

# How Low Can You Go? Reducing Frequency and Time Resolution in Current CNN Architectures for Music Auto-tagging

Andres Ferraro, Dmitry Bogdanov, Xavier Serra  
*Music Technology Group*  
*Universitat Pompeu Fabra*  
Barcelona, Spain  
andres.ferraro@upf.edu

Jay Ho Jeon, Jason Yoon  
*Kakao Corp.*  
Seoul, Republic of Korea  
jason.yoon@kakaocorp.com

**Abstract**—Automatic tagging of music is an important research topic in Music Information Retrieval and audio analysis algorithms proposed for this task have achieved improvements with advances in deep learning. In particular, many state-of-the-art systems use Convolutional Neural Networks and operate on mel-spectrogram representations of the audio. In this paper, we compare commonly used mel-spectrogram representations and evaluate model performances that can be achieved by reducing the input size in terms of both lesser amount of frequency bands and larger frame rates. We use the MagnaTagaTune dataset for comprehensive performance comparisons and then compare selected configurations on the larger Million Song Dataset. The results of this study can serve researchers and practitioners in their trade-off decision between accuracy of the models, data storage size and training and inference times.

**Index Terms**—music auto-tagging, audio classification, convolutional neural networks

## I. INTRODUCTION

Current state-of-the-art systems for music auto-tagging using audio are based on deep learning, in particular convolutional neural networks (CNNs), following two different approaches, one directly using the audio as input (end-to-end models) [1] and the other using the spectrograms as input [2], [3]. Previous works [4] suggest that two approaches can have a comparative performance when they are applied on large datasets.

We can distinguish two architectures for the spectrogram-based CNN solutions, depending on whether they use multiple convolutional layers of small filters [5], [6] or if they use multiple filter shapes [4], [7], [8]. The former is borrowed from the computer vision field (VGG [9]) and gives a good performance without prior domain knowledge, while the latter is based on such a knowledge and employs filters designed to capture information relevant for music auto-tagging such as timbre or rhythm. Commonly mel-spectrograms are used with such architectures although constant-Q [3], [10] and raw short-time Fourier transform (STFT) [11] can be also applied.

In this paper, we compare the performance of two state-of-the-art CNN approaches commonly used for music auto-tagging [4], [6], using different mel-spectrogram represen-

tations as an input. Our research question is whether it is possible for these approaches to operate on reduced data inputs without a significant drop in a prediction accuracy.

To this end, we consider reducing the size of the input spectrograms in terms of both lesser amount of frequency bands and larger frame rates. We show that by reducing the frequency and time resolution we can train the network faster with a small decrease in the performance. The results of this study can help to build faster CNN models as well as reduce the amount of data to be stored and transferred optimizing resources when handling large collections of music.

## II. RELATED WORK

In image processing, there are studies that consider simplifications of CNN architectures by means of reducing network width, depth and input resolution [12]. However, only few previous studies compared different spectrogram representations for CNN architectures in Music Information Retrieval (MIR). Instead, it is common to focus on tuning model hyper-parameters with a fixed chosen input. The choice of the spectrogram input is done empirically and often follows approaches previously reported in literature. Very few information comparing different inputs is available as the authors tend to only report the most successful approaches. Also, as the existing studies on music auto-tagging focus on optimizing accuracy metrics, there is a lack of works that intend to simplify networks and their inputs for computational efficiency and consider practical aspects of the efficient ways to store spectrogram representations.

To the best of our knowledge, there is no systematic comparison of mel-spectrogram representations. The only work we are aware of in this direction has been done by Choi et al. [11], where the authors compare model performances under different pre-processing strategies such as scaling, log-compression, and frequency weighting. The same authors provide an overview of different inputs that can be used for the auto-tagging task in [3]. In relation to mel-spectrograms, they suggest that one can optimize the input to the network by changing some of the signal processing parameters such as

sampling rate, window size, hop size or mel bins resolution. These optimizations can help to minimize data size and train the networks more efficiently, however, no quantitative evaluations are provided.

### III. DATASETS

Researchers in music auto-tagging commonly use the MagnaTagATune dataset [13] to evaluate multiple settings and then repeat some settings on Million Song Dataset [14] to validate differences in performances on a larger scale [2], [4], [6]. It is important to note that both datasets contain unbalanced and noisy and/or weakly-labeled annotations [15] and therefore are challenging to work with, as the reliability of conducted evaluations may be affected [16]. Still, these are the two mostly used datasets for benchmarking due to the availability of audio.

#### A. MagnaTagATune (MTAT)

MagnaTagATune dataset contains multi-label annotations of genre, mood and instrumentation for 25,877 audio segments. Each segment is 30 seconds long, and the dataset contains multiple segments per song. All the audio is in MP3 format with 32 Kbps bitrate and 16 KHz sample rate. The dataset is split into 16 folders, and researchers commonly use the first 12 folders for training, the 13th for validation and the last three for testing. Also, only 50 most frequent tags are typically used for evaluation. These tags include genre and instrumentation labels, as well as eras (e.g., '80s' and '90s') and moods.

#### B. Million Song Dataset (MSD)

The MSD [14] is a large dataset of audio features, expanded by the MIR community with additional information including tags, lyrics and other annotations. It also contains a subset mapped by researchers to 30 seconds audio previews available at 7digital and collaborative tags from Lastfm. This subset contains 241,904 annotated track fragments and is commonly used as another larger scale benchmark for music auto-tagging systems. The tags cover genre, instrumentation, moods and eras. Audio fragments vary in their quality, encoded as MP3 with a bitrate ranging from 64 to 128 Kbps and the sample rates of 22 KHz or 44 KHz.

### IV. BASELINE ARCHITECTURES

In this work, we reproduce two CNN architectures applying them on mel-spectrograms with reduced frequency and time resolution. These architectures are among the best performing according to the existing evaluations on the MTAT and MSD datasets:

- **VGG applied for music (VGG-CNN)** [6]. This architecture contains multiple layers of small-size 2D-filters as it has been adapted from the computer vision field [9]. It is a fully-convolutional network consisting of four convolutional layers with small  $3 \times 3$  filters<sup>1</sup> and max pooling (MP) settings presented in Table I. The network operates on 96-bands mel-spectrograms for 29.1s audio

segments, 12 KHz sample rate, 512 samples frame size and the hop size of 256 samples.

- **Musically-motivated CNN (MUSICNN)** [4]. The architecture contains more filters of different shapes designed with an intention to capture musically relevant information such as timbre ( $38 \times 1$ ,  $38 \times 3$ ,  $38 \times 7$ ,  $86 \times 1$ ,  $86 \times 3$ ,  $86 \times 7$ ) and temporal patterns ( $1 \times 32$ ,  $1 \times 64$ ,  $1 \times 128$ ,  $1 \times 165$ ) in the first layer. The convolution results are concatenated and passed to three additional convolutional layers including residual connections.<sup>2</sup> Original network operates on 96-bands mel-spectrograms computed on smaller 15s audio segments with 16 KHz sample rate, 512 samples frame size and 256 samples hop size.<sup>3</sup> It then averages tag activation scores across multiple segments of the same audio input.

For evaluation on MTAT and MSD, we use batch normalization, Adam [17] as optimization method with a learning rate of 0.001 and binary cross-entropy as loss function for both architectures following their authors.

TABLE I: The baseline VGG CNN model architecture.

Input	Mel-spectrogram ( $96 \times 1366 \times 1$ )
Layer 1	Conv $3 \times 3 \times 128$ MP (2, 4) (output: $48 \times 341 \times 128$ )
Layer 2	Conv $3 \times 3 \times 384$ MP (4, 5) (output: $24 \times 85 \times 384$ )
Layer 3	Conv $3 \times 3 \times 768$ MP (3, 8) (output: $12 \times 21 \times 768$ )
Layer 4	Conv $3 \times 3 \times 2048$ MP (4, 8) (output: $1 \times 1 \times 2048$ )
Output	$50 \times 1$ (sigmoid)

### V. MEL-SPECTROGRAMS

We computed mel-spectrograms using typical setting for the MTAT dataset in the state of the art [4], [6]. The most common settings are 12 KHz or 16 KHz sample rate, frame and hop size of 512 and 256 samples, respectively, and Hann window function. Commonly, 96 or 128 mel bands are used, covering all frequency range below Nyquist (6 KHz and 8 KHz, respectively) and computed using Slaney's mel scale implementation [18]. To normalize the mel-spectrograms we considered two log-compression alternatives denominated as "dB" for  $10 \cdot \log_{10}(x)$  [11] and "log" for  $\log(1 + 10000 \cdot x)$  [2].

Starting with these settings, we then considered different variations in frequency and time resolutions (smaller number of mel bands and larger hop sizes). Table II shows all different spectrogram configurations that we evaluated on the MTAT dataset. Each configuration results in a different dimension of the resulting feature matrix (the number of mel bands  $\times$  the number of frames). An audio segment of 29.1 seconds corresponds to 1366 and 1820 frames in the case of no temporal reduction ( $\times 1$ ) and the 12 KHz and 16kHz sample rate,

<sup>2</sup>We refer the readers to the original paper for all architecture details.

<sup>3</sup>Frame and hop size settings are confirmed in personal communication with the author.

<sup>1</sup>Number of mel bands  $\times$  number of frames.

TABLE II: Mel-spectrograms configurations evaluated on the MTAT dataset. Hop sizes are reported relative to the reference hop size of 256 samples (e.g.,  $\times 5$  stands for a 5 times longer hop size).

sample rate	# mel	hop size	log type
12 KHz	128	$\times 1, \times 2, \times 3, \times 4 \times 5, \times 10$	log, dB
12 KHz	96	$\times 1, \times 2, \times 3, \times 4 \times 5, \times 10$	log, dB
12 KHz	48	$\times 1, \times 2, \times 3, \times 4 \times 5, \times 10$	log, dB
12 KHz	32	$\times 1$	log, dB
12 KHz	24	$\times 1$	log, dB
12 KHz	16	$\times 1$	log, dB
12 KHz	8	$\times 1$	log, dB
16 KHz	128	$\times 1, \times 2, \times 3, \times 4 \times 5, \times 10$	log, dB
16 KHz	96	$\times 1, \times 2, \times 3, \times 4 \times 5, \times 10$	log, dB
16 KHz	48	$\times 1, \times 2, \times 3, \times 4 \times 5, \times 10$	log, dB
16 KHz	32	$\times 1$	log, dB
16 KHz	24	$\times 1$	log, dB
16 KHz	16	$\times 1$	log, dB
16 KHz	8	$\times 1$	log, dB

respectively. In turn, the maximum reduction we considered ( $\times 10$ ) results in 137 and 182 frames.

All spectrograms were computed using Essentia<sup>4</sup> music audio analysis library [19]. It was configured to reproduce mel-spectrograms from another analysis library used by the state of the art, LibROSA,<sup>5</sup> for compatibility. As a matter of interest, to have a better understanding of what information these spectrograms are able to capture, we provide a number of examples sonifying the resulting mel-spectrograms for all considered frequency and time resolutions online.<sup>6</sup>

## VI. BASELINE ARCHITECTURE ADJUSTMENTS

In this section we explain the changes introduced to the original model architectures presented in Section IV.

### A. VGG-CNN

We try to preserve the original architecture defined in [6] in terms of the size and number of filters in each layer, but we need to adjust max pooling settings since we are reducing the dimensions of the mel-spectrogram input. We report all such modifications for the VGG-CNN architecture in Table III. It reports the sizes of square max-pooling windows in each layer selected accordingly to the number of mel bands and the hop size. We prioritize changes in max pooling in the latter layers when possible. We adjust the pooling size to match the input dimensions when possible, otherwise padding is applied. In the case of 16 KHz sample rate, more adjustments to VGG-CNN are necessary because, having a fixed reference hop size of 256 samples, the higher sample rate implies better temporal resolution and the larger mel-spectrograms (1820 frames).

It is important to note that if we change the resolution of the input, the  $3 \times 3$  filters in VGG-CNN capture different ranges of frequency and temporal information. For example, they cover twice the mel-frequency range and a doubled time interval

when using 48 mel bands and  $\times 2$  hop size. This can be an advantage, because it reduces the amount of information that the network needs to learn.

TABLE III: Adjusted sizes for max-pooling windows (time and frequency) in the four consecutive layers of the VGG CNN model with respect to the hop size, sample rate and the number of mel bands. The original sizes are highlighted in bold.

hop size	max-pooling size (time)	
	12 KHz	16 KHz
$\times 1$	<b>4, 5, 8, 8</b>	4, 5, 9, 10
$\times 2$	4, 5, 8, 4	4, 5, 9, 5
$\times 3$	4, 5, 8, 2	4, 5, 9, 3
$\times 4$	4, 5, 8, 2	4, 5, 9, 2
$\times 5$	4, 5, 8, 1	4, 5, 9, 2
$\times 10$	4, 5, 4, 1	4, 5, 9, 1
# mel	max-pooling size (frequency)	
128	2, 4, 4, 4	
96	<b>2, 4, 3, 4</b>	
48	2, 4, 3, 2	
32	2, 2, 3, 2	
24	2, 2, 3, 2	
16	2, 2, 2, 2	
8	2, 2, 2, 1	

### B. MUSICNN

In the original model, timbre filters' sizes in frequency are computed relative to the number of mel bands (90% and 40%). We preserve the same relation when we change this number. In our implementations we modified the segment size to 3 seconds, as we obtained slightly better results in our preliminary evaluation.<sup>7</sup> We keep the temporal dimension of the filters (the number of frames) intact for all considered mel-spectrograms settings.

## VII. EVALUATION METRICS

CNN models for auto-tagging output continuous activation values within  $[0, 1]$  for each tag, and therefore we can study the performance of binary classifications under different activation thresholds. To this end, following previous works [4], [6], [20] we use Receiver Operating Characteristic Area Under Curve (ROC AUC) averaged across tags as our performance metric. We also report Precision-Recall Area Under Curve (PR AUC), because previous studies [21] have shown that ROC AUC can give over-optimistic scores when the data is unbalanced, which is our case. Both ROC AUC and PR AUC are single value measures characterizing the overall performance, which allows to easily compare multiple systems.

To measure the computational cost of models' training and inference we use an estimate of the number of multiply-accumulate operations required by a network to process one batch (1 GMAC is equal to 1 Giga MAC operations). This metric is related to the time a model requires for training and inference. We use an online tool<sup>8</sup> to compute approximate MAC values for our architectures.

<sup>4</sup><https://essentia.upf.edu>

<sup>5</sup><https://librosa.github.io>

<sup>6</sup><https://andrebola.github.io/EUSIPCO2020/demos>

<sup>7</sup>Similar to suggestions by other researchers reproducing this model.

<sup>8</sup><https://dgschwend.github.io/netscope/quickstart.html>

## VIII. RESULTS

We evaluated the considered mel-spectrogram settings on the adjusted CNN models. Full results for all evaluated configurations are available online.<sup>9</sup> In Figure 1 we show the results of the evaluation for VGG-CNN on the MTAT dataset, repeated three times for each configuration. The first two plots show the ROC AUC results for the 12 KHz and 16 KHz sample rate using the log and dB scaling. Similarly, the third and fourth plots show the PR ROC results under the same conditions. The last plot shows GMAC.

The results show that using some of the settings we can reduce the size of the input in frequency and time without affecting much the performance of VGG-CNN on the MTAT dataset. For example, if we reduce the frequency resolution from 96 to 48 mel bands we can reduce the MAC operations near 50% without affecting the performance in all configurations. Similarly, we can also reduce time resolution by 50% without affecting performance, and in this case we also reduce the MAC operations by 50% in all configurations. We can further reduce the number of operations by the cost of some performance decrease. This can be especially useful for applications requiring lightweight models, as we can get a model  $\times 10$  faster by sacrificing between 1.4 and 1.8% of the performance depending on the configuration. Interestingly enough, both ROC AUC and PR AUC slightly improve when using 48 mel bands compared to 96 bands in most of the cases, however no statistically significant difference was found ( $P > 0.08$  for all corresponding configurations in an independent samples t-test).

For the MUSICNN model, we have tested some of the configurations reported in Table IV. We only considered the frequency resolution reduction to 48 mel bands and no hop size increments due to significantly slower training time (see Section IV). The results show comparable performance of 96- and 48-band mel-spectrograms and are consistent with the above mentioned findings for the VGG-CNN model. Overall, using 128 mel bands resolution provided the best performance. Also, according to the results, the MUSICNN architecture outperforms VGG-CNN, which is consistent with the reports from the authors.

To check how our findings scale, we selected a number of configurations and re-evaluated the models on the MSD dataset. The results are reported in Table V. In the case of VGG-CNN the performance of the baseline architectures is slightly superior to the ones working with lower-resolution mel-spectrograms, which comes by cost of a significantly larger computational effort. For example, for the 12 KHz sample rate,  $\times 1$  hop size and dB compression settings, reducing the number of mel bands from 96 to 48 results in the decrease is 0.16% in the ROC AUC performance and 50% reduction in GMACs. For a similar 16 KHz/dB case the reduced model has the same performance with the benefit of twice as low computational speed. In the case of the MUSICNN architecture we see a reduction of the performance

of 0.19% if we compare 96 vs 48 mel bands using 12 KHz sample rate and 0.11% for 16 KHz.

TABLE IV: ROC AUC and PR AUC of the MUSICNN model on the MTAT dataset for a selection of configurations using dB log-compression and the reference hop size ( $\times 1$ ).

# mels	sample rate	ROC AUC	PR AUC
128	12 KHz	90.40	<b>38.54</b>
96	12 KHz	<b>90.50</b>	37.70
48	12 KHz	90.33	37.80
128	16 KHz	<b>90.83</b>	<b>38.92</b>
96	16 KHz	90.60	38.09
48	16 KHz	90.50	37.70

TABLE V: ROC AUC and PR AUC of the models on the MSD dataset for a selection of configurations using dB log-compression.

# mels	hop size	sample rate	ROC AUC	PR AUC
128	$\times 1$	12 KHz	86.48	27.56
96	$\times 1$	12 KHz	<b>86.67</b>	<b>27.70</b>
48	$\times 1$	12 KHz	86.53	27.27
128	$\times 2$	12 KHz	86.28	27.24
96	$\times 2$	12 KHz	86.18	26.93
48	$\times 2$	12 KHz	85.86	26.42
128	$\times 1$	16 KHz	<b>86.84</b>	<b>28.10</b>
96	$\times 1$	16 KHz	86.71	28.06
48	$\times 1$	16 KHz	86.73	27.78
128	$\times 2$	16 KHz	86.34	27.06
96	$\times 2$	16 KHz	86.63	27.70
48	$\times 2$	16 KHz	86.41	26.83

(a) VGG-CNN

# mels	hop size	sample rate	ROC AUC	PR AUC
128	$\times 1$	12 KHz	87.10	26.97
96	$\times 1$	12 KHz	87.16	<b>27.10</b>
48	$\times 1$	12 KHz	86.99	26.66
128	$\times 1$	16 KHz	<b>87.21</b>	26.91
96	$\times 1$	16 KHz	<b>87.21</b>	26.96
48	$\times 1$	16 KHz	87.10	26.64

(b) MUSICNN

## IX. CONCLUSIONS

In this paper we have studied how different mel-spectrogram representations affect the performance of CNN architectures for music auto-tagging. We have compared the performances of two state-of-the-art models when reducing the mel-spectrogram resolution in terms of amount of frequency bands and frame rates. We used the MagnaTagaTune dataset for comprehensive performance comparisons and then compared selected configurations on the larger Million Song Dataset. The results suggest that it is possible to preserve a similar performance while reducing the size of the input. They can help researchers and practitioners to make trade-off decision between accuracy of the models, data storage size and training and inference time, which are crucial in many applications.

As a future work, other approaches such as quantization of mel-spectrogram values will be considered for the reduction of

<sup>9</sup><https://andrebola.github.io/EUSIPCO2020/results>

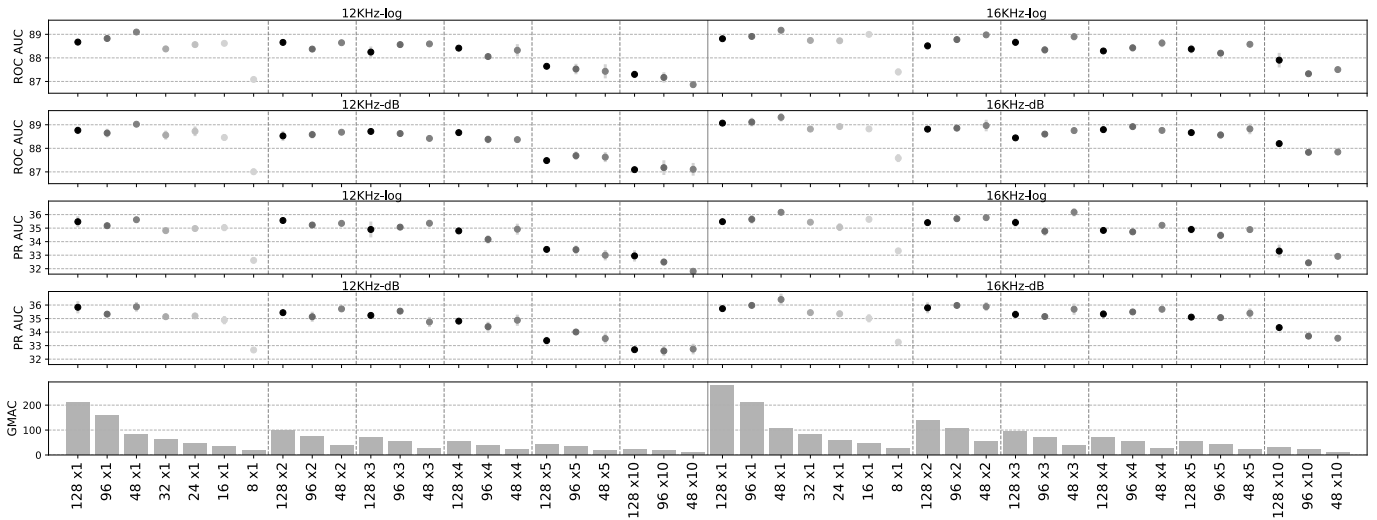


Fig. 1: Mean and standard deviation of ROC AUC and PR AUC of the VGG-CNN model computed on three runs for each mel-spectrogram configuration (# mel, hop size, sample rate, and log type) and the associated GMAC values.

the input data dimensionality. The conducted evaluation can be also extended to other state-of-the-art architectures operating on mel-spectrograms [5], constant-Q [10] and raw waveform approaches [1], [11]. It is also promising to conduct a similar evaluation on other audio auto-tagging tasks. All the code to reproduce this study is open-source and available online.<sup>10</sup>

#### ACKNOWLEDGMENT

This research has been supported by Kakao Corp.

#### REFERENCES

- [1] Jongpil Lee, Jiyoung Park, Keunhyoung Luke Kim, and Juhan Nam, "Sample-level deep convolutional neural networks for music auto-tagging using raw waveforms," in *Proceedings of the 14th Sound and Music Computing Conference (SMC)*, 2017.
- [2] Sander Dieleman and Benjamin Schrauwen, "End-to-end learning for music audio," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 6964–6968.
- [3] Keunwoo Choi, György Fazekas, Kyunghyun Cho, and Mark Sandler, "A tutorial on deep learning for music information retrieval," *arXiv preprint arXiv:1709.04396*, 2017.
- [4] Jordi Pons, Oriol Nieto, Matthew Prockup, Erik M. Schmidt, Andreas F. Ehmman, and Xavier Serra, "End-to-end learning for music audio tagging at scale," in *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR)*, 2018, pp. 637–644.
- [5] Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho, "Convolutional recurrent neural networks for music classification," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2392–2396.
- [6] Keunwoo Choi, George Fazekas, and Mark Sandler, "Automatic tagging using deep convolutional neural networks," in *Proceedings of the 17th International Society for Music Information Retrieval Conference (ISMIR)*, 2016, pp. 805–811.
- [7] Jordi Pons, Olga Slizovskaia, Rong Gong, Emilia Gómez, and Xavier Serra, "Timbre analysis of music audio signals with convolutional neural networks," in *2017 25th European Signal Processing Conference (EUSIPCO)*. IEEE, 2017, pp. 2744–2748.
- [8] Jordi Pons and Xavier Serra, "Designing efficient architectures for modeling temporal features with convolutional neural networks," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2472–2476.
- [9] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [10] Sergio Oramas, Francesco Barbieri, Oriol Nieto, and Xavier Serra, "Multimodal deep learning for music genre classification," *Transactions of the International Society for Music Information Retrieval*. 2018; 1 (1): 4-21., 2018.
- [11] Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho, "A comparison of audio signal preprocessing methods for deep neural networks on music tagging," in *2018 26th European Signal Processing Conference (EUSIPCO)*. IEEE, 2018, pp. 1870–1874.
- [12] Mingxing Tan and Quoc Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning (ICML)*, 2019.
- [13] Edith Law, Kris West, Michael I Mandel, Mert Bay, and J Stephen Downie, "Evaluation of algorithms using games: The case of music tagging," in *Proceedings of the 10th International Society for Music Information Retrieval Conference (ISMIR)*, 2009, pp. 387–392.
- [14] Thierry Bertin-Mahieux, Daniel PW Ellis, Brian Whitman, and Paul Lamere, "The million song dataset," in *Proceedings of the 12th International Society for Music Information Retrieval Conference (ISMIR)*, 2011.
- [15] Keunwoo Choi, György Fazekas, Kyunghyun Cho, and Mark Sandler, "The effects of noisy labels on deep convolutional neural networks for music tagging," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, pp. 139–149, 2018.
- [16] Bob L Sturm, "A survey of evaluation in music genre recognition," in *International Workshop on Adaptive Multimedia Retrieval*. Springer, 2012, pp. 29–66.
- [17] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Malcolm Slaney, "Auditory toolbox," *Interval Research Corporation, Tech. Rep*, vol. 10, no. 1998, 1998.
- [19] Dmitry Bogdanov, Nicolas Wack, Emilia Gómez, Sankalp Gulati, Perfecto Herrera, O. Mayor, Gerard Roma, Justin Salamon, J. R. Zapata, and Xavier Serra, "Essentia: an audio analysis library for music information retrieval," in *Proceedings of the 14th International Society for Music Information Retrieval Conference (ISMIR)*, 2013, pp. 493–498.
- [20] Sergio Oramas, Oriol Nieto, Francesco Barbieri, and Xavier Serra, "Multi-label music genre classification from audio, text and images using deep features," in *Proceedings of the 18th International Society for Music Information Retrieval Conference (ISMIR)*, 2017, pp. 23–30.
- [21] Jesse Davis and Mark Goadrich, "The relationship between precision-recall and roc curves," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 233–240.

<sup>10</sup><https://andrebola.github.io/EUSIPCO2020/>