The Fractal Dimension of Music: Geography, Popularity and Sentiment Analysis

Laura Pollacci[1(✉)], Riccardo Guidotti[2], Giulio Rossetti[1], Fosca Giannotti[2], and Dino Pedreschi[1]

[1] Department of Computer Science, University of Pisa, Pisa, Italy
{laura.pollacci,giulio.rossetti}@di.unipi.it, dino.pedreschi@unipi.it
[2] ISTI-CNR, Pisa, Italy
{riccardo.guidotti,fosca.giannotti}@isti.cnr.it

Abstract. Nowadays there is a growing standardization of musical contents. Our finding comes out from a cross-service multi-level dataset analysis where we study how geography affects the music production. The investigation presented in this paper highlights the existence of a "*fractal*" musical structure that relates the technical characteristics of the music produced at regional, national and world level. Moreover, a similar structure emerges also when we analyze the musicians' popularity and the polarity of their songs defined as the mood that they are able to convey. Furthermore, the clusters identified are markedly distinct one from another with respect to popularity and sentiment.

Keywords: Music data analytics · Hierarchical clustering
Sentiment pattern discovery · Multi-source analytics

1 Introduction

Music has been part of human civilization for centuries: it has been referred to as the *universal language*. Certainly, every culture has given birth to its own music "style": however, as time goes by, the constant collapse of physical barriers as well as the progressive reduction of geographical distances eased by the media and the world wide web caused an overall globalization of music.

During the last decade, the constant growing of on-line streaming services (such as Spotify, Apple Music, Last.fm) has made available to the public the widest choice of music ever. Emerging bands as well as famous ones can obtain a global visibility that was unimaginable only a few years ago. In this rapidly evolving scenario, music seems to have lost its geographical-cultural connotation: is it true that we are observing a growing standardization of music contents? Are there peculiar characteristics able to discriminate the music produced in a given region from the one produced in the country that contains it?

B. Guidi et al. (Eds.): GOODTECHS 2017, LNICST 233, pp. 183–194, 2018.
https://doi.org/10.1007/978-3-319-76111-4_19

In this paper, leveraging a cross-service multi-level dataset, we study how geography affects the music production. Moreover, we study how artists producing a specific type of music are able to reach high popularity and how such popularity is often not related to a specific genre. Our data-driven investigation highlights the existence of a "*fractal*" structure that relates the technical characteristics of the music produced at regional, national and world level. Our structure reflects the well-knowledge characteristics of fractals objects theorized in mathematics since it exhibits a detailed pattern that repeats itself at increasingly big scales. Starting from emergent groups of an Italian region (Tuscany), moving to affirmed Italian artists and finally to a set of world-famous musician, we identify a multi-level set of profiles transversal to the classical concept of genre: we show how such profiles remains stable across the different geographical layers analyzed. Finally, we observed how the mood expressed by artists' songs as well as their popularity vary w.r.t. the multi-level traversal profiles they belong to. Once again, we observed the presence of the same structure that clearly emerges both at different geographical levels and over a multi-level set of profiles.

The rest of the paper is organized as follows. In Sect. 2 is discussed the literature involving music data analysis. Section 3 introduces the datasets used in the current study, and the preprocessing steps performed on them. In Sect. 4 we show and discuss the fractal structures emerging from the datasets and the music profiles that we were able to identify applying an unsupervised learning strategy. Section 5 highlights the relationships of the artists' popularity and music genres with respect to the groups discovered. Following the same approach, in Sect. 6 we analyze the relationships between the song lyrics and the music profiles enriched with sentiment analysis. Finally, Sect. 7 concludes the paper and discusses some future research directions.

2 Related Work

During the last decade, the music world has started receiving increasing attention from the scientific community. Several works [17,18] have analyzed data regarding on-line listening in order to model diffusion of new music genres/artists, as well as to analyze the behaviors and tastes of users.

Moreover, nowadays the on-line platform *Last.Fm* is offering a privileged playground to study different phenomena related to the on-line music consumption. In [19] was proposed a music recommendation algorithm by using multiple social media information and music acoustic-based contents. Hypergraphs developed on Last.fm model the objects and relations, and music recommendation is considered as a ranking problem. In [20] the authors studied the topology of the Last.Fm social graph asking for similarities in taste as well as on demographic attributes and local network structure. Their results suggest that users connect to "on-line" friends, but also indicate the presence of strong "real-life" friendship ties identifiable by the multiple co-attendance of the same concerts. The authors of [21] measured different dimensions of social prominence on a social graph built upon 70k Last.Fm users whose listening were observed for 2 years. In [22] was analyzed the cross-cultural gender differences in the adoption and usage of Last.Fm. Using data from Last.Fm and Twitter the authors of [23], designed a measure that describes diversity of musical tastes and explored its relationship with variables that capture socioeconomic, demographics, and traits such as openness and degree of interest in music. In [24] is shown how to define statistical models to describe patterns of song listening in an on-line music community. Finally, yet using Last.fm data, in [16] the authors defined a Personal Listening Data Model able to capture musical preferences through indicators and patterns, and showed how the employment of such model can provide to the individual users higher levels of self-awareness. They also discovered that all the users are characterized by a limited set of musical preferences, but not by a unique predilection.

In this paper, we are looking for dependencies between different levels of musical artist by developing a musical profile through clustering techniques. Also in [25–27] are used clustering techniques on musical data with different purposes. In [25], is studied the problem of identifying similar artists by integrating features from diverse information sources. In [27] a compression distance has been used to generate the similarities among music files and the paper shows some experimental results using these representations and compares the behavior of various clustering methods. Also in [26] is studied the problem of building clusters of music tracks in a collection of popular music in the presence of constraints and is presented an approach incorporating the constraints for grouping music by similar artists.

To the best of our knowledge, this work is the first attempt in which three different data sources of musical data are analyzed and clustered to find containment patterns and to discover musical dependencies also considering the lyrics of the songs.

3 Datasets and Preprocessing

In this section, we describe the two different types of dataset sources and the preprocessing applied to them.

As proxy of the actual musical scene, we exploit musical dataset having three different level of spatial granularity: world, national and regional. Datasets refer to *world*'s famous musicians, *Italian* musicians, and emerging youth bands in *Tuscany*, respectively.

In Table 1 we describe the details of these datasets. In particular, the `TUSCANY` dataset is referred to emerging artists participating in the "100 Band" contest promoted by "Tuscan Region" and "Controradio" in 2015 [2]. These datasets are built using the Spotify API [3] and are composed by all the songs present on the platform for the selected artists. For each artist, we collect songs titles, album titles, and the number of *followers*. Furthermore, each song is identified by its title, artist, and *popularity* score, that is based on more is played that song. In addition, each track is described by a set of musical features [7]: *acousticness,*

Table 1. Datasets statistics. Within brackets are reported the number of artists for which at least a single song lyric was available.

Dataset	#Artist	#Tracks	#Genres	#Lyrics
`WORLD`	833,197 (19,218)	5,525,222	1,380	79,204
`ITALY`	2,379 (710)	502,582	126	28,582
`TUSCANY`	513 (58)	24,147	28	91

danceability, duration, energy, instrumentalness, liveness, loudness, speechiness, tempo and *valence*. All the features range in [0, 1] with the exception of *duration, loudness* and *tempo*. In the preprocessing phase, we normalize these latter features in order to align all the feature scales.

Moreover, in order to deepen our analysis and consider only the mood transmitted by the players with their sing in terms of lexical content, we integrate these datasets with the lyrics of the songs. In Table 1 are also described the details of the lyrics' datasets. In particular, the `WORLD-lyric` dataset is collected from the Genius [4], the `ITALY-lyric` dataset is collected using SoundCloud API [5], and, finally, the `TUSCANY-lyric` is built extracting texts from the results of a survey. Using Google Form service [6] we gathered musical and personal data about artists who participated at the Tuscany 100 Band contest.

The first problem we decided to deal with is the grouping of musical genres. In fact, all three datasets show a large number genres (both minor and major). Using a list of popular music genres [8,10] we assign each song's minor genres to their major class. Then, we group the collected songs in 12 classes: country, blues, religious, hip hop, latin, electronic, folk, jazz, rock, R&B, pop, a cappella.

In order to process music lyrics with unsupervised methods we first clear the data. Therefore, to obtain normalized texts, we treat lyrics datasets using a rule-based cleaning procedure. Following this method, we obtain standardized music lyrics that can be treated by a general-purpose pipeline of Natural Language Processing (NLP). After that, music lyrics are lemmatized and Part-Of-Speech are tagged with the POS tagger TreeTagger [28,29]. We also reduce the noise selecting only nouns, verbs and adjectives. Doing this, for each text, we obtain only significant words from the sentiment points of view.

4 The Music Scene Fractal Structure

In order to identify the prototypical type of music produced by each artist in the datasets, we describe every performer through his medoid song, i.e., his most representative song identified minimizing the sum of the Euclidean distances between the Spotify features among all his discography. Once built such descriptions of each artist, we move on grouping them together on the basis of the music they produce. Since the available datasets allow us to observe the music phenomenon from three different hierarchical levels (regional, national and world-wide), we perform three different levels. The first one describing the music of regional artists, the second one describing the music of both regional and national artists and the last one describing the music of all the artists observed world-wide.

Through the analysis of the hierarchical clusterings, we aim at understanding if a fractal structure emerges among the type of music produced at different geographical levels. We accomplish this task by employing the *k-means* clustering algorithm [30] to the computed artist profiles [31]. As first step, in order to identify a reasonable value of k, i.e., the most appropriate number of clusters, we calculate the Sum of Squared Error (SSE) distribution for $k \in [2, 18]$. The SSE distributions for the clusterings computed on the three levels - not showed for space reasons - follow a common pattern that identifies optimal values of k in the range $[4, 6]$. After that we have identified such range, we extract the clusters for each value of k in it for the three datasets. Finally, from each cluster of each level, we calculate the medoid of the cluster, i.e., the set of features describing its most representative artist.

In order to understand if our datasets present a fractal structure, we study artists' migration among the clusters when moving from the regional to the world level. We repeat this activity for each k in the identified range. Figure 1 shows the clusters coverage - due to space reasons, only for $k = 5$ - when migrating from `ITALY` to `WORLD` clusters (right figure), and from `TUSCANY` to `ITALY` clusters (left figure). Indeed, these matrices have a strong diagonal prevalence. The same effect can be observed for all the clusters and for all k-values in the range [4, 6]: artists' blocks vary in high percentages from a down-geographical cluster to a top-geographical cluster.

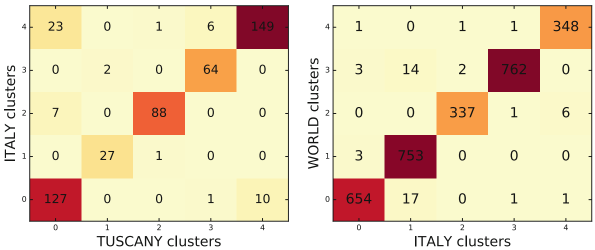

(a) k=5 (b) k=5

Fig. 1. Matrices clusters coverage when migrating from `ITALY` to `WORLD` clusters (left figure), and from `TUSCANY` to `ITALY` clusters (right figure).

We calculated such matrices by matching the pairs of clusters of two different levels with the highest level of coverage, i.e., by maximizing the purity using the cluster identifiers as a label. In fact, we can observe how, for instance, regional artists of a given cluster are re-classified into a cohesive block of the national level that, in turn is re-classified in a single block of the world level.

Since a fractal structure emerges with all the adopted values of k in the following we detail the clustering obtained only for $k = 5$. Moreover, music genres'

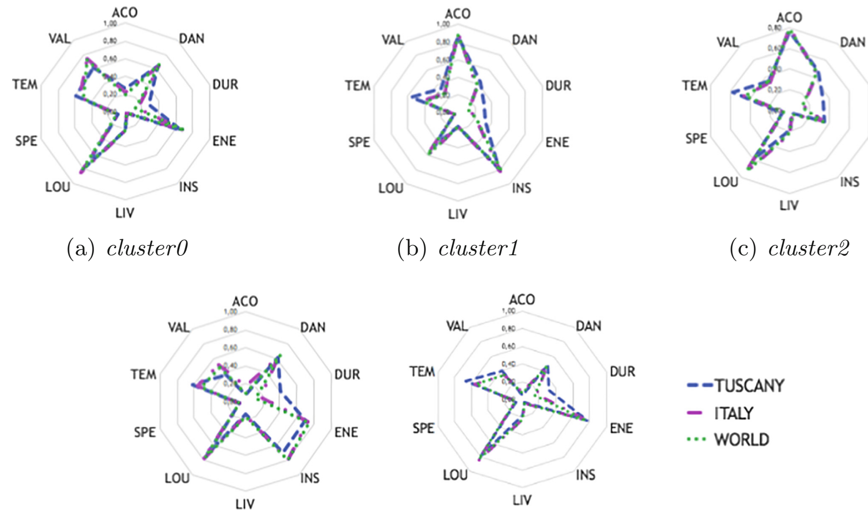

(a) *cluster0* (b) *cluster1* (c) *cluster2*

(d) *cluster3* (e) *cluster4*

Fig. 2. Artists Profiles: `TUSCANY`, `ITALY`, `WORLD` medoids. From left to right: *cluster0, cluster1, cluster2, cluster3, cluster4*

distribution varies within each of the three datasets, consequently affecting their clusterization. Thus, for each level, entire musical sub-genres fall into specific clusters.

In order to better analyze and understand the characteristics of these clusters, we compare the clusters for each dataset by taking advantage of the radar chart in Fig. 2 computed for $k = 5$ (we obtain comparable results for $k = 4$ and for $k = 6$). They describe the medoids of the five clusters identified for each dataset: radar charts underline, one more time, the presence of a fractal structure capturing very similar profiles across the observed hierarchy levels. This demonstrates that, for each different level of observation, musically homogeneous artists are well clustered with artists of the superior step. First of all, we notice how the *fractal* structure is perfectly highlighted also by the radar charts. The spikes of the three different datasets follow the same shape almost for each cluster. At a first glance of Fig. 2, we can observe that some features are more discriminant than others. In fact, features like *speechiness, liveness, loudness*, and *tempo* present similar values in each cluster. Despite this, the others features are determinants for cluster discrimination. Despite some little discrepancy among datasets, we can group clusters by their similarities. We perceive that *cluster0, cluster3* and *cluster4* are expressions of artists in direct opposition to the ones respectively in *cluster1* and *cluster2*:

– *cluster1* and *cluster2* represent melodic, bit danceable without a strong beat, and negative songs. Clusters diverge only for instrumental scores: *cluster1* present low values, on contrary *cluster2* show high values;

– *cluster0*, *cluster3* and *cluster4* represents non-melodic, strongly rhythmic and fairly dancing tracks. *cluster0* and *cluster4* show very low instrumental values, while *cluster3* show high scores. Furthermore, clusters differs for valence values: *cluster0* presents the highest values, on contrary *cluster4* presents the lowest values.

Since a fractal structure able to relate different geographical levels emerges also on the other dimensions, in the following sections we analyze the popularity, followers, genres, and sentiment of the clusters and we will detail only the results observed at the worldwide level.

5 Genres, Popularity and Followers

In this section, we analyze the collected information regarding artist popularity and followers as well as song genres. Starting from a dataset-wide discussion we detail how such dimensions can be used to characterize the identified clusters.

Figure 3(a) shows the overall distributions of both artists and followers with respect to the 12 genres identified after the cleaning stage. Indeed, the most represented genres are rock, electronic and pop, meta-classes. As was foreseeable, these genres also attract a considerable number of followers. However, is interesting noticing that, despite rock ranks a first among the most played genres, it presents a smaller number of followers than pop music.

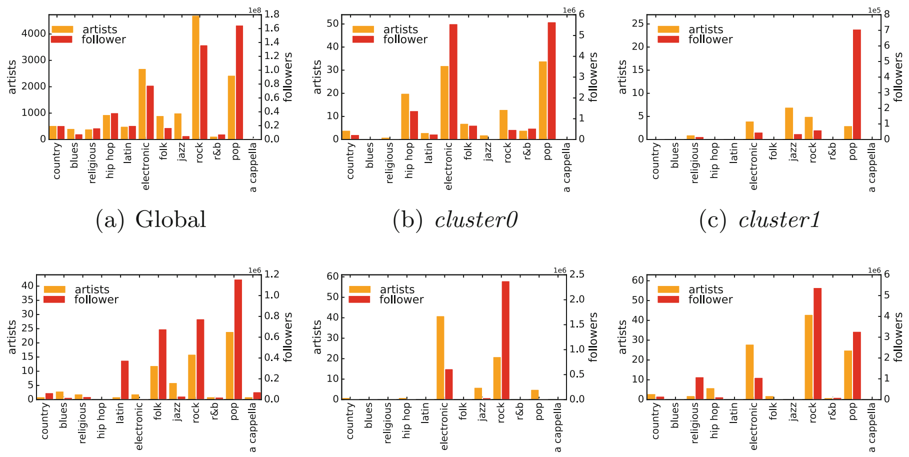

(a) Global (b) *cluster0* (c) *cluster1*

(d) *cluster2* (e) *cluster3* (f) *cluster4*

Fig. 3. Artists and Followers distribution among clusters.

Figure 4(a) shows the relation between artists' popularity and followers. Starting from such plot three subclasses of artists can be identified:

- *Low popularity.* A large number of artists present a low popularity, between 0 and 40, and a small number of followers: indeed, a large number of artists are followed by few people. Such artists could be emerging artists or are likely to play some kind of niche music.
- *Medium-high popularity.* Artists having medium popularity, between 40 and 70, are followed by a consistent number of users.
- *High popularity.* Artists having the very high popularity and that have a relatively low number of followers.

While the former two classes were somehow expected, the latter one breaks the common intuition that expects very popular artists to be the ones attracting more followers.

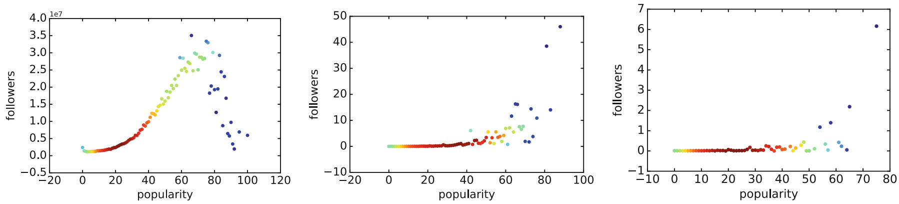

(a) Global (b) *cluster0* (c) *cluster1*

Fig. 4. Relations between popularity and followers.

Once observed the general behaviors of popularity, followers and genre components, we study how they relate with the clustering we obtained in Sect. 4. As a first step, in order to provide a semantic annotation of the identified clusters, we describe them exploiting their main genres as well as their profiles. As we can see from Fig. 3, clusters are strongly heterogeneous since they represent different music genres. Indeed, the clusters obtained are markedly different among each other and each cluster distinctly identifies a subset of genres with specific levels of popularity and number of followers.

Cluster0, Fig. 3(b), can be identified as the electronic/pop/hip hop cluster. It is represented by artists like Laura Pausini and Vanilla Sky. In this cluster fall few genres, however, musicians that belong to this profile are followed by the highest amount of users w.r.t. all clusters. Songs are very suitable for dancing, repetitive, cheerful and sung. In this cluster also fall all few rap artists but probably, they have little influence on medoids' values. However, speechness values are still the highest of the dataset.

Cluster1, Fig. 3(c), is the jazz/rock cluster. The majority of the artists belonging to this cluster, like Stefano Bollani and Doctor Dixie Jazz Band, have few followers. The beat pace and high acoustics make tracks unsuitable for dancing and influence the valence values that are lowest of all others clusters;

Cluster2, Fig. 3(d), is the pop/folk music cluster. These genres are the most represented by artists. In this cluster falls artists like Norah Jones and Lucio Battisti. Tracks are perceived as calm, unsuitable for dancing, primarily vocal and sad, thus leading to valences score that are fairly negative;

Cluster3, Fig. 3(e), is the electronic/rock cluster and it is represented by artists such Calvin Harris and Go!Zilla. The most representative genre is the electronic music, however, users that follow this genre are a few compared by the rock ones. This cluster represents non-acoustic, strongly danceable with a strong beat, like dance, house, and minimal music.

Cluster4, Fig. 3(f), is the pop/rock cluster and in this cluster falls artists like U2 and Green Day. The number of rock and pop followers is moderately high, while the amount of electronic music is quite low. Songs are strongly rhythmic, noisy and energetic but slightly danceable. Often, tracks are perceived as angry, so valences are tendentially negatives.

As final step we study for each cluster the relations between its artists' popularity and their number of followers. Due to space reasons, Fig. 4 shows the relation between artists' popularity and followers only for *cluster0* and for *cluster1*, but same results are showed for all five clusters. All the clusters are characterized by the same trend: most of their artists have medium-low popularity, with scores between 20 and 40, and are followed by a low amount of users. Moreover, as popularity grows also followers increase. However, musicians having both high popularity, over 70, and a high number of followers are few and almost uniformly spread across all the clusters. Moreover, it's interesting to observe that in each cluster there are very few artists with a very high popularity score, over 80, that are followed by a high number of followers: supposedly, these are famous international artists.

6 Sentiment Analysis

As a final analysis, we are interested in observing the correlation between the songs and the feelings they transmit, framing our analysis within our clusters.

We proceed by adopting a lexicon-based approach, exploiting ANEW [32] as seed-lexicon resource. ANEW provides a set of emotional ratings for 1,034 words in terms of *valence*, *arousal*, and *dominance*. In order to determinate artists' polarity, we select the valence values provided from both male and female subjects.

After calculating the weighted average of the words' valences v_{text} as polarity score, we grouped songs by artists, in order to obtain a polarity value for each of them. For each artist, we compute the weighted average among his tracks' polarity. In order to enhance the differences among the levels of the various scores we apply a logistic function [35] $f(x) = \frac{L}{1+e^{-k(x-x_0)}}$; where L is the curve maximum value (we set it equal to 1), $k = 10$ is the steepness of the curve, $x_0 = 0.5$ is the x-value of the sigmoid's midpoint, and $x = v_{text}$.

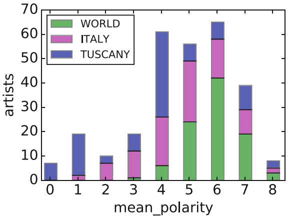

Fig. 5. Artists' polarity score distribution among polarity class.

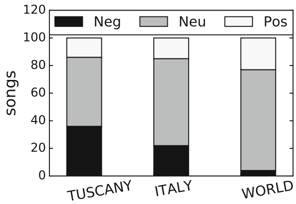

Fig. 6. Artists' polarity score distribution among datasets.

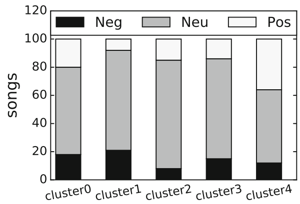

Fig. 7. Artists' polarity score distribution among clusters.

We then apply the aforementioned procedure to each lyrics' dataset. Then we grouping emotionally-tagged tracks for each artist and computing the weighted average among the i-artist's discography. In order to analyze text belonging to `ITALY-lyric` and `TUSCANY-lyric` we translated the ANEW lexicon in Italian by using the Python library Goslate [9].

As result, we obtain a comprehensive list of artists each one of them having a polarity scores in the range [0, 10]. Studying the polarity distributions we note that the most artists have a polarity score in the range [5, 7], as showed in Fig. 5. Anyway, we need to keep in mind that a neutral score could indicate that *(a)* the artist's tracks transmit no strong emotions or, *(b)* they present conflicting emotions. In order to give a better comprehension of the results, we split artists based on the polarity score into three class: *(a) positive*, scores higher than 6; *(b) negative*, scores lower than 4, and *neutral*, scores between 4 and 6. Figure 6 shows the distribution of the artists' polarity among the three datasets using ANEW.

Finally, we analyze how the clusters are affected by the polarity scores. Indeed, we apply the same procedure described above for each of the five clusters. The other clusters are described by a Gaussian curve similar to the one expressed by the complete dataset. In fact, all such distributions are multimodal and define three peaks. The remaining artists reach more extreme polarity scores, with a long tail to negative values.

Going further into the analysis, we split the artists based on their scores into the three polarity class identified above. Figure 7 shows the distribution of the artists' polarity among the three clusters using ANEW. As we can observe, the most represented category still remains the neutral one, while the less represented is the negative one.

7 Conclusion

In this work, we have proposed a data-driven investigation of the music scene. We have fulfilled this task by relying on a composite dataset built upon heterogeneous on-line resources. We compared song technical features, lyrics and artist popularity across three hierarchical geographic layers (world, national and regional). Our analysis identifies the existence of very stable clustering structure able to describe cross-genre music profiles. We highlighted how such clusters describe a *fractal* structure: apparently, disregarding the geographical granularity observed, all the artists observed can be profiled and categorized in a reduced and well defined set of clusters. Moreover, we analyzed artists' popularity and fan base observing how their distributions describe a similar trend in all the identified clusters. Finally, looking to the artists' song lyrics we were able to observe the emotional valence of the identified meta-profiles.

As future works we plan to consolidate (and extend with data coming from other sources) the cross-domain dataset we collected and to build upon it – exploiting the results of the present study – a recommender system that enables its users to discover novel artist on the base of their tastes, mood and locations.

Acknowledgment. This work is partially supported by the European Project *SoBigData: Social Mining & Big Data Ecosystem*, http://www.sobigdata.eu. This work is partially supported by the European Communitys H2020 Program under the funding scheme "INFRAIA-1-2014-2015: Research Infrastructures" grant agreement, http://www.sobigdata.eu, GS501100001809, 654024 "SoBigData: Social Mining & Big Data Ecosystem".

References

1. Spotify. https://www.spotify.com/. Accessed 28 July 2017
2. Toscana100band contest. http://toscana100band.it/. Accessed 28 July 2017
3. Spotify Web API. https://developer.spotify.com/web-api/. Accessed 28 July 2017
4. Genius. https://genius.com/. Accessed 28 July 2017
5. SoundCloud API. https://developers.soundcloud.com/docs/api/guide. Accessed 28 July 2017
6. Google Form service. https://www.google.com/forms/about/. Accessed 28 July 2017
7. Spotify Audio Features Object. https://developer.spotify.com/web-api/get-several-audio-features/. Accessed 28 July 2017
8. AllMusic. http://www.allmusic.com/genres. Accessed 28 July 2017
9. Goslate. http://pythonhosted.org/goslate/. Accessed 28 July 2017
10. List of popular music genres, (n.d.). Wikipedia. http://en.wikipedia.org/wiki/Psychology. Accessed 28 July 2017
11. Clarke, F., Ekeland, I.: Nonlinear oscillations and boundary-value problems for Hamiltonian systems. Arch. Rat. Mech. Anal. **78**, 315–333 (1982)
12. Clarke, F., Ekeland, I.: Solutions périodiques, du période donnée, des équations hamiltoniennes. Note CRAS Paris **287**, 1013–1015 (1978)
13. Michalek, R., Tarantello, G.: Subharmonic solutions with prescribed minimal period for nonautonomous Hamiltonian systems. J. Differ. Equ. **72**, 28–55 (1988)
14. Tarantello, G.: Subharmonic solutions for Hamiltonian systems via a $\mathbb{Z}_p$ pseudoindex theory. Annali di Matematica Pura (to appear)
15. Rabinowitz, P.: On subharmonic solutions of a Hamiltonian system. Commun. Pure Appl. Math. **33**, 609–633 (1980)
16. Guidotti, R., Rossetti, G., Pedreschi, D.: Audio ergo sum. In: Milazzo, P., Varró, D., Wimmer, M. (eds.) STAF 2016. LNCS, vol. 9946, pp. 51–66. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-50230-4_5

17. Rawlings, D., Ciancarelli, V.: Music preference and the five-factor model of the NEO personality inventory. Psychol. Music **25**(2), 120–132 (1997)
18. Rentfrow, P.J., Gosling, S.D.: The do re mi's of everyday life: the structure and personality correlates of music preferences. J. Pers. Soc. Psychol. **84**(6), 1236 (2003)
19. Bu, J., Tan, S., Chen, C., Wang, C., Wu, H., Zhang, L., He, X.: Music recommendation by unified hypergraph: combining social media information and music content. In: Proceedings of the 18th ACM International Conference on Multimedia, pp. 391–400. ACM (2010)
20. Bischoff, K.: We love rock'n'roll: analyzing and predicting friendship links in last. fm. In: Proceedings of the 4th Annual ACM Web Science Conference, pp. 47–56. ACM (2012)
21. Pennacchioli, D., Rossetti, G., Pappalardo, L., Pedreschi, D., Giannotti, F., Coscia, M.: The three dimensions of social prominence. In: Jatowt, A., Lim, E.-P., Ding, Y., Miura, A., Tezuka, T., Dias, G., Tanaka, K., Flanagin, A., Dai, B.T. (eds.) SocInfo 2013. LNCS, vol. 8238, pp. 319–332. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-03260-3_28
22. Putzke, J., Fischbach, K., Schoder, D., Gloor, P.A.: Cross-cultural gender differences in the adoption and usage of social media platforms - an exploratory study of Last.FM. Comput. Netw. **75**, 519–530 (2014). http://dx.doi.org/10.1016/j.comnet.2014.08.027
23. Park, M., Weber, I., Naaman, M., Vieweg, S.: Understanding musical diversity via online social media. In: Ninth International AAAI Conference on Web and Social Media (2015)
24. Zheleva, E., Guiver, J., Mendes Rodrigues, E., Milić-Frayling, N.: Statistical models of music-listening sessions in social media. In: Proceedings of the 19th International Conference on World wide web, pp. 1019–1028. ACM (2010)
25. Li, T., Ogihara, M., Peng, W., Shao, B., Zhu, S.: Music clustering with features from different information sources. IEEE Trans. Multimed. **11**(3), 477–485 (2009)
26. Peng, W., Li, T., Ogihara, M.: Music clustering with constraints. In: ISMIR, pp. 27–32 (2007)
27. González-Pardo, A., Granados, A., Camacho, D., de Borja Rodríguez, F.: Influence of music representation on compression-based clustering. In: 2010 IEEE Congress on Evolutionary Computation (CEC), pp. 1–8. IEEE (2010)
28. Schmid, H.: Improvements in part-of-speech tagging with an application to German. In: Proceedings of the ACL SIGDAT-workshop. Citeseer (1995)
29. Schmid, H.: Part-of-speech tagging with neural networks. In: Proceedings of the 15th Conference on Computational Linguistics, vol. 1, pp. 172–176. Association for Computational Linguistics (1994)
30. Tan, P.-N., Steinbach, M., Kumar, V., et al.: Introduction to Data Mining, vol. 1. Pearson Addison Wesley, Boston (2006)
31. Pollacci, L., Guidotti, R., Rossetti, G.: Are we playing like Music-Stars? Placing emerging artists on the Italian music scene (2016)
32. Bradley, M.M., Lang, P.J.: Affective norms for English words (ANEW): instruction manual and affective ratings. Citeseer, Technical report (1999)
33. Lang, P.J.: Behavioral treatment and bio-behavioral assessment: computer applications (1980)
34. Dodds, P.S., Danforth, C.M.: Measuring the happiness of large-scale written expression: songs, blogs, and presidents. J. Happiness Stud. **11**(4), 441–456 (2010)
35. Verhulst, P.F.: Recherches mathématiques sur la loi d'accroissement de la population. Nouveaux mémoires de l'académie royale des sciences et belles-lettres de Bruxelles **18**, 14–54 (1845)