

REPRESENTATION LEARNING OF MUSIC USING ARTIST LABELS

Jiyoung Park^{1*} Jongpil Lee^{2*} Jangyeon Park¹ Jung-Woo Ha¹ Juhan Nam²

¹ NAVER Corp.

² Graduate School of Culture Technology, KAIST

{j.y.park, jangyeon.park, jungwoo.ha}@navercorp.com, {richter, juhannam}@kaist.ac.kr

ABSTRACT

In music domain, feature learning has been conducted mainly in two ways: unsupervised learning based on sparse representations or supervised learning by semantic labels such as music genre. However, finding discriminative features in an unsupervised way is challenging and supervised feature learning using semantic labels may involve noisy or expensive annotation. In this paper, we present a supervised feature learning approach using artist labels annotated in every single track as objective meta data. We propose two deep convolutional neural networks (DCNN) to learn the deep artist features. One is a plain DCNN trained with the whole artist labels simultaneously, and the other is a Siamese DCNN trained with a subset of the artist labels based on the artist identity. We apply the trained models to music classification and retrieval tasks in transfer learning settings. The results show that our approach is comparable to previous state-of-the-art methods, indicating that the proposed approach captures general music audio features as much as the models learned with semantic labels. Also, we discuss the advantages and disadvantages of the two models.

1. INTRODUCTION

Representation learning or feature learning has been actively explored in recent years as an alternative to feature engineering [1]. The data-driven approach, particularly using deep neural networks, has been applied to the area of music information retrieval (MIR) as well [14]. In this paper, we propose a novel audio feature learning method using deep convolutional neural networks and artist labels.

Early feature learning approaches are mainly based on unsupervised learning algorithms. Lee et al. used convolutional deep belief network to learn structured acoustic patterns from spectrogram [19]. They showed that the learned features achieve higher performance than Mel-Frequency Cepstral Coefficients (MFCC) in genre and artist classification. Since then, researchers have applied various

unsupervised learning algorithms such as sparse coding [12, 24, 29, 31], K-means [8, 24, 30] and restricted Boltzmann machine [24, 26]. Most of them focused on learning a meaningful dictionary on spectrogram by exploiting sparsity. While these unsupervised learning approaches are promising in that it can exploit abundant unlabeled audio data, most of them are limited to single or dual layers, which are not sufficient to represent complicated feature hierarchy in music.

On the other hand, supervised feature learning has been progressively more explored. An early approach was mapping a single frame of spectrogram to genre or mood labels via pre-trained deep neural networks and using the hidden-unit activations as audio features [11, 27]. More recently, this approach was handled in the context of transfer learning using deep convolutional neural networks (DCNN) [6, 20]. Leveraging large-scaled datasets and recent advances in deep learning, they showed that the hierarchically learned features can be effective for diverse music classification tasks. However, the semantic labels that they use such as genre, mood or other timbre descriptions tend to be noisy as they are sometimes ambiguous to annotate or tagged from the crowd. Also, high-quality annotation by music experts is known to be highly time-consuming and expensive.

Meanwhile, artist labels are the meta data annotated to songs naturally from the album release. They are objective information with no disagreement. Furthermore, considering every artist has his/her own style of music, artist labels may be regarded as terms that describe diverse styles of music. Thus, if we have a model that can discriminate different artists from music, the model can be assumed to explain various characteristics of the music.

In this paper, we verify the hypothesis using two DCNN models that are trained to identify the artist from an audio track. One is the basic DCNN model where the softmax output units corresponds to each of artist. The other is the Siamese DCNN trained with a subset of the artist labels to mitigate the excessive size of the output layer in the plain DCNN when a large-scale dataset is used. After training the two models, we regard them as a feature extractor and apply artist features to three different genre datasets in two experiment settings. First, we directly find similar songs using the artist features and K-nearest neighbors. Second, we conduct transfer learning to further adapter the features to each of the datasets. The results show that proposed approach captures useful features for unseen audio datasets

* Equally contributing authors.



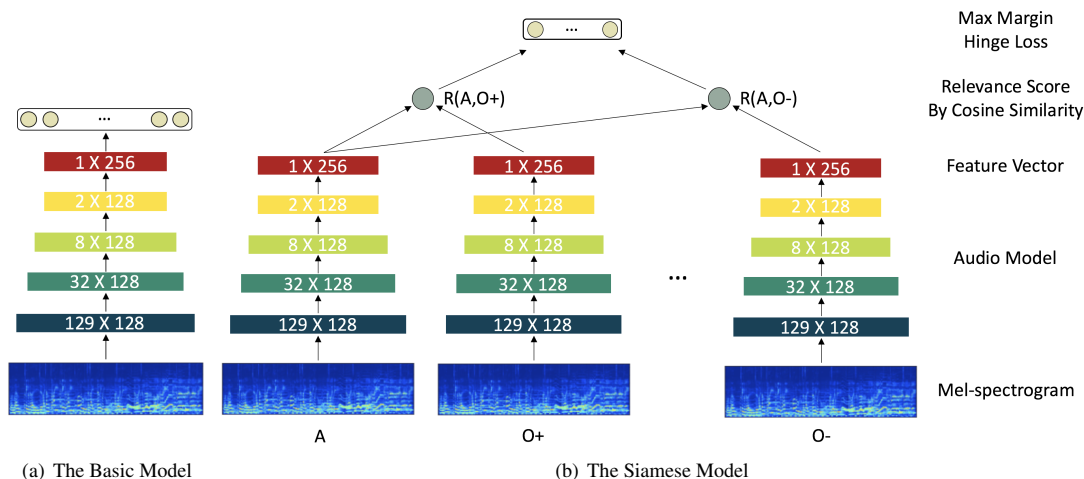


Figure 1. The proposed architectures for the model using artist labels.

and the proposed models are comparable to those trained with semantic labels in performance. In addition, we discuss the advantages and disadvantages of the two proposed DCNN models.

2. LEARNING MODELS

Figure 1 shows the two proposed DCNN models to learn audio features using artist labels. The basic model is trained as a standard classification problem. The Siamese model is trained using pair-wise similarity between an anchor artist and other artists. In this section, we describe them in detail.

2.1 Basic Model

This is a widely used 1D-CNN model for music classification [5, 9, 20, 25]. The model uses mel-spectrogram with 128 bins in the input layer. We configured the DCNN such that one-dimensional convolution layers slide over only a single temporal dimension. The model is composed of 5 convolution and max pooling layers as illustrated in Figure 1(a). Batch normalization [15] and rectified linear unit (ReLU) activation layer are used after every convolution layer. Finally, we used categorical cross entropy loss in the prediction layer.

We train the model to classify artists instead of semantic labels used in many music classification tasks. For example, if the number of artists used is 1,000, this becomes a classification problem that identifies one of the 1,000 artists. After training, the extracted 256-dimensional feature vector in the last hidden layer is used as the final audio feature learned using artist labels. Since this is the representation from which the identity is predicted by the linear softmax classifier, we can regard it as the highest-level artist feature.

2.2 Siamese Model

While the basic model is simple to train, it has two main limitations. One is that the output layer can be excessively large if the dataset has numerous artists. For example, if a

dataset has 10,000 artists and the last hidden layer size is 100, the number of parameters to learn in the last weight matrix will reach 1M. Second, whenever new artists are added to the dataset, the model must be trained again entirely. We solve the limitations using the Siamese DCNN model.

A Siamese neural network consists of twin networks that share weights and configuration. It then provides unique inputs to the network and optimizes similarity scores [3, 18, 22]. This architecture can be extended to use both positive and negative examples at one optimization step. It is set up to take three examples: anchor item (query song), positive item (relevant song to the query) and negative item (different song to the query). This model is often called *triplet networks* and has been successfully applied to music metric learning when the relative similarity scores of song triplets are available [21]. This model can be further extended to use several negative samples instead of just one negative in the triplet network. This technique is called *negative sampling* and has been popularly used in word embedding [23] and latent semantic model [13]. By using this technique, they could effectively approximate the full softmax function when the output class is extremely large (i.e. 10,000 classes).

We approximate the full softmax output in the basic model with the Siamese neural networks using negative sampling technique. Regarding the artist labels, we set up the negative sampling by treating identical artist's song to the anchor song as positive sample and other artists' songs as negative samples. This method is illustrated in Figure 1(b). Following [13], the relevance score between the anchor song feature and other song feature is measured as:

$$R(A, O) = \cos(y_A, y_O) = \frac{y_A^T y_O}{|y_A| |y_O|} \quad (1)$$

where y_A and y_O are the feature vectors of the anchor song and other song, respectively.

Meanwhile, the choice of loss function is important in this setting. We tested two loss functions. One is the softmax function with categorical cross-entropy loss to max-

imize the positive relationships. The other is the max-margin hinge loss to set only margins between positive and negative examples [10]. In our preliminary experiments, the Siamese model with negative sampling was successfully trained only with the max-margin loss function between the two objectives, which is defined as follows:

$$loss(A, O) = \sum_{O^-} \max[0, \Delta - R(A, O^+) + R(A, O^-)] \tag{2}$$

where Δ is the margin, O^+ and O^- denotes positive example and negative examples, respectively. We also grid-searched the number of negative samples and the margin, and finally set the number of negative samples to 4 and the margin value Δ to 0.4. The shared audio model used in this approach is exactly the same configuration as the basic model.

2.3 Compared Model

In order to verify the usefulness of the artist labels and the presented models, we constructed another model that has the same architecture as the basic model but using semantic tags. In this model, the output layer size corresponds to the number of the tag labels. Hereafter, we categorize all of them into *artist-label model* and *tag-label model*, and compare the performance.

3. EXPERIMENTS

In this section, we describe source datasets to train the two artist-label models and one tag-label model. We also introduce target datasets for evaluating the three models. Finally, the training details are explained.

3.1 Source Tasks

All models are trained with the Million Song Dataset (MSD) [2] along with 30-second 7digital¹ preview clips. Artist labels are naturally annotated onto every song, thus we simply used them. For the tag label, we used the Last.fm dataset augmented on MSD. This dataset contains tag annotation that matches the ID of the MSD.

3.1.1 Artist-label Model

The number of songs that belongs to each artist may be extremely skewed and this can make fair comparison among the three models difficult. Thus, we selected 20 songs for each artist evenly and filtered out the artists who have less than this. Also, we configured several sets of the artist lists to see the effect of the number of artists on the model performances (500, 1,000, 2,000, 5,000 and 10,000 artists). We then divided them into 15, 3 and 2 songs for training, validation and testing, respectively for the sets contain less than 10,000 artists. For the 10,000 artist sets, we partitioned them in 17, 1 and 2 songs because once the artists reach 10,000, the validation set already become 10,000 songs even when we only use 1 song from each artist which is already sufficient for validating the model performance.

¹ <https://www.7digital.com/>

We also should note that the testing set is actually not used in the whole experiments in this paper because we used the source dataset only for training the models to use them as feature extractors. The reason we filtered and split the data in this way is for future work².

3.1.2 Tag-label Model

We used 5,000 artists set as a baseline experiment setting. This contains total 90,000 songs in the training and validation set with a split of 75,000 and 15,000. We thus constructed the same size set for tagging dataset to compare the artist-label models and the tag-label model. The tags and songs are first filtered in the same way as the previous works [4, 20]. Among the list with the filtered top 50 used tags, we randomly selected 90,000 songs and split them into the same size as the 5,000 artist set.

3.2 Target Tasks

We used 3 different datasets for genre classification.

- GTZAN (fault-filtered version) [17, 28]: 930 songs, 10 genres. We used a “fault-filtered” version of GTZAN [17] where the dataset was divided to prevent artist repetition in training/validation/test sets.
- FMA small [7]: 8,000 songs, 8 balanced genres.
- NAVER Music³ dataset with only Korean artists: 8,000 songs, 8 balanced genres. We filtered songs with only have one genre to clarify the genre characteristic.

3.3 Training Details

For the preprocessing, we computed the spectrogram using 1024 samples for FFT with a Hanning window, 512 samples for hop size and 22050 Hz as sampling rate. We then converted it to mel-spectrogram with 128 bins along with a log magnitude compression.

We chose 3 seconds as a context window of the DCNN input after a set of experiments to find an optimal length that works well in music classification task. Out of the 30-second long audio, we randomly extracted the context size audio and put them into the networks as a single example. The input normalization was performed by dividing standard deviation after subtracting mean value across the training data.

We optimized the loss using stochastic gradient descent with 0.9 Nesterov momentum with $1e^{-6}$ learning rate decay. Dropout 0.5 is applied to the output of the last activation layer for all the models. We reduce the learning rate when a valid loss has stopped decreasing with the initial learning rate 0.015 for the basic models (both artist-label and tag-label) and 0.1 for the Siamese model. Zero-padding is applied to each convolution layer to maintain its size.

² All the data splits of the source tasks are available at the link for reproducible research <https://github.com/jiyoungpark527/msd-artist-split>.

³ <http://music.naver.com>

Our system was implemented in Python 2.7, Keras 2.1.1 and Tensorflow-gpu 1.4.0 for the back-end of Keras. We used NVIDIA Tesla M40 GPU machines for training our models. Code and models are available at the link for reproducible research⁴.

4. FEATURE EVALUATION

We apply the learned audio features to genre classification as a target task in two different approaches: feature similarity-based retrieval and transfer learning. In this section, we describe feature extraction and feature evaluation methods.

4.1 Feature Extraction Using the DCNN Models

In this work, the models are evaluated in three song-level genre classification tasks. Thus, we divided 30-second audio clip into 10 segments to match up with the model input size and the 256-dimension features from the last hidden layer are averaged into a single song-level feature vector and used for the following tasks. For the tasks that require song-to-song distances, cosine similarity is used to match up with the Siamese model's relevance score.

4.2 Feature Similarity-based Song Retrieval

We first evaluated the models using mean average precision (MAP) considering genre labels as relevant items. After obtaining a ranked list for each song based on cosine similarity, we measured the MAP as following:

$$AP = \frac{\sum_{k \in rel} precision_k}{number\ of\ relevant\ items} \quad (3)$$

$$MAP = \frac{\sum_{q=1}^Q AP(q)}{Q} \quad (4)$$

where Q is the number of queries. $precision_k$ measures the fraction of correct items among first k retrieved list.

The purpose of this experiment is to directly verify how similar feature vectors with the same genre are in the learned feature space.

4.3 Transfer Learning

We classified audio examples using the k-nearest neighbors (k-NN) classifier and linear softmax classifier. The evaluation metric for this experiment is classification accuracy. We first classified audio examples using k-NN to classify the input audio into the largest number of genres among k nearest to features from the training set. The number of k is set to 20 in this experiment. This method can be regarded as a similarity-based classification. We also classified audio using a linear softmax classifier. The purpose of this experiment is to verify how much the audio features of unseen datasets are linearly separable in the learned feature space.

⁴<https://github.com/jongpillee/ismir2018-artist>.

MAP	Artist-label Basic Model	Artist-label Siamese Model	Tag-label Model
GTZAN (fault-filtered)	0.4968	0.5510	0.5508
FMA small	0.2441	0.3203	0.3019
NAVER Korean	0.3152	0.3577	0.3576

Table 1. MAP results on feature similarity-based retrieval.

KNN	Artist-label Basic Model	Artist-label Siamese Model	Tag-label Model
GTZAN (fault-filtered)	0.6655	0.6966	0.6759
FMA small	0.5269	0.5732	0.5332
NAVER Korean	0.6671	0.6393	0.6898

Table 2. KNN similarity-based classification accuracy.

Linear Softmax	Artist-label Basic Model	Artist-label Siamese Model	Tag-label Model
GTZAN (fault-filtered)	0.6721	0.6993	0.7072
FMA small	0.5791	0.5483	0.5641
NAVER Korean	0.6696	0.6623	0.6755

Table 3. Classification accuracy of a linear softmax.

5. RESULTS AND DISCUSSION

5.1 Tag-label Model vs. Artist-label Model

We first compare the artist-label models to the tag-label model when they are trained with the same dataset size (90,000 songs). The results are shown in Table 1, 2 and 3. In feature similarity-based retrieval using MAP (Table 1), the artist-based Siamese model outperforms the rest on all target datasets. In the genre classification tasks (Table 2 and 3), Tag-label model works slightly better than the rest on some datasets and the trend becomes stronger in the classification using the linear softmax. Considering that the source task in the tag-based model (trained with the Last.fm tags) contains genre labels mainly, this result may attribute to the similarity of labels in both source and target tasks. Therefore, we can draw two conclusions from this experiment. First, the artist-label model is more effective in similarity-based tasks (1 and 2) when it is trained with the proposed Siamese networks, and thus it may be more useful for music retrieval. Second, the semantic-based model is more effective in genre or other semantic label tasks and thus it may be more useful for human-friendly music content organization.

5.2 Basic Model vs. Siamese Model

Now we focus on the comparison of the two artist-label models. From Table 1, 2 and 3, we can see that the Siamese model generally outperforms the basic model. However, the difference become attenuated in classification tasks and the Siamese model is even worse on some datasets. Among them, it is notable that the Siamese model is significantly worse than the basic model on the NAVER Music dataset




Figure 2. MAP results with regard to different number of artists in the feature models.




Figure 3. Genre classification accuracy using k-NN with regard to different number of artists in the feature models.




Figure 4. Genre classification accuracy using linear softmax with regard to different number of artists in the feature models.

in the genre classification using k-NN even though they are based on feature similarity. We dissected the result to see whether it is related to the cultural difference between the training data (MSD, mostly Western) and the target data (the NAVER set, only Korean). Figure 5 shows the detailed classification accuracy for each genre of the NAVER dataset. In three genres, ‘Trot’, ‘K-pop Ballad’ and ‘Kids’ that do not exist in the training dataset, we can see that the basic model outperforms the Siamese model whereas the results are opposite in the other genres. This indicates that the basic model is more robust to unseen genres of music. On the other hand, the Siamese model slightly over-fits to the training set, although it effectively captures the artist features.

5.3 Effect of the Number of Artists

We further analyze the artist-label models by investigating how the number of artists in training the DCNN affects the performance. Figure 2, 3 and 4 are the results that show similarity-based retrieval (MAP) and genre classification (accuracy) using k-NN and linear softmax, respectively, according to the increasing number of training artists. They show that the performance is generally proportional to the number of artists but the trends are quite different between the two models. In the similarity-based retrieval, the MAP of the Siamese model is significantly higher than that of the basic model when the number of




Figure 5. The classification results of each genre for the NAVER dataset with only Korean music.

artists is greater than 1,000. Also, as the number of artists increases, the MAP of the Siamese model consistently goes up with a slight lower speed whereas that of the basic model saturates at 2,000 or 5,000 artists. On the other hand, the performance gap changes in the two classification tasks. On the GTZAN dataset, while the basic model is better for 500 and 1,000 artists, the Siamese model reverses it for 2,000 and more artists. On the NAVER dataset, the basic model is consistently better. On the FMA small, the results are mixed in two classifiers. Again, the results may be explained by our interpretation of the models in Section 5.2. In summary, the Siamese model seems

Models	GTZAN (fault-filtered)	FMA small
2-D CNN [17]	0.6320	-
Temporal features [16]	0.6590	-
Multi-level Multi-scale [20]	0.7200	-
SVM [7]	-	0.5482 [†]
Artist-label Basic model	0.7076	0.5687
Artist-label Siamese model	0.7203	0.5673

Table 4. Comparison with previous state-of-the-art models: classification accuracy results. Linear softmax classifier is used and features are extracted from the artist-label models trained with 10,000 artists. [†] This result was obtained using the provided code and dataset in [7].

to work better in similarity-based tasks and the basic model is more robust to different genres of music. In addition, the Siamese model is more capable of being trained with a large number of artists.

5.4 Comparison with State-of-the-arts

The effectiveness of artist labels is also supported by comparison with previous state-of-the-art models in Table 4. For this result, we report two artist-label models trained with 10,000 artists using linear softmax classifier. In this table, we can see that the proposed models are comparable to the previous state-of-the-art methods.

6. VISUALIZATION

We visualize the extracted feature to provide better insight on the discriminative power of learned features using artist labels. We used the DCNN trained to classify 5,000 artists as a feature extractor. After collecting the feature vectors, we embedded them into 2-dimensional vectors using t-distributed stochastic neighbor embedding (t-SNE).

For artist visualization, we collect a subset of MSD (apart from the training data for the DCNN) from well-known artists. Figure 6 shows that artists' songs are appropriately distributed based on genre, vocal style and gender. For example, artists with similar genre of music are closely located and female pop singers are close to each other except Maria Callas who is a classical opera singer. Interestingly, some songs by Michael Jackson are close to female vocals because of his distinctive high-pitched tone.

Figure 7 shows the visualization of features extracted from the GTZAN dataset. Even though the DCNN was trained to discriminate artist labels, they are well clustered by genre. Also, we can observe that some genres such as disco, rock and hip-hop are divided into two or more groups that might belong to different sub-genres.

7. CONCLUSION AND FUTURE WORK

In this work, we presented the models to learn audio feature representation using artist labels instead of semantic labels. We compared two artist-label models and one tag-label model. The first is a basic DCNN consisting of a softmax output layer to predict which artist they belong to out of all artists used. The second is a Siamese-style architecture that maximizes the relative similarity score be-




Figure 6. Feature visualization by artist. Total 22 artists are used and, among them, 15 artists are represented in color.




Figure 7. Feature visualization by genre. Total 10 genres from the GTZAN dataset are used.

tween a small subset of the artist labels based on the artist identity. The last is a model optimized using tag labels with the same architecture as the first model. After the models are trained, we used them as feature extractors and validated the models on song retrieval and genre classification tasks on three different datasets. Three interesting results were found during the experiments. First, the artist-label models, particularly the Siamese model, is comparable to or outperform the tag-label model. This indicates that the cost-free artist-label is as effective as the expensive and possibly noisy tag-label. Second, the Siamese model showed the best performances on song retrieval task in all datasets tested. This can indicate that the pair-wise relevance score loss in the Siamese model helps the feature similarity-based search. Third, the use of a large number of artists increases the model performance. This result is also useful because the artists can be easily increased to a very large number.

As future work, we will investigate the artist-label Siamese model more thoroughly. First, we plan to investigate advanced audio model architecture and diverse loss and pair-wise relevance score functions. Second, the model can easily be re-trained using new added artists because the model does not have fixed output layer. This property will be evaluated using cross-cultural data or using extremely small data (i.e. one-shot learning [18]).

8. ACKNOWLEDGEMENT

This work was supported by Basic Science Research Program through the National Research Foundation of Korea funded by the Ministry of Science, ICT & Future Planning (2015R1C1A1A02036962) and by NAVER Corp.

9. REFERENCES

- [1] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [2] Thierry Bertin-Mahieux, Daniel PW Ellis, Brian Whitman, and Paul Lamere. The million song dataset. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, volume 2, pages 591–596, 2011.
- [3] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a “siamese” time delay neural network. In *Advances in Neural Information Processing Systems (NIPS)*, pages 737–744, 1994.
- [4] Keunwoo Choi, George Fazekas, and Mark Sandler. Automatic tagging using deep convolutional neural networks. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 805–811, 2016.
- [5] Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho. Convolutional recurrent neural networks for music classification. In *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 2392–2396, 2017.
- [6] Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho. Transfer learning for music classification and regression tasks. In *Proc. of the International Conference on Music Information Retrieval (ISMIR)*, pages 141–149, 2017.
- [7] Michaël Defferrard, Kirell Benzi, Pierre Vandergheynst, and Xavier Bresson. Fma: A dataset for music analysis. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 316–323, 2017.
- [8] Sander Dieleman and Benjamin Schrauwen. Multi-scale approaches to music audio feature learning. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 116–121, 2013.
- [9] Sander Dieleman and Benjamin Schrauwen. End-to-end learning for music audio. In *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6964–6968, 2014.
- [10] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems (NIPS)*, pages 2121–2129, 2013.
- [11] Philippe Hamel and Douglas Eck. Learning features from music audio with deep belief networks. In *Proc. of the International Conference on Music Information Retrieval (ISMIR)*, pages 339–344, 2010.
- [12] Mikael Henaff, Kevin Jarrett, Koray Kavukcuoglu, and Yann LeCun. Unsupervised learning of sparse features for scalable audio classification. In *Proc. of the International Conference on Music Information Retrieval (ISMIR)*, pages 681–686, 2011.
- [13] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. Learning deep structured semantic models for web search using clickthrough data. In *Proc. of the 22nd ACM international conference on information & knowledge management*, pages 2333–2338. ACM, 2013.
- [14] Eric Humphrey, Juan Bello, and Yann LeCun. Feature learning and deep architectures: new directions for music informatics. *Journal of Intelligent Information Systems*, 41(3):461–481, Dec 2013.
- [15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pages 448–456, 2015.
- [16] Il-Young Jeong and Kyogu Lee. Learning temporal features using a deep neural network and its application to music genre classification. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 434–440, 2016.
- [17] Corey Kereliuk, Bob L Sturm, and Jan Larsen. Deep learning and music adversaries. *IEEE Transactions on Multimedia*, 17(11):2059–2071, 2015.
- [18] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, volume 2, 2015.
- [19] Honglak Lee, Peter Pham, Yan Largman, and Andrew Y Ng. Unsupervised feature learning for audio classification using convolutional deep belief networks. In *Advances in neural information processing systems (NIPS)*, pages 1096–1104, 2009.
- [20] Jongpil Lee and Juhan Nam. Multi-level and multi-scale feature aggregation using pretrained convolutional neural networks for music auto-tagging. *IEEE Signal Processing Letters*, 24(8):1208–1212, 2017.
- [21] Rui Lu, Kailun Wu, Zhiyao Duan, and Changshui Zhang. Deep ranking: Triplet matchnet for music metric learning. In *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 121–125, 2017.

- [22] Pranay Manocha, Rohan Badlani, Anurag Kumar, Ankit Shah, Benjamin Elizalde, and Bhiksha Raj. Content-based representations of audio using siamese neural networks. In *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [23] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems (NIPS)*, pages 3111–3119, 2013.
- [24] Juhan Nam, Jorge Herrera, Malcolm Slaney, and Julius O. Smith. Learning sparse feature representations for music annotation and retrieval. In *Proc. of the International Conference on Music Information Retrieval (ISMIR)*, pages 565–570, 2012.
- [25] Jordi Pons, Thomas Lidy, and Xavier Serra. Experimenting with musically motivated convolutional neural networks. In *Proc. of the International Workshop on Content-Based Multimedia Indexing (CBMI)*, pages 1–6, 2016.
- [26] Jan Schlüter and Christian Osendorfer. Music Similarity Estimation with the Mean-Covariance Restricted Boltzmann Machine. In *Proc. of the International Conference on Machine Learning and Applications*, pages 118–123, 2011.
- [27] Erik M. Schmidt and Youngmoo E. Kim. Learning emotion-based acoustic features with deep belief networks. In *Proc. of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 65–68, 2011.
- [28] George Tzanetakis and Perry Cook. Musical genre classification of audio signals. *IEEE Transactions on speech and audio processing*, 10(5):293–302, 2002.
- [29] Yonatan Vaizman, Brian McFee, and Gert Lanckriet. Codebook-based audio feature representation for music information retrieval. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 22(10):1483–1493, 2014.
- [30] Jan Wülfing and Martin Riedmiller. Unsupervised learning of local features for music classification. In *Proc. of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 139–144, 2012.
- [31] Chin-Chia Yeh, Li Su, and Yi-Hsuan Yang. Dual-layer bag-of-frames model for music genre classification. In *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 246–250, 2013.