

Multi-Level and Multi-Scale Feature Aggregation Using Pretrained Convolutional Neural Networks for Music Auto-Tagging

Jongpil Lee and Juhan Nam, *Member, IEEE*

Abstract—Music auto-tagging is often handled in a similar manner to image classification by regarding the two-dimensional audio spectrogram as image data. However, music auto-tagging is distinguished from image classification in that the tags are highly diverse and have different levels of abstraction. Considering this issue, we propose a convolutional neural networks (CNN)-based architecture that embraces multi-level and multi-scaled features. The architecture is trained in three steps. First, we conduct supervised feature learning to capture local audio features using a set of CNNs with different input sizes. Second, we extract audio features from each layer of the pretrained convolutional networks separately and aggregate them altogether giving a long audio clip. Finally, we put them into fully connected networks and make final predictions of the tags. Our experiments show that using the combination of multi-level and multi-scale features is highly effective in music auto-tagging and the proposed method outperforms the previous state-of-the-art methods on the MagnaTagATune dataset and the Million Song Dataset. We further show that the proposed architecture is useful in transfer learning.

Index Terms—Convolutional neural networks, feature aggregation, music auto-tagging, transfer learning.

I. INTRODUCTION

MUSIC auto-tagging is a task that predicts descriptive words of music from the audio signals. Recently, as convolutional neural network (CNN) has become the de-facto standard in image classification, the deep learning approach has been actively explored in music auto-tagging as well as by using the spectrogram and its variants as input data and so recasting it as a multi-label classification task on the two-dimensional (2-D) time-frequency images [1]–[4].

However, music auto-tagging is distinguished from image classification in that the tags consist of words that are highly diverse and have different levels of abstraction. For example,

some words, particularly instrument-related ones, such as *female vocalist*, *guitar*, and *saxophone* are objective descriptions of specific sound sources. They tend to be local and repetitive within an audio clip and are basically predicted from the physical properties of the sound sources. On the other hand, other words such as *rock*, *happy*, and *80s* are discriminative descriptions of the overall content in terms of genre, mood, and years. They are more global and comprehensive, requiring longer audio segments to predict them.

While the tags are positioned in different levels or time-scales in a hierarchy, the majority of previous work predicted them from the same level or scale of features as typically done in image classification. There are a few that handled this issue explicitly by comparing or combining multi-layer or multi-scale audio features. In terms of feature level, Lee *et al.* [5] used convolutional restricted Boltzmann machine to learn hierarchical features in an unsupervised manner. They compared each layer of features and their combination for music genre and artist classification. However, the evaluation was not sufficiently comprehensive as they used small datasets. Hamel and Eck [6] applied deep neural networks pretrained with deep belief networks to learn hierarchical features and compared each layer of features. However, the learned features were obtained from single frames of spectrogram, which is too local to capture musically meaningful and rich patterns. In terms of time-scale, Hamel *et al.* [7] investigated combining different resolutions of spectrograms, and Dieleman and Schrauwen [8] improved the approach further using Gaussian and Laplacian pyramids. However, they conducted the multi-scaling and feature concatenation only on the input layer, focusing on the spectrograms.

The general consensus from the previous work is that the individual tags have different performance sensitivity to different time-scales and levels of features, and so combining all of them provides the best results. With this lesson, we propose a CNN-based architecture that handles multi-level and multi-scale of audio features more comprehensively. The architecture is trained in three steps: local feature learning, feature aggregation, and global classification. The local feature learning is carried out using a set of CNNs. They are trained in a supervised manner with the tag labels, taking different sizes of input frames and accordingly more hidden layers. The feature aggregation step extracts local features from all layers and time-scales using the pretrained CNNs and summarizes them into a single feature vector. The last stage performs final predictions of tags from the aggregated feature vector using a fully connected neural network. By nature of the separated steps, this architecture is capable of transfer learning, that is, by conducting local feature

Manuscript received January 16, 2017; revised April 1, 2017; accepted May 30, 2017. Date of publication June 8, 2017; date of current version June 28, 2017. This work was supported in part by the Korea Advanced Institute of Science and Technology (Project G04140049) and in part by the National Research Foundation of Korea (Project 2015R1C1A1A02036962). The associate editor coordinating the review of this manuscript and approving it for publication was Dr. James E. Fowler. (*Corresponding author: Juhan Nam.*)

The authors are with the Graduate School of Culture Technology, Korea Advanced Institute of Science and Technology, Daejeon, 34141, South Korea (e-mail: richter@kaist.ac.kr; juhanam@kaist.ac.kr).

Color versions of one or more of the figures in this letter are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LSP.2017.2713830

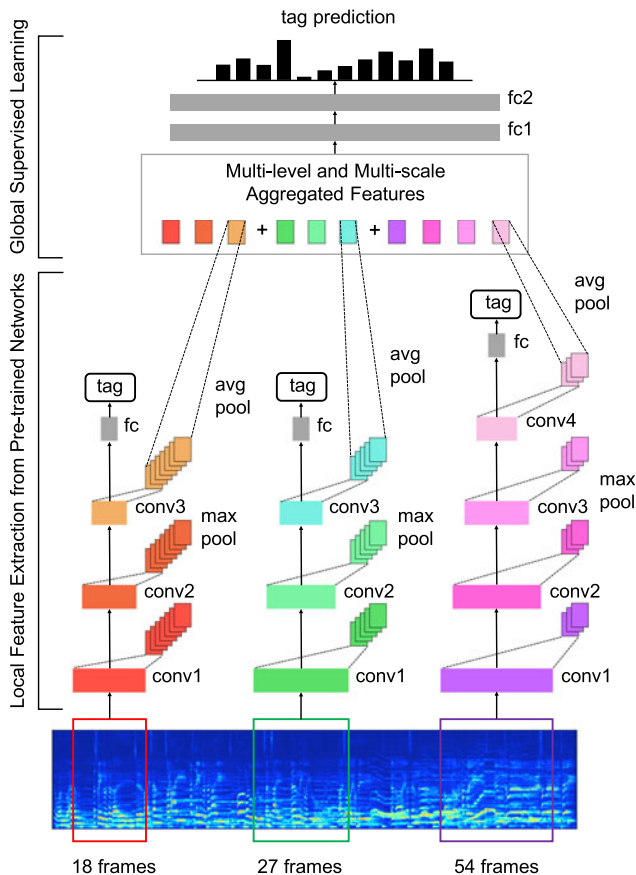


Fig. 1. Proposed architecture for music auto-tagging. “fc” stands for “fully-connected layer.” The CNNs capture local audio features by supervised learning and are used as multi-level and multi-scale feature extractors. Given a long audio clip, the extracted features are aggregated. Finally, the global classifier makes final predictions of tags on them.

learning with one large dataset, and then feature aggregation and global classification with another dataset.

We evaluate the proposed architecture with two popularly used datasets primarily for music auto-tagging and also in transfer learning settings where music genre classification is performed using pretrained CNNs. Our experiments show how different combinations of multi-layer and multi-time-scale features improve the accuracy and also how the architecture outperforms the previous state-of-the-art architectures.

II. PROPOSED METHOD

The overall architecture that we propose is illustrated in Fig. 1. The following describes the three steps to train it.

A. Local Feature Learning

In the first step, we perform supervised feature learning with a set of CNNs. We choose the segment sizes such that the hidden layers capture multi-level audio features within one to several musical beats for different beats per minute (BPM). For example, 18, 27, and 54 frames correspond to 420, 630, and 1260 ms. They take care of at least one beat long in songs with 48 to 143 BPM, which is the tempo range that covers the majority of popular music. The CNNs are configured to conduct one-dimensional (1-D) convolution in all layers, assuming that

low-frequency and high-frequency content do not share the weights (as opposed to images) and so the whole frequency range is under the receptive field of the filters. We determined the filter width “3” in the convolutional layers by referring to the VGG net [9]. We first built 27 frames model, which is composed of 3 convolutional layers [(3,128)-(3)-(3,128)-(3)-(3,256)-(3)], 1 fully connected layer (256), and final prediction layer (50). Using this as a reference configuration, for the 18 frames model, the last convolutional layer was replaced by (2,256)-(2) and for models with more than 54 frames as input, we simply added (2,256)-(2) layers according to the input size. Zero-padding is applied to each convolution layer to preserve the size. The convolution stride is fixed to 1 and max-pooling stride is set to the same as the pooling length. We perform the back-propagation with tag labels from the dataset. Each of the models can be actually used to predict the tags for a long audio clip by averaging the local predictions. We call them “local models”

B. Multi-level and Multi-scale Feature Aggregation

The pretrained CNNs can be viewed as feature extractors. Since a single CNN model can extract different levels of features given the input size and we train them with different input sizes, we can extract multi-level and multi-scale features from them. In order to handle the long audio clips (typically, about 30 s), we aggregate them into a single large feature vector and use them as a song-level representation. For example, the output shape of each convolution layer of one segment in a 27-frame model is (27,128)-(9,128)-(3,256). After extracting features for all segments in 30 s audio, the song-level feature dimensions become (46,27,128)-(46,9,128)-(46,3,256). In order to extract the most representative feature, we apply max-pooling over each segment separately. We then summarize them into single feature vectors by average pooling over the long audio clip separately for each layer. This scheme, that is, max-pooling followed by average pooling, was used as an effective means to summarize local features [10], [11]. As a result, the dimensionality of the concatenated feature vector will be determined by the sum of the numbers of filters from all layers. For instance, the 27 frames model will have $128 + 128 + 256$ dimensional feature vector. This is repeated for all different input sizes and they are finally concatenated into a large feature vector.

C. Global Classification

In this step, we make final predictions of tags from the aggregated multi-level and multi-scale features. We train another classifier, which is a neural network with two fully connected hidden layers with 512 or 1024 units depending on the datasets. Since the feature aggregation and global classification steps are separated from the local feature learning, we can conduct transfer learning, which has been explored effectively as well for music audio data [12], [13], using pretrained CNNs with a large dataset. In our experiment, we evaluate the transfer learning setting for several different datasets.

III. EXPERIMENTS

A. Datasets

To evaluate the proposed architecture, we primarily used the MagnaTagATune (MTAT) dataset [14] and Million Song Dataset

(MSD) annotated with the Last.fm tags [15]. We filtered out the tags and used most frequently labeled 50 tags for both MTAT and MSD, following the previous work [2], [4].¹ Also, all songs in the two datasets were trimmed to 29.1 s long. We used area under receiver operating characteristic (AUC) as a primary evaluation metric for music auto-tagging. In addition, we conducted genre classification tasks, GTZAN [16] (10 genres, fault-filtered split that is designed to avoid the repetition of artist across training/validation/test list [17]) and Tagtraum genre annotations on MSD (15 genres, stratified split with 80% training data of CD2C version)[18], in a transfer learning setting where the pretrained CNNs with MSD are used as feature extractors.

B. Training Details

Mel-frequency spectrogram with 128 bins is used as the input representation. The parameters are set to 22.05 kHz sampling rate (by resampling when necessary), 512 samples of hop size, 1024 samples of Hanning window, and magnitude compression with a nonlinear curve, $\log(1 + C|A|)$ where A is the magnitude and C is set to 10. As a result, each clip has 1250 frames and is divided into 69, 46, 23, 11, and 5 segments for the corresponding 18, 27, 54, 108, and 216 frames models, respectively. The detailed parameters to train the networks are as follows: sigmoid activation for output layer with binary cross entropy loss, batch normalization [19] and ReLU activation for every intermediate layer, 0.5 of dropout for hidden fully connected layers, and stochastic gradient descent with 0.9 Nesterov momentum. Also, we conducted the input normalization simply by dividing standard deviation after subtracting mean value of the entire input data. We used Keras [20] and Theano [21] framework running on GPUs. Training of local models with small input size such as 18 frames model on MSD have taken about 5 days in total.

C. Compared Models

A typical approach for music auto-tagging is to take about 3-sec long audio segments as input and average the outputs to make final predictions for an audio clip (e.g., [2]). Here we call them “local” models as already denoted in Section II-A. On the other hand, we call our proposed architecture “global” models as it aggregates features from local models and makes final predictions directly from the audio clip. In our experiment, we evaluate the two models with various combinations of input sizes and feature levels.

IV. RESULTS AND DISCUSSION

A. Comparison of Local and Global Models

Fig. 2 shows the evaluation results for the local and global models on MTAT and MSD for different input sizes. From the local models, the AUC reaches the highest level when the input size is 108 frames (about 2.5 s), indicating that taking 3 s as the input size is actually a reasonable choice. However, the proposed global models consistently outperform the local models and the performance increment is more vivid on MSD.

B. Effect of Multi-level Features

Fig. 3 dissects the effect of multi-level features further in the global model. When a single-level feature is used, higher level

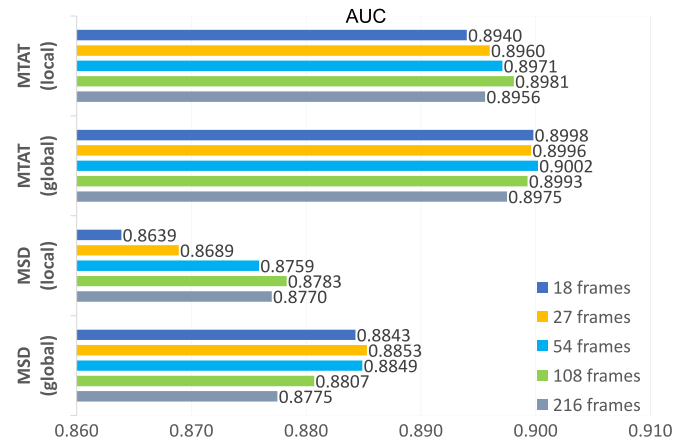


Fig. 2. Comparison of local and global models for different input sizes.

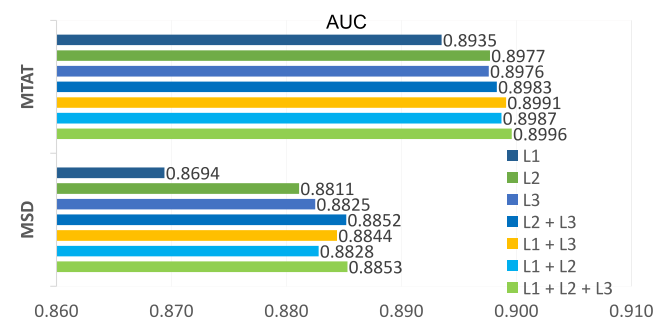


Fig. 3. Comparison of various combinations of feature levels in the 27 frames model. L1, L2, and L3 denote the hidden layer activations in each layer in the CNN model.

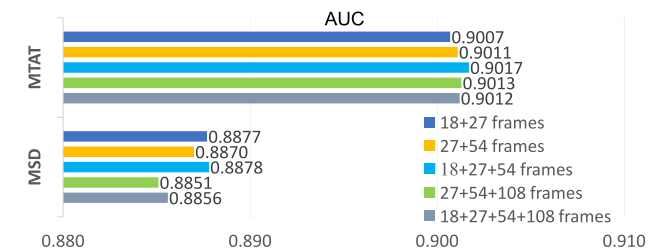


Fig. 4. Comparison of various combinations of multi-scale features.

features (L3) apparently work better than lower ones (L1 and L2). When multi-level features are concatenated, the AUC levels consistently increase on both datasets. One interesting result is that each layer has different importance. For example, on MTAT, the absence of L1 features decreases the AUC most. On the other hand, on MSD, the absence of L3 features makes the highest drop. This may attribute to difference in tag words between the two datasets. For example, MTAT contains more instrument-related tags than genre or mood tags, compared to MSD.

C. Effect of Multi-Scaled Features

We now discuss the effect of multi-scale features. Fig. 4 shows the results for different combinations of varying input size in the global models. Compared to multi-level features, the performance gain is not strong but the use of multi-scaled features is definitely helpful. The best result is achieved in both MSD and

¹https://github.com/jongpillee/music_dataset_split

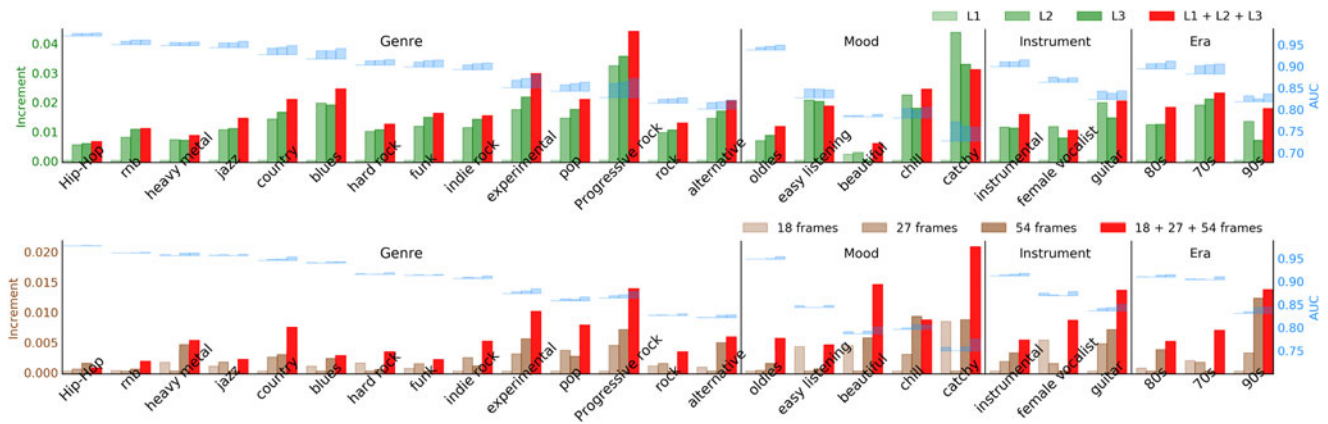


Fig. 5. AUC results of the global model on MSD for selected tags. The top compares multi-level features (L1, L2, L3) and their concatenation when the input size is 27 frames. The bottom compares multi-scale features (18, 27, 54 frames) and their concatenation. The right blue bar-graphs scaled by the y-axis on the right side represents the absolute AUCs by which the tags are sorted. The left green (top) and brown (bottom) bar-graph scaled by the y-axis on the left-side represents the performance increments given the the smallest level on individual tags.

TABLE I
GLOBAL CLASSIFICATION RESULTS USING FEATURES EXTRACTED FROM THE CNN PRETRAINED WITH MSD

Dataset	Task	Measurement	
MTAT	Autotagging	AUC	0.9021
Tagtraum	Genre classification	Accuracy	0.766
GTZAN (fault-filtered)	Genre classification	Accuracy	0.720

MTAT when 18, 27, and 54 frames models are combined. This may be inferred from Fig. 2.

D. Performances Visualization of Individual Tags

We investigate the global model even further by comparing the performance sensitivity of individual tags to different feature levels and time-scales, as illustrated in Fig. 5. In the multi-level comparison (top), since supervised training is performed with the tags in the local feature learning stage, gradual increment is expected as the layer goes up. This trend, however, does not work consistently for every single tag. For example, some tags including *blues*, *chill*, *guitar*, and *80s* favor L2 features more than L3. The nongradual increment is observed in the multi-scale comparison (bottom) as well. Some tags including *heavy metal*, *experimental*, *progressive rock*, *alternative*, *chill*, *guitar*, and *90s* favor 54 frames whereas others including *hard rock*, *easy listening*, *female vocalist*, and *70s* work better on 18 frames. Overall, we can ensure that the best AUCs in almost all tags are achieved when the multi-level and multi-scale features are concatenated.

E. Transfer Learning and Comparison to State-of-the-Art Results

In Table I, we show the performance of the proposed architecture in the transfer learning settings where MSD is used to pretrain the CNNs as a feature extractor and other datasets are for the final classification. Interestingly, the auto-tagging performance on MTAT is even greater than those using local models trained from the MTAT dataset itself. Also, it shows the

TABLE II
COMPARISON OF OUR MODELS TO PRIOR STATE-OF-THE-ART MODELS

Model	MTAT	MSD	GTZAN (fault-filtered)
Bag of multi-scaled features [8]	0.898	–	–
1-D CNN [2]	0.8815	–	–
Transfer learning [13]	0.88	–	–
Persistent CNN [22]	0.9013	–	–
Time–frequency CNN [23]	0.9007	–	–
2-D CNN [4]	0.894	0.851	–
CRNN [24]	–	0.862	–
2-D CNN [17]	–	–	0.632
Temporal features [25]	–	–	0.659
Local model	0.8981	0.8783	–
Global model with multi-level features	0.9002	0.8853	–
Global model with both multi-features	0.9017	0.8878	–
Global model with both multi-features (pretrained with MSD)	0.9021	–	0.720

music genre classification results on fault-filtered GTZAN and Tagtraum genre annotations on MSD. To our knowledge, we report the performance on the Tagtraum for the first time. From Table II, the accuracy on the fault-filtered GTZAN is greater than previously reported ones. Table II also compares the best results in the local and global models to those from the previous state-of-the-art models on MTAT and MSD. They show that our proposed architecture is highly effective.

V. CONCLUSION AND FUTURE WORK

In this letter, we presented a CNN-based architecture, which is designed considering different levels of abstraction in music tags. We showed the effectiveness of the architecture by evaluating different combinations of the multi-level and multi-scale features and also by applying it to transfer learning settings. Finally, we showed that our proposed architecture achieves the best results on the three popularly used datasets. As future work, we plan to train the architecture in a multi-task learning manner by optimizing the local CNNs and global aggregated networks simultaneously.

REFERENCES

- [1] P. Hamel, S. Lemieux, Y. Bengio, and D. Eck, "Temporal pooling and multiscale learning for automatic annotation and ranking of music audio," in *Proc. 12th Int. Conf. Music Inf. Retrieval*, 2011, pp. 729–734.
- [2] S. Dieleman and B. Schrauwen, "End-to-end learning for music audio," in *Proc. 2014 IEEE Int. Conf. Acoust., Speech Signal Process.* 2014, pp. 6964–6968.
- [3] J. Pons, T. Lidy, and X. Serra, "Experimenting with musically motivated convolutional neural networks," in *Proc. 2016 14th Int. Workshop Content-Based Multimedia Indexing*. 2016, pp. 1–6.
- [4] K. Choi, G. Fazekas, and M. Sandler, "Automatic tagging using deep convolutional neural networks," in *Proc. 17th Int. Conf. Music Inf. Retrieval*, 2016, pp. 805–811.
- [5] H. Lee, Y. Largman, P. Pham, and A. Y. Ng, "Unsupervised feature learning for audio classification using convolutional deep belief networks," in *Proc. 22nd Int. Conf. Neural Inf. Process. Syst.*, 2009, pp. 1096–1104.
- [6] P. Hamel and D. Eck, "Learning features from music audio with deep belief networks," in *Proc. 11th Int. Conf. Music Inf. Retrieval*. Utrecht, The Netherlands, 2010, pp. 339–344.
- [7] P. Hamel, Y. Bengio, and D. Eck, "Building musically-relevant audio features through multiple timescale representations," in *Proc. 13th Int. Conf. Music Inf. Retrieval*, 2012, pp. 553–558.
- [8] S. Dieleman and B. Schrauwen, "Multi-scale approaches to music audio feature learning," in *Proc. 14th Int. Soc. Music Inf. Retrieval Conf. Pontificia Universidade Católica do Paraná*, 2013, pp. 116–121.
- [9] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd Int. Conf. on Learning Representations*, 2015.
- [10] J. Nam, J. Herrera, M. Slaney, and J. O. Smith, "Learning sparse feature representations for music annotation and retrieval," in *Proc. 13th Int. Conf. Music Inf. Retrieval*, 2012, pp. 565–570.
- [11] J. Nam, J. Herrera, and K. Lee, "A deep bag-of-features model for music auto-tagging," arXiv:1508.04999, 2015.
- [12] P. Hamel, M. E. P. Davies, K. Yoshii, and M. Goto, "Transfer learning in MIR: Sharing learned latent representations for music audio classification and similarity," in *Proc. 14th Int. Conf. Music Inf. Retrieval*, 2013, pp. 9–14.
- [13] A. Van Den Oord, S. Dieleman, and B. Schrauwen, "Transfer learning by supervised pre-training for audio-based music classification," in *Proc. Conf. Int. Soc. Music Inf. Retrieval*, 2014, pp. 29–34.
- [14] 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. Conf. Int. Soc. Music Inf. Retrieval*, 2009, pp. 387–392.
- [15] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere, "The million song dataset," in *Proc. 10th Int. Conf. Music Inf. Retrieval*, 2009, pp. 387–392.
- [16] G. Tzanetakis and P. Cook, "Musical genre classification of audio signals," *IEEE Trans. Speech Audio Process.*, vol. 10, no. 5, pp. 293–302, Jul. 2002.
- [17] C. Kereliuk, B. L. Sturm, and J. Larsen, "Deep learning and music adversaries," *IEEE Trans. Multimedia*, vol. 17, no. 11, pp. 2059–2071, Nov. 2015.
- [18] H. Schreiber, "Improving genre annotations for the million song dataset," in *Proc. 16th Int. Conf. Music Inf. Retrieval*, 2015, pp. 241–247.
- [19] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. 32nd Int. Conf. Machine Learning*, 2015, pp. 448–456.
- [20] F. Chollet, "Keras: Deep learning library for theano and tensorflow," 2015. [Online]. Available: <https://github.com/fchollet/keras>
- [21] R. Al-Rfou *et al.*, "Theano: A python framework for fast computation of mathematical expressions," arXiv:1605.02688, 2016.
- [22] J.-Y. Liu, S.-K. Jeng, and Y.-H. Yang, "Applying topological persistence in convolutional neural network for music audio signals," arXiv:1608.07373, 2016.
- [23] U. Güçlü, J. Thielen, M. Hanke, M. van Gerven, and M. A. van Gerven, "Brains on beats," in *Proc. Advances Neural Inf. Process. Syst.*, 2016, pp. 2101–2109.
- [24] K. Choi, G. Fazekas, M. Sandler, and K. Cho, "Convolutional recurrent neural networks for music classification," in *Proc. 2017 IEEE Int. Conf. Acoust., Speech Signal Process.* 2017, pp. 2392–2396.
- [25] I.-Y. Jeong and K. Lee, "Learning temporal features using a deep neural network and its application to music genre classification," in *Proc. 17th Int. Conf. Music Inf. Retrieval*, 2016, pp. 434–440.