



# The Sound Demixing Challenge 2023 – Music Demixing Track

OVERVIEW ARTICLE

GIORGIO FABBRO 

STEFAN UHLICH 

CHIEH-HSIN LAI 

WOOSUNG CHOI

MARCO MARTÍNEZ-RAMÍREZ

WEIHSIANG LIAO

IGOR GADELHA

GERALDO RAMOS

EDDIE HSU

HUGO RODRIGUES

FABIAN-ROBERT STÖTER

ALEXANDRE DÉFOSSEZ

YI LUO

JIANWEI YU

DIPAM CHAKRABORTY

SHARADA MOHANTY 

ROMAN SOLOVYEV 

ALEXANDER STEMPOVSKIY 

TATIANA HABRUSEVA 

NABARUN GOSWAMI 

TATSUYA HARADA 

MINSEOK KIM

JUN HYUNG LEE 

YUANLIANG DONG 

XINRAN ZHANG 

JIAFENG LIU 

YUKI MITSUFUJI 



CORRESPONDING AUTHOR:

**Giorgio Fabbro**

Sony Europe B.V., Stuttgart,  
Germany

Giorgio.Fabbro@sony.com

KEYWORDS:

Music Source Separation; Deep Learning; Neural Networks; Robust Training; Sound; Signal Processing

TO CITE THIS ARTICLE:

Fabbro, G., Uhlich, S., Lai, C.-H., Choi, W., Martínez-Ramírez, M., Liao, W., Gadelha, I., Ramos, G., Hsu, E., Rodrigues, H., Stöter, F.-R., Défossez, A., Luo, Y., Yu, J., Chakraborty, D., Mohanty, S., Solovyev, R., Stempkovskiy, A., Habruseva, T., Goswami, N., Harada, T., Kim, M., Lee, J. H., Dong, Y., Zhang, X., Liu, J., & Mitsufuji, Y. (2024). The Sound Demixing Challenge 2023 – Music Demixing Track. *Transactions of the International Society for Music Information Retrieval*, 7(1), 63–84. DOI: <https://doi.org/10.5334/tismir.171>

\*Author affiliations can be found in the back matter of this article

## ABSTRACT

This paper summarizes the music demixing (MDX) track of the Sound Demixing Challenge (SDX'23). We provide a summary of the challenge setup and introduce the task of robust music source separation (MSS), i.e., training MSS models in the presence of errors in the training data. We propose a formalization of the errors that can occur in the design of a training dataset for MSS systems and introduce two new datasets that simulate such errors: *SDXDB23\_LabelNoise* and *SDXDB23\_Bleeding*.<sup>1</sup> We describe the methods that achieved the highest scores in the competition. Moreover, we present a direct comparison with the previous edition of the challenge (the Music Demixing Challenge 2021): the best performing system achieved an improvement of over 1.6dB in signal-to-distortion ratio over the winner of the previous competition, when evaluated on MDXDB21. Besides relying on the signal-to-distortion ratio as objective metric, we also performed a listening test with renowned producers and musicians to study the perceptual quality of the systems and report here the results. Finally, we provide our insights into the organization of the competition and our prospects for future editions.

## 1. INTRODUCTION

Audio source separation has a long history in research, partially motivated by the many applications it enables. Thanks to source separation, music producers and artists are able to make use of material that was created decades ago or under limited recording conditions. Movie production studios have now the possibility to revive old classics,<sup>2</sup> for which only single-channel master tracks exist, and bring them back to the theatre taking advantage of newer innovations, such as spatial audio (Petermann et al., 2022). Moreover, people who are experiencing hearing difficulties, a condition which makes it challenging for them to communicate in loud and noisy surroundings, can be supported by this technology and successfully engage in conversations.

In the context of research, the more recent success of audio source separation should be attributed also to the presence of benchmarks that allowed different methods to be effectively compared: one example is the MUSDB18 dataset (Rafii et al., 2017), introduced during the Signal Separation Evaluation Campaign (SiSEC) in 2018 (Stöter et al., 2018). The dataset includes 150 different songs, along with corresponding separate recordings for each musical instrument. These tracks were grouped into four classes for consistency across songs: *vocals*, *bass*, *drums*, and a final class named *other* for any other instrument present. This allowed participants to carry out training and evaluation of their models on a standardized and consistent pool of data.

In 2021 we continued the tradition of SiSEC competitions with the Music Demixing Challenge (MDX'21) (Mitsufuji et al., 2022). We decided to keep MUSDB18 as the reference dataset for training, but introduced a new benchmark for testing based on data that could not be accessed by the participants. The submissions to the MDX'21 challenge were evaluated on a new hidden dataset called MDXDB21 (Mitsufuji et al., 2022), which contained 30 songs produced by Sony Music Entertainment Japan for the purpose of the challenge.

Two years later, we organized a new edition of the challenge, named Sound Demixing Challenge 2023 (SDX'23): while MDX'21 focused exclusively on music source separation (MSS), SDX'23 presented two independent tracks, one for music source separation and one for cinematic sound separation (Uhlich et al., 2024). In the music track, we again used MDXDB21 as test benchmark. This allowed a comparison of the submissions across the two editions of the challenge.

While we wanted to offer the prospective participants a familiar research playground (e.g., keeping the usual four instrument classes introduced by MUSDB18), we complemented that with a novel aspect of the source separation problem: *robust music source separation*. In robust MSS the trained system needs to be able to handle errors and inconsistencies in the training data. We

decided to focus one part of SDX'23 on this topic after we analyzed the outcome of the previous edition of the challenge.

At the end of MDX'21 it was evident that the volume of data accessible for training can have a significant effect on the quality of the final source separation system. At the same time, since state-of-the-art methods for source separation are still dominated by supervised learning approaches, the availability of appropriate data is limited. This formulation requires individual tracks, each containing the signal of a single instrument, which are difficult to obtain, especially in large quantities. Simulation of the recordings using MIDI and virtual instruments (e.g., the Slakh dataset (Manilow et al., 2019)) can be a compromise: the size of the dataset grows very easily, but the results lack realism and have some critical limitations (e.g., they often do not contain vocals).

In the context of the challenge, rather than tackling the problem of the scarcity of the data, we assumed that such data would be available and looked further down the road: is having more songs enough to obtain better models? How can we best leverage and curate a large corpus of tracks to improve the quality of the separation model?

A set of internal experiments revealed that training source separation models on a large volume of high-quality data does not guarantee better network convergence. The convergence behavior of our models was dramatically affected by label errors in the data. These were mostly related to the ground truth labels used when training (i.e., the identity of the instruments present in each audio recording). For example, tracks labeled as *vocals* actually contained the signal of a guitar. Intuitively, the impact of such label errors should decrease proportionally to the amount of data used to train the model. Against our expectations, increasing the amount of data was causing the models to stop converging. Only through an expensive activity of data cleaning we were able to make the model converge again.

While the relative amount of erratic data definitely impacts whether this phenomenon occurs or not, our experiment showed how the difference in performance caused by the errors in the data justifies a closer look into the phenomenon. There are no aspects of our experiments that prevent us from generalizing this finding to tasks beyond source separation, or even outside the audio domain. Nevertheless, confirming the generalization of this finding would be outside the scope of this work. We leave such activity for future work. What can be drawn from our experiments is that an activity that ensures that the data is clean has a visible impact on the final performance of the model.

The process of cleaning a dataset is very expensive and does not scale easily with its size: manually

checking the annotations for a large number of tracks is not feasible. While automatic methods can provide an initial solution (e.g., audio classification and tagging models: García et al., 2021; Mahanta et al., 2021; Kong et al., 2020), some errors are intrinsically difficult to fix: if a single track erroneously contains the signal of two different instruments, simply changing its label will not solve the problem. Source separation methods can also be used to exclude noisy samples in a large dataset, to ensure that the new model is trained only on clean data (Rouard et al., 2023).

A more interesting solution would be to make the training of the network invariant to these errors. Ideally, if the learning process is robust to such inconsistencies, adding new data upon availability becomes easier, as we can avoid a cleaning activity that is expensive, likely incomplete and potentially ineffective. The research community has proposed some robust training methods to address label noise in classification problems (Han et al., 2018; Li et al., 2020; Cheng et al., 2020; Wang et al., 2022) or handle out-of-distribution samples (Lai et al., 2019, 2023; Hendrycks et al., 2021; Mukherjee et al., 2021). However, to the best of our knowledge no existing approach explicitly tackles audio source separation.<sup>3</sup> With the SDX'23 challenge we aimed to bring the attention of the research community to the topic of robust MSS and provided the participants with an environment that presented the issues outlined above, even when working with small datasets.

In this paper we summarize the music track of SDX'23: we show the challenge setup in Sec. 2, we introduce the topic of robust MSS in Sec. 3, we outline the parts of the challenge related to MSS without robustness constraints in Sec. 4, we present the challenge results, together with the descriptions of some of the winning approaches in Sec. 5 and 6 and we elaborate on the technical challenges in the organization of the competition in Sec. 7.

## 2 MDX CHALLENGE SETUP

In the following, we summarize the structure of the competition. Similarly to the previous edition, the challenge was hosted on AICrowd.<sup>4</sup>

### 2.1 TASK DEFINITION

Participants in the music track (MDX) of the Sound Demixing Challenge 2023 were asked to submit systems to separate a stereo song  $\mathbf{x}(n) \in \mathbb{R}^2$  into one stereo track for *vocals* ( $\mathbf{s}_v(n) \in \mathbb{R}^2$ ), one for *bass* ( $\mathbf{s}_b(n) \in \mathbb{R}^2$ ), one for *drums* ( $\mathbf{s}_d(n) \in \mathbb{R}^2$ ) and one for *other* ( $\mathbf{s}_o(n) \in \mathbb{R}^2$ ), where the song can be obtained as:

$$\mathbf{x}(n) = \mathbf{s}_v(n) + \mathbf{s}_b(n) + \mathbf{s}_d(n) + \mathbf{s}_o(n), \quad (1)$$

and  $n \in \mathcal{Z}$  denotes the time index. All signals are sampled at 44100 Hz.

### 2.2 LEADERBOARDS

The previous edition of the MDX challenge (Mitsufuji et al., 2022) focused on music source separation, similarly to the SiSEC competitions (Ono et al., 2015; Liutkus et al., 2017; Stöter et al., 2018). This year, we devoted two of three leaderboards to the issue of training source separation models with data containing errors. We elaborate on this task in Section 3.

On top of that, we provided a third leaderboard that was devoted to the music source separation task, without any constraint on the training data. By allowing training under any condition, we are interested in tracking the progress of the source separation community. Unlike the previous edition of the challenge, we did not offer any leaderboard for models trained exclusively on MUSDB18 (Rafii et al., 2017).

In summary, the submissions were categorized under the following three leaderboards:

- Leaderboard *LabelNoise* was designated for models trained on data suffering from *label noise*,
- Leaderboard *Bleeding* was designated for models trained on data suffering from *bleeding*,
- Leaderboard *Standard* was designated for models trained on any data.

For the definitions of *label noise* and *bleeding*, we refer the reader to Sec. 3.2.

It is important to note that the leaderboards *LabelNoise* and *Bleeding* did not allow the usage of any external trained resource, as that would have been equivalent to the usage of external (possibly clean) data.

### 2.3 RANKING METRIC

The evaluation of the systems followed the same strategy as in MDX'21: we used the global *signal-to-distortion ratio* (SDR) as metric, which is defined for one song as

$$\text{SDR} = \frac{1}{4}(\text{SDR}_v + \text{SDR}_b + \text{SDR}_d + \text{SDR}_o), \quad (2)$$

with

$$\text{SDR}_c = 10 \log_{10} \frac{\sum_n \|\mathbf{s}_c(n)\|^2}{\sum_n \|\mathbf{s}_c(n) - \hat{\mathbf{s}}_c(n)\|^2}, \quad (3)$$

where  $\mathbf{s}_c(n) \in \mathbb{R}^2$  and  $\hat{\mathbf{s}}_c(n) \in \mathbb{R}^2$  denote the stereo target and estimate for source  $c \in \{V, B, D, O\}$ . Finally, the global SDR of (2) is averaged over all songs in the hidden test dataset to obtain the final score.<sup>5</sup>

In addition to the objective metric above, this edition of the challenge also featured a subjective evaluation based on a listening test carried out on the estimates of the top three systems in the leaderboard *Standard* (ranked by their SDR score). We refer the reader to Sec. 4.2 and Sec. 6 for details on the subjective evaluation.

The total prize pool for the competition was \$32,000. \$2,000 was reserved for the listening test winner, while the rest was equally distributed among the leaderboards.

## 2.4 TIMELINE

The challenge featured two rounds. Phase I started on January 23<sup>rd</sup> 2023, while Phase II started on March 6<sup>th</sup> 2023. Due to the submission system experiencing difficulties in handling the surge in the number of submissions towards the end of the challenge, the end date of Phase II was extended by one week. Originally scheduled to conclude on May 1st, 2023, the challenge was extended to May 8, 2023, to ensure a fair competition for all teams. A *warm-up round* was also organized, which began on December 8<sup>th</sup> 2022 and lasted until the beginning of Phase I. During this round, participants could get acquainted with the submission system and prepare their submissions for the challenge.

To carry out the evaluation of the submissions, we used the same approach as in the previous edition of the challenge. The hidden test set (MDXDB21) contains 30 songs, of which we held out three songs: we used two of them to give feedback to the participants about their performance (i.e., they could listen to the separations of their model on them), while we excluded the third one because the bass track is silent. The remaining 27 songs were used to compute the scores displayed in the leaderboards. During Phase I, the scores available to the participants were computed on one third of them (nine songs). During Phase II, nine more songs were added (i.e., the scores displayed on the leaderboards were computed over 18 songs). These two subsets were selected randomly (we kept the same random subsets used in MDX'21). Once the competition ended and the participants could not submit anymore to be eligible for prizes, the final scores computed on the whole test set (27 songs) were displayed on the leaderboards.

We chose to carry out the evaluation in this way in order to prevent participants from implicitly adapting their submissions to the test set throughout the course of the challenge. Please note how, at the end of the MDX'21 challenge, all submissions from Phase II were evaluated on the 27 songs to produce the scores on the final leaderboards. This encouraged participants to maximize the number of submissions during Phase II, so that they would increase the likelihood of one of them achieving a high score on the full test set. For this reason, at the end of Phase II of SDX'23, we asked each participant in every leaderboard to *manually* select three candidate submissions that would move on to the final evaluation: only these submissions were then evaluated on the whole test set. The best out of three was then displayed as the final entry for each team.

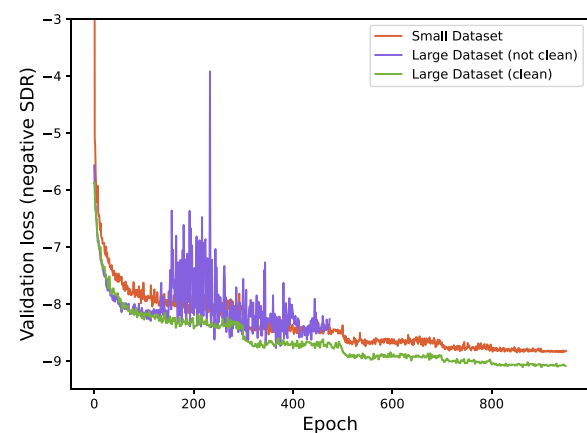
## 3 ROBUST MUSIC SEPARATION

From the results of the last edition of the challenge (Mitsufuji et al., 2022) it became clear to the source separation community that the performance of a model very often correlates with the amount of data used to train it. One example was the participant *defossez*: his model<sup>6</sup> (Rouard et al., 2023) trained only on MUSDB18 achieved an average SDR score of 7.32dB,<sup>7</sup> while training on additional data improved its performance by almost 1dB.<sup>8</sup>

It would be safe to assume that increasing the amount of data used for training increases the performance of the model, but this is not necessarily the case. While performing experiments using internal data, we experienced that a model whose validation loss used to converge when trained on a small high-quality (i.e., clean) dataset was not converging anymore when trained on a high-quality dataset one order of magnitude larger (see Figure 1). After careful inspection, we realized that label errors in this large pool of data were responsible for this phenomenon. Please note that the general practice of increasing the amount of training data to increase performance still applies. Once the large pool of data had been cleaned, we achieved a higher performance than with the smaller dataset.

### 3.1 WHY ROBUST MUSIC SEPARATION?

Intuitively, the more data we have, the lower the impact of any incorrect recording should be. However, increasing the amount of data also increases the number of incorrect recordings and the likelihood that globally the dataset presents inconsistencies. Examples of such errors can be recordings where the recorded instrument does not match the instrument label, or individual recordings that contain the signal of more than one instrument.



**Figure 1:** Comparison of validation loss when training the same model on a small dataset (red), a large dataset with errors (purple) and the same large dataset once the errors have been corrected (green). All experiments were evaluated on the same validation set.

As [Figure 1](#) shows, only through an expensive activity of *data cleaning* we were able to make those models converge again, confirming that those label errors were the cause of the issue.

Given how time-consuming and resource-expensive a data cleaning activity is, we decided to focus the research community on this issue during the Sound Demixing Challenge 2023. Our intention was to constrain the participants in the data at their disposal, so as to make them devise strategies that would allow training models *robust* to inconsistencies in the training data.

### 3.2 FORMALIZING ERRORS IN THE TRAINING DATA

We aimed at providing a common framework to characterize errors and inconsistencies in the training data from the perspective of audio source separation. We assumed that the training of a source separation model requires individual stems and does not rely exclusively on mixtures. Under these assumption, the integrity of the individual recordings becomes paramount. We identified two categories of errors that can occur:

- the recorded instrument and the label identifying it do not agree,
- the recording contains signals belonging to more than one instrument.

We named the first category **label noise**, since a (possibly) random change in the label of the recording is responsible for it; we named the second category **bleeding**, as it occurs when the signal of a second instrument *bleeds* into the recording of the first one. Two leaderboards in the challenge were dedicated to robust MSS, one for each of the two error categories (*LabelNoise* and *Bleeding*).

**Label Noise** Being a creative process, music production does not follow conventional workflows. Different music producers adhere to different conventions for their procedures, which has direct consequences in how their final deliverable is organized. Often producers are asked to deliver not only the final mix of their song, but also individual stems: these tracks usually group multiple recordings that belong to a single identifiable instrument (e.g., all microphones for the drum set). The choice of which stems to deliver and how to name them is typically influenced by the preferences and background of the producer and by the style of the song. In the context of music production there is no immediate benefit in following a strict convention when naming the stems.

When we repurpose these recordings for training source separation models, we need to collect stems coming from different music producers. For one stem, the only information we require is the identity of the recorded instrument. Typically, the only label we possess that indicates the instrument identity in a recording is

its file name. Some producers follow more organized workflows and might keep some extra metadata, but this does not hold for all of them. The stem file name is chosen manually: a process that is prone to errors and not conforming to any naming convention. The result is a very large collection of valuable audio recordings and a chaotic set of instrument names.

Even if we were able to efficiently collapse all the variations of one instrument name (e.g., *el\_guitar*, *electric\_guitar*, *el\_gtr*, etc.), we would still have some which are intrinsically ambiguous. Some examples are: *lead* (which could refer to vocals, guitar, or any other instrument playing a leading part), *choir* (which could be a human choir or a synthesizer imitating a choir), *bells* (which could refer to church bells, chime bells or a synthesizer sound), attributes used as names (such as *clean*, *dark*, *bright*), *sfx* (which could potentially refer to any sound), and others.

**Bleeding** When producing music in a studio, the priority is given to the performance. Each performer must be in the condition of delivering the perfect rendition of the artist's vision. One example is for the whole band to play together, instead of recording one instrument at a time. This improves the musical interpretation of the song, as each musician can directly react to small changes in the performance of the others: this leads to increased interactions that translate into a more truthful and lively recording. Recording studios are designed to maximize acoustic isolation between different rooms and booths, but there are limits. For example, low frequencies are notoriously difficult to isolate, due to their long wavelength: a bass amplifier produces sound that can overcome the isolation barriers in the studio and reach the microphones devoted to other instruments. If those instruments are recorded at the same time, some signal from the bass amplifier will *bleed* into their tracks. This situation arises also when recording performances with an orchestra ([Prätzlich et al., 2015](#)). This is not an issue in the context of the song production, as those signals will be summed anyway to produce the final mix. But if these recordings are now repurposed as training material for a source separation system, this bleeding becomes problematic.

### 3.3 CREATING DATASETS WITH CORRUPTIONS

In the context of the challenge, we wanted participants to tackle the issues above while ensuring that the overall competition remained fair when systems are evaluated against a common test set. Ideally, we would force the participants to train their systems using the same datasets, which have been corrupted using label noise and bleeding. On top of that, we would need to prevent access to an error-free version of those datasets.

For both categories of errors we have the option of simulating them on existing recordings. Although simulation might result in a loss of realism, it preserves

the conditions required during the setup of the challenge and does not change the underlying task that needs to be solved. For this reason, we decided to simulate both label noise and bleeding starting from error-free recordings.


Although clean recordings are easy to find in the community, we needed to avoid participants having access to the source material, as it would have given them an unfair advantage. For this reason, we made use of MoisesDB (Pereira et al., 2023), which was not released until after the end of the challenge. It contains 240 individual tracks sourced from 45 diverse artists, spanning twelve distinctive musical genres. Each track is broken down into its constituent audio components and is categorized within a hierarchical, two-tier taxonomy of stems (Manilow et al., 2020).

We selected 203 songs from MoisesDB to be the source data for our corrupted datasets. Then, we generated two versions of such data, one containing label noise (*SDXDB23\_LabelNoise*), another containing bleeding (*SDXDB23\_Bleeding*), and shared them with the participants. The new datasets follow the same structure and design of MUSDB18 and MDXDB21, where each song is composed of four stems: *vocals*, *bass*, *drums* and *other*. The original song can be recovered as the summation of the four stems.<sup>9</sup> We make both datasets available for download.<sup>10</sup>

When simulating the errors, our objective was to cause a degradation in the training loss of Open-Unmix (Stöter et al., 2019) (upon convergence) of approximately 1dB SDR, when trained individually on label noise and bleeding, compared to training the same model on error-free data. Moreover, we chose not to split the datasets into a training and validation part to give the participants more freedom to decide how (and whether) to perform validation.

**Label Noise** The simulation of label noise is based on randomly changing the label of a stem. This is applied on a subset of the stems in the clean dataset: this allowed us to effectively simulate label noise as occurs in real life. We chose to apply label noise to 20% of the stems in the source data, independently of their label. If a stem was subject to label noise, the new label was sampled from a distribution that reflected the label noise we found in our internal datasets. To do this, we collected statistics over the frequency of corrections we performed during our data cleaning activity: if a stem initially labeled as  $c_{src}$  was corrected to  $c_{dst}$ , we increased the likelihood that during the simulation of label noise we would change  $c_{dst}$  into  $c_{src}$ . These statistics are reported in Figure 2.

Please note that our statistics refer to a taxonomy of ten musical instruments, while the final dataset contains stems for four: this means that some corruptions we perform will not have an effect in the final version of the dataset. For example, changing *drums* into *bass* represents an error in the final dataset, changing *guitar* to *piano* does not (as they both belong to the class *other*).



**Figure 2:** Statistics collected during our internal data cleaning activity. The values in the rows are normalized so that they sum to 1. For example, in all the errors we found in our internal data, the chance that a guitar was labeled as bass is 32%.

In the end, 34% of the stems in *SDXDB23\_LabelNoise* were affected by label noise.<sup>11</sup>

**Bleeding** Simulating bleeding was a more elaborate process. The amount of bleeding in a recording is usually low (in terms of signal-to-noise ratio). In order to achieve the target impact of 1dB SDR on the model convergence, we decided to apply corruptions to *every single stem* in *SDXDB23\_Bleeding*. More specifically, *every stem bleeds into every other stem* in the same song.

The bleeding component in a stem was obtained as a copy of another stem, where we applied gain reduction and filtering. The scaled and filtered signal was then summed to the stem that would contain the corruption. We randomly scaled each bleeding signal in the range  $[-7, -12]$ dB and applied either a low-pass filter or a band-pass filter (we randomly selected for every stem). The order of the filter was randomly chosen in the range  $[3, 10]$ . When we applied a low-pass filter, the cut-off frequency was randomly chosen in the range  $[900, 9000]$ Hz. When we applied a band-pass filter, we randomly chose the low and high cut-off frequencies in the range  $[200, 600]$ Hz and  $[8, 10]$ kHz, respectively. All random choices were drawn from uniform distributions. The ranges were designed empirically to compromise between bleeding realism and desired impact on the model convergence.

### 3.4 ROBUST BASELINE MODEL

We present here a simple baseline model for the task of robust MSS, which can be used for both label noise and bleeding. This method is invariant to the choice of network architecture: in our experiments we used Open-Unmix (Stöter et al., 2019).

We first trained the model on the full noisy dataset  $D_1$ , without any data cleaning, and obtained a system that achieved suboptimal performance: we name this  $UMX^{(1)}$ . We then ran inference of  $UMX^{(1)}$  on the same data it was trained upon. In other words, we used  $UMX^{(1)}$  to remove

some of the errors present in  $D_1$ : although the model had only suboptimal performance, we could expect it to be able to remove part of the errors in the data. This step created a new improved version of  $D_1$ , which we name  $D_2$ . We then trained a new model  $UMX^{(2)}$ , from scratch, on  $D_2$ ; we expected this model to achieve better performance than  $UMX^{(1)}$ . We could then repeat this iterative refinement of the training data  $N$  times and train the final model  $UMX^{(N)}$ . The maximum number of iterations  $N$  was found empirically: to realize our baseline, we used  $N = 2$ . Please note how this method can be interpreted as a distillation approach, where the current model  $UMX^{(i)}$  acts as a teacher during the training of the next model  $UMX^{(i+1)}$  (the student) (Hinton et al., 2015; Wang et al., 2021; Luo and Yu, 2023).

The creation of the improved dataset at every iteration can be performed in two ways: we name them *filtered* and *redistributed*. Every song in the dataset  $D_i$  at iteration  $0 < i \leq N$  is composed of stems  $\mathbf{s}_c^{(i)}(n), \forall c \in C$ , where  $c$  is the label of the stem and  $c = \{V, B, D, O\}$  is the set of all sources. These stems are used as input to the current model  $UMX^{(i)}$  to create  $D_{i+1}$ . We denote the separation of instrument  $c \in C$  with the current model as  $UMX_c^{(i)}(\cdot)$ .

In the *filtered* method we define a new stem with label  $c$  as the following:

$$\mathbf{s}_c^{(i+1)}(n) = UMX_c^{(i)}(\mathbf{s}_c^{(i)}(n)), \quad (4)$$

that is, we only consider as input the stem of the current instrument  $c$ . Any contribution of the current instrument  $c$  present in the other stems of the same song is discarded (i.e., we discard all contributions of instrument  $c$  present in  $\mathbf{s}_{\bar{c}}^{(i)}(n), \forall \bar{c} \neq c$ ).

In the *redistributed* method we define a new stem with label  $c$  as the following:

$$\mathbf{s}_c^{(i+1)}(n) = \sum_{\bar{c} \in C} UMX_{\bar{c}}^{(i)}(\mathbf{s}_{\bar{c}}^{(i)}(n)), \quad (5)$$

that is, we sum together all the estimates for one instrument  $c$  given as input the stems in the current

song for every instrument  $\bar{c} \in C$ . The contributions of the current instrument  $c$  present in the wrong input stems (i.e., when  $\bar{c} \neq c$ ) are preserved in the new dataset and placed in the correct stem. We provide a visual comparison of the *filtered* and *redistributed* methods in Figure 3.

Please note that the model described above was not provided to the participants as a baseline during the challenge. To supply the participants with reference scores during the competition, we trained models separately on *SDXDB23\_LabelNoise* and *SDXDB23\_Bleeding* without using any robust method: we provided two baselines based on Open-Unmix (Stöter et al., 2019), two based on Hybrid Demucs (Défossez, 2021) and two based on MDX-Net (Kim et al., 2021).

## 4. STANDARD MUSIC SEPARATION

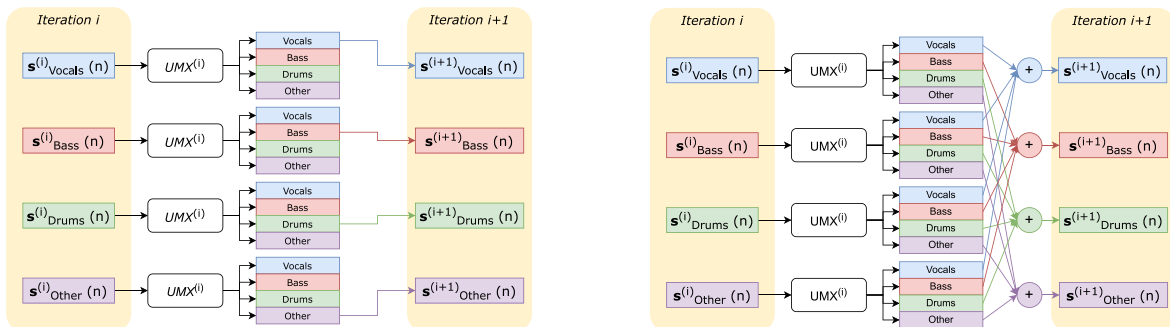
Part of the success of the last edition of the challenge was due to the fact that not only researchers in academia but also industry players participated and submitted their systems. This was made possible by the presence of a leaderboard where no constraint was given: any model of any complexity, trained on any data, could participate.<sup>12</sup>

For this edition, we offered the same in the leaderboard *Standard*. Our intention was to measure the improvement in the performance of the submitted systems with respect to the same benchmark in 2021.

### 4.1 BASELINE MODELS

We provided several baselines to the participants also for the leaderboard *Standard*. The role of a baseline model was not only to appear in the leaderboard as a reference for existing methods; each participant could choose also to use it as starting point for their submissions.

We provided Open-Unmix (Stöter et al., 2019) (the large model, *UMX-L*), Band-Split RNN (BSRNN) (Luo and



(a) The *filtered* method: each stem in the cleaned dataset (i.e., at iteration  $i + 1$ ) contains only contributions of the corresponding stem in the noisy dataset (i.e., at iteration  $i$ ).

(b) The *redistributed* method: each stem in the cleaned dataset (i.e., at iteration  $i + 1$ ) is the sum of contributions coming from all the stems of the corresponding song in the noisy dataset (i.e., at iteration  $i$ ).

**Figure 3:** The process of cleaning the stems of one song in the noisy dataset using the proposed robust baseline model. We propose two different methods: *filtered* and *redistributed*.

Yu, 2023) and Cross-Net-Open-Unmix (Sawata et al., 2021) (X-UMX-M). The former was already provided for the previous edition of the challenge, while the other two were new.

As described in the original paper (Luo and Yu, 2023), to create the BSRNN baseline we trained 4 separate BSRNN models for the four instruments in the MUSDB18 dataset (Rafii et al., 2017). We used the same band-split schemes for the four tracks, but used slightly smaller model sizes. More specifically, we used 10 BSRNN blocks instead of 12, and set the feature dimension to 80 instead of 128. The full pipeline as well as the model weights are available online.<sup>13</sup>

The X-UMX-M model is an extension of Open-Unmix (Stötter et al., 2019) where the activations of the four sources are averaged before and after the recurrent layers. This enables the exchange of information across the activations for different sources (hence “Cross-Net”). The baseline we provide corresponds to the one trained on 20k songs (hence the “M” in the name), as described by Sawata et al. (2023).

## 4.2 LISTENING TEST

For non-robust MSS only, in addition to the objective evaluation based on SDR, we organized a subjective evaluation based on a listening test where professionals in the music industry rated the separations of the best systems. The main goal of the listening test was to assess the performance of various source separation models via comparative AB sampling. We included in the listening test one model for each of the three teams that achieved the highest SDR scores on the *Standard* leaderboard (SAMI-ByteDance, ZFTurbo and kimberley\_jensen). During the listening test, each assessor was in charge of judging segments of either the output of a source separation model or its residual (i.e., we subtract the output of the model from the input mixture). The latter reveals how well the model can suppress a source, highlighting potential residues of the signal of interest that are not correctly suppressed.

The listening test was carried out online, using a web interface developed specifically for this occasion, by a panel of seven assessors who bring experience from various domains of music. This panel comprised award-winning singers, songwriters, composers, music producers, sound engineers, and an educator. The assessors were adequately trained to understand and identify common issues associated with source separation, such as distortion and artifacts. The online interface featured a brief explanation of the purpose of the test, a description of the interface and recommendations to use high-grade audio equipment.<sup>14</sup> Each assessor was expected to complete a minimum of 72 comparisons in total (possibly in multiple sessions), equivalent to three comparisons for each pair of models, across the four instrument classes plus their residuals.

To generate the evaluation data for the test, we selected ten songs from those in MoisesDB not used to create *SDXDB23\_LabelNoise* and *SDXDB23\_Bleeding*. We intentionally selected songs belonging to six different genres, to maximize their diversity. We applied the three candidate models to each song and obtained the separated signals. In order to keep the duration of the listening test feasible, we did not have the assessors evaluate the complete separated songs. Rather, we identified in each song four segments of three musical bars each, where the audio energy was sufficiently high, and we used only the separated signals of those segments. We then created the residual signals of the separations. Finally, we paired separations of the same segment obtained with different models and randomly assigned them to the assessors.

## 5. CHALLENGE RESULTS AND WINNING APPROACHES

In this section we provide the results of the MDX track of the challenge. Tables 1, 2 and 3 show the final scores on the challenge leaderboards. For each leaderboard, we

Rank	Participant	Prize	Global SDR (dB)					Submissions to <i>LabelNoise</i>	
			Mean	Bass	Drums	Other	Vocals	1st phase	2nd phase
<i>Submissions</i>									
1.	CCOM	1 <sup>st</sup>	<b>7.46</b>	<b>8.12</b>	<b>7.99</b>	<b>5.34</b>	<b>8.37</b>	7	50
2.	subatomicseer	2 <sup>nd</sup>	6.60	6.67	7.03	4.61	8.07	65	33
3.	kuielab	3 <sup>rd</sup>	6.51	6.71	6.71	4.82	7.82	99	25
4.	aim-less		6.44	6.75	7.19	4.56	7.28	10	22
5.	yang_tong		6.33	6.29	7.46	3.94	7.65	-	2
<i>Baselines</i>									
	UMX		3.01	3.77	2.84	1.62	3.83		
	Demucs		4.84	5.55	5.68	2.89	5.23		
	MDX-Net		3.49	4.26	2.84	2.42	4.42		

**Table 1:** Final *LabelNoise* leaderboard (models trained only on *SDXDB23\_LabelNoise*; top-5).



Rank	Participant	Prize	Global SDR (dB)					Submissions to <i>Bleeding</i>	
			Mean	Bass	Drums	Other	Vocals	1st phase	2nd phase
<i>Submissions</i>									
1.	kuielab	1 <sup>st</sup>	<b>6.58</b>	<b>6.98</b>	6.65	<b>4.96</b>	<b>7.74</b>	99	13
2.	ZFTurbo	2 <sup>nd</sup>	6.38	6.94	<b>6.86</b>	4.62	7.12	32	4
3.	subatomicseer	3 <sup>rd</sup>	6.31	6.33	<b>6.86</b>	4.59	7.47	65	11
4.	CCOM		6.20	6.34	6.32	4.28	7.87	7	17
5.	alina_porechina		5.87	6.01	6.10	4.09	7.30	99	118
<i>Baselines</i>									
	UMX		3.61	3.90	3.85	2.50	4.17		
	Demucs		5.33	5.90	5.56	3.69	6.19		
	MDX-Net		3.56	4.00	2.30	2.65	5.29		

**Table 2:** Final *Bleeding* leaderboard (models trained only on *SDXDB23\_Bleeding*; top-5).

Rank	Participant	Prize	Global SDR (dB)					Submissions to <i>Standard</i>	
			Mean	Bass	Drums	Other	Vocals	1st phase	2nd phase
<i>Submissions</i>									
1.	SAMI-ByteDance		<b>9.97</b>	<b>11.15</b>	<b>10.27</b>	<b>7.08</b>	<b>11.36</b>	13	5
2.	ZFTurbo	1 <sup>st</sup>	9.26	9.94	9.53	7.05	10.51	32	24
3.	kimberley_jensen	2 <sup>nd</sup>	9.18	10.06	9.47	6.80	10.40	86	134
4.	kuielab	3 <sup>rd</sup>	8.97	9.72	9.43	6.72	10.01	99	54
5.	alina_porechina		8.63	9.92	9.29	6.23	9.07	99	172
<i>Baselines</i>									
	UMX-L		6.52	6.62	6.84	4.89	7.73		
	BSRNN		6.14	5.63	6.53	4.43	7.98		
	X-UMX-M		6.30	5.85	6.87	4.42	8.04		

**Table 3:** Final *Standard* leaderboard (models trained on any data; top-5).

also report which participants received a prize. This does not correspond to the final ranking as, in order to receive a prize, the participant needed to disclose their solution. We report in section 5.1 the results of our iterative refinement baseline for robust MSS.

After the conclusion of the challenge, we contacted the top four teams in each leaderboard and invited them to contribute to this manuscript with a description of their approaches. Sections 5.2, 5.3, 5.4 and 5.5 report the systems of those teams who accepted our invitation and are written by the corresponding teams. Please note that, despite the selected teams participating in all leaderboards, we choose to describe below only their most significant contributions. For example, when one team proposed an interesting approach for robust MSS, we mainly describe their solutions for leaderboards *LabelNoise* and *Bleeding*, while for teams that achieved very high scores on non-robust MSS, we report their solution for leaderboard *Standard* and neglect their solutions to the other two leaderboards. This way, we aim to provide more insight into the scientific choices that allowed the team to achieve higher performance. At the beginning of each section we therefore indicate what kind of contribution is reported.

## 5.1 ITERATIVE REFINEMENT BASELINE

We report in Table 4 the performance of our iterative refinement baseline. First of all, we highlight the impact that the errors in the data have on the performance of the model: training on *SDXDB23\_LabelNoise* degrades the average separation quality by 1.42dB, while training on *SDXDB23\_Bleeding* degrades it by 0.83dB.

We then use these models to improve the training data, first by using our *redistributed* approach. If we train a new model on this improved dataset, average performance on *SDXDB23\_LabelNoise* increases by 0.43dB in SDR, while we experience virtually no change for *SDXDB23\_Bleeding*.

If we use our *filtered* approach to improve the training data, we see an average improvement of 0.89dB for *SDXDB23\_LabelNoise* and of 0.50dB for *SDXDB23\_Bleeding*.

Please note that we also report scores for when we apply our iterative approach to the clean data only (*MoisesDB*): we experience a loss of 0.16dB only when using the *redistributed* strategy, likely due to distortions and artifacts introduced by the model during the first iteration.

	Global SDR (dB)				
	Mean	Bass	Drums	Other	Vocals
<i>MoisesDB (203 songs)</i>					
Original dataset	4.43	4.65	5.06	<b>3.02</b>	<b>5.00</b>
Improved dataset (redistributed)	4.27	4.68	4.93	2.72	4.75
Improved dataset (filtered)	<b>4.46</b>	<b>5.07</b>	<b>5.16</b>	2.77	4.86
<i>SDXDB23_LabelNoise</i>					
Original dataset	3.01	3.76	2.83	1.62	3.82
Improved dataset (redistributed)	3.44	4.00	3.81	1.86	4.08
Improved dataset (filtered)	<b>3.90</b>	<b>4.57</b>	<b>4.57</b>	<b>2.22</b>	<b>4.25</b>
<i>SDXDB23_Bleeding</i>					
Original dataset	3.60	3.90	3.84	2.50	4.17
Improved dataset (redistributed)	3.59	3.73	4.07	2.40	4.17
Improved dataset (filtered)	<b>4.09</b>	<b>4.65</b>	<b>4.76</b>	<b>2.52</b>	<b>4.44</b>

**Table 4:** Results of our iterative refinement baseline. We use a source separation algorithm trained on corrupted data to improve the dataset: training the same model on the improved data increases the separation quality.

## 5.2 TEAM ZFTURBO (ROMAN SOLOVYEV, ALEXANDER STEPPOVSKIY, TATIANA HABRUSEVA)

We describe below our system for leaderboard *Standard*, which ranked 2<sup>nd</sup> place, with 9.26dB SDR (mean over all instruments).

**Ensemble Models for Better Separation** Nowadays the open-source community has already collected several models for music source separation and made them available for usage (Stöter et al., 2019; Défossez, 2021; Rouard et al., 2023; Hennequin et al., 2020). These models perform relatively well without the need for training them further. We argue that, in order to create a model for source separation able to score well in a competition, we do not need to train new models and can instead reuse existing ones and combine them to increase their performance.

The practice of creating ensembles of models is not new (Uhlich et al., 2017), but no clear guidelines on how to create an effective ensemble exist. Typically, the simplest solution is to find an ensemble of models that performs well on a small selection of songs and assume that the same performance will be achieved on all other songs. Our approach aims to make the process of selecting and combining models easier, faster and more effective.

For this purpose, we created an open benchmark (Solovyev et al., 2023) for sound demixing that we use to run objective evaluations using a standardized interface. This system allows us to quickly add more models and more evaluation data, with the final objective of finding the best ensemble of models available. This evaluation system is available online.<sup>15</sup> At the time of writing, our evaluation system features two datasets: a *synthetic* one, where small segments of vocals and instrumental songs are mixed together, even if they are not musically coherent, and a realistic one, called *MultiSong dataset*,

which comprises 100 one-minute tracks of varying genre. Although the data at our disposal for the evaluation still represent a relatively small subset of the possible scenarios a separation model will encounter during inference, our approach is already a major improvement with respect to the practice of manually selecting ensemble weights based on subjective listening of a small selection of songs.

**Details of the Proposed Ensemble** In the following, we describe the ensemble model that we submitted to the challenge. We found the coefficients to ensemble the individual models using our evaluation system. It is important to note how we use different combinations of models for extracting different sources. The best ensemble for one instrument likely is not the best ensemble for another instrument. This requires us to run independent evaluations on every single source to be extracted.

To extract the vocals, we combine three pretrained models: UVR-MDX1,<sup>16</sup> UVR-MDX2<sup>17</sup> (these models are part of the Ultimate Vocal Remover project<sup>18</sup>) and HTDemucs (Rouard et al., 2023) (we use the fine-tuned version available on GitHub<sup>19</sup>). Please note that UVR-MDX2 is a model trained to estimate the instrumental part of a song (i.e., to remove the vocals), so we subtract its output from the original mixture to obtain an estimate for the vocals. HTDemucs allows one to estimate the separated signals several times by shifting the input mixture over time by a small random amount. The estimations are then aggregated after compensating for the shift. We found that increasing the number of estimations (each with a different shift) increases the performance, with the improvement becoming negligible for very high values. Moreover, HTDemucs splits the input signal in overlapping windows, runs independent predictions on all of them and combines the estimations at the end.

We can choose how much overlap to have between consecutive windows. We found that maximizing the overlap increases the separation quality, at the cost of slower inference. Also, for high overlap values the gain in performance can become negligible. We report the improvements we observed using these two features in Table 5. Additionally, we found beneficial to separate the song twice: once using the original signal, once inverting its phase (this is simply achieved by multiplying the waveform by  $-1$ ). The two estimations are then averaged together after compensating for the change in phase. The outputs of UVR-MDX1, UVR-MDX2 and HTDemucs are combined with a weighted sum of the time-domain signals.

In order to extract the remaining three instruments, we found beneficial to compute the instrumental track and use it as input of the separation models instead of the original mixture. We compute the instrumental track by subtracting the estimation of the vocals from the original song. Intuitively, this simplifies the task of the remaining models, as the vocals do not need be removed anymore. We use four different variants of HTDemucs among those available on GitHub<sup>20</sup> (*demucs\_ft*, *demucs*, *demucs\_6s*, *demucs\_mmi*) and produce one estimate of each instrument for each of them. Please note that we use the phase inversion method described above for each model. We then combine the four outputs (one per model) using specific weights for each instrument. To produce the final estimates, we further combine the estimation of models for different instruments, to further make sure that the estimations are as robust/stable as possible.

Table 6 reports the SDR scores for our final model ensemble on MDXDB21 (the test set used in the challenge) and on our MultiSong dataset. For all details about our approach, the source code is publicly available on GitHub.<sup>21</sup>

### 5.3 TEAM SUBATOMICSEER (NABARUN GOSWAMI, TATSUYA HARADA)

We focus on our approaches for robust MSS, submitted to leaderboards *LabelNoise* and *Bleeding*, which ranked 2<sup>nd</sup> and 3<sup>rd</sup> place respectively. Please note that we

Shifts	Overlap Ratio	Global SDR (dB)				
		Mean	Bass	Drums	Other	Vocals
2	0.5	9.43	12.15	11.35	5.81	8.40
4	0.75	9.47	12.22	11.40	5.84	8.41
1	0.95	9.48	12.24	<b>11.41</b>	5.84	<b>8.43</b>
10	0.95	<b>9.49</b>	<b>12.25</b>	<b>11.41</b>	<b>5.85</b>	<b>8.43</b>

**Table 5:** (Team *ZFTurbo*) Separation performance varying the number of shifts and overlap during the inference of HTDemucs. Increasing both lead to higher performance, with marginal improvements for very high parameter values.

Dataset	Global SDR (dB)				
	Mean	Bass	Drums	Other	Vocals
MultiSong MVSep	<b>10.11</b>	<b>12.68</b>	<b>11.68</b>	6.67	9.62
MDXDB21 (18 songs)	9.41	9.87	9.52	<b>7.43</b>	<b>10.81</b>
MDXDB21 (27 songs)	9.25	9.94	9.53	7.05	10.51

**Table 6:** (Team *ZFTurbo*) SDR scores for the final ensemble on our MultiSong dataset (Solovyev et al., 2023) and on MDXDB21. We report separately the scores visible during the competition (only on 18 songs) and at the end (on 27 songs).

use different models for the two leaderboards. Then, we briefly present the system we submitted to the leaderboard *Standard*, which ranked 7<sup>th</sup> place.

**Discrete Wavelet Transform in MSS** Signals of different musical instruments have distinct features at several resolutions. We argue that exploiting the resolutions of such features can be helpful in carrying out tasks such as source separation. For this reason, we propose to take advantage of the multi-resolution analysis capabilities of the discrete wavelet transform (DWT) and introduce two new models for source separation that make use of it. Moreover, we propose a noise-robust training scheme to train models for robust MSS.

The existing HTDemucs (Rouard et al., 2023) processes the signals in two separate branches, one for time-domain waveforms and one for spectrograms. The two branches communicate using a cross-transformer. We propose to add a third branch, which operates on the DWT of the signals. This branch is based on a stack of encoders, decoders and cross-transformers. The branch that processes spectrograms acts as a residual bridge between the branch that processes the DWT and the one that processes time-domain waveforms. We call this architecture *Wavelet HTDemucs* (WHTDemucs).

We propose also a second model, based on a simpler architecture, which we name *DWT Transformer UNet* (DTUNet). It is based on two separate branches: one for the time-domain waveform, one for the DTW. A single cross-transformer allows communication between the two branches. The outputs of the two branches are element-wise summed. We also apply to the final output a source-independent post-filter, composed of a stack of three convolutional layers with GELU activations.

**Robust Training** To train models for robust MSS, we use the following strategy. From the start of the training, in order to fix the source permutation at the output of the model, we employ a reconstruction objective that is not permutation invariant. We employ the L1 loss, which has been shown to work well for supervised audio source separation training (Braun and Tashev, 2020; Rouard et al., 2023; Luo and Yu, 2023). Throughout the training, we rely on the Mixture Consistency loss (Wisdom et al., 2019) to prevent the model from collapsing to outputting silence.

Additionally, we employ the unsupervised MixIT loss (Wisdom et al., 2020). Having four sources, we group them in two pairs (randomly). The signals in each pair are mixed together. These two mixes represent the separation targets, given the original mix as input. Moreover, we incorporate the Mean Teacher loss (Tarvainen and Valpola, 2017), where a student is trained to estimate the output of a teacher model (in our case, the exponential moving average (EMA) of the model parameters). We experiment with two variants of the Mean Teacher loss. In version 1 (V1) we compute the mean teacher target by mixing the outputs of the EMA model for each noisy input stem. In version 2 (V2) we only use the output of the EMA model on the input mixture. When using V1, we apply it after 30k iterations; when using V2, we apply it after 75k iterations, using a smaller batch size.

In order to address the task of leaderboard *LabelNoise*, we also design a method to remove the noisy stems from the training data. We train a DTUNet model using V2 of the Mean Teacher loss. We then separate all the individual stems in the training data and compute the SDR between them and the model output corresponding to the stem label. If the SDR value is higher than 9dB, we consider the stem potentially clean. We then manually check a subset of the stems whose separation exceeds 9dB and only keep those effectively clean. Finally, we train a BSRNN model (Luo and Yu, 2023) (one per instrument) on this clean data. The final submission is an ensemble of a WHTDemucs (trained on the full *SDXDB23\_LabelNoise* with V1), a DTUNet (trained on the full *SDXDB23\_LabelNoise* with V2) and a BSRNN (trained on the clean data). We select the ensemble weights based on the performance of the individual models on the leaderboard and by manually inspecting the separations. Table 7 reports the performance of the two proposed models and of the final ensemble for the Label Noise task.

Manually selecting a clean subset of the dataset as described above could only be done for the training set with label noise, since all the stems in *SDXDB23\_Bleeding* were corrupted. Therefore, our submission for leaderboard *Bleeding* comprises an ensemble of a WHTDemucs (trained on *SDXDB23\_Bleeding* with V1) and a DTUNet (trained on *SDXDB23\_Bleeding* with V2). We select the ensemble weights based on the performance

Model (Mean Teacher loss)	Global SDR (dB)				
	Mean	Bass	Drums	Other	Vocals
WHTDemucs (V1)	5.93	6.41	5.73	4.42	7.17
DTUNet (V2)	5.93	5.84	6.71	4.10	7.08
Blend	<b>6.60</b>	<b>6.70</b>	<b>7.03</b>	<b>4.61</b>	<b>8.07</b>

**Table 7:** (Team *subatomicseer*) Our scores on the *LabelNoise* leaderboard.

of the individual models in the leaderboard. Table 8 reports the performance of the two proposed models and of the final ensemble for the Bleeding task.

**Standard Leaderboard** For the leaderboard *Standard*, we train a DTUNet and a BSRNN on all the data at our disposal: MUSDB18 (Rafii et al., 2017), MedleyDB-V1 (Bittner et al., 2014), and stems from our private collection, for a total of 347 songs. Our final submission is an ensemble of those two models and HTDemucs (the default model available on GitHub<sup>22</sup>). We select the ensemble weights that maximize the score on our validation set. Table 9 reports the results of the three individual models on our validation set.

#### 5.4 TEAM *CCOM* (YUANLIANG DONG, XINRAN ZHANG, JIAFENG LIU)

We focus on our approach for leaderboard *LabelNoise*, which ranked 1<sup>st</sup> place, with 7.46dB SDR (mean over all instruments).

**Robust Training** Our method to train a model for robust MSS is composed of two steps: first, we train the model by making the loss robust to noisy ground truth stems, then we use the trained model to detect and exclude noisy stems in the training data and use the new dataset to train a better model.

First, we make the training loss robust to noisy stems using the *loss truncation* technique. The idea of *loss truncation* was introduced by Kang and Hashimoto (2020). Suppose an oracle model exists that can perfectly separate each stem (i.e., with infinite SDR). With this model, correct samples in the training data would have zero loss, while noisy samples would have a loss greater than zero. We can consider this oracle model a perfect classifier for the classes *clean labels* and *noisy labels*, using the (quantile of the) loss as classification criterion. In practice, in a batch of samples, we can sort the

Model (Mean Teacher)	Global SDR (dB)				
	Mean	Bass	Drums	Other	Vocals
WHTDemucs (V1)	5.86	5.90	5.61	<b>4.68</b>	7.25
DTUNet (V2)	5.62	5.37	6.18	3.92	7.00
Blend	<b>6.31</b>	<b>6.33</b>	<b>6.86</b>	4.59	<b>7.47</b>

**Table 8:** (Team *subatomicseer*) Our scores on the *Bleeding* leaderboard.

Model (Training Songs)	Global SDR (dB)				
	Mean	Bass	Drums	Other	Vocals
DTUNet (347)	8.79	8.75	10.65	6.76	8.99
BSRNN (347)	8.65	8.06	<b>10.80</b>	6.38	<b>9.37</b>
HTDemucs (800)	<b>9.19</b>	<b>9.68</b>	10.76	<b>7.17</b>	9.15

**Table 9:** (Team *subatomicseer*) Performance of the individual models of our ensemble on our validation set. Please note that HTDemucs is trained with more data than our internal models.

samples by their loss value in descending order, calculate some quantile of the losses as a threshold and discard the samples above such threshold. Such an oracle model does not exist in practice, and the rules of the challenge prevent us from using an existing source separation model trained on clean data as an approximation. Therefore, we assume that a source separation model trained on noisy data is good enough at approximating the oracle method and train HTDemucs on *SDXDB23\_LabelNoise* using *loss truncation* from the beginning (i.e., the model being trained and the one approximating the oracle are the same).

In order to further improve the separation quality, we design a process to remove the stems with corrupted labels from the dataset. We start from the assumption that the HTDemucs model trained with *loss truncation* already has the ability to discriminate between a stem with a noisy label and one with a correct label. We then use this model to separate every individual stem in the dataset into the four instruments and measure the energy of each estimate. For example, by processing a stem labeled *vocals*, we obtain four estimates (one per instrument). If the stem is clean (i.e., if the label is correct), only the estimate of the vocals should have positive energy. On the other hand, if the stem is not clean, some other estimate will have higher energy. In practice, we compare the energy of the estimate corresponding to the stem label and, if it is at least 20dB higher than the energy of the other estimates, we consider the stem clean. After this procedure, we create a cleaned dataset of 517 stems (out of 812 stems). We repeat this process another time, by retraining HTDemucs on the cleaned dataset (without *loss truncation*), and create a new set of 519 clean stems. This amounts to excluding approximately 36% of all the stems. Finally, we use these 519 stems to fine-tune the previous HTDemucs and further improve the results. Due to limitations in time and computational resources, we only conduct two iterations of this process. In principle, with more iterations we could further clean the dataset and yield models that perform better.

In general, since we are using a model on the same data we trained it upon, the model could estimate the wrong separation as it might have *memorized* the wrong stem during the training procedure. We implicitly address this since the model processes individual stems, instead of mixtures (during training, the model always experienced mixtures at the input). In addition to that, we choose to explicitly address this by changing the parameters for the augmentations between training and inference. During training, we apply random pitch-shifting on the audio by  $\pm 2$  semitones and randomly change the tempo by up to 12%. Instead, during inference, we transpose the input stems up by 6 semitones and double the speed of the original signal.

We report in Table 10 the results of our approach, including all intermediate steps of our procedure.

Training Setup	Global SDR (dB)				
	Mean	Bass	Drums	Other	Vocals
Baseline	4.96	5.07	5.76	3.14	5.85
With <i>loss truncation</i>	6.26	6.94	6.62	4.45	7.09
With filtered data (1 <sup>st</sup> )	6.89	7.34	7.58	4.88	7.74
With filtered data (2 <sup>nd</sup> )	<b>7.46</b>	<b>8.12</b>	<b>7.99</b>	<b>5.34</b>	<b>8.37</b>

**Table 10:** (Team CCOM) Performance of HTDemucs using our approach. The baseline is trained on *SDXDB23\_LabelNoise*, then we train a model using *loss truncation* only. We use this model to filter the dataset (denoted with 1<sup>st</sup> in the Table) and train a new model. Finally, we repeat the dataset filtering (denoted with 2<sup>nd</sup>) and fine-tune the model to obtain the best performance.

## 5.5 TEAM KUIELAB (MINSEOK KIM, JUN HYUNG LEE)

We focus on our approaches for robust MSS, submitted to leaderboards *LabelNoise* and *Bleeding*, which ranked 3<sup>rd</sup> and 1<sup>st</sup> place respectively. Please note that we use different models for the two leaderboards. Then, we briefly present the system we submitted to the leaderboard *Standard*, which ranked 4<sup>th</sup> place.

Our participation in the previous edition of the challenge (Mitsufuji et al., 2022) resulted in the creation of an improved TFC-TDF-UNet model (named v2) (Kim et al., 2021). In the context of this edition of the challenge, we propose TFC-TDF-UNet v3, which achieves higher separation quality than its predecessor and reduces the inference time. Additionally, we also propose a training strategy to handle errors in the training data, for models trained for robust MSS.

**Improvements to the TFC-TDF-UNet Model** TFC-TDF-UNet is a music source separation model first introduced by Choi et al. (2020): it uses a UNet architecture that processes the complex spectrogram of the mixture and estimates the complex spectrogram of the source. Unlike Open-Unmix (Stöter et al., 2019), this model does not estimate a mask for the spectrogram, but directly estimates all components of the separated source. The first version employed time-frequency convolutions (TFC) together with a time-distributed fully-connected (TDF) block as fundamental building blocks. The TFC-TDF-UNet v2 (Kim et al., 2021) was proposed in the context of the MDX'21 challenge, as a way to reduce the computational complexity of the first version, while achieving a higher separation quality.

The new version of TFC-TDF-UNet (v3), which we submitted to the SDX'23 challenge, features several architectural improvements. First, the architecture of the model is now based on the ResUNet structure (Zhang et al., 2017), where a TDF block is added in every residual path. We employ a channel-wise sub-band filter-bank (Liu et al., 2020) while also increasing the size of the frequency dimension in our spectrograms, to improve the spectral resolution of the signals. We train a single

model that estimates all sources, instead of training one model for each source. We found particularly effective to concatenate the input spectrogram to the model activation before the final convolutional layer. We replace batch normalization and ReLU activations with instance normalization and GELU activations respectively. We employ L2 on the time-domain waveform as training loss, instead of the L1 loss.

**Robust Training** For the models trained for robust MSS, we employ a simple strategy to make the training procedure less sensitive to the errors in the ground truth data, based on the idea of *loss truncation* (Kang and Hashimoto, 2020). Please note how this approach is the same as that used by the CCOM team, described in Section 5.4. We assume that, in the presence of a relatively good model, an error in the ground truth data will generate a high loss value, even if the estimation made by the model is good. In this case, the gradients computed on this data will not necessarily improve the quality of the model, so we choose to exclude them. In other words, when we compute the loss for a minibatch during training, we discard the samples with a higher loss value and only update the model parameters using the gradients for the samples with a low loss value. This approach is sufficient when training on data corrupted with label noise, but not with bleeding, as bleeding occurs in all stems in the dataset. In this case, we found it beneficial to discard not only samples along the batch dimension, but also along the time dimension. In other words, only the time instants of a sample with a low loss value contribute to the learning procedure.

For leaderboards *LabelNoise* and *Bleeding*, we trained TFC-TDF-UNet v3 only on *SDXDB23\_LabelNoise* and *SDXDB23\_Bleeding* respectively. In both trainings, we use the strategy above to make the optimization robust to noisy stems. Table 11 reports an ablation study on the strategy we propose to enable noise-robust training. The results show that our strategy is an effective method to handle the errors in the training data, especially for label noise.

**Standard Leaderboard** For leaderboard *Standard*, our final submission is an ensemble of TFC-TDF-UNet v3, HDemucs<sup>23</sup> (commonly known as *Demucs v3*) and

Task	Loss Truncation	Global SDR (dB)				
		Mean	Bass	Drums	Other	Vocals
<i>LabelNoise</i>	No	5.05	5.31	5.31	3.45	6.12
<i>LabelNoise</i>	Yes	<b>6.26</b>	<b>6.43</b>	<b>6.38</b>	<b>4.64</b>	<b>7.58</b>
<i>Bleeding</i>	No	5.80	6.11	5.86	4.36	6.87
<i>Bleeding</i>	Yes	<b>6.22</b>	<b>6.58</b>	<b>6.20</b>	<b>4.69</b>	<b>7.41</b>

**Table 11:** (Team *kuielab*) Ablation study on loss truncation. Please note that these are the scores of an individual TFC-TDF-UNet v3 model, not of the final ensemble.

HTDemucs (also known as *Demucs v4*). For all details on the models used in the submission, please refer to Kim and Lee (2023)

Table 12 shows a comparison of the new version of TFC-TDF-UNet (v3) with the previous model (v2). Not only do we achieve higher SDR score on all instruments, but we also increase the inference speed on GPU.

## 6. LISTENING TESTS RESULTS

In this section, we report the results of the listening test. Only the teams who achieved the three highest SDR scores in the leaderboard *Standard* (*SAMI-ByteDance*, *ZFTurbo* and *kimberley\_jensen*) participated in the listening test. The best performing model of each team was used to produce the audio clips used in the test.

The assessors were expected to complete at least 72 comparisons each. At the end of the test, they performed a total of 583 comparisons. Each comparison involved the separated outputs produced by two of the three models under evaluation and the original mixture to serve as reference.

First, we quantified the interactions of each assessor with the audio samples. From the data gathered on the Moises.ai testing platform, we found that the assessors spent on average  $27.95 \pm 17.88$  seconds to perform one comparison. The assessors performed an average of  $3.13 \pm 1.93$  switches between the segments of each comparison.

We report the global results of the listening test in Figure 4. The table in Figure 4 shows the number of times a model in each row won an evaluation against a model in each column. We also report a second table (Figure 4b) showing the same results normalized by the number of evaluations for each pair of models.

In order to detect potential biases as a result of our choice of assessors, we group them into two categories: *Producer* and *Musician-Educator* (we consider musicians and educators in the same category, since our assessor panel only includes one educator, who is also a performer). We report the results of the listening test independently for the two categories in Figure 5. Figure 5a shows the results for *Producer*, while Figure 5b shows the results for *Musician-Educator*. We observe that assessors in the category *Producer* showed a higher

Model	Global SDR (dB)					
	Mean	Bass	Drums	Other	Vocals	Speed
v2	7.03	6.85	6.87	5.44	8.96	12.8x
v3	<b>7.90</b>	<b>7.36</b>	<b>8.81</b>	<b>6.19</b>	<b>9.22</b>	<b>15.0x</b>

**Table 12:** (Team *kuielab*) Comparison of TFC-TDF-UNets v2 and v3 on the MUSDB18-HQ benchmark. *Speed* denotes the relative GPU inference speed with respect to real-time on the challenge evaluation server.

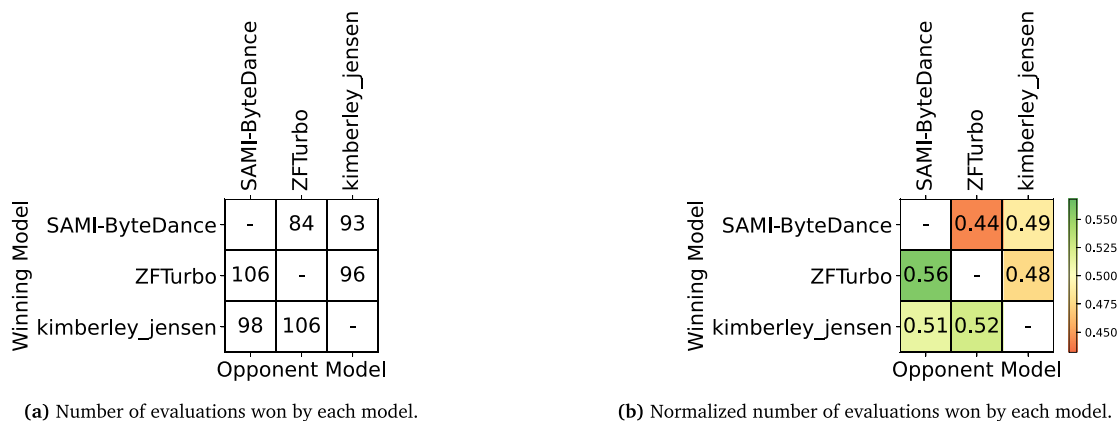


Figure 4: Results of the listening test.




Figure 5: Results of the listening test by assessor category.

preference for *kimberley\_jensen*, while those in *Musician-Educator* preferred *SAMI-ByteDance*. Given the limited number of assessors that took part in the listening test, it remains unclear whether this preference is the result of some statistical anomaly, or whether people with different backgrounds do really prefer one model over the other. We leave this investigation to future work.

The audio segments used in the test were obtained by either extracting one of the four instruments, or by

removing it. Figures 6, 7, 8 and 9 show the results on extracting and removing bass, drums, other and vocals respectively.

### 6.1 TRUESKILL RATINGS

In order to generate a valid global ranking of the three models, we employed the TrueSkill ranking system (Herbrich et al., 2007) to summarize the results of our test. TrueSkill generates a ranking based on a series




Figure 6: Results of the listening test on bass removal and extraction.




Figure 7: Results of the listening test on drum removal and extraction.




Figure 8: Results of the listening test on other removal and extraction.




Figure 9: Results of the listening test on vocal removal and extraction.

of matches (i.e., the comparisons in our listening test) between pairs of players (i.e., the models under evaluation). Each match refines the estimation of a mean value ( $0 < \mu < 50$ ) and a standard deviation ( $\sigma$ ) which represent the perceived skill of the competitor (the higher the better) and a confidence on the estimation of such skill (the lower the more confident), respectively. The procedure starts with  $\mu = 25$  and  $\sigma = 8.333$ . The final ranking is shown in Table 13. We see that the differences between the three models are very small, and that the confidence is very high (due to a low  $\sigma$ ) This might be due to the models having similar performance, but also to the relatively contained size of the listening test.

Finally, TrueSkill enables us to compute the probability of a hypothetical draw between any two models. For the match SAMI-ByteDance vs ZFTurbo the draw probability is 0.981, for the match ZFTurbo vs kimberley\_jensen it is 0.980 and for the match SAMI-ByteDance vs kimberley\_jensen it is 0.975. Given how high the draw probabilities between the teams are, we decided to equally split the prize pool of the listening test among the three teams.


## 6.2 HUMAN PREFERENCE VS OBJECTIVE METRIC

Figure 10 shows how often the human judgments agree with the objective scores based on SDR. Each point refers



	Model	$\mu$	$\sigma$	SDR (Mean)
1	<i>kimberley_jensen</i>	<b>24.793</b>	0.779	9.18
2	<i>ZFTurbo</i>	24.362	0.779	9.26
3	<i>SAMI-ByteDance</i>	24.011	0.779	<b>9.97</b>

**Table 13:** Final ranking obtained with TrueSkill. We used the default parameters for each player ( $\mu = 25$  and  $\sigma = 8.33$ ). We report the average SDR score on leaderboard *Standard* as reference.



**Figure 10:** Correlation between the SDR scores and the results of the listening test.

to a single evaluation of the listening test. We encode in the horizontal axis the SDR difference between the two competing models in that evaluation. In the vertical axis we encode whether the model with higher SDR was also selected by the assessor. As we have multiple human judgements for the same combination of model pair, song and instrument extraction or separation, we average those results.

The figure does not show a clear correlation between objective scores and the judgements made by the listeners (this is in line with the findings of Torcoli et al., 2021). The quality of all the three models participating in the listening test is high and this makes it difficult to confidently choose the best one. Nevertheless, we find that the model by *kimberley\_jensen* achieved first place in the final ranking, despite being third in the leaderboard obtained using SDR.

## 7. ORGANIZING THE CHALLENGE AND FUTURE EDITIONS

At the end of MDX'21, we stated (Mitsufuji et al., 2022) that source separation can still bring benefits to many application and research areas, and this motivated the need for future editions of the competition. For this reason, we organized the Sound Demixing Challenge (SDX'23). Our aim was to provide once more the familiar benchmark that participants knew already, but also expand its formulation in various directions. At the

time of MDX'21, we designed the competition so that researchers both new and experienced on this topic would be able to evaluate their models against a new test set. We tried to keep the competition fair, by hiding the test data from the participants, and we maximized the visibility of the event by hosting it on AICrowd.<sup>24</sup>

In 2023, the principles we used in organizing the challenge have not changed. Rather, we built upon them and expanded the scope of the whole competition, to include more application areas. Including cinematic sound separation (specifically, dedicating a whole challenge track to it, Uhlich et al., 2024)) and introducing robust MSS are two directions that we wanted to highlight as promising items for research. Among many reasons, this expansion has been successful also thanks to an enlarged pool of organizers and entities involved: Sony, Moises and Mitsubishi Electric Research Labs joined forces and shared efforts and resources in order to realize a bigger competition.

### 7.1 ROBUST MUSIC SOURCE SEPARATION

Our objective of having participants develop solutions for robust training of separation models has been met only partially. The provision of two new training sets (one for label noise and one for bleeding) has been positively received. Given how low the availability of training data is in the research world, the participants saw this as an opportunity for more experimentation and this contributed to the success of leaderboards A and B.

On the other hand, the necessarily small size of the datasets we provided (203 songs) allowed the participants to resort to manual cleaning strategies (at least in the case of label noise): some participants spent valuable time identifying which stems were corrupted and excluded them from the dataset. From the perspective of the challenge, this is acceptable; at the same time, though, it undermined our initial motivation for exploring such a topic, as such an activity would not scale with the amount of data and, as a consequence, would not be realistic. Some participants reported that bleeding was clearly more difficult to tackle than label noise: we believe this to be related to the fact that *SDXDB23\_Bleeding* contained exclusively problematic stems, while *SDXDB23\_LabelNoise* only had approximately 34% of problematic stems. For bleeding, no manual selection of the data was possible (i.e., the settings were more realistic): as a consequence, it was perceived as more difficult to solve.

During the course of the challenge, many participants asked us whether we would allow the usage of pretrained models available on the Internet. We explicitly forbade the usage of such pretrained models, as that would have enabled the participants to clean the data with e.g. existing high quality source separation models. This would have again undermined our initial motivation for exploring robust training of source separation models.

On top of that, although using a pretrained model is a practically viable solution for corrupted data, it would have threatened the aspects related to fairness in the competition, enabling different teams to use different tools. For this reason, we prevented participants from potentially achieving higher scores, but ensured that everyone would use the same resources to develop their models.

## 7.2 STANDARD MUSIC SOURCE SEPARATION

Overall, the major trend in the solutions still remains blending multiple models (Uhlich et al., 2017). What the participants provided were in general not individual neural networks, but many, whose outputs were combined together to create the final estimates. We observed the same phenomenon also in MDX'21: a possibility is that, by providing trained baseline models to the participants, we implicitly encouraged this behavior. Although in SDX'23 we provided significantly less baselines than in the previous edition, the winning models from MDX'21 were made open source after the challenge, so they represented a viable alternative to our baselines.

We noticed a larger degree of interest towards the *Standard* leaderboard than towards the other two: source separation without robustness constraints was still the most competitive playground we offered. This allowed us to compare the evolution of the systems between the two editions of the challenge. In MDX'21, the highest score on the final leaderboard was achieved by the team *Audioshake*, with an average SDR score of 8.33dB. In SDX'23, the highest score on the final leaderboard was achieved by the team *SAMI-ByteDance*, with an average SDR score of 9.97dB. In other words, the highest public score achieved on our benchmark has increased by approximately 1.64dB over the course of two years.

Finally, running a full listening test on the top-performing submissions allowed us to get an alternative source of evaluation for the separation quality, besides the SDR score. This made possible the involvement of professionals in the music industry, who represent potential users of the technology: their feedback is therefore a very important signal to take into consideration when judging the quality of a system.

## 8. CONCLUSIONS

This paper summarized the music demixing (MDX) track of the Sound Demixing Challenge 2023. We provided a description of the challenge setup, we presented the topic of robust MSS and formalized the errors that can occur in a training dataset for source separation: label noise and bleeding. We explained how we realized two new datasets for robust MSS: *SDXDB23\_LabelNoise* and *SDXDB23\_Bleeding*. Then, we described the outcome of

the challenge and reported the final results, together with a description of the winning approaches. We detailed how the evaluation has taken place, in particular with the introduction of a listening test specifically carried out with professional figures in the music industry. We believe that the SDX'23 challenge has given benefits to the source separation community and hope that we will continue organizing a long series of competitions in the future.

## NOTES

- 1 The datasets are available for download at <https://developer.moises.ai/research#datasets>.
- 2 <https://www.youtube.com/watch?v=jcWINJxnw70>.
- 3 One paper (Koo et al., 2023) was published during the preparation of this article which already makes use of our new proposed dataset *SDXDB23\_LabelNoise*. Their approach is not based on training the model on the noisy data, but on automatically correcting the labels in the data before training the model.
- 4 <https://www.aicrowd.com/challenges/sound-demixing-challenge-2023>.
- 5 Please note that this evaluation metric computes the mean over all time samples and over all songs. This is different from the evaluation metric used in the SiSEC competitions, where the SDR score of a model is the median over the scores on all songs in the test set, and the SDR score on a single song is computed as the median of the SDR scores over segments of one second.
- 6 <https://github.com/facebookresearch/demucs/tree/v3>.
- 7 [https://www.aicrowd.com/challenges/music-demixing-challenge-ismir-2021/leaderboards?challenge\\_leaderboard\\_extra\\_id=868&challenge\\_round\\_id=886](https://www.aicrowd.com/challenges/music-demixing-challenge-ismir-2021/leaderboards?challenge_leaderboard_extra_id=868&challenge_round_id=886).
- 8 [https://www.aicrowd.com/challenges/music-demixing-challenge-ismir-2021/leaderboards?challenge\\_round\\_id=886](https://www.aicrowd.com/challenges/music-demixing-challenge-ismir-2021/leaderboards?challenge_round_id=886).
- 9 Please note that summing the four stems will not yield the exact mixture for *SDXDB23\_Bleeding* as we simulate it by adding bleeding components to the original stems.
- 10 <https://developer.moises.ai/research#datasets>.
- 11 Please note how, despite applying label noise only on 20% of the stems in the source data, since several stems were grouped into one (*other*) for every song, the relative amount of corrupted stems increased to 34%.
- 12 We only imposed a limit on the model's inference time on a GPU.
- 13 <https://gitlab.aicrowd.com/Tomasysu/sdx-2023-music-demixing-track-starter-kit>.
- 14 Please note that, since the test was run online, we did not have control over the equipment that the assessor used to perform the test.
- 15 [https://mvsep.com/quality\\_checker/](https://mvsep.com/quality_checker/).
- 16 Checkpoint "Kim\_Vocal\_1.onnx" available at [https://github.com/TRvlvr/model\\_repo/releases/download/all\\_public\\_uvr\\_models/Kim\\_Vocal\\_1.onnx](https://github.com/TRvlvr/model_repo/releases/download/all_public_uvr_models/Kim_Vocal_1.onnx).
- 17 Checkpoint "UVR-MDX-NET-Inst\_HQ\_2.onnx" available at [https://github.com/TRvlvr/model\\_repo/releases/download/all\\_public\\_uvr\\_models/UVR-MDX-NET-Inst\\_HQ\\_2.onnx](https://github.com/TRvlvr/model_repo/releases/download/all_public_uvr_models/UVR-MDX-NET-Inst_HQ_2.onnx).
- 18 <https://github.com/Anjok07/ultimatevocalremovergui>.
- 19 Checkpoint *htdemucs\_ft* available at <https://github.com/facebookresearch/demucs>.
- 20 <https://github.com/facebookresearch/demucs>.
- 21 <https://github.com/ZFTurbo/MVSEP-MDX23-music-separation-model>.
- 22 <https://github.com/facebookresearch/demucs>.
- 23 <https://github.com/facebookresearch/demucs/tree/v3>.
- 24 <https://www.aicrowd.com/challenges/music-demixing-challenge-ismir-2021>.

## ACKNOWLEDGEMENTS

The authors would like to thank Sony Music Entertainment Japan for the creation of MDXDB21. The authors also thank the assessors who took part in the listening test.

## FUNDING INFORMATION


The work of Nabarun Goswami and Tatsuya Harada was partially supported by JST Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015, and Basic Research Grant (Super AI) of the Institute for AI and Beyond of the University of Tokyo. The work of Minseok Kim and Jun Hyung Lee was supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) (No. NRF-2021R1A2C2011452).

## COMPETING INTERESTS

The authors have no competing interests to declare.

## AUTHOR AFFILIATIONS

**Giorgio Fabbro**  [orcid.org/0000-0002-6302-9033](https://orcid.org/0000-0002-6302-9033)  
Sony Europe B.V., Stuttgart, Germany

**Stefan Uhlich**  [orcid.org/0000-0003-3158-4945](https://orcid.org/0000-0003-3158-4945)  
Sony Europe B.V., Stuttgart, Germany

**Chieh-Hsin Lai**  [orcid.org/0009-0009-3059-929X](https://orcid.org/0009-0009-3059-929X)  
Sony Group Corporation, Tokyo, Japan

**Woosung Choi**  
Sony Group Corporation, Tokyo, Japan

**Marco Martínez-Ramírez**  
Sony Group Corporation, Tokyo, Japan

**Weihsiang Liao**  
Sony Group Corporation, Tokyo, Japan

**Igor Gadelha**  
Moises.ai, João Pessoa, Brazil

**Geraldo Ramos**  
Moises.ai, Salt Lake City, USA

**Eddie Hsu**  
Moises.ai, João Pessoa, Brazil

**Hugo Rodrigues**  
Moises.ai, Salt Lake City, USA

**Fabian-Robert Stöter**  
AudioShake, San Francisco, USA


**Alexandre Défossez**  
Meta AI, Paris, France


**Yi Luo**  
Tencent AI Lab, Shenzhen, China


**Jianwei Yu**  
Tencent AI Lab, Shenzhen, China

**Dipam Chakraborty**  
AICrowd, Bengaluru, India

**Sharada Mohanty**  [orcid.org/0000-0003-4578-0170](https://orcid.org/0000-0003-4578-0170)  
AICrowd, Lausanne, Switzerland

**Roman Solovyev**  [orcid.org/0000-0003-0312-452X](https://orcid.org/0000-0003-0312-452X)  
Institute for Design Problems in Microelectronics of Russian Academy of Sciences, Moscow, Russian Federation


**Alexander Stempkovskiy**  [orcid.org/0000-0001-5585-014X](https://orcid.org/0000-0001-5585-014X)  
Institute for Design Problems in Microelectronics of Russian Academy of Sciences, Moscow, Russian Federation


**Tatiana Habruseva**  [orcid.org/0000-0003-3940-8639](https://orcid.org/0000-0003-3940-8639)  
Independent researcher, Cork, Ireland


**Nabarun Goswami**  [orcid.org/0000-0002-3960-5627](https://orcid.org/0000-0002-3960-5627)  
The University of Tokyo, Japan


**Tatsuya Harada**  [orcid.org/0000-0002-3712-3691](https://orcid.org/0000-0002-3712-3691)  
The University of Tokyo, Japan; RIKEN, Japan


**Minseok Kim**  
Korea University, Seoul, South Korea

**Jun Hyung Lee**  [orcid.org/0009-0008-2740-2006](https://orcid.org/0009-0008-2740-2006)  
Korea University, Seoul, South Korea

**Yuanliang Dong**  [orcid.org/0000-0001-5208-9447](https://orcid.org/0000-0001-5208-9447)  
Central Conservatory of Music, Beijing, China

**Xinran Zhang**  [orcid.org/0000-0001-9055-061X](https://orcid.org/0000-0001-9055-061X)  
Central Conservatory of Music, Beijing, China

**Jiafeng Liu**  [orcid.org/0009-0005-5530-8744](https://orcid.org/0009-0005-5530-8744)  
Central Conservatory of Music, Beijing, China

**Yuki Mitsufuji**  [orcid.org/0000-0002-6806-6140](https://orcid.org/0000-0002-6806-6140)  
Sony Group Corporation, Tokyo, Japan

## 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 Wang, H., Yang, Y., and Lee, J. H., editors, *Proceedings of the 15th International Society for Music Information Retrieval Conference (ISMIR)*, pages 155–160.
- Braun, S. and Tashev, I. J.** (2020). A consolidated view of loss functions for supervised deep learning-based speech enhancement. *2021 44th International Conference on Telecommunications and Signal Processing (TSP)*, pages 72–76. DOI: <https://doi.org/10.1109/TSP52935.2021.9522648>
- Cheng, H., Zhu, Z., Li, X., Gong, Y., Sun, X., and Liu, Y.** (2020). Learning with instance-dependent label noise: A sample sieve approach. *arXiv preprint arXiv:2010.02347*.
- 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 of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, pages 192–198.
- Défossez, A.** (2021). Hybrid spectrogram and waveform source separation. In *Proceedings of the ISMIR 2021 Workshop on Music Source Separation*.
- García, H. F., Aguilar, A., Manilow, E., and Pardo, B.** (2021). Leveraging hierarchical structures for fewshot musical instrument recognition. In Lee, J. H., Lerch, A., Duan, Z., Nam, J., Rao, P., van Kranenburg, P., and Srinivasamurthy,

- A., editors, *Proceedings of the 22nd International Society for Music Information Retrieval Conference (ISMIR)*, pages 220–228.
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., and Sugiyama, M.** (2018). Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in Neural Information Processing Systems*, 31.
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.** (2021). The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349. DOI: <https://doi.org/10.1109/ICCV48922.2021.00823>
- Hennequin, R., Khlif, A., Voituret, F., and Moussallam, M.** (2020). Spleeter: A fast and efficient music source separation tool with pre-trained models. *Journal of Open Source Software*, 5(5):2154. DOI: <https://doi.org/10.21105/joss.02154>
- Herbrich, R., Minka, T., and Graepel, T.** (2007). Trueskilltm: A Bayesian skill rating system. In Scholkopf, B., Platt, J. C., and Hoffman, T., editors, *Advances in Neural Information Processing Systems 19 (NIPS-06)*, pages 569–576. MIT Press. DOI: <https://doi.org/10.7551/mitpress/7503.003.0076>
- Hinton, G. E., Vinyals, O., and Dean, J.** (2015). Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531.
- Kang, D. and Hashimoto, T. B.** (2020). Improved natural language generation via loss truncation. *ArXiv*, abs/2004.14589. DOI: <https://doi.org/10.18653/v1/2020.acl-main.66>
- 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 of the ISMIR 2021 Workshop on Music Source Separation*.
- Kim, M. and Lee, J. H.** (2023). Sound Demixing Challenge 2023 – Music Demixing Track technical report. *arXiv preprint arXiv:2306.09382*.
- Kong, Q., Cao, Y., Iqbal, T., Wang, Y., Wang, W., and Plumbley, M. D.** (2020). PANNs: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2880–2894. DOI: <https://doi.org/10.1109/TASLP.2020.3030497>
- Koo, J., Chae, Y., Jeon, C.-B., and Lee, K.** (2023). Self-refining of pseudo labels for music source separation with noisy labeled data. *arXiv preprint arXiv:2307.12576*.
- Lai, C.-H., Zou, D., and Lerman, G.** (2019). Robust subspace recovery layer for unsupervised anomaly detection. *arXiv preprint arXiv:1904.00152*.
- Lai, C.-H., Zou, D., and Lerman, G.** (2023). Robust variational autoencoding with Wasserstein penalty for novelty detection. In *International Conference on Artificial Intelligence and Statistics*, pages 3538–3567. PMLR.
- Li, J., Socher, R., and Hoi, S. C.** (2020). Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*.
- Liu, H., Xie, L., Wu, J., and Yang, G.** (2020). Channelwise subband input for better voice and accompaniment separation on high resolution music. In *Interspeech*. DOI: <https://doi.org/10.21437/Interspeech.2020-2555>
- Liutkus, A., Stöter, F.-R., Rafii, Z., Kitamura, D., Rivet, B., Ito, N., Ono, N., and Fontcave, J.** (2017). The 2016 Signal Separation Evaluation Campaign. In Tichavsky, P., Babaie-Zadeh, M., Michel, O. J., and Thirion-Moreau, N., editors, *Latent Variable Analysis and Signal Separation*, pages 323–332. Cham. Springer International Publishing. DOI: [https://doi.org/10.1007/978-3-319-53547-0\\_31](https://doi.org/10.1007/978-3-319-53547-0_31)
- Luo, Y. and Yu, J.** (2023). Music source separation with band-split RNN. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:1893–1901. DOI: <https://doi.org/10.1109/TASLP.2023.3271145>
- Mahanta, S. K., Khilji, A. F. U. R., and Pakray, P.** (2021). Deep neural network for musical instrument recognition using MFCCs. *Computacion y Sistemas*, 25(2):351–360. DOI: <https://doi.org/10.13053/cys-25-2-3946>
- Manilow, E., Wichern, G., and Roux, J. L.** (2020). Hierarchical musical instrument separation. In *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*.
- Manilow, E., Wichern, G., Seetharaman, P., and Le Roux, J.** (2019). Cutting music source separation some Slakh: A dataset to study the impact of training data quality and quantity. In *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 45–49. IEEE. DOI: <https://doi.org/10.1109/WASPAA.2019.8937170>
- Mitsufuji, Y., Fabbro, G., Uhlich, S., Stoter, F.-R., Defossez, A., Kim, M., Choi, W., Yu, C.-Y., and Cheuk, K.-W.** (2022). Music Demixing Challenge 2021. *Frontiers in Signal Processing*, 1:18. DOI: <https://doi.org/10.3389/frsip.2021.808395>
- Mukherjee, D., Guha, A., Solomon, J. M., Sun, Y., and Yurochkin, M.** (2021). Outlier-robust optimal transport. In *International Conference on Machine Learning*, pages 7850–7860. PMLR.
- Ono, N., Rafii, Z., Kitamura, D., Ito, N., and Liutkus, A.** (2015). The 2015 Signal Separation Evaluation Campaign. In Vincent, E., Yeredor, A., Koldovský, Z., and Tichavský, P., editors, *Latent Variable Analysis and Signal Separation*, pages 387–395. Cham. Springer International Publishing. DOI: [https://doi.org/10.1007/978-3-319-22482-4\\_45](https://doi.org/10.1007/978-3-319-22482-4_45)
- Pereira, I., Araújo, F., Korzeniowski, F., and Vogl, R.** (2023). MoisesDB: A dataset for source separation beyond 4-stems. In *Proceedings of the 24th International Society for Music Information Retrieval Conference (ISMIR)*, pages 619–626.
- Petermann, D., Wichern, G., Wang, Z.-Q., and Roux, J. L.** (2022). The cocktail fork problem: Threestem audio separation for real-world soundtracks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. DOI: <https://doi.org/10.1109/ICASSP43922.2022.9746005>

- Prätzlich, T., Müller, M., Bohl, B. W., and Veit, J.** (2015). Freischutz Digital: Demos of audio-related contributions. In *Demos and Late Breaking News of the International Society for Music Information Retrieval Conference (ISMIR)*.
- Rafii, Z., Liutkus, A., Stöter, F.-R., Mimitakis, S. I., and Bittner, R.** (2017). The MUSDB18 corpus for music separation. <https://doi.org/10.5281/zenodo.1117372>.
- Rouard, S., Massa, F., and Défossez, A.** (2023). Hybrid transformers for music source separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. DOI: <https://doi.org/10.1109/ICASSP49357.2023.10096956>
- Sawata, R., Takahashi, N., Uhlich, S., Takahashi, S., and Mitsufuji, Y.** (2023). The whole is greater than the sum of its parts: Improving DNNbased music source separation. *arXiv preprint arXiv:2305.07855*.
- Sawata, R., Uhlich, S., Takahashi, S., and Mitsufuji, Y.** (2021). All for one and one for all: Improving music separation by bridging networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 51–55. DOI: <https://doi.org/10.1109/ICASSP39728.2021.9414044>
- Solovyev, R., Stempkovskiy, A., and Habruseva, T.** (2023). Benchmarks and leaderboards for sound demixing tasks. *arXiv preprint arXiv:2305.07489*.
- Stöter, F.-R., Liutkus, A., and Ito, N.** (2018). The 2018 Signal Separation Evaluation Campaign. In Deville, Y., Gannot, S., Mason, R., Plumbley, M. D., and Ward, D., editors, *Latent Variable Analysis and Signal Separation*, pages 293–305, Cham. Springer International Publishing. DOI: [https://doi.org/10.1007/978-3-319-93764-9\\_28](https://doi.org/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. *Journal of Open Source Software*, 4(41):1667. DOI: <https://doi.org/10.21105/joss.01667>
- Tarvainen, A. and Valpola, H.** (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems*, 30.
- Torcoli, M., Kastner, T., and Herre, J.** (2021). Objective measures of perceptual audio quality reviewed: An evaluation of their application domain dependence. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1530–1541. DOI: <https://doi.org/10.1109/TASLP.2021.3069302>
- Uhlich, S., Fabbro, G., Masato, H., Takahashi, S., Wichern, G., Le Roux, J., Chakraborty, D., Mohanty, S., Li, K., Luo, Y., Yu, J., Gu, R., Solovyev, R., Stempkovskiy, A., Habruseva, T., Sukhovei, M., and Mitsufuji, Y.** (2024). The Sound Demixing Challenge 2023 – Cinematic Demixing Track. *Transactions of the International Society for Music Information Retrieval*, 7(1):44–62. DOI: <https://doi.org/10.5334/tismir.172>
- Uhlich, S., Porcu, M., Giron, F., Enekl, M., Kemp, T., Takahashi, N., and Mitsufuji, Y.** (2017). Improving music source separation based on deep neural networks through data augmentation and network blending. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 261–265. DOI: <https://doi.org/10.1109/ICASSP.2017.7952158>
- Wang, H., Xiao, R., Dong, Y., Feng, L., and Zhao, J.** (2022). Promix: Combating label noise via maximizing clean sample utility. *arXiv preprint arXiv:2207.10276*.
- Wang, Z., Giri, R., Isik, U., Valin, J.-M., and Krishnaswamy, A.** (2021). Semi-supervised singing voice separation with noisy self-training. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 31–35. DOI: <https://doi.org/10.1109/ICASSP39728.2021.9413723>
- Wisdom, S., Hershey, J. R., Wilson, K., Thorpe, J., Chinen, M., Patton, B., and Saurous, R. A.** (2019). Differentiable consistency constraints for improved deep speech enhancement. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 900–904. DOI: <https://doi.org/10.1109/ICASSP.2019.8682783>
- Wisdom, S., Tzinis, E., Erdogan, H., Weiss, R., Wilson, K., and Hershey, J.** (2020). Unsupervised sound separation using mixture invariant training. *Advances in Neural Information Processing Systems*, 33:3846–3857.
- Zhang, Z., Liu, Q., and Wang, Y.** (2017). Road extraction by deep residual U-Net. *IEEE Geoscience and Remote Sensing Letters*, 15:749–753. DOI: <https://doi.org/10.1109/LGRS.2018.2802944>

---

**TO CITE THIS ARTICLE:**

Fabbro, G., Uhlich, S., Lai, C.-H., Choi, W., Martínez-Ramírez, M., Liao, W., Gadelha, I., Ramos, G., Hsu, E., Rodrigues, H., Stöter, F.-R., Défossez, A., Luo, Y., Yu, J., Chakraborty, D., Mohanty, S., Solovyev, R., Stempkovskiy, A., Habruseva, T., Goswami, N., Harada, T., Kim, M., Lee, J. H., Dong, Y., Zhang, X., Liu, J., & Mitsufuji, Y. (2024). The Sound Demixing Challenge 2023 – Music Demixing Track. *Transactions of the International Society for Music Information Retrieval*, 7(1), 63–84. DOI: <https://doi.org/10.5334/tismir.171>

**Submitted:** 22 August 2023    **Accepted:** 13 February 2024    **Published:** 18 April 2024

**COPYRIGHT:**

© 2024 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

*Transactions of the International Society for Music Information Retrieval* is a peer-reviewed open access journal published by Ubiquity Press.