

Transfer Learning of Artist Group Factors to Musical Genre Classification

Kim, Jaehun; Won, Minz; Serra, Xavier; Liem, Cynthia C. S.

DOI

[10.1145/3184558.3191823](https://doi.org/10.1145/3184558.3191823)

Publication date

2018

Document Version

Final published version

Published in

WWW'18 Companion Proceedings of the The Web Conference 2018

Citation (APA)

Kim, J., Won, M., Serra, X., & Liem, C. C. S. (2018). Transfer Learning of Artist Group Factors to Musical Genre Classification. In *WWW'18 Companion Proceedings of the The Web Conference 2018* (pp. 1929-1934). International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3184558.3191823>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Transfer Learning of Artist Group Factors to Musical Genre Classification

Jaehun Kim*
Delft University of Technology
Delft, Netherlands
j.h.kim@tudelft.nl

Xavier Serra
Universitat Pompeu Fabra
Barcelona, Spain
xavier.serra@upf.edu

Minz Won†
Universitat Pompeu Fabra
Barcelona, Spain
minz.won@upf.edu

Cynthia C. S. Liem
Delft University of Technology
Delft, Netherlands
c.c.s.liem@tudelft.nl

ABSTRACT

The automated recognition of music genres from audio information is a challenging problem, as genre labels are subjective and noisy. Artist labels are less subjective and less noisy, while certain artists may relate more strongly to certain genres. At the same time, at prediction time, it is not guaranteed that artist labels are available for a given audio segment. Therefore, in this work, we propose to apply the transfer learning framework, learning artist-related information which will be used at inference time for genre classification. We consider different types of artist-related information, expressed through artist group factors, which will allow for more efficient learning and stronger robustness to potential label noise. Furthermore, we investigate how to achieve the highest validation accuracy on the given FMA dataset, by experimenting with various kinds of transfer methods, including single-task transfer, multi-task transfer and finally multi-task learning.

CCS CONCEPTS

• **Information systems** → **Music retrieval**; • **Computing methodologies** → **Multi-task learning**; **Transfer learning**; **Neural networks**;

KEYWORDS

music information retrieval; multi-task learning; transfer learning; neural network

ACM Reference Format:

Jaehun Kim, Minz Won, Xavier Serra, and Cynthia C. S. Liem. 2018. Transfer Learning of Artist Group Factors to Musical Genre Classification. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23-27, 2018, Lyon, France*, Jennifer B. Sartor, Theo D'Hondt, and Wolfgang De Meuter (Eds.). ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3184558.3191823>

*This is the corresponding author

†This research was partially conducted during the author's internship at Kakao Corp.

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion, April 23-27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3191823>

1 INTRODUCTION

Learning to Recognize Musical Genre from Audio is a challenge track of *The Web Conference 2018*. The main goal of the challenge is to predict musical genres of unknown audio segments correctly, by utilizing the FMA dataset [10] as a training set. The challenge therefore focuses on a classification task.

In machine learning, many classification tasks, such as visual object recognition, consider objective and clearly separable classes. In contrast, music genres consider subjective, human-attributed labels. These may be inter-correlated (e.g. a *rock* song may also be considered *pop*, many *classical* works are also *instrumental*) and dependent of a user's context (e.g., a *French rock* song is not *International* to a French listener). Generally, no universal genre taxonomy exists, and even the definition of 'genre' itself is problematic: what is usually understood as 'genre' in Music Information Retrieval would rather be characterized as 'style' in Musicology [17]. This makes genre classification a challenging problem. In our work, considering the given labels in the challenge, we consider a musical genre to be a category that consists of songs sharing certain aspects of musical characteristics.

Commonly, music tracks are released with explicit mentioning of titles and artists. The identity of the artist does not suffer from semantic taxonomy problems, and can thus be considered as a more objective label than the genre label. At the same time, songs from the same artist tend to share prominent musical characteristics. Considering that an artist is commonly mapped into one or multiple specific genres, but not the whole universe of possible genres, and that the other way around, sets of artists can be seen as exemplars for certain music genres, the musical characteristics that identify an artist may also be key features of certain musical genres.

Therefore, it will be beneficial to exploit artist-related information in a genre classification task. At the same time, learning a direct mapping from artist identity to genre label would not be practical. First of all, for an unknown audio segment for which a genre classification should be performed, the artist label may also not be available. Secondly, artist labels may not always be informative to a system, especially when an artist is newly introduced, so no previous history on the artist exists. Finally, an artist may have been active in multiple genres at once, but not be equally representative for all these genres. Given such constraints, we wish to employ a learning framework which only requires artist labels

at training time, but not at prediction time, and that will allow for the inclusion of newly introduced artists, for whom not much extra information is available beyond their songs.

In this work, we therefore present a multi-task transfer framework for using artist labels to improve a genre classification model. Assuming that artist labels are given for each track in the training set, these labels are used as side information, allowing a model to learn the mapping between audio and artists, while capturing patterns that might as well be useful for genre prediction.

It has been shown that music representations learned from raw artist labels can effectively transfer to other music-related tasks [21]. However, learning more than thousands of artists as individual classes is not efficient for at least two reasons:

- Due to data sparsity, only a few tracks are assigned per class;
- Despite the uniqueness of each artist, it can be beneficial to group them into clusters of similar artists, avoiding learning bottlenecks caused by large numbers of classes.

To overcome these potential problems, we therefore apply a label pre-processing step, obtaining Artist Group Factors (AGF) as learning targets, rather than individual artist identities.

Finally, we train Deep Convolutional Neural Networks (DCNNs) employing different learning setups, ranging from targeting genre and various types of AGFs with individual networks, to employing a shared architecture as introduced in multiple previous Multi-Task Learning (MTL) works [2, 3, 6, 14, 16, 18, 24, 25].

In the remainder of this paper, we first discuss an initial data exploration leading to our choice for AGFs (Section 2). Subsequently, we will give a detailed description of the proposed approach (Section 3), followed by a discussion of experimental settings (Section 4). Finally, we will present our results (Section 5), followed by a short discussion and conclusion (Section 6).

2 INITIAL DATA EXPLORATION

In the beginning of the challenge, we first explored the training data, and investigated a conventional data-driven approach using a DCNN for music genre classification, with genre labels as targets.

First of all, we had some concerns about the reliability of the genre annotations. As they were provided by users who uploaded the content, the users did not have access to a single genre taxonomy and unified annotation strategy. Thus, user-contributed annotations are expected to show more variability than annotations by experts. Furthermore, the dataset included 25,000 tracks from 5,152 unique albums. For 5,028 out of these 5,152 albums, genre annotations were made at the album level. While all tracks in an album can belong to a single genre, this is not always true. Indeed, we could discover examples of the case in which different tracks on the same album would belong to different genres, as well as multiple misannotations. Given these reliability issues, it is not guaranteed that by targeting these annotations only, generalized model performance for genre classification can be achieved.

To this end, while we will consider performance for direct (main top-)genre labels as targets (which we will denote as learning task category g in the remainder of this paper), in order to obtain more generalizable results obtained on more objective and consistent labeling data, we propose a multi-task transfer framework, introducing an Artist Group (AG) prediction task targeting AGFs.

3 METHODOLOGY

3.1 Artist Group Factors

The main idea of extracting AGFs is to cluster artists based on meaningful feature sets that allow for aggregation at (and beyond) the artist level. For instance, one can collect genre labels from songs belonging to each artist, and then construct a Bag-of-Word (BoW) artist-level feature vector. Each dimension of the vector represents a genre, with the magnitude of the vector indicating genre frequency among a song collection. Alternatively, a BoW feature vector can be constructed by counting latent ‘terms’ belonging to each artist, which can be obtained by a dictionary learned from song-level or frame-level features through K-means clustering [20] or the Sparse Coding [9] method.

Once artist-level BoW feature vectors are constructed, standard clustering methods such as K-Means, or more sophisticated topic modeling algorithms such as Latent Dirichlet Allocation (LDA) [4] can be applied to find a small number of latent groups of artists: the AGFs for this particular feature set. This 2-step cascading pipeline is illustrated in Figure 1.

In this work, we exploit four feature sets, which reflect different levels of musical and acoustical aspects of songs. From these feature sets, we obtain artist-level BoW vectors. Subsequently, LDA is applied to transform artist-level BoW vectors into dedicated AGF representations for the particular feature set. We will both consider these artist group prediction tasks and the main genre classification task within our learning framework: an overview summary is given in Table 1.

3.1.1 MFCCs. Mel-Frequency Cepstral Coefficients (MFCCs), which are known to be efficient low-level descriptors for timbre analysis, were used as features of the artist grouping. The coefficients are initially calculated for short-time audio frames. Considering the coefficients over all audio frames of tracks for all artists, we build an universal dictionary of features using K-Means clustering. AGFs resulting from this feature set will belong to learning task category m .

3.1.2 dMFCCs. Along with MFCCs, we also use time-deltas of MFCCs (first-order differences between subsequent frames), to consider the temporal dynamics of the timbre for the artist grouping. AGFs resulting from this feature set will be denoted by d .

3.1.3 Essentia. We use song-level feature vectors from Essentia [5], which is a music feature extraction library. It extracts descriptors ranging from low-level features, such as statistics of spectral characteristics, to high-level features, including danceability [12] or semantic features learned from the data. After filtering descriptor entries that include missing values or errors, we obtained a 4374-dimensional feature vector per track. Before training a dictionary, we apply quantile normalization: a rank-based normalization process that transforms the distribution of the given features to follow a target distribution [1], which we set to be a normal distribution in this case. AGFs resulting from this feature set will belong to learning task category e .

3.1.4 Subgenres. We also use the 150 genre labels, including sub-genres, as a pre-defined dictionary for semantic description. For these, we directly build artist-level BoW vectors by aggregating

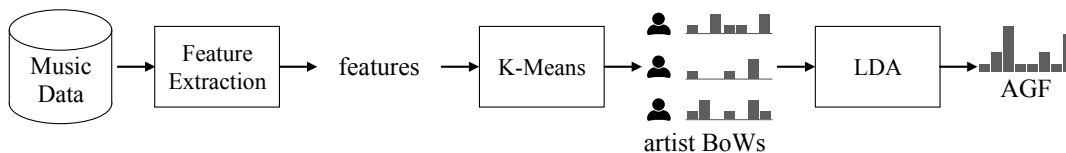


Figure 1: Artist group factor extraction pipeline.

Table 1: Details of Learning Targets

id	Category	Source	Dictionary	Dimension
g	Main	Genre	N / A	16
m		MFCC		25
d	AGF	dMFCC	K-means	25
e		Essentia [5]		4374
s		Subgenre	N / A	150

Table 2: Proposed DCNN structure

Layers	Output shape
Input layer	$128 \times 43 \times 1$
Conv 5×5 , ELU	$128 \times 43 \times 1$
MaxPooling 2×1	$64 \times 43 \times 16$
Conv 3×3 , BN, ELU	$64 \times 43 \times 32$
MaxPooling 2×2	$32 \times 21 \times 32$
Dropout (0.1)	$32 \times 21 \times 32$
Conv 3×3 , ELU	$32 \times 21 \times 64$
MaxPooling 2×2	$16 \times 10 \times 64$
Conv 3×3 , BN, ELU	$16 \times 10 \times 64$
MaxPooling 2×2	$8 \times 5 \times 64$
Dropout (0.1)	$8 \times 5 \times 64$
Conv 3×3 , ELU	$8 \times 5 \times 128$
MaxPooling 2×2	$4 \times 2 \times 128$
Conv 3×3 , ELU	$4 \times 2 \times 256$
Conv 1×1 , BN, ELU	$4 \times 2 \times 256$
GlobalAveragePooling, BN	256
Dense, BN, ELU	256
Dropout (0.5)	256
Output layer 16 or 40	16 or 40

all the genre labels from tracks by an artist. AGFs resulting from this feature set will belong to learning task category s.

3.2 Network Architectures

The architecture of the proposed system can be divided into two parts, as shown in Figure 2. We first train multiple DCNNs, targeting the various categories of learning targets (genres or various AGFs). Subsequently, transfer takes place: a multilayer perceptron (MLP) for the final genre classification is trained, utilizing features that were derived from the previously trained DCNNs.

3.2.1 DCNN. We adapted DCNN models to obtain transferable features for genre classification (Table 2). The input size of the input layer is 128×43 , which is the size of a spectrogram with 128 mel bins and 43 samples (1 second of audio). After the input layer, there are

seven convolutional layers followed by a max-pooling layer, except for the last two layers. The first convolutional layer has 5×5 kernels and the last convolutional layer has 1×1 kernels. Except for those two layers, all convolutional layers have 3×3 kernels. Outputs of the last convolutional layer are subsampled by global-average-pooling. Finally, they are connected to two dense layers for predicting AGF clusters or genres. Batch normalization [13] and dropouts [22] are sparsely used to prevent overfitting. Exponential Linear Unit (ELU) [8] is used as an activation function for the convolutional layers and Softmax is used for the output layer.

3.2.2 Shared Architecture. Considering that lower layers of DCNNs usually capture lower-level features such as edges from images or spectrograms, we hypothesized that sharing lower layers among the various DCNNs can be effective under the scenario where multiple learning sources are available. With this approach, one can expect that it not only ensures sufficient specialization on task-specific upper layers, but also benefits from regularization effects on lower layers[14]. Joint learning of multiple tasks with shared layers can prevent the shared layer to overfit for a specific task, instead learning underlying factors that have commonalities required across tasks [6, 19].

Throughout the experiment, we used the shared architecture that shares only the first convolutional block. It consists of the first convolutional and the max-pooling layer. For brevity, for the remainder of the paper, we use Single-Task Nets (STNs) and an Multi-Task Net (MTN) to refer to the non-shared networks and shared networks respectively.

3.2.3 Transfer method. The proposed system learns and predicts a genre of an input spectrogram by transferring pre-trained features from Section 3.2.1. We trained an MLP with a single hidden layer; the size of the hidden layer was 1024. ELU non-linearity was used for the hidden layer and Softmax was used for the output layer. Dropouts of 50% were applied for the input layer and a hidden layer.

Note that for both the feature learning phase and the transfer learning phase, we keep using a segment-wise learning approach. Only at the final inference step, we aggregate all the segment-level predictions, by taking the average of each segment's predicted probability for the genres.

3.2.4 Training. At training time, we iteratively update the model parameters with the mini-batch stochastic gradient descent method using the Adam algorithm [15]. For data augmentation, we randomly crop 1-second excerpts from the entire track included in the mini-batch. We use 64 samples per batch and set the learning rate to 0.001 across the experiments.

For comparison between methods, experiments are run with a fixed number of epochs. We set 1000 epochs for an MTN and

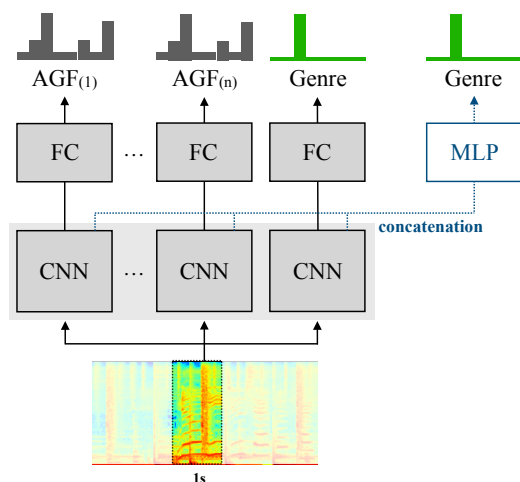


Figure 2: Illustration for the transfer learning scenario. Dotted lines indicate the setup for the multilayer perceptron for performing final genre classification.

200 for STNs. Since we took a similar stochastic update algorithm to [18] for the shared architecture, for the number of updates for task-specific layers in a shared network, the number of epochs used for training non-shared networks should be multiplied with the number of involved learning tasks. For the transfer learning phase, we also set the number of epochs to train the MLP to 50.

3.3 Pre-processing

We use mel spectrograms as the input representation for the neural networks. We extract 128-dimensional mel spectra for audio frames of 46ms, with 50% overlap with adjacent frames. To enhance lower-intensity levels of input mel spectrograms at higher frequencies, we take dB-scale log amplitudes of each mel spectrum.

3.4 Implementation Details

The experiments were run on GPU-accelerated hardware and software environments. We used Lasagne [11], Theano [23] and Keras [7] as main experimental frameworks¹. We used a number of different GPUs, including NVIDIA GRID-K2, NVIDIA GTX 1070, NVIDIA TITAN X.

4 EXPERIMENTS

To investigate the effectiveness of various types of AGFs for transfer learning, we trained all 31 possible combinations of given learning tasks, including AGFs (m, d, e, s) and main top-genre labels (g). For each run, to investigate the optimal feature architecture, we tested both shared networks and separate networks for each learning task. This leads to a total number of 62 cases, including all the combinations of learning tasks per network architecture.

However, in all cases in which multiple tasks are considered, the networks have a larger number of parameters compared to the case

¹The main code for the experiment can be found in <https://github.com/eldrin/Lasagne-MultiTaskLearning>

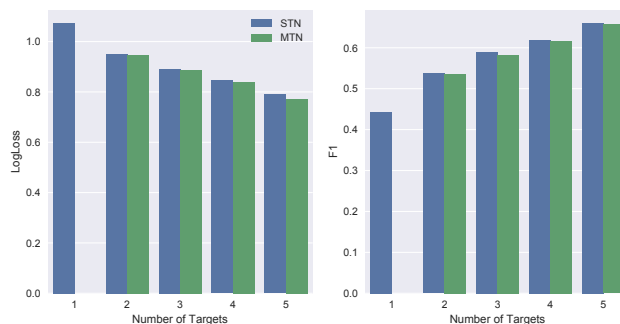


Figure 3: Average performance for the number of tasks involved in feature learning

Table 3: Comparison of the average performance with or without the main task

	LogLoss		F1	
	STN	MTN	STN	MTN
without g	1.0079	0.9618	0.4932	0.5168
with g	0.8540	0.8486	0.6154	0.6155

in which a network focuses on a single task. With a subsequent experiment, we therefore tried to verify the effect of more parameters and larger networks vs. the effect of using more tasks. To this end, we train wide Single Task Networks (wSTNs), targeting only genre, but having an equal number of parameters to the MTNs/STNs targeting multiple tasks. Finally, with respect to the number of tasks involved, we compare the best performance of MTNs/STNs to the performance of wSTNs with the same number of parameters.

As for the AGFs using song-level or frame-level features, we trained K-means algorithms employing 2048 clusters. We observed that lower numbers of clusters (e.g. 1024) can cause artists with few tracks to get a zero vector as artist-level BoW representation, due to data sparsity. Throughout the experiments, we used a fixed number of latent artist groups, set to 40.

Finally, for the internal evaluation, we divided the given training dataset employing a stratified random 85/15 split.

5 RESULTS

5.1 Multiple Learning Tasks in STN vs. MTN

In general, we observe that the number of learning tasks has a positive effect on both performance metrics. As shown in Table 3, it also is found that cases in which the main top-genre classification are included yield better results in comparison to other combinations of tasks.

Considering STN vs. MTN, on the log loss metric, MTN shows better results, but in the case of the f1-measure, the opposite is shown. Generally, considering the number of learning tasks and absolute magnitude of differences, the difference observed between the two methods cannot be deemed significant; more experiments with additional datasets and multiple splits would be needed to

Table 4: The performance of various combinations of AGFs and the top-level main genre target as a feature learning task.

	STN		MTN	
	LogLoss	F1	LogLoss	F1
g	0.8891	0.5963		
m	1.1812	0.3581		
d	1.0987	0.3967	N/A	N/A
e	1.2542	0.3437		
s	0.9404	0.5218		
gs	0.8606	0.6114	0.8578	0.6190
ge	0.8811	0.5953	0.8792	0.5996
gd	0.8845	0.5898	0.8803	0.5955
gm	0.8874	0.5957	0.8813	0.6037
se	0.9124	0.5537	0.9079	0.5502
sd	0.9191	0.5601	0.9146	0.5412
sm	0.9260	0.5581	0.9283	0.5458
ed	1.0557	0.4433	1.0422	0.4399
em	1.1186	0.4244	1.1060	0.4376
dm	1.0583	0.4373	1.0704	0.4280
gse	0.8361	0.6255	0.8335	0.6277
gsd	0.8579	0.6280	0.8519	0.6150
gsm	0.8486	0.6289	0.8541	0.6153
ged	0.8528	0.6051	0.8601	0.6067
gem	0.8645	0.5988	0.8701	0.6056
gdm	0.8773	0.5985	0.8845	0.5941
sed	0.8965	0.5818	0.8867	0.5640
sem	0.9104	0.5834	0.8889	0.5668
sdm	0.9211	0.5629	0.9109	0.5572
edm	1.0359	0.4879	1.0365	0.4675
gsed	0.8211	0.6343	0.8132	0.6328
gsem	0.8264	0.6352	0.8172	0.6284
gsdm	0.8407	0.6379	0.8288	0.6170
gedm	0.8466	0.6053	0.8450	0.6152
sedm	0.8906	0.5856	0.8875	0.5870
gsedm	0.7894	0.6599	0.7727	0.6571

assess whether statistically significant differences between STN vs. MTN approaches can be obtained.

For both STN and MTN, the best performance we achieved uses all the learning tasks, as shown in the last row of Table 4.

5.2 Networks for Multiple Learning Tasks vs. Large Network on a Single Task

We also compared the performance between the best STNs and MTNs for a given number of learning tasks, versus the performance of a wSTN that has equal model capability to these multi-task setups in terms of parameters and architecture, but only is trained on direct main top-genre classification. The corresponding results are shown in Table 5. It can be seen that MTN representations yield better performance on the log loss metric when all 5 learning tasks (all AGFs and the main top-genre) are used, although at the same time, wSTN performs better when considering the f1-measure for

Table 5: Comparison between wSTN (single genre classification task) and STN/MTN setups (multiple tasks) learning setups. The reported performances of STN and MTN consider the task combinations for which the best performance was obtained, given the mentioned number N of tasks.

N	LogLoss			F1		
	wSTN	STN	MTN	wSTN	STN	MTN
2	0.8688	0.8606	0.8578	0.6071	0.6114	0.6190
3	0.8546	0.8361	0.8335	0.6629	0.6289	0.6277
4	0.8278	0.8211	0.8132	0.6451	0.6352	0.6328
5	0.8290	0.7893	0.7727	0.6528	0.6599	0.6571

the case in which 2 learning tasks are used. In other cases, differences between the setups appear marginal; further experiments would be needed to assess whether STNs/MTNs will give significant performance boosts in case a larger set of tasks would be considered.

6 DISCUSSION & CONCLUSION

In this work, we proposed including several categories of low-rank AGFs, expressing artist-level information, into the task of classifying music genre based on musical audio. Our experimental results support the hypothesis that by targeting different categories of AGFs, deep networks can learn features from musical audio that can meaningfully support genre classification. The inclusion of multiple parallel learning tasks considering different AGF categories, and the inclusion of both genre- and AGF-based tasks in a multi-task setup, also both seem beneficial, although further work will need to be done to assess whether observed effects are truly significant. For this, other datasets will have to be included for training and testing; furthermore, alternative cluster algorithms and clustering parameters should be investigated to achieve the most robust AGF-based features.

ACKNOWLEDGMENTS

This work was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative. And this work is partially supported by the Maria de Maeztu Programme (MDM-2015-0502). We further acknowledge the computing support of Kakao Corporation.

REFERENCES

- [1] Dhammika Amarantunga and Javier Cabrera. 2001. Analysis of Data From Viral DNA Microchips. *J. Amer. Statist. Assoc.* 96, 456 (2001), 1161–1170. <https://doi.org/10.1198/016214501753381814> arXiv:<https://doi.org/10.1198/016214501753381814>
- [2] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35, 8 (2013), 1798–1828.
- [3] Joachim Bingel and Anders Søgaard. 2017. Identifying beneficial task relations for multi-task learning in deep neural networks. *arXiv preprint arXiv:1702.08303* (2017).
- [4] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.
- [5] Dmitry Bogdanov, Nicolas Wack, Emilia Gómez Gutiérrez, Sankalp Gulati, Perfecto Herrera Boyer, Oscar Mayor, Gerard Roma Trepas, Justin Salamon, José Ricardo Zapata González, and Xavier Serra. 2013. Essentia: An audio analysis library for music information retrieval. In *Britto A, Gouyon F, Dixon S, editors. 14th Conference of the International Society for Music Information Retrieval (ISMIR)*;

- 2013 Nov 4-8; Curitiba, Brazil.[place unknown]: ISMIR; 2013. p. 493-8. International Society for Music Information Retrieval (ISMIR).
- [6] Rich Caruana. 1998. Multitask learning. In *Learning to learn*. Springer, 95–133.
- [7] François Chollet. 2015. keras. <https://github.com/fchollet/keras>. (2015).
- [8] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2015. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289* (2015).
- [9] Adam Coates and Andrew Y Ng. 2011. The Importance of Encoding Versus Training with Sparse Coding and Vector Quantization. (2011).
- [10] Michaël Defferrard, Kirell Benzi, Pierre Vandergheynst, and Xavier Bresson. 2017. FMA: A Dataset for Music Analysis. In *18th International Society for Music Information Retrieval Conference*.
- [11] Sander Dieleman, Jan Schlüter, Colin Raffel, Eben Olson, Søren Kaae Sønderby, Daniel Nouri, Daniel Maturana, Martin Thoma, Eric Battenberg, Jack Kelly, Jeffrey De Fauw, Michael Heilman, diogo149, Brian McFee, Hendrik Weideman, takacs84, peterderivaz, Jon, instagibbs, Dr. Kashif Rasul, CongLiu, Britefury, and Jonas Degraeve. 2015. Lasagne: First release. (Aug. 2015). <https://doi.org/10.5281/zenodo.27878>
- [12] Perfecto Herrera and Sebastian Streich. 2005. Detrended Fluctuation Analysis of Music Signals: Danceability Estimation and further Semantic Characterization. In *Audio Engineering Society Convention 118*. Audio Engineering Society.
- [13] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*. 448–456.
- [14] Jaehun Kim, Julián Urbano, Cynthia Liem, and Alan Hanjalic. 2018. One Deep Music Representation to Rule Them All?: A comparative analysis of different representation learning strategies. *arXiv preprint arXiv:1802.04051* (2018).
- [15] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16] Sijin Li, Zhi-Qiang Liu, and Antoni B Chan. 2014. Heterogeneous multi-task learning for human pose estimation with deep convolutional neural network. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 482–489.
- [17] Cynthia C. S. Liem, Andreas Rauber, Thomas Lidy, Richard Lewis, Christopher Raphael, Joshua D. Reiss, Tim Crawford, and Alan Hanjalic. 2012. Music Information Technology and Professional Stakeholder Audiences: Mind the Adoption Gap. In *Multimodal Music Processing*, Meinard Müller, Masataka Goto, and Markus Schedl (Eds.). Dagstuhl Follow-Ups, Vol. 3. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 227–246.
- [18] Wu Liu, Tao Mei, Yongdong Zhang, Cherry Che, and Jiebo Luo. 2015. Multi-task deep visual-semantic embedding for video thumbnail selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3707–3715.
- [19] Xiaodong Liu, Jianfeng Gao, Xiaodong He, Li Deng, Kevin Duh, and Ye-Yi Wang. 2015. Representation learning using multi-task deep neural networks for semantic classification and information retrieval. (2015).
- [20] Stuart Lloyd. 1982. Least squares quantization in PCM. *IEEE transactions on information theory* 28, 2 (1982), 129–137.
- [21] Jiyoun Park, Jongpil Lee, Jangyeon Park, Jung-Woo Ha, and Juhan Nam. 2017. Representation Learning of Music Using Artist Labels. *arXiv preprint arXiv:1710.06648* (2017).
- [22] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [23] Theano Development Team. 2016. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints* abs/1605.02688 (May 2016). <http://arxiv.org/abs/1605.02688>
- [24] Wenlu Zhang, Rongjian Li, Tao Zeng, Qian Sun, Sudhir Kumar, Jieping Ye, and Shuiwang Ji. 2016. Deep model based transfer and multi-task learning for biological image analysis. *IEEE Transactions on Big Data* (2016).
- [25] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2014. Facial landmark detection by deep multi-task learning. In *European Conference on Computer Vision*. Springer, 94–108.