

A STEM-AGNOSTIC SINGLE-DECODER SYSTEM FOR MUSIC SOURCE SEPARATION BEYOND FOUR STEMS

Karn N. Watcharasupat Alexander Lerch

Music Informatics Group, Georgia Institute of Technology, Atlanta, GA, USA
{kwatcharasupat, alexander.lerch}@gatech.edu

ABSTRACT

Despite significant recent progress across multiple sub-tasks of audio source separation, few music source separation systems support separation beyond the four-stem vocals, drums, bass, and other (VDBO) setup. Of the very few current systems that support source separation beyond this setup, most continue to rely on an inflexible decoder setup that can only support a fixed pre-defined set of stems. Increasing stem support in these inflexible systems correspondingly requires increasing computational complexity, rendering extensions of these systems computationally infeasible for long-tail instruments. We propose Banquet, a system that allows source separation of multiple stems using just one decoder. A bandsplit source separation model is extended to work in a query-based setup in tandem with a music instrument recognition PaSST model. On the MoisesDB dataset, Banquet — at only 24.9 M trainable parameters — performed on par with or better than the significantly more complex 6-stem Hybrid Transformer Demucs. The query-based setup allows for the separation of narrow instrument classes such as clean acoustic guitars, and can be successfully applied to the extraction of less common stems such as reeds and organs.

1. INTRODUCTION

Music Source Separation (MSS) is the task of separating a musical audio mixture into its constituent components, commonly referred to as stems. The releases of DSD100 [1] and MUSDB18 [2, 3], both being four-stem MSS datasets, have defined a de-facto standard, with nearly every major work since relying on the four-stem *vocals, bass, drum*, and *others* (VDBO) setup [4–19]. While this has significantly improved the comparability and reproducibility of the task, it has also disproportionately favored the VDBO setup. Very few works have tackled MSS beyond the VDBO setup, each relying on datasets with significant limitations: Wang et al. [20] relied on MedleyDB [21, 22], whose stem ontology is somewhat unfriendly to source separation, Manilow et al. [23] relied on the syn-

thetically generated Slakh dataset [24], and others relied on proprietary data inaccessible to other research groups [11, 18], limiting reproducibility. The recently released MoisesDB [25], a multitrack source separation dataset, attempts to address these limitations, particularly in terms of stem availability and taxonomy. This aims at broadening the task beyond VDBO based on publicly available data. However, to the best of our knowledge, while MoisesDB was used in the 2023 Sound Demixing Challenge (SDX) [26], no published system has utilized MoisesDB for source separation beyond VDBO yet.

In this work, we propose Banquet,¹ a query-based source separation model that can separate an arbitrary number of stems using just one set of stem-agnostic encoder and decoder, and a pre-trained feature extractor [27]. Our model was adapted from the cinematic audio source separation Bandit model [28], which was in turn adapted from the music source separation Bandsplit RNN model [17]. Bandit significantly reduces the complexity of Bandsplit RNN by adopting a common-encoder approach with stem-specific decoders. In this work, we take the complexity reduction further by switching to a query-based setup, using only one decoder shared amongst all possible stems. Performance evaluation on MoisesDB demonstrated separation performance above oracle for drum and bass, state-of-the-art for guitar and piano, and at least 7.4 dB SNR for vocals. Our system additionally provided support for fine-level stem extraction currently available only in a few MSS systems.

2. RELATED WORK

Nearly every major MSS works since 2017 have relied on the VDBO setup. Early systems [4, 6, 29], including OpenUnmix [8], were usually Time-Frequency (TF) masking models with LSTM forming the core of the systems, with some experimenting with densely-connected convolutional systems [5, 12]. Beginning with Wave-U-Net [7], the U-Net architecture became a popular choice for MSS, with notable models such as Demucs [9, 10, 14, 18], Spleeter [11], ByteSep [13], and KUIELab-MDX-Net [15] all being some variations of a U-Net. More recently, Bandsplit RNN [17] became one of the few state-of-the-art systems to not rely on a U-Net setup. This was followed by the Bandsplit RoPE Transformer model [19] topping the



© K. N. Watcharasupat, and A. Lerch. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).
Attribution: K. N. Watcharasupat, and A. Lerch, “A Stem-Agnostic Single-Decoder System for Music Source Separation Beyond Four Stems”, in *Proc. of the 25th Int. Society for Music Information Retrieval Conf.*, San Francisco, United States, 2024.

¹ Banquet is a portmanteau of **Q**uery-based **B**andit. Code available at github.com/kwatcharasupat/query-bandit. Last accessed 24 July 2024.

leaderboard of SDX 2023 [26]. Of existing open-source systems, very few offer separation functionality beyond the VDBO setup. Spleeter [11] supports 5-stem separation with VDBO and piano. HT-Demucs [18] supports a 6-stem setup with VDBO, piano, and guitar.

2.1 Conditional source separation

The systems mentioned above were mostly designed with either stem-specific models, stem-specific decoders, or a shared decoder with predetermined outputs. As a result, these systems are not particularly amenable to the addition of new stems, especially if these new stems have limited data availability. Below we review some of the common approaches for conditional source separation that may be useful for extending existing systems beyond VDBO.

Meseguer-Brocal and Peeters [30] were likely amongst the first to attempt a conditioned U-Net for source separation using a single decoder. They used multiple feature-wise linear modulation (FiLM) [31] layers within the encoder to perform MSS in a VDBO setup. Slizovskaia et al. [32] used a similar setup with FiLMs either throughout the encoder, at the bottleneck layer, or at the final decoder layer. The systems in [32] were tested on the 13-instrument URMP dataset [33], with up to 4 active instruments in any recording, but all performed poorly in terms of mean signal-to-distortion ratio (SDR). Lin et al. [34] proposed a joint separation-transcription U-Net system, which performed well for string and brass instruments in URMP, but struggled on woodwind instruments. The system in [34] used FiLMs throughout the encoder with a query embedding from another convolutional model, and across all skip connections with transcription embeddings.

Lee et al. [35] proposed a U-Net with two methods of less aggressive conditioning with examples beyond VDBO, but only provided objective results for a VDBO setup on MUSDB18. Wang et al. [20] also proposed a U-Net, with FiLM conditioning only at the bottleneck layer. The system in [20] was able to support a substantial number of stems beyond VDBO with the caveat that its reported performance is significantly below contemporary models for VDBO stems. Gfeller et al. [36] utilized a FiLM-conditioned wave-to-wave U-Net to perform one-shot conditional audio filtering. Similar approaches were also adopted in Choi et al. [37] and Jeong et al. [38] for MSS, in Chen et al. [39] for source activity-queried separation, in Kong et al. [40] for universal source separation, and in Liu et al. [41, 42] for language-queried source separation. These works [36–42] applied FiLM or generalizations thereof to nearly every single layer of the network, significantly increasing the computational complexity of the system. We surmise that the apparent need for multiple conditioning in a U-Net is probably due to the nature of its information flow [43], which may require a significant number of information streams to be conditioned to achieve acceptable performance.

In a different direction, source separation systems relying on audio embedding “distances” have also been developed, notably with Le Roux et al. in [23, 44, 45]. In 2018,

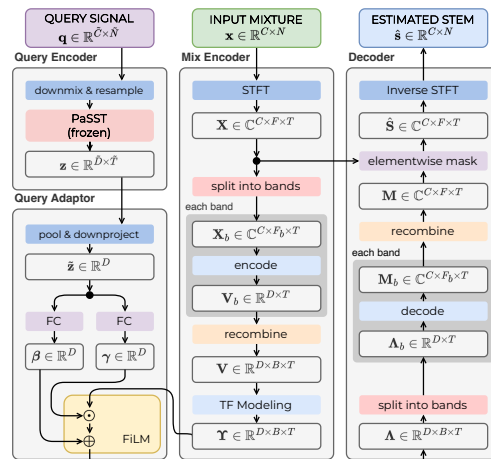


Figure 1. Overview of the Banquet System.

Kumar et al. [46] presented an early work using Euclidean audio embedding distance from a “query” embedding to inform music source separation. A similar system using a Gaussian mixture posterior in lieu of standard distance was proposed in [44]. Hierarchical masking [23] was later utilized to allow the extraction of stems at multiple levels of specificity. More recently, source separation systems with audio embedding in a low-dimensional hyperbolic space have been developed to allow music [45] and speech [47] source separation with some degrees of control on the specificity of the extraction. Uniquely, Samuel et al. [48] proposed a network-generating network approach for instrument-conditioned source separation.

3. PROPOSED SYSTEM

The overview of the proposed Banquet system is shown in Fig. 1. The system is a single-encoder single-decoder adaptation of Bandit [28], that takes in a mixture signal x and a query signal q , and extracts a stem estimate \hat{s} from the mixture signal of the “same” stem type as the query signal using a complex-valued TF mask. This is done by (i) encoding the mixture into a subband-level time-varying embedding tensor Υ , (ii) encoding the query into a single-vector representation \tilde{z} , (iii) adapting the mixture embedding, conditioned on the query, into a stem-specific embedding $\Lambda = Q(\Upsilon; \tilde{z})$, then (iv) decoding the Λ to a TF mask M that is used to obtain the source estimate.

3.1 Bandit encoder

The encoder module of the system used in this work is the *musical* variant of the Bandit encoder, with $B = 64$ bands. Specifically, given an input mixture $x \in \mathbb{R}^{C \times N}$ with C channels and N samples, a short-time Fourier transform (STFT) of x is computed to obtain $X \in \mathbb{C}^{C \times F \times T}$ with a frame size of $N_{FFT} = 2(F - 1) = 2048$ and 75% overlap. The STFT is then split into overlapping subbands as detailed in [28]. Each of the subbands is then viewed as a real-valued tensor in $\mathbb{R}^{2CF \times T}$, passed through a layer norm and an affine transformation with $D = 128$ neurons to obtain $V_b \in \mathbb{R}^{D \times T}$. These tensors are then stacked to

obtain $\mathbf{V} \in \mathbb{R}^{D \times B \times T}$. TF modeling is then applied on \mathbf{V} to obtain \mathbf{Y} using 8 pairs of residual gated recurrent units (GRUs), the first of each pair operating along the time axis and the second along the band axis.

Note that this TF modeling is the only part of the model that is recurrent across either the time or the subband axes. The rest of the encoder and the decoder operate in a subband-wise manner identically for any time frame.

3.2 Query encoding

To obtain the query embedding, a PaSST model [27] trained on the OpenMIC-2018 dataset [49] is used. The 20 instruments in OpenMIC span all coarse-level classes of MoisesDB, except *other*. For compatibility, each query signal is downmixed to mono and downsampled to 32 kHz before being fed to PaSST. Although the query feature extractor could, in theory, be jointly trained with the rest of the system, preliminary experiments showed that this can result in considerable instability during training, especially if the query feature extractor is not at least pretrained. Due to the size and complexity of PaSST, the query feature extractor is fully frozen in this work. The embedding from the PaSST variant used is a time series with a feature dimension of $\tilde{D} = 784$. The embedding is averaged over time and linearly down-projected to obtain $\tilde{\mathbf{z}} \in \mathbb{R}^D$.

3.3 Query-based adaptation

In the original Bandit system [28], each stem was estimated through a dedicated decoder. As a result, \mathbf{Y} typically contains information from all stems, with most of the “separation” occurring within each of the decoders. This is evident in the fact that the encoder of a Bandit system trained on the cinematic audio Divide and Remaster (DnR) dataset [?] could be successfully used in a 4-stem MSS on the MUSDB18-HQ dataset [2] with separation quality on par with Open-Unmix [28].

In this work, only a single decoder is responsible for mask estimation for any stem. As a result, the query-based adaptation $\mathcal{Q}: (\mathbb{R}^{D \times B \times T}, \mathbb{R}^D) \mapsto \mathbb{R}^{D \times B \times T}$ has an important role in filtering out irrelevant information from \mathbf{Y} , or at least “hinting” to the decoder the nature of the target stem. A single FiLM layer is used to map from the mixture embedding to the stem-specific embedding, that is,

$$\mathbf{A}[d, b, t] = \gamma[d] \cdot \mathbf{Y}[d, b, t] + \beta[d], \quad \forall d, b, t, \quad (1)$$

where modulating variables $\gamma, \beta \in \mathbb{R}^D$ are obtained from a two-layer nonlinear affine map of $\tilde{\mathbf{z}}$. This is similar to the conditioning method used in [20].

Crucially, note that the modulating variables are not subband-specific. Due to the nature of the TF modeling module within the encoder, features of \mathbf{Y} are already aligned across subbands and time frames. Moreover, BSRNN-like models only contain one stream of information flow, with a clear bottleneck, thus lending itself to the global conditioning mechanism significantly more than, for example, U-Net-style models in [20, 41, 42].

The use of embedding-based query, as opposed to one-hot class-based query, provides significant practical flexibility in adding new instruments as data become available or in adjusting the level of specificity in the querying, as these can be done via finetuning with no architectural changes to the model. Moreover, class-based query can be emulated in an embedding-based system but not vice versa.

3.4 Bandit decoder

The decoder used is identical in structure to that in [28]. The major difference is that there is only one stem-agnostic decoder. Given a conditioned embedding tensor \mathbf{A} , the embedding tensor is split into subband-level representation $\mathbf{A}_b = \mathbf{A}[:, b, :]$. Each \mathbf{A}_b is passed through a layer norm and a gated linear unit (GLU) to obtain a real-valued tensor $\mathbb{R}^{2CF_b \times T}$ which is then viewed as a complex-valued tensor $\mathbf{M}_b \in \mathbb{C}^{C \times F_b \times T}$. Frequency-domain overlap-add is then applied to obtain the full-band mask using

$$\mathbf{M}[c, f, t] = \sum_{b=0}^{B-1} \frac{\mathbf{W}[b, f] \cdot \mathbf{M}_b[c, f - \min \mathfrak{F}_b, t]}{\sum_{k=0}^{B-1} \mathbf{W}[k, f]} \quad (2)$$

Finally, the source estimates are then obtained using elementwise masking $\hat{\mathbf{S}} = \mathbf{X} \circ \mathbf{M}$.

3.5 Loss function

The loss function used in this work is the multichannel version of the L1SNR loss proposed in [28]. The contribution for each sample of the loss function is given by

$$\mathcal{L}(\hat{\mathbf{s}}; \mathbf{s}) = \mathcal{D}(\hat{\mathbf{s}}; \mathbf{s}) + \mathcal{D}(\Re \hat{\mathbf{S}}; \Re \mathbf{S}) + \mathcal{D}(\Im \hat{\mathbf{S}}; \Im \mathbf{S}), \quad (3)$$

$$\mathcal{D}(\hat{\mathbf{y}}; \mathbf{y}) = 10 \log_{10} \frac{\|\text{vec}(\hat{\mathbf{y}} - \mathbf{y})\|_1 + \epsilon}{\|\text{vec}(\mathbf{y})\|_1 + \epsilon}, \quad (4)$$

where $\hat{\mathbf{s}} = \text{iSTFT}(\hat{\mathbf{S}})$, \mathbf{s} and \mathbf{S} are defined similarly for the ground truth, $\text{vec}(\cdot)$ is the vectorization operator, and $\epsilon = 10^{-3}$ for stability.

4. DATA AND EXPERIMENTAL SETUP

This work utilizes the MoisesDB dataset [25], which consists of 240 songs from 47 artists, in stereo format at 44.1 kHz. MoisesDB defined their stem ontology with more than 30 fine-level classes, which are then grouped into 11 coarse-level classes [25, Table 2]. Due to the lack of official splits for MoisesDB, we performed a five-fold split² on the dataset stratified by genres. The first three splits are used as the training set, the fourth as the validation set, and the last as the test set.

4.1 Query extraction

For each possible stem of each song, a 10-second chunk of the clean audio of the same stem is extracted as the query signal. This is done by computing a time series of onset strength for each stem and then aggregating the mean onset

² The splits are available in the repository. Note that not all stems contain a sufficient number of data points to be split into a five-fold validation setup. As a result, some stems are only present in a subset of folds.

strength for each 10-second sliding window with a hop size of 512 samples. The 10-second window with the strongest average onset is taken as the query signal. A t-SNE plot of the query embedding is shown in Fig. 2. While clusters can be clearly seen amongst related stems, it can also be seen that there are varying degrees of non-separability of the embedding between fine-level stems.

4.2 Training

Each model was trained using an NVIDIA H100 GPU (80 GB) for up to 150 epochs, unless otherwise stated. A training epoch consists of 8192 mixture-query pairs, with a batch size of 4. We used Adam optimizer with an initial learning rate of 10^{-3} and a decay factor of 0.98 per epoch.

In the default sampling strategy, a random song is chosen, a random trainable stem for that song is chosen as the target stem, then a random chunk of 6 s is chosen. If the current target chunk has an RMS below -36 dBFS, a new random chunk is chosen for up to 10 more trials. Otherwise, the threshold is dropped to -48 dBFS for another 10 trials. If a suitable chunk is still not found, the next random chunk is chosen regardless of RMS. A pre-extracted query of the same stem is then randomly chosen from the available pool of songs, including the song of the mixture.

4.3 Testing and inference

During testing and inference, each track is split into 6-s segments with a hop size of 0.5 s, as per [17]. The estimated stems were then reconstructed into a full track using time-domain overlap-add with a Hann window. The Banquet models are tested in two scenarios: one using a query from a different song, and another using a query from the same song (SSQ). In different-song querying, the query song for each stem is randomly chosen from another song within the test split that contains the stem. When possible, the query song is chosen so that it is from the same genre as the mixture song but from a different artist. Otherwise, a song from any genre with a different artist is chosen.

4.4 Evaluation metric

In this work, we report the full-track multichannel signal-to-noise ratio (SNR)³ as the main metric. Specifically, for a test signal \hat{s} and a reference signal s , both in $\mathbb{R}^{C \times N}$, the SNR is computed by

$$\text{SNR}(\hat{y}; y) = 10 \log_{10} \left(\frac{\|s\|_F^2}{\|\hat{s} - s\|_F^2} \right). \quad (5)$$

5. RESULTS AND DISCUSSION

In this section, we provide the results and discussion of our experiments. Section 5.1 discusses pretraining of the Bandit/Banquet encoder. Section 5.2 trials the use of the query-based setup on a subset of vocals, drums, and bass

³ Signal-to-interference ratio (SIR) and signal-to-artifact ratio (SAR) were not computed as the number of the constituent stems can be large, making the required subspace projection intractable and/or unreliable. It is also unclear if coarse-level ground truth or fine-level ground truth should be used for such a projection. See [50–52] for background.




Figure 2. A t-SNE plot of the PaSST embeddings of the query signals. Stems from the same coarse-level grouping, as defined by [25], share the same symbol.

stems. Section 5.3 extends the system to include fine-level stems from *guitar* and *piano* families. Finally, Section 5.4 attempts to perform extraction on all possible fine-level stems with sufficient data.

5.1 Encoder pretraining

Preliminary experiments indicated that encoder pretraining is an important step to stabilize the training of the query-based model, especially as the number of query stems grows. The encoder pretraining is done with a common-encoder multi-decoder setup similar to [28] with a VDBO setup for 100 epochs. The VDBO decoders were discarded and the encoder was used for subsequent experiments. The performance of the pretrained model is shown in Table 1, with performance above oracle ideal ratio mask (IRM) for drums and bass, and on par with HT-Demucs for vocals.⁴

5.2 Learning to separate from queries

As a first step to verify the query-based ability of the model, a Banquet model is trained to extract only *lead female singer*, *lead male singer*, *drums*, and *bass* stems, referred to as the Q:VDB setup. We experimented with training from scratch, using a frozen pretrained encoder (FE), and using a trainable pretrained encoder (TE). While the frozen-encoder setup did not demonstrate any sign of overfitting during the training, the trainable-encoder system demonstrated (very slight) overfitting. As a result, an additional setup with data augmentation (DA) was attempted with the trainable encoder setup, using simple stem-wise within-song random gain (up to ± 6 dB), random time shifting, polarity inversion, and channel swapping.

The results are shown in Table 2. All three variants with pretrained encoder provided better performance than the model trained from scratch, except for drums in the trainable-encoder model without DA being 0.1 dB lower. Thus, for all subsequent experiments, the encoder is always pretrained. Without DA, there was no clear benefit to unfreezing the encoder. However, in a trainable en-

⁴ All coarse-level results for oracle methods, HT-Demucs, and Spleeter were recomputed only on the test set using song-wise results from github.com/moises-ai/moises-db. The song-wise results were missing for five of the songs (as of 6 April 2024), two of these belong in the test set, thus the aggregates were computed over 46 songs instead of 48 songs.

Table 1. Median SNR of the models trained on the VDBO setup, evaluated on the test set of MoisesDB.

Model	Median SNR (dB):	Vocals	Drums	Bass	Other
Bandit [28]		9.1	9.9	10.6	6.4
HT-Demucs [18]		9.1	11.0	12.2	7.3
Spleeter [11]		7.4	6.6	6.8	5.0
Oracle IRM		10.3	9.2	8.8	7.6

Table 2. Median SNR of Banquet models on the Q:VDB setup, evaluated with different-song queries.⁶

Pretrained Enc.	FE	DA	Female Vox	Male Vox	Drums	Bass
N	N	N	8.3	7.2	9.4	9.4
Y	Y	N	9.8	7.6	9.9	10.2
Y	N	N	9.8	8.0	9.3	9.8
Y	N	Y	10.2	8.0	10.1	10.8

coder system with DA, slight to moderate improvements were observed across all but the male vocal stem. Note, however, that allowing full-model training significantly increases the number of trainable parameters from 13.5 M to 24.9 M thus the computational cost and training time also increases accordingly. The performances of the drums and bass stems are on par or better than the dedicated-stem setup in Table 1. Generally, the models perform better on female vocals than on male vocals.

5.3 Extending to guitar and piano

Amongst systems that tackled MSS beyond four stems, the next two stems beyond VDBO are usually guitar and piano, due to their high prevalence within pop/rock music. The set of possible queries is thus extended from Q:VDB to also include *acoustic guitar*, *clean electric guitar*, *distorted electric guitar*, *grand piano*, and *electric piano*. This is referred to as the Q:VDBGP setup. Due to the significantly lower number of available training data for guitar and piano stems, we also experimented with a balanced sampling (BS) strategy. In this strategy, a random stem is first chosen as the target stem, then a random song containing that stem is chosen. The remainder of the sampling process is the same as the default. This strategy ensures that every stem has a similar number of training pairs, but distorts the “natural” distribution of stem occurrences.

For comparability with existing systems, the inference outputs of fine-level stems in this setup were added together to form their respective coarse-level predictions.⁷ Coarse-level results are shown in Table 3. Fine-level results for trainable-encoder models are shown in Table 4.

At the coarse level, most variants of Banquet continue to perform above the oracle IRM for drums and bass. With the default-sampling trainable encoder systems, the Banquet performed better than HT-Demucs on guitar and piano. Without DA, balanced sampling generally did not lead to consistent improvements for guitar and piano. With balanced sampling and DA on a trainable-encoder model,

⁶ Median results for the same-song query and different-song query are within 0.2 dB of each other.

⁷ The ground truth signals for are the full coarse-level tracks, e.g. *vocals* ground truth include contributions from *background vocals* even if we do not have *background vocals* in the predictions.

Table 3. Coarse-level performance of the Banquet models with different-song queries on the Q:VDBGP setup

Model	FE	DA	BS	Vox	Lead Vox	Drums	Bass	Guitar	Piano
Banquet	Y	N	N	8.0	7.9	9.8	10.5	2.3	0.8
			Y	7.9	7.7	9.6	10.5	2.2	0.9
	N	N	N	7.4	8.0	9.6	10.6	3.0	2.3
			Y	7.6	7.7	9.3	10.2	2.9	2.5
	Y	N		7.8	7.9	10.1	10.9	3.2	2.2
		Y		7.6	7.9	9.5	11.0	3.3	2.5
HT-Demucs (VDBGPO)				8.9	—	11.6	12.4	2.4	1.7
Spleeter (VDBPO)				7.0	—	6.9	6.7	—	0.7
Oracle IRM				10.0	—	9.6	7.8	5.2	5.0

Bold: best Banquet model **and/or** best non-oracle model.

however, slight gains in median SNRs of guitar and piano were observed, albeit at the cost of vocals and drum SNRs.

At the fine level, the model performance follows a similar trend to that of the coarse level. Drums and bass continue to perform above the oracle IRM, while both lead vocals performed close to the IRM. Guitar and piano performances are still well below IRM. Interestingly, it appears that querying with excerpts from the same or different track did not affect the model performance for most stems except for electric piano. This is likely due to both the small sample size of electric piano limiting generalizability, and the highly diverse set of possible timbres thus the intertwined nature of both the query embedding and the target audio with other keyboard instruments. The ability of the model to query with stems from different tracks is a double-edged sword, however, since this also means that the model is somewhat insensitive to fine differences in timbre between different renditions of the “same” instruments. This could potentially limit its usefulness when applied to a scenario where multiple target stems have very similar timbres.

5.4 Extending beyond guitar and piano

The results for the Q:VDBGP setup demonstrated that the model is able to learn to extract 5 additional stems. In this experiment, we extend the set of possible queries to include all remaining stems with at least one data point per fold: *effects*, *pitched percussion*, *organs & electronic organs*, *synth pad*, *synth lead*, *string section*, *brass*, and *reeds*. Additionally, *bass* is now broken up into *bass guitar* and *bass synth*. This is referred to as the Q:ALL setup. Although these are all fine-level stems as defined by MoisesDB, some of these classes are more specific than others. For example, *brass* is a fine-level stem despite possibly including trumpets, trombones, horns, and tuba. The experimental setups are similar to that of Setup B.⁸

The same-song query results⁹ for the models trained in

⁸ BS and DA models for Q:ALL were significantly more unstable during training than for the Q:VDBGP setup, despite being identical architecturally. When this happens, we discard the collapsed model and restart the training from scratch until we have a model that completes the entire training run with nonsilent output for most stems. No TE+DA+BS system was stable enough to finish the training run without collapse.

⁹ Note that when the FE and the TE+DA systems have SNR concentrated at 0 dB for the long-tail stems, these are indicators of the model outputting very soft, practically silent output. In general, a model yielding negative SNR for a particular stem might be more desirable than a model that has collapsed for a particular stem.

Table 4. Model performance on the Q:VDBGP setup fine-level stems.

FE	DA	BS	SSQ	Female Vox			Male Vox			Drums			Bass			Acoust. Gtr.			Clean E. Gtr.			Dist. E. Gtr.			Grand Piano			E. Piano		
				Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3	Q1	Q2	Q3
N	N	N	N	5.5	9.6	13.2	6.7	7.9	10.0	8.0	9.6	11.6	7.9	9.9	12.0	0.9	1.8	3.6	0.2	0.7	2.4	0.9	2.4	5.3	0.7	2.3	2.9	0.0	0.6	0.7
			Y	5.6	9.6	13.2	6.7	7.9	10.0	8.0	9.6	11.6	7.9	9.9	12.0	0.9	1.8	3.7	0.2	0.9	2.6	0.9	2.4	5.3	0.7	2.2	3.0	0.0	0.8	1.5
Y	N	N	N	6.1	9.6	13.1	6.8	7.7	9.7	7.8	9.3	11.3	7.6	10.0	11.5	0.8	1.8	3.6	0.2	0.8	2.5	1.0	2.5	5.4	0.8	2.5	3.1	-0.1	0.7	0.8
			Y	6.1	9.6	13.1	6.8	7.7	9.7	7.8	9.3	11.3	7.6	10.0	11.5	0.8	1.8	3.7	0.0	0.9	2.7	1.2	2.5	5.4	0.8	2.5	3.1	-0.6	0.8	1.8
Y	N	N	N	5.5	10.1	13.0	6.9	7.9	10.2	8.5	10.1	12.3	8.4	10.7	13.2	1.2	1.7	4.5	0.2	0.9	3.0	0.9	2.8	4.7	0.8	2.8	3.2	0.1	0.5	0.9
			Y	5.5	10.1	13.1	6.9	7.9	10.2	8.5	10.1	12.3	8.4	10.7	13.2	1.2	1.7	4.6	0.2	1.1	2.7	0.9	2.8	4.7	0.8	2.4	3.1	-0.1	0.6	0.9
Y	N	N	N	5.5	10.1	13.5	6.5	7.8	10.0	8.3	9.5	11.8	8.4	10.3	12.1	1.1	1.7	3.9	0.0	0.4	2.7	0.9	3.0	4.9	0.8	2.6	3.2	0.2	0.5	0.9
			Y	5.5	10.1	13.5	6.5	7.8	10.0	8.3	9.5	11.8	8.4	10.3	12.1	1.0	1.7	3.9	0.3	0.6	2.7	0.6	3.0	4.8	0.8	2.5	3.2	0.6	0.9	2.1

FE: frozen encoder, DA: data augmentation, BS: balanced sampling, SSQ: same-song query, Q1: lower quartile, Q2: median, Q3: upper quartile




Figure 3. Performance of the Banquet models with same-song queries on Q:ALL fine-level stems

Q:ALL are shown in Fig.3. The performances of the model trained on Q:ALL on the stems from the Q:VDBGP setup are similar to those in Table 4, with the exception of the significant drop in performance for the balanced-sampled trainable-encoder model. Amongst the newly added stems, there are significant variations in performance, but they are all still very weak in terms of SNR, with no sample performing above 5 dB SNR. For organs, background vocals, and both synth stems, the trainable-encoder model yielded the better upper quartile and maximum performance, but is also very unreliable. Unfortunately, balanced sampling on a trainable encoder model only worsened the performance. DA on a trainable-encoder model with default sampling slightly improved the lower quartile performance, but is also accompanied by lower maximum and upper quartile performance. Frozen-encoder system collapsed for most long-tail stems in default sampling, but balanced sampling interestingly was more stable and performed the best for bass synth, pitched percussion, reeds, and brass. Evidently, the classical tradeoffs are at play here; allowing the model more flexibility with a trainable encoder also comes with a higher risk of model collapse or unreliable

performance. More surprisingly, the fact that even a frozen encoder trained on a VDBO setup was able to function at all beyond Q:VDB indicates that the embedding space of a Bandit encoder already contains information that is partially generalizable beyond VDBO, as also observed in [28].

The results of the long-tail stems are somewhat unsurprising given that the genre distribution in MoisesDB skewed heavily toward pop, rock, and singer-songwriter. In addition to the low track counts, these long-tail instruments also tend to have infrequent active segments and relatively softer levels within a song. In fact, of the long-tail stems, reeds and pitched percussion are the only ones with median RMS above -35 dBFS. Analysis of the SNR distribution shows that the model performance is quite correlated to the track-level RMS of the target signal (Spearman’s ρ between 0.78 and 0.81). This is likely due to a combination of low data availability and the inherent difficulty associated with cleanly extracting these “supporting” stems when there are significant spectral overlaps from more prominent co-occurring stems. In light of the recently published analysis in [53], we may have been too conservative with our DA setup. In particular, we made a conscious choice to only perform gain augmentation close to the original levels, instead of significantly amplifying softer stems. Whether the latter may improve the result at all will have to be addressed in future work. Moreover, given that [34] saw partial success with the predominantly classical instrumentation of URMP, there may also be an opportunity for a much more aggressive cross-dataset DA.

6. CONCLUSION

In this work, Banquet, a stem-agnostic single-decoder query-based source separation system was proposed to address MSS beyond the VDBO stems. At 24.9 M trainable parameters, this highly modularized model with a single stream of information flow provided strong performance for vocals, drums, and bass; outperformed significantly more complex HT-Demucs on guitar and piano; and provided a proof-of-concept for extractions of additional long-tail and/or fine-grained stems at no additional complexity. While there remains room for improvements for long-tail stems with low data availability, this work demonstrated the opportunity for further research on single-decoder systems toward supporting a large and diverse set of stems.

7. ACKNOWLEDGMENTS

This work was supported by a Cyber-Infrastructure Resource Award from the Institute for Data Engineering and Science (IDEaS), Georgia Institute of Technology.

K. N. Watcharasupat was separately supported by the American Association of University Women (AAUW) International Fellowship, and the IEEE Signal Processing Society Scholarship Program.

The authors would like to thank Yiwei Ding and Chih-Wei Wu for their assistance with the project.

8. ETHICS STATEMENT

Machine learning-based systems are inherently data-dependent. Our models rely both directly and indirectly on datasets with inherent imbalance, not only in terms of instrument and genre distribution, but also in terms of cultural origins. As a result, our system inevitably inherit some bias and may not perform on music that have not been well represented in the training data. The authors acknowledge this important limitation and are committed to continue exploring approaches to correct these biases, both in terms of data acquisition and algorithmic development.

9. REFERENCES

- [1] A. Liutkus, F.-R. Stöter, Z. Rafii, D. Kitamura, B. Rivet, N. Ito, N. Ono, and J. Fontecave, “The 2016 Signal Separation Evaluation Campaign,” in *Proceedings of the 13th International Conference on Latent Variable Analysis and Signal Separation*. Grenoble, France: Springer International Publishing, 2017, pp. 323–332.
- [2] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimitakis, and R. Bittner, “The MUSDB18 corpus for music separation,” Dec. 2017.
- [3] —, “MUSDB18-HQ - an uncompressed version of MUSDB18,” Aug. 2019.
- [4] S. Uhlich, M. Porcu, F. Giron, M. Enenkl, T. Kemp, N. Takahashi, and Y. Mitsufuji, “Improving music source separation based on deep neural networks through data augmentation and network blending,” in *Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing*. New Orleans, LA: IEEE, Mar. 2017, pp. 261–265.
- [5] N. Takahashi and Y. Mitsufuji, “Multi-Scale multi-band denisenets for audio source separation,” in *Proceedings of the 2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. New Paltz, NY: IEEE, Oct. 2017, pp. 21–25.
- [6] F. R. Stöter, A. Liutkus, and N. Ito, “The 2018 Signal Separation Evaluation Campaign,” in *Proceedings of the 14th International Conference on Latent Variable Analysis and Signal Separation*. Guildford, United Kingdom: Springer International Publishing, 2018, pp. 293–305.
- [7] D. Stoller, S. Ewert, and S. Dixon, “Wave-U-Net: A Multi-Scale Neural Network for End-to-End Audio Source Separation,” in *Proceedings of the 19th International Society for Music Information Retrieval Conference*. Paris, France: ISMIR, 2018, pp. 334–340.
- [8] F.-R. Stöter, S. Uhlich, A. Liutkus, and Y. Mitsufuji, “Open-Unmix - A Reference Implementation for Music Source Separation,” *Journal of Open Source Software*, vol. 4, no. 41, p. 1667, Sep. 2019.
- [9] A. Défossez, N. Usunier, L. Bottou, and F. Bach, “Demucs: Deep Extractor for Music Sources with extra unlabeled data remixed,” 2019.
- [10] —, “Music Source Separation in the Waveform Domain,” Nov. 2019.
- [11] R. Hennequin, A. Khelif, F. Voituret, and M. Moussallam, “Spleeter: A fast and efficient music source separation tool with pre-trained models,” *Journal of Open Source Software*, vol. 5, no. 50, p. 2154, Jun. 2020.
- [12] N. Takahashi and Y. Mitsufuji, “D3Net: Densely connected multidilated DenseNet for music source separation,” Mar. 2021.
- [13] Q. Kong, Y. Cao, H. Liu, K. Choi, and Y. Wang, “Decoupling Magnitude and Phase Estimation with Deep ResUNet for Music Source Separation,” in *Proceedings of the 22nd International Society for Music Information Retrieval Conference*, Online, 2021, pp. 342–349.
- [14] A. Défossez, “Hybrid Spectrogram and Waveform Source Separation,” in *Proceedings of the 2021 Music Demixing Workshop at the 22nd International Society for Music Information Retrieval Conference*. Online: ISMIR, Nov. 2021.
- [15] M. Kim, W. Choi, J. Chung, D. Lee, and S. Jung, “KUIELab-MDX-Net: A Two-Stream Neural Network for Music Demixing,” in *Proceedings of the 2021 Music Demixing Workshop at the 22nd International Society for Music Information Retrieval Conference*. Online: ISMIR, Nov. 2021.
- [16] Y. Mitsufuji, G. Fabbro, S. Uhlich, F.-R. Stöter, A. Défossez, M. Kim, W. Choi, C.-Y. Yu, and K.-W. Cheuk, “Music Demixing Challenge 2021,” *Frontiers in Signal Processing*, vol. 1, p. 808395, Jan. 2022.
- [17] Y. Luo and J. Yu, “Music Source Separation With Band-Split RNN,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 1893–1901, 2023.
- [18] S. Rouard, F. Massa, and A. Défossez, “Hybrid Transformers for Music Source Separation,” in *Proceedings of the 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*. Rhodes Island, Greece: IEEE, May 2023.
- [19] W.-T. Lu, J.-C. Wang, Q. Kong, and Y.-N. Hung, “Music Source Separation with Band-Split RoPE Transformer,” in *Proceedings of the 2024 IEEE International Conference on Acoustics, Speech and Signal Processing*. Seoul, Korea, Republic of: IEEE, Sep. 2023, pp. 481–485.
- [20] Y. Wang, D. Stoller, R. M. Bittner, and J. Pablo Bello, “Few-Shot Musical Source Separation,” in *Proceedings of the 2022 IEEE International Conference on Acoustics, Speech and Signal Processing*. Singapore, Singapore: IEEE, May 2022, pp. 121–125.
- [21] R. Bittner, J. Salamon, M. Tierney, M. Mauch, C. Cannam, and J. Bello, “MedleyDB: A Multitrack Dataset for Annotation-Intensive MIR Research,” in *Proceedings of the 15th Conference of the International Society for Music Information Retrieval*. Taipei, Taiwan: ISMIR, 2014, pp. 155–160.
- [22] R. M. Bittner, J. Wilkins, H. Yip, and J. P. Bello, “MedleyDB 2.0 : New Data and a System for Sustainable Data Collection,” in *Extended Abstracts for the Late-Breaking Demo Session of the 17th International Society for Music Information Retrieval Conference*. New York City, USA: ISMIR, 2016.
- [23] E. Manilow, G. Wichern, and J. Le Roux, “Hierarchical Musical Instrument Separation,” in *Proceedings of the 21st International Society for Music Information Retrieval Conference*. Montréal, Canada: ISMIR, 2020, pp. 376–383.

- [24] E. Manilow, G. Wichern, P. Seetharaman, and J. Le Roux, "Cutting Music Source Separation Some Slakh: A Dataset to Study the Impact of Training Data Quality and Quantity," in *Proceedings of the 2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. New Paltz, NY, USA: IEEE, Oct. 2019, pp. 45–49.
- [25] I. Pereira, F. Araújo, F. Korzeniowski, and R. Vogl, "MoisesDB: A Dataset for Source Separation Beyond 4-Stems," in *Proceedings of the 24th International Society for Music Information Retrieval Conference*, Milan, Italy, 2023, pp. 619–626.
- [26] G. Fabbro, S. Uhlich, C.-H. Lai, W. Choi, M. Martínez-Ramírez, W. Liao, I. Gadelha, G. Ramos, E. Hsu, H. Rodrigues, F.-R. Stöter, A. Défossez, Y. Luo, J. Yu, D. Chakraborty, S. Mohanty, R. Solovyev, A. Stempkovskiy, T. Habruseva, N. Goswami, T. Harada, M. Kim, J. Hyung Lee, Y. Dong, X. Zhang, J. Liu, and Y. Mitsufuji, "The Sound Demixing Challenge 2023 – Music Demixing Track," *Transactions of the International Society for Music Information Retrieval*, vol. 7, no. 1, pp. 63–84, Apr. 2024.
- [27] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, "Efficient Training of Audio Transformers with Patchout," in *Proceedings of the 23rd Annual Conference of the International Speech Communication Association*. Incheon, Korea: ISCA, Sep. 2022, pp. 2753–2757.
- [28] K. N. Watcharasupat, C.-W. Wu, Y. Ding, I. Orife, A. J. Hipple, P. A. Williams, S. Kramer, A. Lerch, and W. Wolcott, "A Generalized Bandsplit Neural Network for Cinematic Audio Source separation," *IEEE Open Journal of Signal Processing*, vol. 5, pp. 73–81, 2023.
- [29] N. Takahashi, N. Goswami, and Y. Mitsufuji, "MMDenseLSTM: An Efficient Combination of Convolutional and Recurrent Neural Networks for Audio Source Separation," in *Proceedings of the 16th International Workshop on Acoustic Signal Enhancement*. Tokyo, Japan: IEEE, Sep. 2018, pp. 106–110.
- [30] G. Meseguer-Brocal and G. Peeters, "Conditioned-U-Net: Introducing a Control Mechanism in the U-Net for Multiple Source Separations," in *Proceedings of the 20th International Society for Music Information Retrieval Conference*. Delft, Netherlands: ISMIR, Nov. 2019, pp. 159–165.
- [31] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. Courville, "FiLM: Visual Reasoning with a General Conditioning Layer," in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. New Orleans, LA, USA: AAAI, Dec. 2017.
- [32] O. Slizovskaia, G. Haro, and E. Gómez, "Conditioned Source Separation for Music Instrument Performances," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2083–2095, 2021.
- [33] B. Li, X. Liu, K. Dinesh, Z. Duan, and G. Sharma, "Creating a Multitrack Classical Music Performance Dataset for Multimodal Music Analysis: Challenges, Insights, and Applications," *IEEE Transactions on Multimedia*, vol. 21, no. 2, pp. 522–535, Feb. 2019.
- [34] L. Lin, Q. Kong, J. Jiang, and G. Xia, "A Unified Model for Zero-shot Music Source Separation, Transcription and Synthesis," in *Proceedings of the 22nd International Society for Music Information Retrieval Conference*. Online: ISMIR, Aug. 2021, pp. 381–388.
- [35] J. H. Lee, H.-S. Choi, and K. Lee, "Audio query-based music source separation," in *Proceedings of the 20th International Society for Music Information Retrieval Conference*. Delft, Netherlands: ISMIR, 2019.
- [36] B. Gfeller, D. Roblek, and M. Tagliasacchi, "One-Shot Conditional Audio Filtering of Arbitrary Sounds," in *Proceedings of the 2021 IEEE International Conference on Acoustics, Speech and Signal Processing*. Toronto, Canada: IEEE, Jun. 2021, pp. 501–505.
- [37] W. Choi, M. Kim, J. Chung, and S. Jung, "LaSAFT: Latent Source Attentive Frequency Transformation For Conditioned Source Separation," in *Proceedings of the 2021 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, Jun. 2021, pp. 171–175.
- [38] Y.-S. Jeong, J. Kim, W. Choi, J. Chung, and S. Jung, "LightSAFT: Lightweight Latent Source Aware Frequency Transform for Source Separation," in *Proceedings of the 2021 Music Demixing Workshop at the 22nd International Society for Music Information Retrieval Conference*. Online: ISMIR, 2021.
- [39] K. Chen, X. Du, B. Zhu, Z. Ma, T. Berg-Kirkpatrick, and S. Dubnov, "Zero-shot Audio Source Separation through Query-based Learning from Weakly-labeled Data," in *Proceedings of the 36th AAAI Conference on Artificial Intelligence*. Online: AAAI, Feb. 2022.
- [40] Q. Kong, K. Chen, H. Liu, X. Du, T. Berg-Kirkpatrick, S. Dubnov, and M. D. Plumbley, "Universal Source Separation with Weakly Labelled Data," May 2023.
- [41] X. Liu, H. Liu, Q. Kong, X. Mei, J. Zhao, Q. Huang, M. D. Plumbley, and W. Wang, "Separate What You Describe: Language-Queried Audio Source Separation," in *Proceedings of the 23rd Annual Conference of the International Speech Communication Association*. Incheon, Korea: ISCA, Sep. 2022, pp. 1801–1805.
- [42] X. Liu, Q. Kong, Y. Zhao, H. Liu, Y. Yuan, Y. Liu, R. Xia, Y. Wang, M. D. Plumbley, and W. Wang, "Separate Anything You Describe," Oct. 2023.
- [43] S. Lee and I. V. Bajic, "Information Flow Through U-Nets," in *Proceedings of the 18th IEEE International Symposium on Biomedical Imaging*. Nice, France: IEEE, Apr. 2021, pp. 812–816.
- [44] P. Seetharaman, G. Wichern, S. Venkataramani, and J. Le Roux, "Class-conditional Embeddings for Music Source Separation," in *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing*. Brighton, United Kingdom: IEEE, May 2019, pp. 301–305.
- [45] D. Petermann, G. Wichern, A. Subramanian, and J. L. Roux, "Hyperbolic Audio Source Separation," in *Proceedings of the 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*. Rhodes Island, Greece: IEEE, Jun. 2023.
- [46] R. Kumar, Y. Luo, and N. Mesgarani, "Music Source Activity Detection and Separation Using Deep Attractor Network," in *Proceedings of the 19th Annual Conference of the International Speech Communication Association*. Hyderabad, India: ISCA, Sep. 2018, pp. 347–351.
- [47] D. Petermann and M. Kim, "Hyperbolic Distance-Based Speech Separation," in *Proceedings of the 2024 IEEE International Conference on Acoustics, Speech and Signal Processing*. Seoul, Korea, Republic of: IEEE, Apr. 2024, pp. 1191–1195.
- [48] D. Samuel, A. Ganeshan, and J. Naradowsky, "Meta-Learning Extractors for Music Source Separation," in *Proceedings of the 2020 IEEE International Conference on Acoustics, Speech and Signal Processing*. Barcelona, Spain: IEEE, May 2020, pp. 816–820.

- [49] E. J. Humphrey, S. Durand, and B. McFee, "OpenMIC-2018: An open dataset for multiple instrument recognition," in *Proceedings of the 19th International Society for Music Information Retrieval Conference*. Paris, France: ISMIR, 2018, pp. 438–444.
- [50] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio Speech and Language Processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [51] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR - Half-baked or Well Done?" in *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing*. Brighton, United Kingdom: IEEE, 2019, pp. 626–630.
- [52] R. Scheibler, "SDR - Medium Rare With Fast Computations," in *Proceedings of the 2022 IEEE International Conference on Acoustics, Speech and Signal Processing*. Singapore, Singapore: IEEE, 2022, pp. 701–705.
- [53] C.-B. Jeon, G. Wichern, F. G. Germain, and J. Le Roux, "Why does music source separation benefit from cacophony?" in *Proceedings of the 1st Workshop on Explainable Machine Learning for Speech and Audio at the 2024 IEEE International Conference on Acoustics, Speech, and Signal Processing*. Seoul, Korea, Republic of: IEEE, Feb. 2024.