

Pre-training Music Classification Models via Music Source Separation

Christos Garoufis^{1,2,3}, Athanasia Zlatintsi^{1,2,3}, and Petros Maragos^{2,3}

¹*Institute of Language and Speech Proc., Athena Research Center, Athens, Greece*

²*Robotics Institute, Athena Research Center, Athens, Greece*

³*School of ECE, National Technical University of Athens, Athens, Greece*

{christos.garoufis,nancy.zlatintsi}@athenarc.gr, maragos@cs.ntua.gr

Abstract—In this paper, we study whether music source separation can be used as a pre-training strategy for music representation learning, targeted at music classification tasks. To this end, we first pre-train U-Net networks under various music source separation objectives, such as the isolation of vocal or instrumental sources from a musical piece; afterwards, we attach a classification network to the pre-trained U-Net and jointly finetune the whole network. The features learned by the separation network are also propagated to the tail network through a convolutional feature adaptation module. Experimental results in two widely used and publicly available datasets indicate that pre-training the U-Nets with a music source separation objective can improve performance compared to both training the whole network from scratch and using the tail network as a standalone in two music classification tasks, music auto-tagging and music genre classification. We also show that our proposed framework can be successfully integrated into both convolutional and Transformer-based backends, highlighting its modularity.

Index Terms—music source separation, transfer learning, music auto-tagging, music genre classification

I. INTRODUCTION

The recent upsurge in the size of available datasets, as well as their utilization through non-fully supervised learning schemes, is causing a paradigm shift from task-specific approaches to task-agnostic models and representations, suitable for tackling a multitude of inter-related tasks. This trend has transferred to music signal processing as well, with the developed approaches roughly divided in two categories: those that apply transfer learning through models trained in larger datasets, in either similar [1] or not directly related [2] tasks, and those employing a self-supervised pre-training procedure, using either a single [3], [4] or multiple [5]–[7] modalities.

A task that is underexplored for music representation learning, or in conjunction with other tasks in music information retrieval, is music source separation. Music source separation has proven beneficial for artist identification [8], [9], and employed within a joint framework with automatic music transcription [10], [11], instrument activity detection [12], and key estimation [11]. The association between automatically separated audio sources has also been successfully utilized for audio representation learning [13], [14].

This research was supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “3rd Call for H.F.R.I. Research Projects to support Post-Doctoral Researchers” (Project Acronym: i-MRePlay, Project Number: 7773).

A family of music source separation architectures [10], [15]–[18] makes use of a U-Net [19] structure, which involves an encoder that gradually reduces the input resolution to produce low-dimensional features, and a decoder that recovers the isolated sources by iteratively combining these low-dimensional features. This multi-scale structure has motivated the utilization of U-Nets into neural network architectures tackling other tasks, such as music transcription [20] and music representation learning [21], where stacking a U-Net prior to the rest of the network has resulted in performance improvement. Apart from the encoder-decoder structure of U-Nets, their suitability for music source separation implies their potential utilization as an either pre-trained, or able to be finetuned, feature extractor for other tasks. Indeed, spectral, timbral, or high-level attributes of musical pieces can be better captured using either excerpts of isolated sources [14], [22], or recombined low resolution features, containing meaningful high-level information pertaining to these sources [21].

In this paper¹, we investigate whether music source separation can be used for pre-training purposes for other music classification tasks. To this end, driven by the recent success of source separation architectures operating on the Short-Time Fourier Transform (STFT) magnitude [17], [23], we pre-train U-Net networks with a source separation objective, and finetune them jointly with a classification backend (*tail*) in downstream tasks, i.e. music auto-tagging and music genre classification. Our approach bears similarities to a recently introduced waveform-domain architecture for self-supervised music representation learning, TUne+ [21]; however, it operates in a supervised setting in the time-frequency domain, employing a novel pre-training strategy based on music source separation. Experimental results in two datasets, Magna-Tag-A-Tune (MTAT) [24] for music auto-tagging and FMA [25] for music genre classification, with two different classification backends, a VGG-like convolutional network and an Audio Spectrogram Transformer (AST) [26] indicate that training from scratch the classification network with the pre-stacked U-Net does not improve the performance of the classification backends. On the other hand, using a separation objective to pre-train the U-Net can lead to improved results in spite of

¹Code, pretrained models and qualitative results are available at <https://github.com/cgaroufis/MSSPT>

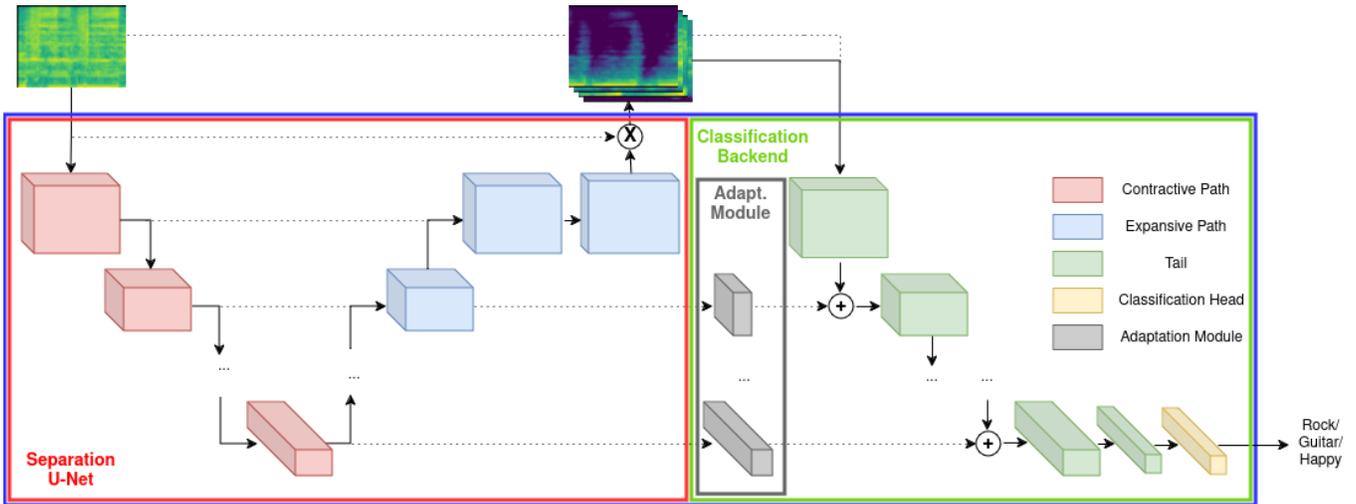


Fig. 1. An overview of our proposed framework. We first train a U-Net architecture (red rectangle) with a music source separation objective; then we attach a classification backend (green rectangle) to the pre-trained separation task network, using an adaptation module (grey rectangle) for feature adaptation, and train the complete network (blue rectangle) to the desired downstream task, using spectrograms from both the original audio excerpt and the isolated sources.

the smaller size of the pre-training dataset [27], indicating the generalizability of our approach to various music classification tasks. We also discuss, for each classification network, the necessity of pre-training it independently to the target downstream task before finetuning it jointly with the U-Net, while showing that the proposed training framework guides the networks towards learning properties of musical signals tied to the musical source used during U-Net pre-training.

II. METHODOLOGY

A. Network Architecture

The framework utilized, an overview of which is depicted in Fig. 1 for the case of a convolutional backend, is inspired by TUne+, which was recently introduced by Vasquez et al. [21] in the context of self-supervised waveform-domain music representation learning. In short, the architecture of TUne+ consists of: a) a *contractive* path (encoder), which gradually downsamples the input through a series of convolutional layers, producing thus multi-resolution features, b) an *expansive* path (decoder), which re-instates the feature map to its initial dimensionality by combining and upsampling the learned contractive features, and c) the *tail*, which produces the final learned embeddings. Since TUne+ was trained within a contrastive learning framework, the output of the tail was used as a representation for downstream tasks. It is also worth noting that feature maps in the expansive path are connected to compatible, dimensionality-wise, feature maps in both the contractive path and the tail through skip connections.

Contractive and Expansive Paths: Both the encoder and the decoder of our network are modeled after the baseline U-Net presented in [17], with the caveat of reducing the filter capacity at half of the original. The encoder receives STFT magnitude spectrograms and consists of 6 blocks, each of which contains 2 2D-convolutional sub-blocks², followed by a

²Throughout the networks, all convolutional kernels are of dimensionality 3×3 , while pooling and upsampling operations are performed at a factor of 2 in each dimension, unless stated otherwise.

max pooling operation. Conversely, the decoder is built symmetrically to the encoder and consists of 6 blocks, comprising series of transposed convolutional layers and convolutional sub-blocks, followed by a final convolutional sub-block, which applies soft multiplicative masks to the input STFT magnitude. Each decoder block receives the outputs of both the previous decoder block and its symmetric encoder block.

Tail and Classification Head: Regarding the tail, we experiment with two different network architectures, a) a convolutional network resembling the Short-Chunk CNN audio backend developed in [28] and b) a Transformer-based backend, based on AST [26]. In both cases, the tail receives the STFT magnitudes of both the original input and the output of the U-Net encoder-decoder, and transforms them to the mel scale, using 128 mel bands, after applying log compression, resulting in a more compact representation. The convolutional backend is built symmetrically to the encoder and the decoder of the network, with 7 2D-convolutional blocks each incorporating 2 convolutional sub-blocks and a max pooling operation; the number of filters in each block is set to the half of [28]. On the other hand, the AST backend follows the typical architecture of a Transformer encoder, consisting of 12 Transformer blocks with an embedding dimension of 768, an internal dimension of 3072 and 12 attention heads per block. Before being inserted into the first Transformer block, the input is split into 16×16 patches, without overlap. Each patch further undergoes a linear projection into an 1D embedding, which is then flattened, summed with a learnable positional embedding and concatenated with a classification token ($[CLS]$). Following the literature [26], [28], the classification head is implemented as a linear layer for the AST, receiving the final representation of the classification token, whereas it further includes an intermediate layer of 512 neurons in the case of the CNN.

Feature Adaptation Module: Similar to [21], the features of the expansive path are propagated towards the classification tail. To this end, the dimensions of the feature maps of the expansive path are first aligned with those of the classification

tail. After alignment, a 2D-convolution is applied, using an equal number of filters to the channels of the representation in the tail, and the two representations are finally summed.

In the case of the CNN backend, since the tail is built symmetrically to the expansive path, its feature maps are paired with those of the tail with the same temporal resolution. The spectral dimensions are aligned through a strided 1D max pooling operation, whereas the convolutional layer operates as a linear projection layer, using 1×1 kernels. On the other hand, the resolution of the embedding sequence remains constant throughout the AST, with the ordering of the embeddings being derived from patchifying the two-dimensional input into an $N_r \times N_c$ grid, which is then flattened. Thus, to produce an embedding sequence with a similar semantic ordering, the 2D convolution is applied with appropriately sized non-overlapping kernels, to generate an $N_r \times N_c$ grid. In this case, the processed feature maps from the expansive path are inserted into the AST at every second layer, after flattening.

B. Training Procedure

We devise a three-stage training scheme for training the network. During the first stage, the U-Net composed by the *contractive* and *expansive* paths (red rectangle in Fig. 1) is trained with a source separation objective, where isolated vocal or instrumental sources of the input audio segment are extracted via applying multiplicative soft masks on the input STFT magnitudes. As the loss function, we utilize the time-domain \mathcal{L}_1 loss, using the mixture phase, between the ground truth and estimated sources. Then, during the (optional) second stage, the classification network (*tail*) is trained independently in the targeted downstream task until convergence. Finally, the classification network is connected to the trained U-Net via the feature adaptation module, and the complete network (blue rectangle in Fig. 1) is jointly finetuned. During the second and third stages, a supervised loss is applied at the output of the network, whereas the separation loss is discarded.

We note that for architectures involving the AST, we follow the commonly used practice [26], [29] of initializing the model parameters from an Imagenet-pretrained checkpoint instead of random weight initialization. Similarly, after conducting preliminary experiments, the convolutional filters of the feature adaptation module were initialized using the Imagenet-pretrained weights of the patch embedding layer, after resizing them by bilinear interpolation to fit the respective kernels.

III. EXPERIMENTAL SETUP

Data & Preprocessing: For source separation pre-training, we utilized the musdb18 dataset [27], a widely used open-access dataset for music source separation. It consists of a total of 150 songs, sampled at 44.1 kHz, with a total duration approximating 10 hours, as well as separate stems for the vocals, drums, bass, and the rest of the instrumental accompaniment (“other”). On the other hand, Magna-Tag-A-Tune (MTAT) [24], as well as the *medium* subset of the Free Music Archive (FMA) dataset [25], were employed as downstream datasets. MTAT consists of 25,863 29-sec song

clips, each associated with a number of tags and sampled at 16 kHz, and has a total duration of approximately 210 hours. Similar to previous work [14], [28], [30] we consider the filtered subset of MTAT, which includes only the portion of the initial dataset labelled with at least one of the top-50 tags. The *medium* subset of FMA comprises 25,000 30-sec song segments, at a sampling rate of 44.1 kHz, each grouped into one out of 16 root genres. It reaches a total duration of approximately 208 hours. For all datasets, either the default [25], [27] or the commonly-used [30] split between training, validation and testing data was adopted.

As preprocessing, all audio excerpts in musdb18 and FMA were resampled, for compatibility purposes, at 16 kHz. The STFT magnitude of all audio segments was computed using a window length of 512 samples and a hop length of 160.

Training and Evaluation Protocol: Regarding the source separation pre-training, U-Nets were pre-trained for all four uni-source cases, i.e. the extraction of the vocals, bass, drums, or the rest of the accompaniment (“other”), as well as the multi-source case, where source estimates for all four aforementioned sources are provided. The initial learning rate was adjusted separately for each source [17], whereas additive data augmentation was employed for all sources with the exception of bass. The Adam [31] optimizer was used with a batch size equal to 8, whereas the learning rate was scheduled after [17].

During fine-tuning, an appropriately sized segment for each backend (3.85 sec for the CNN, 7.70 sec for the AST) was randomly sampled from each song clip at each epoch and fed to the network. As the optimizer, in the case of the CNN we followed the setup proposed in [32], using a batch size of 16 and applying early stopping via the validation loss, with a patience of 10 epochs. On the other hand, for the architecture incorporating the AST, we use Adam, a batch size of 4 and an initial learning rate of $5e-05$, halving it every 2nd epoch after the 5th one; after 15 epochs, we finetune the networks for two additional epochs using stochastic gradient descent.

During inference, song-level predictions are obtained by splitting the clips into slices, with overlap equal to half their length, and then averaging the per-slice predictions. For music auto-tagging, the binary cross-entropy was used as the loss function, and the per-tag average ROC-AUC and PR-AUC scores are employed as metrics. On the other hand, for genre classification, the categorical cross-entropy was used as the loss function, and the weighted accuracy (WA) as the metric.

IV. RESULTS AND DISCUSSION

Main Results: In Table I, we present the results for each case of source-wise pre-training, for both backends. As baselines, we consider the performance of the *tail* models, without receiving the input, or the features, of the pre-trained separator (first row in Table I), as well as the complete models, with the *contractive* and *expansive* paths randomly initialized (second row in Table I). For all cases, we report the mean value and standard error over 5 runs. An initial conclusion is that in contrast to self-supervised learning [21], the inclusion of a randomly initialized U-Net before the network tail does

TABLE I

PERFORMANCE OF THE TRAINED NETWORKS, ACCORDING TO THE TARGET SOURCE FOR PRE-TRAINING THE U-NET, IN MUSIC AUTO-TAGGING (MTAT DATASET) AND MUSIC GENRE CLASSIFICATION (FMA DATASET). WE ALSO COMPARE TO DISCARDING THE U-NET AND USING THE BARE BACKEND (TOP ROW), AS WELL AS RANDOMLY INITIALIZING THE U-NET (SECOND ROW). BOLD DENOTES IMPROVEMENT OVER THE TAIL BASELINE; *DENOTES STATISTICALLY SIGNIFICANT IMPROVEMENT ($p < 0.05$) AFTER APPLICATION OF BONFERRONI-CORRECTED T-TESTS.

U-Net	Pre-Training	Source(s)	CNN Backend			AST Backend		
			MTAT		FMA	MTAT		FMA
			ROC-AUC	PR-AUC	WA (%)	ROC-AUC	PR-AUC	WA (%)
\times	-	-	91.47 \pm 0.03	46.50 \pm 0.11	66.30 \pm 0.31	91.65 \pm 0.05	46.82 \pm 0.20	67.22 \pm 0.25
\checkmark	\times	-	91.38 \pm 0.08	46.32 \pm 0.35	66.53 \pm 0.13	91.58 \pm 0.04	46.77 \pm 0.12	67.10 \pm 0.61
\checkmark	\checkmark	Bass	91.45 \pm 0.08	46.48 \pm 0.18	66.58 \pm 0.24	91.69 \pm 0.08	47.00 \pm 0.24	67.89 \pm 0.31
\checkmark	\checkmark	Drums	91.47 \pm 0.07	46.68 \pm 0.20	66.80 \pm 0.19	91.57 \pm 0.03	46.57 \pm 0.20	66.11 \pm 0.46
\checkmark	\checkmark	Other	91.59 \pm 0.07	46.98 \pm 0.13*	67.14 \pm 0.12*	91.78 \pm 0.12	47.49 \pm 0.23	67.09 \pm 0.33
\checkmark	\checkmark	Vocals	91.85 \pm 0.03*	47.16 \pm 0.17*	66.10 \pm 0.41	91.89 \pm 0.08*	47.21 \pm 0.34	66.77 \pm 0.14
\checkmark	\checkmark	Multiple	91.50 \pm 0.02	46.65 \pm 0.12	66.52 \pm 0.21	91.87 \pm 0.08*	47.31 \pm 0.13*	67.40 \pm 0.93

TABLE II

IMPACT OF PRE-TRAINING THE CLASSIFICATION BACKEND BEFORE FINE-TUNING FOR THE CASE OF ACCOMPANIMENT SEPARATION AND BOTH CNN AND AST-BASED BACKENDS, IN MUSIC AUTO-TAGGING (MTAT).

Backend	Pre-Training	ROC-AUC	PR-AUC
CNN	\times	91.35 \pm 0.15	46.15 \pm 0.44
CNN	\checkmark	91.59 \pm 0.07	46.98 \pm 0.13
AST	\times	91.78 \pm 0.12	47.49 \pm 0.23
AST	\checkmark	91.47 \pm 0.08	46.58 \pm 0.16

not necessarily guarantee improved performance. On the other hand, pre-training the U-Net in music source separation, despite using a dataset smaller than the downstream ones for pre-training, can prove beneficial for music classification tasks. In particular, most separation objectives lead to improved performance in music auto-tagging, irrespective of the classification backend, with the largest PR-AUC increase achieved for vocal separation in the case of the CNN, and accompaniment separation in the case of AST; this improvement approaches a relative increase of 1.5% over the bare backends, whose performance is close to the one reported in the literature [28], [29]. Regarding music genre classification, evaluation of the CNN-based architecture leads to better results for the majority of the pre-training objectives; this improvement is statistically significant for accompaniment separation. However, the trend is reversed for the AST backend, where only the bass tracks lead to consistent performance increase as pre-training targets.

Regarding the model comparison according to the integrated backend, we observe that, in agreement with the recent literature [29], the variants combining the pre-trained U-Net with AST slightly outperform, for most source-task configurations, the fully convolutional ones. We also note that, as we will expand on later, the best results (reported in Table I) were obtained through the complete three-stage scheme for the CNN, but without the optional second stage for the AST.

Concerning the performance of individual sources, the results indicate a consistently positive effect of pre-training the U-Net on separating the melodic accompaniment of the musical mixture. On the other hand, while setting the vocals as the target source during pre-training results in worse performance than the tail baselines in music genre classification, it leads to improved performance in the task of music auto-tagging. This result can be correlated to the efficacy of vocal excerpts in identifying artist-specific features, as shown in [8], [9], and as will be further discussed, is linked to the nature of the

tags in MTAT. Finally, in contrast to contrastive representation learning [14], using a multi-source pre-training objective does not necessarily outperform single-source pre-training; it is noteworthy, though, that it appears to benefit more from the higher capacity of the AST backend in both downstream tasks.

Impact of Backend Pre-Training: In Table II, we present the ROC-AUC and PR-AUC scores achieved in MTAT for both classification backends, depending on whether the classification backend was independently pre-trained before finetuning, for the case of accompaniment separation pre-training. The results indicate that while in the case of the CNN, the network benefits from independent pre-training of the classification backend, this technique proves detrimental for the AST backend. We hypothesize that the relative lack of inductive biases in Transformer architectures, compared to CNNs, leads to easier network overfitting; additionally, the use of Imagenet-pretrained weights provides a good starting point for the AST, enabling its smooth generalization into other domains.

Qualitative Analysis: In Fig. 2, we display the tag-wise PR-AUC score difference between the tail baseline and our proposed framework (average between the CNN and AST configurations) for the case of vocal separation pre-training, for the top-10 and bottom-10 MTAT tags in terms of performance differential; we also display the performance achieved with the other separation pre-training targets in the same tags. We observe that vocal separation pre-training succeeds in steering the model towards tags related to the vocals, as well as semantic information about them; for some of these tags, application of the paired t-test showed a statistically significant improvement ($p < 0.01$). On the other hand, the tags where the vocally pre-trained model performs worse than the baseline largely concern instruments or instrumental genres.

V. CONCLUSIONS

In this work, we investigated whether music source separation can be utilized as a pre-training objective for music classification tasks. Experimental results in two different tasks, music auto-tagging and music genre classification, using both convolutional and Transformer-based backends indicate that our proposed approach can lead to improved performance in both tasks, depending on the target source during pre-training, highlighting its potential applicability in diverse music classification tasks as well as its flexibility.

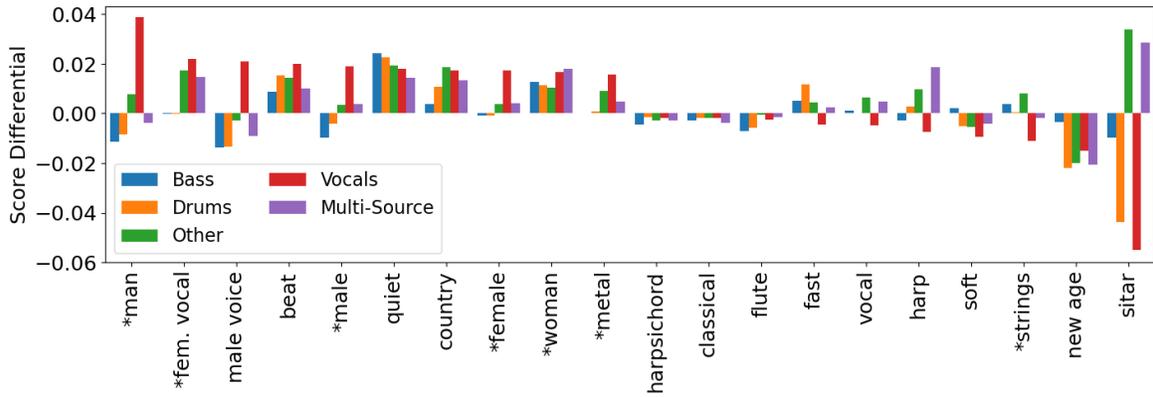


Fig. 2. Relative difference (PR-AUC scores, averaged across CNN-based and AST-based models) between the baseline and our proposed framework in selected MTAT tags, according to the pre-training objective; * denotes statistically significant ($p < 0.01$) difference between the baseline and the vocals.

In the future, we would like to exploit more diverse source separation techniques for the pre-training stage [33], as well as alleviate the computational overhead introduced by network stacking by means of feature-level supervision [34].

Acknowledgments: We sincerely thank C. Papaioannou for his constructive comments on the content of the paper.

REFERENCES

- [1] K. Choi, G. Fazekas, M. Sandler, and K. Cho, "Transfer Learning for Music Classification and Regression Tasks," in *Proc. ISMIR 2017*, Suzhou, China, 2017.
- [2] R. Castellon, C. Donahue, and P. Liang, "Codified Audio Language Modeling Learns Useful Representations for Music Information Retrieval," in *Proc. ISMIR 2021*, online, 2021.
- [3] J. Spijkervet and J. A. Burgoyne, "Contrastive Learning of Musical Representations," in *Proc. ISMIR 2021*, online, 2021.
- [4] H. Zhao, C. Zhang, B. Zhu, Z. Ma, and K. Zhang, "S3T: Self-Supervised Pre-training with Swin Transformer for Music Classification," in *Proc. ICASSP 2022*, Singapore, Singapore, 2022.
- [5] Q. Huang, A. Jansen, J. Lee, R. Ganti, J. Y. Li, and D. P. Ellis, "MuLan: A Joint Embedding of Music Audio and Natural Language," in *Proc. ISMIR 2022*, Bengaluru, India, 2022.
- [6] I. Manco, E. Benetos, E. Quinton, and G. Fazekas, "Learning Music Audio Representations via Weak Language Supervision," in *Proc. ICASSP 2022*, Singapore, Singapore, 2022.
- [7] K. Avramidis, S. Stewart, and S. Narayanan, "On the Role of Visual Context in Enriching Music Representations," in *Proc. ICASSP 2023*, Rhodes, Greece, 2023.
- [8] B. Sharma, R. K. Das, and H. Li, "On the Importance of Audio-Source Separation for Singer Identification in Polyphonic Music," in *Proc. Interspeech 2019*, Graz, Austria, 2019.
- [9] K. Kim, J. Lee, S. Kum, and J. Nam, "Learning a Cross-Domain Embedding Space of Vocal and Mixed Audio with a Structure-Preserving Triplet Loss," in *Proc. ISMIR 2021*, online, 2021.
- [10] L. Lin, Q. Kong, J. Jiang, and G. Xia, "A Unified Model for Zero-Shot Music Source Separation, Transcription and Synthesis," in *Proc. ISMIR 2021*, online, 2021.
- [11] K. W. Cheuk, K. Choi, Q. Kong *et al.*, "Jointist: Simultaneous Improvement of Multi-instrument Transcription and Music Source Separation via Joint Training," *arXiv preprint arXiv:2302.00286*, 2023.
- [12] Y.-N. Hung and A. Lerch, "Multitask Learning for Instrument Activation Aware Music Source Separation," in *Proc. ISMIR 2020*, online, 2020.
- [13] E. Fonseca, A. Jansen, D. P. Ellis, S. Wisdom *et al.*, "Self-Supervised Learning from Automatically Separated Sound Scenes," in *Proc. WAS-PAA 2021*, New Waltz, USA, 2021.
- [14] C. Garoufis, A. Zlatintsi, and P. Maragos, "Multi-Source Contrastive Learning from Musical Audio," in *Proc. SMC 2023*, Stockholm, Sweden, 2023.
- [15] A. Jansson, E. Humphrey, N. Montecchio, R. Bittner, A. Kumar, and T. Weyde, "Singing Voice Separation with Deep U-Net Convolutional Networks," in *Proc. ISMIR 2017*, Suzhou, China, 2017.
- [16] D. Stoller, S. Ewert, and S. Dixon, "Wave-U-Net: A Multi-Scale Neural Network for End-to-End Audio Source Separation," in *Proc. ISMIR 2018*, Paris, France, 2018.
- [17] Q. Kong, Y. Cao, H. Liu, K. Choi, and Y. Wang, "Decoupling Magnitude and Phase Estimation with Deep Res-U-Net for Music Source Separation," in *Proc. ISMIR 2021*, online, 2021.
- [18] C. Garoufis, A. Zlatintsi, and P. Maragos, "HTMD-Net: A Hybrid Masking-Denoising Approach to Time-Domain Monaural Singing Voice Separation," in *Proc. EUSIPCO 2021*, online, 2021.
- [19] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proc. MICCAI 2015*, Munich, Germany, 2015.
- [20] F. Pedersoli, G. Tzanetakis, and K. M. Yi, "Improving Music Transcription by Pre-Stacking a U-Net," in *Proc. ICASSP 2020*, online, 2020.
- [21] M. Vázquez and J. Burgoyne, "Tailed U-Net: Multi-Scale Music Representation Learning," in *Proc. ISMIR 2022*, Bengaluru, India, 2022.
- [22] J. De Berardinis, A. Cangelosi, and E. Coutinho, "The Multiple Voices of Musical Emotions: Source Separation for Improving Music Emotion Recognition Models and their Interpretability," in *Proc. ISMIR 2020*, online, 2020.
- [23] 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.
- [24] E. Law, K. West, M. I. Mandel, M. Bay, and J. S. Downie, "Evaluation of Algorithms using Games: The Case of Music Tagging," in *Proc. ISMIR 2009*, Kobe, Japan, 2009.
- [25] M. Defferrard, K. Benzi, P. Vandergheynst, and X. Bresson, "FMA: A Dataset for Music Analysis," in *Proc. ISMIR 2017*, Suzhou, China, 2017.
- [26] Y. Gong, Y.-A. Chung, and J. Glass, "AST: Audio Spectrogram Transformer," in *Proc. Interspeech 2021*, Brno, Czechia, 2021.
- [27] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimilakis, and R. Bittner, "The MUSDB18 Corpus for Music Separation," <https://doi.org/10.5281/zenodo.1117372>, 2017.
- [28] M. Won, A. Ferraro, D. Bogdanov, and X. Serra, "Evaluation of CNN-Based Automatic Music Tagging Models," in *Proc. SMC 2020*, online, 2020.
- [29] C. Papaioannou, E. Benetos, and A. Potamianos, "From West to East: Who Can Understand the Music of the Others Better?" in *Proc. ISMIR 2023*, Milan, Italy, 2023.
- [30] J. Lee, J. Park, K. L. Kim, and J. Nam, "SampleCNN: End-to-end Deep Convolutional Neural Networks using Very Small Filters for Music Classification," *Applied Sciences*, vol. 8, no. 1, p. 150, 2018.
- [31] D. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proc. ICLR 2015*, San Diego, CA, USA, 2015.
- [32] M. Won, S. Chun, and X. Serra, "Toward Interpretable Music Tagging with Self-Attention," *arXiv preprint arXiv:1906.04972*, 2019.
- [33] 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 *Proc. AAAI 2022*, Vancouver, BC, Canada, 2022.
- [34] Y.-N. Hung and A. Lerch, "Feature-Informed Embedding Space Regularization for Audio Classification," in *Proc. EUSIPCO 2022*, Belgrade, Serbia, 2022.