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Toronto, ON, June 6–11, 2021, 171–175. doi:10.1109/icassp39728.2021.9413896
- Défossez, A., Usunier, N., Bottou, L., and Bach, F. (2021). *Music Source Separation in the Waveform Domain*. arXiv. arXiv preprint arXiv:1911.13254.
- Défossez, A. (2021). “Hybrid Spectrogram and Waveform Source Separation,” in Proceedings: The ISMIR 2021 Workshop on Music Source Separation, November 12, 2021.
- Hendrycks, D., and Gimpel, K. (2016). *Gaussian Error Linear Units (GELUs)*. arXiv. arXiv preprint arXiv:1606.08415.
- Hennequin, R., Khlif, A., Voituret, F., and Moussallam, M. (2020). Spleeter: a Fast and Efficient Music Source Separation Tool with Pre-trained Models. *Joss* 5, 2154. doi:10.21105/joss.02154
- Isik, U., Giri, R., Phansalkar, N., Valin, J.-M., Helwani, K., and Krishnaswamy, A. (2020). “PoCoNet: Better Speech Enhancement with Frequency-Positional Embeddings, Semi-supervised Conversational Data, and Biased Loss,” in Proceedings: Annual Conference of the International Speech Communication Association (INTERSPEECH), Shanghai, China, October 25–29, 2020, 2487–2491.
- Kim, M., Choi, W., Chung, J., Lee, D., and Jung, S. (2021). “KUIELab-MDX-Net: A Two-Stream Neural Network for Music Demixing,” in Proceedings: ISMIR 2021 Workshop on Music Source Separation, November 12, 2021.
- Le Roux, J., Wisdom, S., Erdogan, H., and Hershey, J. R. (2019). “SDR—half-baked or Well Done,” in Proceedings: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brighton, United Kingdom, May 12–17, 2019, 626–630. doi:10.1109/ICASSP.2019.8683855
- Liutkus, A., Stöter, F.-R., Rafii, Z., Kitamura, D., Rivet, B., Ito, N., et al. (2017). “The 2016 Signal Separation Evaluation Campaign,” in Proceedings: International Conference on Latent Variable Analysis and Signal Separation (LVA/ICA), Grenoble, France, February 21–23, 2017, 323–332. doi:10.1007/978-3-319-53547-0_31
- Nugraha, A. A., Liutkus, A., and Vincent, E. (2016). “Multichannel Music Separation with Deep Neural Networks,” in Proceedings: 24th European Signal Processing Conference (EUSIPCO), Budapest, Hungary, August 29–September 2, 2016, 1748–1752. doi:10.1109/eusipco.2016.7760548
- Ono, N., Rafii, Z., Kitamura, D., Ito, N., and Liutkus, A. (2015). “The 2015 Signal Separation Evaluation Campaign,” in Proceedings: International Conference on Latent Variable Analysis and Signal Separation (LVA/ICA), Liberec, Czech Republic, August 25–28, 2015, 387–395. doi:10.1007/978-3-319-22482-4_45
- Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., et al. (2018). “Image Transformer,” in Proceedings: 35th International Conference on Machine Learning (ICML), Stockholm, Sweden, July 10–15, 2018, 4055–4064.
- Rafii, Z., Liutkus, A., Stöter, F.-R., Mimilakis, S. I., and Bittner, R. (2017). The MUSDB18 Corpus for Music Separation. [Dataset]. doi:10.5281/zenodo.1117372
- Rafii, Z., Liutkus, A., Stöter, F.-R., Mimilakis, S. I., and Bittner, R. (2019). MUSDB18-HQ - an Uncompressed Version of MUSDB18. [Dataset]. doi:10.5281/zenodo.3338373
- Ronneberger, O., Fischer, P., and Brox, T. (2015). “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in Proceedings: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Munich, Germany, October 5–9, 2015, 234–241. doi:10.1007/978-3-319-24574-4_28
- Sawata, R., Uhlich, S., Takahashi, S., and Mitsufuji, Y. (2021). “All for One and One for All: Improving Music Separation by Bridging Networks,” in Proceedings: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Toronto, ON, June 6–11, 2021, 51–55. doi:10.1109/icassp39728.2021.9414044
- Stöter, F.-R., Liutkus, A., and Ito, N. (2018). “The 2018 Signal Separation Evaluation Campaign,” in Proceedings: International Conference on Latent Variable Analysis and Signal Separation (LVA/ICA), Guildford, United Kingdom, July 2–5, 2018, 293–305. doi:10.1007/978-3-319-93764-9_28
- Stöter, F.-R., Uhlich, S., Liutkus, A., and Mitsufuji, Y. (2019). Open-Unmix - a Reference Implementation for Music Source Separation. *Joss* 4, 1667. doi:10.21105/joss.01667
- Takahashi, N., and Mitsufuji, Y. (2017). “Multi-scale Multi-Band Densenets for Audio Source Separation,” in Proceedings: IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA), New Paltz, NY, October 15–18, 2017, 21–25. doi:10.1109/wasppaa.2017.8169987
- Takahashi, N., and Mitsufuji, Y. (2021). “Densely Connected Multi-Dilated Convolutional Networks for Dense Prediction Tasks,” in Proceedings: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 19–25, 2021, 993–1002. doi:10.1109/cvpr46437.2021.00105
- Uhlich, S., Giron, F., and Mitsufuji, Y. (2015). “Deep Neural Network Based Instrument Extraction from Music,” in Proceedings: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), South Brisbane, QLD, April 19–24, 2015, 2135–2139. doi:10.1109/icassp.2015.7178348
- Uhlich, S., Porcu, M., Giron, F., Enenkl, M., Kemp, T., Takahashi, N., et al. (2017). “Improving Music Source Separation Based on Deep Neural Networks through Data Augmentation and Network Blending,” in Proceedings: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), New Orleans, LA, March 5–9, 2017, 261–265. doi:10.1109/icassp.2017.7952158
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention Is All You Need,” in Proceedings: Advances in Neural Information Processing Systems, Long Beach, CA, December 4–9, 2017, 5998–6008.
- Vincent, E., Sawada, H., Bofill, P., Makino, S., and Rosca, J. P. (2007). “First Stereo Audio Source Separation Evaluation Campaign: Data, Algorithms and Results,” in Proceedings: International Conference on Independent Component Analysis and Signal Separation (ICA), London, United Kingdom, September 9–12, 2007, 552–559. doi:10.1007/978-3-540-74494-8_69
- Vincent, E., Araki, S., Theis, F., Nolte, G., Bofill, P., Sawada, H., et al. (2012). The Signal Separation Evaluation Campaign (2007–2010): Achievements and Remaining Challenges. *Signal. Process.* 92, 1928–1936. doi:10.1016/j.sigpro.2011.10.007
- Wu, Y., and He, K. (2018). “Group Normalization,” in The European Conference on Computer Vision (ECCV), Munich, Germany, September 8–14, 2018, 3–19. doi:10.1007/978-3-030-01261-8_1
- Yoshida, Y., and Miyato, T. (2017). *Spectral Norm Regularization for Improving the Generalizability of Deep Learning*. arXiv. arXiv preprint arXiv:1705.10941.
- Yu, C.-Y., and Cheuk, K.-W. (2021). “Danna-Sep: Unite to Separate Them All,” in Proceedings: The ISMIR 2021 Workshop on Music Source Separation, November 12, 2021.

Conflict of Interest: YM is employed by Sony Group Corporation. GF, SU, and F-RS are employed by Sony Europe B.V. AD is employed by Facebook AI Research.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Mitsufuji, Fabbro, Uhlich, Stöter, Défossez, Kim, Choi, Yu and Cheuk. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.