Learning to Recognize Musical Genre from Audio

Challenge Overview

Michaël Defferrard
EPFL, Lausanne, Switzerland
michael.defferrard@epfl.ch

Sean F. Carroll
EPFL, Lausanne, Switzerland
sean.carroll@epfl.ch

Sharada P. Mohanty
EPFL, Lausanne, Switzerland
sharada.mohanty@epfl.ch

Marcel Salathé
EPFL, Lausanne, Switzerland
marcel.salathe@epfl.ch

ABSTRACT

We here summarize our experience running a challenge with open data for musical genre recognition. Those notes motivate the task and the challenge design, show some statistics about the submissions, and present the results.

KEYWORDS

Music Information Retrieval (MIR); ML Challenge; Open Data

ACM Reference Format:

Michaël Defferrard, Sharada P. Mohanty, Sean F. Carroll, and Marcel Salathé. 2018. Learning to Recognize Musical Genre from Audio: Challenge Overview. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23-27, 2018, Lyon, France*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3184558.3192310>

1 INTRODUCTION

Like never before, the web has become a place for sharing creative work — such as music — among a global community of artists and art lovers. While music and music collections predate the web, the web has enabled much larger scale collections. Whereas people used to own a handful of vinyl records or CDs, they nowadays have instant access to the whole of published musical content via online platforms such as Spotify, iTunes, Youtube, FMA, Jamendo, Bandcamp, etc. Such dramatic increases in the size of music collections has created two challenges: (i) the need to automatically organize a collection (as users and publishers cannot manage them manually anymore), and (ii) the need to automatically recommend new songs to a user knowing their listening habits. An underlying task in both those challenges is to be able to group songs in semantic categories.

Music genres are categories that have arisen through a complex interplay of cultures, artists, and market forces to characterize similarities between compositions and organize music collections. Yet the boundaries between genres still remain fuzzy, making the problem of music genre recognition (MGR) a nontrivial task [5]. While its utility has been debated, mostly because of its ambiguity and cultural definition, it is widely used and understood by end-users who find it useful to discuss musical categories [3].

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/10.1145/3184558.3192310>

The task of this challenge, one of the four Web Conference's challenges, was to recognize the musical genre of a piece of music of which only a recording is available. Genres are broad, e.g. *pop* or *rock*, and each song only has one target genre. Other metadata, e.g. the song title or artist name, were not to be used for the prediction.

The data for this challenge comes from the recently published FMA dataset [1], a dump of the Free Music Archive (FMA).¹ The dataset is a collection of 917 GiB and 343 days of Creative Commons-licensed audio from 106,574 tracks from 16,341 artists and 14,854 albums, arranged in a hierarchical taxonomy of 161 genres. It provides full-length and high-quality audio, pre-computed features, together with track- and user-level metadata, tags, and free-form text such as biographies.

2 THE CHALLENGE

To avoid overfitting and cheating, we organized the challenge in two rounds. The final ranking was based on results from the second round. The training data for both rounds consisted of the FMA medium subset, which is composed of 25,000 clips of 30 seconds, categorized in 16 genres. The categorization is unbalanced with 21 to 7,103 clips per genre. As the data is public, we collected new test data for the second round to prevent access to the test set.

In the first round, participants were provided a test set of 30,000 clips of 30 seconds each and had to upload the predicted genre for each of these clips. The platform readily evaluated those predictions and ranked the participants upon each submission. A subset of these clips were sampled from the FMA large dataset, while ensuring that none overlaps with any clip provided in the training set. The other subset was sampled from songs in the FMA full dataset which are not present in the medium subset.

For the second round, the participants had to provide their models as git repositories which contained the prediction code and the trained model along with an executive summary of their approach. Docker containers were built out of those repositories.² We then ran them against a new unseen test set which was sampled from new contributions to the Free Music Archive.

Both rounds used the same evaluation metric. The primary score was the mean log loss and the secondary score was the mean F1 score. The mean log loss is defined by

$$L = -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C y_{nc} \ln(p_{nc}), \quad (1)$$

¹<https://freemusicarchive.org>

²<https://github.com/jupyter/repo2docker>

Figure 1: Joint distribution of the F1 score and the log loss (if smaller than 5) of all submissions. The red point represents the baseline prepared by the organizers.

where $N = 35000$ is the number of examples in the test set, $C = 16$ is the number of genres, y_{nc} is a binary value indicating if the n -th instance belongs to the c -th label, p_{nc} is the probability according to a submission that the n -th instance belongs to the c -th label, and $\ln$ is the natural logarithm. The mean F1 score is given by

$$F_1 = \frac{2}{C} \sum_{c=1}^C \frac{p^c r^c}{p^c + r^c}, \quad p^c = \frac{tp^c}{tp^c + fp^c}, \quad r^c = \frac{tp^c}{tp^c + fn^c}, \quad (2)$$

where p^c and r^c are the precision and recall for class c , and tp^c , fp^c , fn^c refers to the number of true positives, false positives, and false negatives.

The challenge was hosted on crowdAI, a public platform for open challenges. Instructions on how to participate, access to training and test data, graded submissions, and the leaderboard were available on the challenge page.³ Moreover, we developed a starter kit with code to handle the data and make a submission.⁴ It also featured some examples and a baseline. Finally, participants were encouraged to review the FMA paper [1] for a detailed description of the data as well as the GitHub repository for Jupyter notebooks showing how to use the data, explore it, and train baseline models.⁵

3 RESULTS

At the end of the first round, we had engaged a total of 246 participants who either made a submission, downloaded the datasets, or contributed to the forums. From those 246 participants, 38 made at least one submission, with some of the top participants making as many as 110 submissions. A total of 671 submissions were made in the first round. Of these, 77 were invalid and 576 were successfully graded. From those 576 submissions, 364 had a score better than the baseline provided by the organizers. Figure 1 shows the distribution of F1 scores and log losses. The current best solution has an F1 score of 0.909 and a log loss of 0.330. Figure 2 shows how the participants progressed through the first round.

Moreover, we reviewed and accepted two papers. In [4], the authors compared the following approaches: (i) ConvNet on spectrograms, and (ii) deep neural net, (iii) ExtraTrees, and (iv) XGBoost on higher-level features extracted by Essentia. They found that

³<https://www.crowdai.org/challenges/www-2018-challenge-learning-to-recognize-musical-genre>

⁴<https://github.com/crowdAI/crowdai-musical-genre-recognition-starter-kit>

⁵Code and data available at <https://github.com/mdeff/fma>. The rc1 version was used.

Figure 2: Score progression of all participants through the first round of the challenge. Each line represents an active participant. The dotted line represents the baseline.

ensemble methods outperformed neural networks, with XGBoost performing best. In [2], the authors argued that genres are subjective and noisy labels, whereas artists are more objective labels. As an artist is commonly part a subset of genres, and that sets of artists can be seen as exemplars for genres, they hypothesized that musical characteristics which identify an artist may also be key features of certain genres. As such, they proposed to train a multi-task neural network to jointly predict artist group and genre. Results showed that features learned for artist recognition were indeed useful for MGR. They thus achieved transfer learning, though keeping the main MGR task as one of the multiple subtasks was crucial.

4 CONCLUSION

The outcomes of the challenge were multiple. First, the accepted papers presented new perspectives and introduced new methods. Then, all the participants to the second round had to share their code as open-source. We hope that those implementations will be useful to the community, for example to serve as baselines, to be scrutinized, or to be improved upon. Finally, the challenge introduced participants to the new FMA dataset and was an opportunity for them to get familiar with it.

That challenge was part of a wider effort to promote open evaluation in machine learning for music data, of which the release of the open FMA dataset was the first step [1]. The goal of this initiative is to establish a reference benchmark based on open data. MIR research has historically suffered from the lack of publicly available benchmark datasets, which stem from the commercial interest in music by record labels, and therefore imposed rigid copyright. The FMA's solution was to aim for tracks which license permits redistribution. All data and code produced during the project and challenge are released under the CC BY 4.0 and MIT licenses.

REFERENCES

- [1] Michaël Defferrard, Kirell Benzi, Pierre Vandergheynst, and Xavier Bresson. 2017. FMA: A Dataset for Music Analysis. In *ISMIR*.
- [2] Jaehun Kim, Minz Won, Xavier Serra, and Cynthia C. S. Liem. 2018. Transfer Learning of Artist Group Factors to Musical Genre Classification. In *WWW*.
- [3] C McKay and I Fujinaga. 2006. Musical genre classification: Is it worth pursuing and how can it be improved?. In *ISMIR*.
- [4] Benjamin Murauer and Günther Specht. 2018. Detecting Music Genre Using Extreme Gradient Boosting. In *WWW*.
- [5] N Scaringella, G Zoia, and D Mlynek. 2006. Automatic genre classification of music content: a survey. *IEEE Signal Processing Magazine* (2006).