Modelling Moral Traits with Music Listening Preferences and Demographics

Vjosa Preniqi[1], Kyriaki Kalimeri[2], and Charalampos Saitis[1]

[1] Centre for Digital Music, Queen Mary University of London, London UK
[2] ISI Foundation, Turin, Italy
v.preniqi@qmul.ac.uk

Abstract. Music has always been an integral part of our everyday lives through which we express feelings, emotions, and concepts. Here, we explore the association between music genres, demographics and moral values employing data from an ad-hoc online survey and the Music Learning Histories Dataset. To further characterise the music preferences of the participants the generalist/specialist (GS) score employed. We exploit both classification and regression approaches to assess the predictive power of music preferences for the prediction of demographic attributes as well as the moral values of the participants. Our findings point out that moral values are hard to predict (.62 $AUROC_{avg}$) solely by the music listening behaviours, while if basic sociodemographic information is provided the prediction score rises to 4% on average (.66 $AUROC_{avg}$), with the Purity foundation to be the one that is steadily the one with the highest accuracy scores. Similar results are obtained from the regression analysis. Finally, we provide with insights on the most predictive music behaviours associated with each moral value that can inform a wide range of applications from rehabilitation practices to communication campaign design.

1 Introduction

Music played a fundamental role in the evolution of societies being tightly related to communication, bonding, and cultural identity development [14]. Influencing a wide range of cognitive functions such as reasoning, problem-solving, creativity, and mental flexibility [17], musical taste is also known to be strongly related to personality [7] and political orientation [6]. Musical sophistication is also shown to be related to personality traits regardless of demographics or musicianship level [10].

More recently, scientists aside from the traditional self-reported surveys [6], employed digital data and in particular online music streaming [2] and social media [20] data to assess music preferences. Employing data from the myPersonality Facebook project, Nave et al. [20], found that both people's reactions to unfamiliar music samples and "likes" for music artists predicted personality traits. Krismayer et al. [13] studied the Last.fm platform showing that the music listening behaviours can predict demographics, including age, gender, and nationality. More recently, Anderson et al. [2] presented evidence about the connection between personalities and music listening preferences studying Spotify music streaming data.

Building on comparable interactionist theories, we set to explore the less attended relation between moral values and music preferences. We operationalise morality according to the Moral Foundations Theory (MFT) [9], which defines five moral traits, namely *Care/Harm, Fairness/Cheating, Loyalty/Betrayal, Authority/Subversion*, and *Purity/Degradation*. These can further collapse into two superior moral foundations: of *Individualising*, compounded by fairness and care, that asserts that the basic constructs of society are the individuals and hence focuses on their protection and fair treatment, and of *Binding*, that summarises purity, authority and loyalty, and is based on the respect of leadership and traditions.

Moral values are considered to be higher psychological constructs than the more commonly investigated personality traits yet they have attracted less attention from music scientists. In recent literature there are indications that negative emotions enforced by types of music can worsen moral judgement [3] although that study did not rely on a psychometrically validated theory like the MFT. Kalimeri et al. [12] demonstrated the predictability of moral foundations from a variety of digital data including smartphone usage and web browsing. Their results showed that moral traits and human values are indeed complex, and thus harder to predict compared to demographics, nevertheless, they provide a realistic dimension of the possibilities of modelling moral traits for delivering better targeted and more effective interventions.

Here, we train classification and regression models which infer on self-reported survey data regarding the music preferences. We thoroughly assess the rapresentativity of our data, not only in terms of sociodemographic attributes but also from music behavioural patterns comparing against the open access dataset of music learning histories dataset (MLHD). Our results show that moral values are indeed predictable from music preference information and in line with the findings of the related literature. Further, we discuss the most predictive music behaviours, contributing to an in-depth understanding of the moral profiles. Such insights are fundamental to the broader picture since moral values are a key element in the decision making process on several societal issues [11, **?**]. Modelling moral values from music represents a great opportunity for improving recommendation systems; designing online streaming applications with user well-being in focus [18]; increasing engagement to communication campaigns for social good applications.

2 Data Collection and Feature Engineering

Here, we employ data from a third-party survey administered online for a general scope marketing project. The survey consists of 2,003 participants (51% females) from 12 different regions in Canada. The participants filled in, among other items, information about basic demographic attributes, including age, gender, education, and political views. They also completed the validated Moral Foundations questionnaire [9], while stated their preferences on 13 music genres (on a 5-point Likert scale where 1 = strongly dislike and 5 = strongly like). The considered music genres were: alternative pop/rock, christian, classical, country, folk, heavy metal, rap/hip-hop, jazz, latin, pop, punk, R&B, and rock. These genres were set from the survey creators and were not further described to the respondents. Even so, they are commonly used to define general musical tastes among

Table 1. Summary of the survey dataset (cleaned) with major demographic attributes utilised for this research work.

Attributes	**Demographics**	**Sample size** (N = 1062)
Age	18-24	80 (7.5%)
	25-34	154 (14.5%)
	35-44	205 (19.3%)
	45-54	205 (21.9%)
	55-64	187 (17.6%)
	65+	203 (19.1%)
Gender	Male	474 (44.6%)
	Female	588 (55.3%)
Education	Less than High School	35 (3.2%)
	High school graduate	195 (18.3%)
	Some College	154 (14.5%)
	Trade or professional school	115 (10.8%)
	College Graduate	349 (32.8%)
	Post Graduate work or degree	205 (19.3%))
Political Party	Conservative	328 (30.8%)
	Liberal	279 (26.2%)
	NDI (New Democratic Party)	184 (17.3%)
	Green Party	66 (6.2%)
	Party Quebecois	56 (5.2%)
	I don't vote	149 (14%)

non-musician respondents. To justify these genres and observe if there is any affiliation between survey reported preferences and digital music listening patterns, we explored digital data of 1062 Canadian listeners extracted from the Music Learning Histories Dataset (MLHD) [22] with a similar age and gender distribution to our survey.

Moving on to our survey data, to make sure that participants were paying attention to the survey questions, two "catch questions" were included, which we later used to filter the data. After excluding these users we were left with 1,062 participants (55% females), a sample size substantially higher than previous survey-based studies [7, 6]. Table 1 summarises the demographic features of our dataset.

We then applied a factor analysis using principal axis factoring with promax rotation to identify the major dimensions of participants' music preferences. A 5-factor solution was retained, which explained 67% of total data variance: {jazz, classical, latin}, {punk, heavy metal, rap/hip-hop}, {pop, R&B}, {country, Christian, folk}, and {rock, alternative pop/rock} (genres ordered by decreasing factor loading). These factors are in line with the ones obtained in related studies [7].

To quantify the respondents' diversity in music preferences, we employed an adapted version of the generalist-specialist (GS) score, inspired by the work of Anderson et al. [1]. The projections of the 13 genres onto the five factors were considered as that genre's vector representation in the "preference space". Intuitively, generalists versus specialists

Table 2. Detailed list of the experiments we performed with the list of features employed as predictors in each one of them.

ID	Features Employed as Predictors
EX1	13 Music Genres
EX2	5 factors
EX3	GS score
EX4	13 Music Genres, Age, Gender
EX5	13 Music Genres, Age, Gender, Education
EX6	13 Music Genres, Age, Gender, Education, Political Views

will have genre vectors spread apart versus close together in the preference space. We calculate the user centroid $\overrightarrow{ct_i}$ of genre vectors representing the loadings of genres on the 5 factors $\overrightarrow{l_j}$, weighted by the number of genre scores rated by each respondent w_j. The GS score is the cosine similarity between a genre vector and the preference-weighted average of a users' genre vectors:

$$GS(u_i) = \frac{1}{\sum w_j} \cdot \sum w_j \frac{\overrightarrow{l_j} \cdot \overrightarrow{ct_i}}{\parallel \overrightarrow{l_j} \parallel \cdot \parallel \overrightarrow{ct_i} \parallel} \quad , \qquad \overrightarrow{ct_i} = \frac{1}{\sum w_j} \cdot \sum w_j \overrightarrow{l_j}$$

3 Experiments and Results

Exploratory Analysis. As a first step we assess the correlation between musical genres' preferences, demographics, political views and moral traits. We observed a positive Spearman correlation of age with Christian music, classical, country and folk music genres (ρ_s = {0.18, 0.21, 0.20, 0.25}), while heavy metal, hip-hop/rap and punk were more preferred by younger ages, whereas older people expressed their dislike towards these genres (ρ_s = {-0.22, -0.38, -0.38 }). Education was positively related with classical music, jazz and latin music (ρ_s = {0.22, 0.13, 0.13}), indicating that people with higher education preferred these genres. Loyalty, authority and purity were positively correlated with Christian music (ρ_s = {0.18, 26, 38}) and country music (ρ_s = {0.17, 20, 21}). Looking at the political views of the respondents, conservatives were positively correlated with Christian genre and country (ρ_s = {0.12, 0.12}) and negatively correlated to hip-hop/rap and punk (ρ_s = {-0.17, -0.15}).

Further, we assessed whether the obtained self-reported responses of the questionnaire are in line with digital music listening data. From the MLHD dataset [22] we extracted artists' genres using MusicBrainz identifiers. From the survey data we discerned that the top 10 most preferred genres were: rock, pop, alternative pop-rock, classical, r&b, country, jazz, folk, latin, and hip-hop/rap. Similar trends were encountered in the music listening histories of Canadian users in MLHD where the 10 most frequently listened genres were: rock, alternative rock, pop-rock, pop, electronic, folk, punk, jazz, heavy metal, and hip-hop.

Moral Values Classification. Our main research question is whether we can predict peoples' moral values from their music preferences. To answer this question, we postulate

Table 3. Moral traits classification with XGBoost, average weighted AUROC and standard deviation over 5-fold cross-validation (baseline is .50).

	EX1	EX2	EX3
Care	**.57 (3.7)**	.54 (2.1)	.52 (1.5)
Fairness	**.56 (2.9)**	.52 (1.1)	.48 (2.7)
Authority	**.63 (0.8)**	.60 (1.1)	.49 (1.7)
Purity	**.69 (2.8)**	.65 (3.0)	.57 (2.3)
Loyalty	**.61 (2.4)**	.56 (1.9)	.48 (3.1)
Individ.	**.55 (3.5)**	.51 (0.8)	.50 (1.6)
Binding	**.67 (2.4)**	.63 (2.2)	.52 (1.9)

Table 4. Moral traits classification with XGBoost for different predictors (see Table 2 where they are defined). Models evaluated based on AUROC and standard deviation over 5-fold cross-validation (baseline is 50).

	EX1	EX4	EX5	EX6
Care	.57 (3.7)	.62 (3.2)	.62 (3.0)	**.63 (2.3)**
Fairness	.56 (2.9)	.58 (2.5)	.57 (2.3)	**.62 (4.3)**
Authority	.63 (0.8)	.64 (1.6)	.65 (2.0)	**.66 (1.6)**
Purity	.69 (2.8)	.71 (3.0)	**.71 (1.4)**	.71 (1.6)
Loyalty	.61 (2.4)	**.67 (3.5)**	.66 (2.2)	.66 (2.9)
Individ.	.55 (3.5)	.59 (2.4)	.59 (3.3)	**.61 (1.8)**
Binding	.67 (2.4)	.71 (3.2)	.70 (2.2)	**.72 (2.9)**

the task as a supervised classification one, developing a series of experiments to assess the predictive power of different variables (see Table 2). We assign the class label "high" to individuals with moral scores higher than the population median for the specific foundation, and "low", otherwise. We perform 5-fold cross-validation on shuffled data (to avoid dependencies in successive data points), with 70% of training and 30% testing data. We opt for the gradient boosting algorithm XGBoost (XGB) as it performed better than Random Forest (RF) and Support Vector Machine (SVM) in this task.

To take into account the effect of unbalanced class labels in the performance metric, we evaluate our models with the area under the receiver operating characteristic (AUROC) metric which is a performance measure for binary classifiers that employs a discrimination threshold to differentiate between a high and a low class [12]. The best model is then chosen as the one that maximized the weighted area under receiver operating characteristic (AUROC) statistic.

Initially, we compared the predictive power of the genre information against the features engineered by us (EX1, EX2, and EX3). We trained one model per moral foundation, and we present the cross validated results in Table 3. We notice that the information obtained directly about the music preferences (EX1) outperforms the features we developed. When comparing the scenarios, we observe that the 5 factors, and the GS score accounting only for part of the variance in the data, did not manage to outperform the explicit information on music preferences. A question that emerges naturally, is

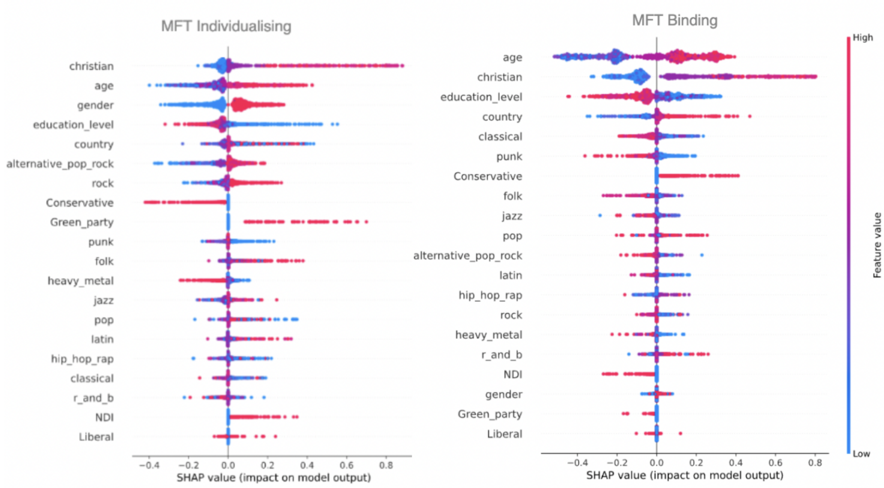

Fig. 1. Feature contributions (via SHAP values). The higher the SHAP value, the more the feature contributes to the moral prediction.

.

whether including knowledge regarding the participants' basic demographic features (i.e. age, political views, education level) will improve the prediction of their moral values. Table 4 summarises the results when age, gender, education and political views are incorporated in the design. As expected, the more information we have about the participants the more precise our predictions become, however, the improvement is minimum. This shows us the importance of music behaviours alone in explaining the variability of our moral values.

Further, we employed SHAP (SHapley Additive exPlanations), a game theory approach developed to explain the contribution of each feature to the final output of any machine learning model [15]. SHAP values provide both global and local interpretability, meaning that we can assess both how much each predictor and each observation, respectively, contribute to the performance of the classifier. SHAP's output helps to understand the general behaviour of our model by assessing the impact of each input feature in the final decision, thus enhancing the usefulness of our framework (Figure 1).

Moral Values Regression. Data binning is a common way to aggregate information and facilitate the classification tasks. However, there are known issues to dichotomisation of variables which often lead to misleading results [16]. Here, to ensure that the most predictive features as emerged from the classification process are indeed descriptive of the respective moral trait, we conducted a regression analysis. At this point, the aim is to understand whether we can estimate the original moral scores (predicting the quantity) based on our explanatory variables in disposition (i.e., music genres ad demographics).

To do so, we trained an XGBoost Regressor for each moral foundation. We maintained the same experimental designs and settings as in the classification task. For

Table 5. Mean Absolute Error (MAE) and standard deviation over 5-fold cross-validation for XGBoost regression on music preference features (see Table 2).

	EX1	EX2	EX3
Care	3.86 (13.2)	**3.72 (10.9)**	3.89 (7.0)
Fairness	**3.27 (11.1)**	3.28 (9.6)	3.55 (8.7)
Authority	**4.19 (23.3)**	4.20 (16.7)	4.47 (13.9)
Purity	**4.86 (19.7)**	4.99 (25.0)	5.35 (21.0)
Loyalty	4.46 (12.1)	**4.33 (19.4)**	4.64 (11.6)
Individ.	3.23 (9.5)	**3.17 (8.5)**	3.35 (9.9)
Binding	3.86 (15.1)	**3.79 (6.3)**	4.22 (18.5)

Table 6. Mean Absolute Error (MAE) and standard deviation over 5-fold cross-validation for XGBoost regression on music preference and demographic features (see Table 2).

	EX1	EX4	EX5	EX6
Care	3.86 (13.2)	3.72 (6.2)	3.71 (9.3)	**3.60 (8.0)**
Fairness	3.27 (11.1)	3.25 (8.2)	3.19 (10.5)	**3.12 (13.4)**
Authority	4.19 (23.3)	4.14 (15.9)	4.10 (9.3)	**4.09 (11.0)**
Purity	4.86 (19.7)	4.86 (20.4)	4.74 (18.9)	**4.71 (15.8)**
Loyalty	4.46 (12.1)	**4.19 (22.8)**	4.20 (18.7)	4.21 (18.7)
Individ.	3.23 (9.5)	3.17 (14.4)	3.17 (8.0)	**3.0 (8.9)**
Binding	3.86 (15.1)	3.80 (11.0)	3.76 (13.1)	**3.74 (5.4)**

evaluation, we used Mean Absolute Error (MAE). These options allow for a direct comparison of the most predictive features with the ones emerged from the classification task (Table 5). We noticed that as in the classification task, when adding information to the models the MAE decreases indicating that the model fits the data better. Also in this case the gain of adding more information is relatively small with respect to the music genres alone.

We visualised the most predictive features using again the Shap values (see Figure 2). Interestingly, the christian music genre appears again as the most important predictor for both the Binding and Individualising traits. The feature importance for the output of the XGboost regressor, is in line with the feature significance obtained with the classification approach. The same holds for all the moral foundations which are not depicted here for spacing issues.

4 Discussions and Conclusions

Henry Wadsworth Longfellow wrote, "Music is the universal language of mankind." Contemporary research has found converging evidence that people listen to music that reflects their psychological traits and needs and help express emotions, cultures, values and personalities. In this paper, we analysed the less explored links between musical preferences, demographics (age, gender, political views, and education level) and Moral

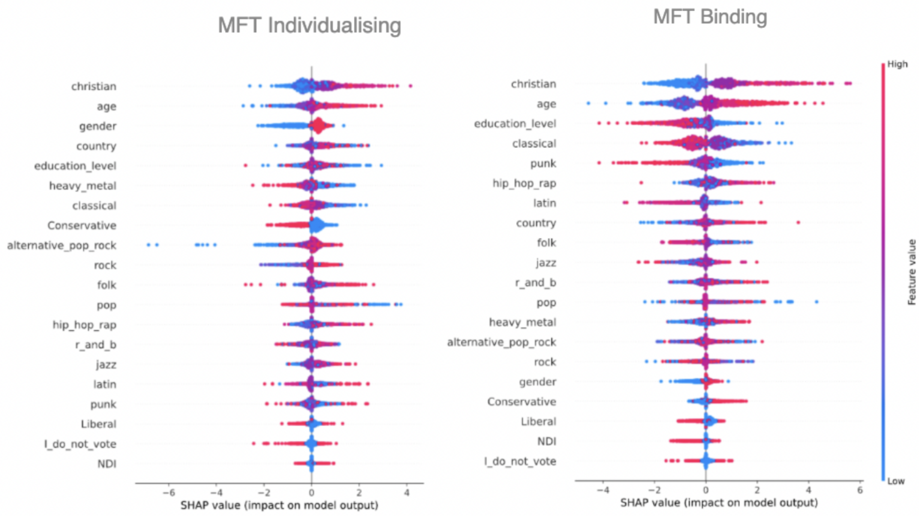

Fig. 2. Features with the most impact on the XGBoost regression model output.

Foundations (MFT [9]). We applied both classification and regression models for moral traits prediction. From classification results, it was inferred that MFT Binding was best predicted with AUROC score 72%, whereas MFT individualising showed weaker results with AUROC score 61%. While, for the regression task the lowest MAE was 3.0 for the Individualising and 3.74 for Social Binding. In both approaches, the most impactful features on inferring morality were christian music and age.

Moral foundations are strongly tied to political views; despite that, the musical features are more predictive than political leanings. Social binding is related to conservative political views [8] - and in fact is predicted by christian, and country music. We notice that people naturally express their moral values through the music they listen to. We instinctively *categorize* objects, symbols, but also people, creating a notion of *social identity*. According to the social identity theory members of a group will seek to find negative aspects to other groups thus enhancing their self-image [21]. Such reasoning reflects on a broad range of attitudes related to stereotype formations [19] but also as we notice here to musical preferences. For instance, people higher in social binding foundations tend to listen to country music which often expresses notions of patriotism. Christian music is also a predictor of this superior foundation, which again fosters the notion of belonging to a group. Across all experiments, Christian music emerged as the most predictive genre. On the other hand, genres such as punk, and hip hop are known to challenge the traditional values and the status quo, hence are preferred by people who strongly value these aspects. Our findings suggest that musical preferences are quite informative of deeper psychological attributes; still there is space for improvement. For instance, we noticed that the care, fairness, and loyalty foundations are harder to predict. To this end we aim to explore musical content analysis, for instance, incorporating linguistic cues, and the moral valence scores as proposed by Araque et al. [4, 5] on lyrics to further improve the performance.

In future work we aim to delve deeper into the relation between music and morality, and between music and other universal human values, by using passively collected digital traits of music listening behaviours outside a laboratory setting and over a period of time [2], while using self-reported surveys as a solid groundtruth. We will further investigate the association between music listening preferences other psychological aspects such as human values and emotions. Developing data-informed models will help unlock the potential of personalised, uniquely tailored digital music experiences and communication strategies [12, 1]. Predicting the moral values from listening behaviours can provide noninvasive insights on the values or other psychological aspects of populations at a large scale.

5 Acknowledgements

This work was supported by the QMUL Centre for Doctoral Training in Data-informed Audience-centric Media Engineering (2021–2025) as part of a PhD studentship awarded to VP. KK acknowledges support from the "Lagrange Project" of the ISI Foundation, funded by the CRT Foundation. We would like to thank Dr. Robert Raleigh for providing the survey data, and the two anonymous reviewers for their thoughtful comments.

References

1. Anderson, A., Maystre, L., Anderson, I., Mehrotra, R., Lalmas, M.: Algorithmic effects on the diversity of consumption on spotify. In: Proc. Web Conf. 2020. pp. 2155–2165 (2020)
2. Anderson, I., Gil, S., Gibson, C., Wolf, S., Shapiro, W., Semerci, O., Greenberg, D.M.: "just the way you are": Linking music listening on spotify and personality. Soc. Psychol. Personal. Sci. **12**(4), 561–572 (2021)
3. Ansani, A., D'Errico, F., Poggi, I.: 'you will be judged by the music i hear': A study on the influence of music on moral judgement. In: Web Intel. vol. 17, pp. 53–62 (2019)
4. Araque, O., Gatti, L., Kalimeri, K.: Moralstrength: Exploiting a moral lexicon and embedding similarity for moral foundations prediction. Knowl. Based Syst. **191**, 105184 (2020)
5. Araque, O., Gatti, L., Kalimeri, K.: The language of liberty: A preliminary study. In: Companion Proc. Web Conf. 2021. pp. 623–626 (2021)
6. Devenport, S.P., North, A.C.: Predicting musical taste: Relationships with personality aspects and political orientation. Psychol. Music **47**(6), 834–847 (2019)
7. Gardikiotis, A., Baltzis, A.: 'rock music for myself and justice to the world!': Musical identity, values, and music preferences. Psychol. Music **40**(2), 143–163 (2012)
8. Graham, J., Haidt, J., Nosek, B.A.: Liberals and conservatives rely on different sets of moral foundations. J. Personal. Soc. Psychol. **96**(5), 1029–1044 (2009)
9. Graham, J., Nosek, B., Haidt, J., Iyer, R., Koleva, S., Ditto, P.H.: Mapping the moral domain. J. Personal. Soc. Psychol. **101**(2), 366–385 (2011)
10. Greenberg, D.M., Müllensiefen, D., Lamb, M.E., Rentfrow, P.J.: Personality predicts musical sophistication. J. Res. Pers. **58**, 154–158 (2015)
11. Haidt, J., Joseph, C.: Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. Daedalus **133**(4), 55–66 (2004)

12. Kalimeri, K., Beiró, M.G., Delfino, M., Raleigh, R., Cattuto, C.: Predicting demographics, moral foundations, and human values from digital behaviours. Comput. Hum. Behav. **92**, 428–445 (2019)
13. Krismayer, T., Schedl, M., Knees, P., Rabiser, R.: Predicting user demographics from music listening information. Multimed. Tools. Appl. **78**(3), 2897–2920 (2019)
14. Loersch, C., Arbuckle, N.L.: Unraveling the mystery of music: Music as an evolved group process. J. Personal. Soc. Psychol. **105**(5), 777–798 (2013)
15. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Proc. 31st Int. Conf. Neural. Inf. Process. Syst. pp. 4768–4777 (2017)
16. MacCallum, R.C., Zhang, S., Preacher, K.J., Rucker, D.D.: On the practice of dichotomization of quantitative variables. Psychol. Methods **7**(1), 19–40 (2002)
17. MacDonald, R., Kreutz, G., Mitchell, L.: Music, health, and wellbeing. Oxford University Press (2013)
18. Mejova, Y., Kalimeri, K.: Effect of values and technology use on exercise: implications for personalized behavior change interventions. In: Proc. ACM Conf. User Model. Adapt. Personaliz. (UMAP) 2019. pp. 36–45 (2019)
19. Miller, S.L., Maner, J.K., Becker, D.V.: Self-protective biases in group categorization: Threat cues shape the psychological boundary between “us” and “them”. J. Personal. Soc. Psychol. **99**(1), 62–77 (2010)
20. Nave, G., Minxha, J., Greenberg, D.M., Kosinski, M., Stillwell, D., Rentfrow, J.: Musical preferences predict personality: evidence from active listening and facebook likes. Psychol. Sci. **29**(7), 1145–1158 (2018)
21. Tajfel, H., Turner, J.C., Austin, W.G., Worchel, S.: An integrative theory of intergroup conflict. Org. Id.: A reader **56**(65), 33–47 (1979)
22. Vigliensoni, G., Fujinaga, I.: The music listening histories dataset. In: Proc. Int. Soc. Music Inf. Retr. Conf. (ISMIR). pp. 96–102 (2017)