Know your big data: De-biasing subsamples of large datasets for personality research
using importance sampling and kNN matching

Fricke, Kai R.

Department of Personality Psychology and Psychological Assessment,
Helmut-Schmidt-University/University of the German Federal Armed Forces Hamburg,
Hamburg, Germany

Herzberg, Philipp Yorck

Department of Personality Psychology and Psychological Assessment,
Helmut-Schmidt-University/University of the German Federal Armed Forces Hamburg,
Hamburg, Germany

Author Note

Corresponding author: Kai R. Fricke <fricke@hsu-hh.de>

Abstract

Large behavioral datasets collect ecologically valid data on user behavior, but only seldom provide more data on the users. To gain further insights on the participants, small subsamples can be collected, assessing variables of interest, such as personality characteristics. However, such samples are often biased through self-selection of participation. Removal of this bias is often hard because no demographic data is present and only distributions of (multiple) continuous target variables in the population are known from the large dataset. In two studies, we examined de-biasing methods with respect to arbitrary continuous variables. In the first study, we examined an artificial dataset with full information on the population and found that sequential importance sampling and single-nearest-neighbour matching proved successful in removing bias from subsamples. In the second study, we took this method to practice and examined the relationship of personality characteristics and music preferences in a sample of Spotify users with respect to the musical preference distributions of one million users of the Million Song Dataset. Our results show that sequential importance sampling is a promising way to remove bias from samples when the distribution of relevant variables in the population is known.

Keywords: importance sampling, nearest neighbour matching, big data, personality psychology

Know your big data: De-biasing subsamples of large datasets for personality research using importance sampling and kNN matching

Introduction

The collection of huge digital data sets on user behavior provides great opportunities to examine individual differences between these users. However, current research on big data seldom relates their results to personality characteristics, mostly because such data is either not collected, or not freely available for researchers. Deeper knowledge of the biographic and psychological characteristics of the users included in such large datasets is needed to generalize insights to the general population. But how can we gain these insights if the data is not available?

Large behavioral data sets are utilized for the improvement of methods, such as speech recognition, or natural language generation, as well as predictive analytics, i.e. predicting certain outcome variables (such as likelihood of sales) from various input variables (Gandomi & Haider, 2015). Research with personality as the outcome variable (i.e. on the prediction of personality characteristics from digital traces) was largely guided by research on the myPersonality dataset (Kosinski, Stillwell, & Graepel, 2013). In this dataset, tens of thousands of users provided insights on their demographic information, their personality profiles (obtained using established instruments), and their digital behavior (i.e. data produced within the Facebook social network service). The advantages of this dataset were compelling: The access to a large body of behavioral data, preference statements ("likes"), and related personality data from psychological instruments gave host to much research and intriguing insights. For instance, it was shown that personality predictions from Facebook likes achieved the same average accuracy as personality ratings from family members with just 100 likes (Youyou, Kosinski, & Stillwell, 2015).

The significance of insights gained from the myPersonality dataset cannot be overstated. However, data collection for the dataset ceased in 2012. Further, it is currently not available for new research projects. In the last seven years, digital behavior changed. While Facebook remains tremendously popular with over 1.5 billion

daily active users (Facebook Inc., 2018), some market researchers suggested that user interaction with the Facebook service declined, as did the willingness to create and share personal data (Perrin, 2018; Peters, 2018).

However, there are many datasets available that did not collect any data on the user personality at all. For instance, there is behavioral data available on music listening (Bertin-Mahieux, Ellis, Whitman, & Lamere, 2011; Vigliensoni & Fujinaga, 2018), movie ratings (Harper & Konstan, 2016), or millions of user comments of the Reddit service (Baumgartner, 2015). These datasets provide a huge amount of data points, but don't reveal much on the user behind the service. For example, the *Million Song Dataset* (Bertin-Mahieux et al., 2011) provides song listening events for over one million users, but no further data other than an anonymous user ID. The *Music Listening Histories Dataset* (Vigliensoni & Fujinaga, 2018) provides over 27 billion listening events, but only includes age, gender, and location (when stated by the user) in the dataset. Thus, it is currently often not possible to relate these ecological valid behavioral datasets to individual differences, or more so personality characteristics.

There is much value to be gained from examining these datasets from a personality research perspective. One way to gain additional insights on a large sample is to draw a subsample, assess variables of interest (e.g. personality characteristics) on this sample, and then infer these insights to the larger sample. In a way, this approach is similar to the regular research process aiming to generate knowledge about the (national) population. Such smaller samples are, however, biased such that only a (self-)selected subset of users answers these personal questions. This bias has to be accounted for in analysis. However, the data available to remove such biases (e.g. demographic variables) is often not present in the large data set. In that case, other methods to remove the bias have to be reviewed: Since we know about the distribution of target variables (e.g. preferences) in the large sample, we can try to remove the bias from our subsample with respect to these variables.

This research examines methods to conduct such analyses. In the first of two studies, we create an artificial dataset, from which we draw a biased sample and review

methods of de-biasing this sample. We then show that at least one of the methods is able to create valid insights on the artificial population. A notable difference to other de-biasing methods is that we deal with multiple continuous target variables which are to be balanced – most balancing methods only deal with binary or categorical variables. In the second study we take this method to practice as we apply it to a newly collected sample of users of the Spotify music streaming service. We use this sample to gain insights on the demographic and personality characteristics of the one million users reflected in the Million Song Dataset, for which no such data is available at this time (Bertin-Mahieux et al., 2011; Fricke, Greenberg, Rentfrow, & Herzberg, 2019).

Big data in personality research

Fortunately, some large datasets were specifically created to examine personality characteristics. The myPersonality dataset (Kosinski et al., 2013) originated from an eponymous application for the Facebook social network, created in 2007. Until 2012, over six million people participated in various questionnaires on the web application, and 575,380 unique participants provided their results to 100 items of the IPIP personality item pool (Goldberg et al., 2006; Stillwell & Kosinski, 2012; Johnson, 2014). Because of its vast size and large amount of scales available for researchers, numerous articles based on the myPersonality dataset have been published. For instance, it was found that personality traits can be predicted from Facebook likes (Kosinski et al., 2013), that Facebook likes are more accurate in personality predictions than ratings from close friends and relatives (Youyou et al., 2015), and that written text from Facebook comments predicts personality characteristics (Park et al., 2015). Compared to other datasets the main advantage of the myPersonality dataset is the combination of ecologically valid behavior clues (e.g. Facebook likes, comments, profile pictures, friend networks, and alike) and explicit personality statements and questionnaire responses. No other datasets of this size that we know of provide such integrated information. However, data collection for the myPersonality dataset ceased in 2012, and since 2018 the dataset is not shared to researchers for new research projects anymore.

Other large personality datasets exist. Two datasets by Johnson (2014) collected personality data from close to one million users total: The IPIP-NEO measure of 300 IPIP items has been completed by over 300,000 participants, and a shorter 120 subset has been completed by over 600,000 participants. While these scales provide an accurate measurement of personality, only the demographic variables sex, age and country of origin were assessed additionally.

From 2009 to 2011, 588,014 UK citizens participated in an online survey assessing self-ratings on various personality-related statements (the *BBC* dataset) (Rentfrow, Jokela, & Lamb, 2015). The resulting dataset is freely available, and was for instance used to examine regional personality differences in Great Britain (Rentfrow et al., 2015). Similarly, de Vries, Gosling, and Potter (2011) examined regional differences of Agreeableness in the United States, using data from the ongoing Gosling–Potter Internet Personality Project with over 650,000 participants (Srivastava, John, Gosling, & Potter, 2003). This non-exhaustive list of available datasets shows that most datasets assess personality profiles in addition to demographic variables, providing a good basis for structural analysis of personality characteristics. For instance, Gerlach, Farb, Revelle, and Amaral (2018) examined the myPersonality dataset, two of the Johnson datasets, and the BBC dataset and found four robust personality types reflected in each of the datasets. However, with the exception of the myPersonality dataset, most datasets are restricted to a specific set of assessed variables. With a lack of linked behavioral data, no insights on the relationship to other data or distant constructs can be examined.

However, many datasets are available that collect behavioral data on a host of different subjects, but in turn don't collect much information on the user. In the case of the Million Song Dataset, the authors specifically stripped the dataset such that the highest degree of anonymity is preserved, and no links to other variables, even demographic, can be made (Bertin-Mahieux et al., 2011).

Large datasets are compelling because they reflect a large part of the population. When variables of interest are missing from these datasets, an approach to fill these

voids is to collect this data from a small subsample, and infer from the results to the larger sample. This method is similar to the regular psychological research approach, in which a sample is drawn to generate insights on the general population. Most studies reflect on their generalizability, often discussing the sample collection procedure and the distribution of demographic variables in contrast to those of the general population. However, if such data is not available in the large dataset, we need another approach to examine the generalizability of our findings, or possibly even correct for the (sub)sample selection bias.

Sample selection bias

In most cases when a sample from a population is drawn, the resulting sample is biased towards specific variables, reflecting which part of the population is more likely to take part in the assessment. For instance, online assessments are biased towards users who regularly browse the web and are interested in specific topics, and thus are more likely to be exposed to adverts for such assessments. In overt assessments, samples are generally biased towards people who are willing to share their data. These self-selection biases can lead to non-representative samples of the population: For instance, Marcus and Schütz (2005) found that non-responders to a survey request for website owners were rated to have less Extraversion and Openness than responders. Similarly, Rogelberg et al. (2003) found non-responders to an organizational survey to be less agreeable and conscientious. Fan and Yan (2010) pointed out that participant self-selection appears at various stages, and that many factors might influence the self-selection bias, such as length of the study, participation incentives, and reminder mails, whose influence might also be related to specific personality characteristics of the respondent.

Dealing with imbalance. Most psychological studies report on their demographics and discuss the generalizability of their findings. Matching strategies are often employed when sample sizes are small, e.g. when examining rare diseases. Other strategies that deal with the imbalance between a minority and a majority class use

techniques such as oversampling, in which entities of the minority class are sampled more than once to reach balance. The Synthetic Minority Over-sampling Technique (SMOTE) (Chawla, Bowyer, Hall, & Kegelmeyer, 2002) is similar, but creates artificial (synthetic) sample points to avoid overfitting of classification algorithms to cofounder variables. The ADASYN algorithm (He, Bai, Garcia, & Li, 2008) builds on the SMOTE algorithm and extends it with specific weights such that samples that are harder to learn are represented more. While these algorithms were originally developed to reduce bias in test-training-sample settings (i.e. binary classifications), they have been extended to deal with multiple categories, and also with continuous variables (e.g. SMOTER; Torgo, Ribeiro, Pfahringer, & Branco, 2013).

Another possibility to achieve balance between two classes is to undersample the majority or the minority class, i.e. leaving out data points in order to gain balance. Undersampling the majority class alters the distribution in the target variables and is thus not suitable for our task of obtaining information on the population. Undersampling the minority class thins out the data on which our results are based and is thus only suitable when dealing with large subsamples.

Aim of the study. In our problem domain we often deal with missing information on demographic variables, but usually have information on other characteristics, such as preferences. These characteristics are often represented by continuous variables, e.g. a preference indicator, or likelihoods. Thus, for our use case we need to achieve balance on multiple continuous variables. We consider two main approaches for de-biasing strategies: First, we try importance sampling, in which we sample some cases of the biased dataset more than once so that we achieve a similar distribution in our target variables as in the main dataset. Since importance sampling depends on the weights that are given to the biased examples, we compare four methods of generating these weights: Single variable distribution ratio, sequential variable distribution ratios, combined variable distribution ratios, and direct density ratio estimation. Second, we use a nearest-neighbour approach. In the first approach, we sample the single nearest neighbour for each instance of the main dataset from the

sub dataset. In the second approach, we examine the k-nearest-neighbours (kNN) and create synthetic samples of the minority class. This approach is inspired by the SMOTER algorithm, but omits the need for an importance threshold hyperparameter. Instead, a synthetic sample is created for each instance of the main dataset.

Study 1: Simulation study with full information

In the first study, we created an artificial dataset to compare different methods of importance sampling, and to show that the technique leads to interpretable results concerning the population. To that end, the dataset contained multiple variables, of which some were correlated and some independent. From this population, we drew a biased sample. This biased sample was then administered to four oversampling methods and two nearest-neighbour matching methods, leading to new datasets in which the bias was supposedly removed. Since in the simulation we have access to the true distributions of the variables in the population, we can estimate how much our sampling method leads to correct estimates of the population distribution and hence removal of the bias.

Method

Artificial dataset. The artificial dataset contains four Gaussian-distributed continuous variables, $v_1, v_2, v_3, v_4 \sim \mathcal{N}$. v_2 is correlated to v_1 with $\rho \approx .67$ and v_4 is correlated to v_3 with $\rho \approx .44$. The population correlations resulted from adding variables to each other, i.e. $v_2 \sim \mathcal{N} + 0.9 \cdot v_1$ and $v_4 \sim \mathcal{N} + 0.5 \cdot v_3$. The factors were chosen so that a moderate/large and moderate correlation are present. All other correlations are approximately zero. The first three continuous variables act as the target variables, for which we aim to restore balance.

Three factors exist in the dataset, created from continuous variables, of which f_1 is correlated with v_1 ($\rho \approx .81$) and f_2 is correlated with v_3 ($\rho \approx -.77$). Depending on the random sampling, this leads to approximately 11 factor levels for f_1 and f_2 . f_3 is an independently drawn factor with five levels.

The population size is 10,000 persons. A biased sample of 1,000 cases is drawn from this population, with sample probabilities $p = (5 + v_1)^8$, resulting in a right-shift of the biased sample compared to the population with respect to v_1 . The R script creating and analysing the artificial dataset can be obtained from Appendix A. The R scripts for the analysis can be obtained from the supplemental material.

Importance sampling

To achieve balance with respect to the target variables we can subsample from the biased sample such that a specific distribution is reached. Some samples of the biased sample might end up multiple times in the de-biased sample, while some may be omitted completely. This approach is generally known as weighted random oversampling, or importance sampling.

We examine four importance sampling methods aiming to restore balance to the variables v_1, v_2 , and v_3 . In all cases, we used density ratios to re-weight the sampling probabilities. Intuitively, adjusting the sampling weights to the density ratio of the population D_P and the sample D_S with $DR = \frac{D_P}{D_S}$ removes the bias towards the sample density by dividing and shifts the samples towards the population density. Our methods differ in how the density ratios are estimated.

Single-variable distribution ratio. In the first case, we only seek to restore balance to the v_1 variable. For this, we estimate a one-dimensional density kernel using the bandwidth estimation described by Sheather and Jones (1991). We estimate the density functions for the population and the biased sample. For each datapoint in the sample, we then calculate the likelihood under P and the likelihood under S and calculate the sample weight for this case through division.

Sequential importance sampling. In the second case, we match on all three variables sequentially. Thus, we first sample a de-biased sample with respect to the v_1 variable (like above), from which we sample w.r.t. the v_2 variable, from which we sample w.r.t. the v_3 variable.

Combined distribution ratios. In the third case, we use one-dimensional density kernels for each of the variables, D_{p1}, D_{p2}, D_{p3} and D_{s1}, D_{s2}, D_{s3} . We then

calculate the density ratios for each of the pairs and compute the product. The sample weight for each data point is thus given by $\lambda(x) = \frac{D_{p1} \cdot D_{p2} \cdot D_{p3}}{D_{s1} \cdot D_{s2} \cdot D_{s3}}$.

Direct (3D) density ratio estimation. In the last case, we use direct density ratio estimation, a method which has been shown to be applicable to importance estimation in shift adaptation before (Sugiyama, Nakajima, Kashima, Buenau, & Kawanabe, 2008). For our analyses, we use the implementation from the `densratio` R package (Makiyama, 2019). Here, the density ratio is directly estimated w.r.t. to the three continuous matching variables, v_1 , v_2 , and v_3 .

Nearest-neighbour sampling

We also examined two methods of nearest-neighbour sampling. Both methods calculate the k-nearest-neighbours (kNN) of each case in the population set from the biased set. In other words, for each case in the large sample, the closest matches in the biased sample are found. For estimating the distance between two elements, we calculated the euclidean distance.

Single nearest-neighbour matching. In the first method, the single nearest neighbour is sampled. Thus, similar to oversampling, the resulting de-biased sample consists only of real collected examples, but might omit some of these.

Synthetic kNN sampling. In the second method, the k nearest neighbours are sampled, and a weighted average for all variables between these neighbours is then calculated. Thus, the resulting data points are artificial. This procedure follows the idea of SMOTE and SMOTER, in which a similar approach is taken. The relative weights used for averaging the values are calculated as $p = \frac{1}{d+\epsilon}$, where d is the distance between the two elements and ϵ is a small scalar value to avoid division through zero when a perfect match is found. In practice, a perfect match leads to a very high relative weight, returning the matched case with virtually no alteration. For categorical variables, the same probability is used for sampling the factor level from the nearest neighbours. In our experiments, we considered the $k = 10$ nearest neighbours.

Analysis

For our analyses, we used each of the six methods to create a de-biased sample from the biased sample. We then evaluate if the approaches yield a good approximation of the population parameters.

First, we compare the variable distributions between the real dataset, biased dataset, and de-biased datasets visually with distribution plots. For numerical analysis, we rescale the density functions such that the total area sums to 1, and then calculate the Kullback-Leibler divergence, $KL(Q, R) = \int_{-\infty}^{\infty} D_R(x) \cdot \log \frac{D_Q(x)}{D_R(x)} dx$. Because the KL-divergence is asymmetrical, we report the mean KL-divergence of both directions, $KLD = \frac{KL(P,S)+KL(S,P)}{2}$. Further, we calculate the area of the overlay of the two density curves, $A = \int_{-\infty}^{\infty} \min(D_P(x), D_S(x)) dx$. The area is a scalar value between 0 and 1 and represents the percentage of overlap between the two density functions, with 1 being a perfect fit. Lastly, we report the effect size for the mean differences between the datasets. The effect size is calculated using Cohen's d, $d = \frac{\mu_P - \mu_S}{\sqrt{\frac{1}{2} \cdot (\sigma_P^2 + \sigma_S^2)}}$.

We then compare the factor distributions. Again, we calculate the KL-divergence and the area of overlap of the two densities, like in the continuous case, but with the discrete equivalents. For an estimate of the effect size, we calculate Cohen's w, with $w = \sqrt{\frac{\chi^2}{n}}$, where χ^2 is the Chi-squared test statistic.

Lastly, we examine the correlation patterns within the population and the samples. For numerical analysis, we first transform the correlation between two variables into z values using Fisher's z transformation $z = \frac{1}{2} \cdot (\log(1 + r) - \log(1 - r))$, and then take the difference, resulting in Cohen's q with $q = z_P - z_S$.

Please note that we do not report the probability statistics (i.e. significance levels) from mean difference tests, distribution tests, or the χ^2 tests, as these are highly dependent on the sample size, which we can artificially inflate or deflate. Instead we report effect sizes so that differences can be interpreted in terms of relative magnitude rather than (bogus) statistical significance.

The importance sampling methods rely on random sampling and thus return different results each run. To estimate the variance in our results, we ran bootstrapped

the samples $n = 100$ times and report the mean and SD of the results. Since our biased sample is fixed, the nearest-neighbours remain the same, and hence no variance between different iterations exists.

Results

Continuous variable distributions. In the first variable v_1 , the biased dataset shows a large divergence from the original sample ($KLD = .81$), a moderate overlap ($A = .54$) and a large effect size for the mean difference ($d = -1.26$). All six proposed de-biasing methods improved upon all metrics, as seen in Table 1. The single variable distribution ratio, sequential importance sampling, and single nearest-neighbour matching showed the best results: The effect sizes of the mean difference became small (mean $d < -.19$) while the overlap area increased (mean $A > .81$). However, the Kullback-Leibler divergence for the single variable and sequential importance sampling was much less than for the single nearest-neighbour matching method (mean $KLD = .08$ and mean $KLD = .17$ vs. $KLD = .61$).

— please insert Table 1 about here —

For the other variables, almost all methods showed good results. The combined distribution ratios method showed the worst improvement in v_2 with a mean difference effect size of $d = .56$ and an overlap of only $A = .72$. The single variable and sequential importance sampling which performed best for v_1 were slightly worse than single nearest-neighbour and synthetic kNN matching, but still showed small effect sizes for the mean differences, satisfactory overlap and small KL-divergences.

Figure 1 shows a visual comparison of the density plots for the sequential importance sampling, single nearest-neighbour matching, and synthetic kNN matching methods on the v_1 variable. The sequential importance sampling method is best able to produce a similar distribution plot. The single nearest-neighbour method is a bit jagged, which improves with the smoothing from the synthetic kNN matching.

Overall, the combined matching and direct density ratio estimation methods performed worst, while the other methods showed satisfactory de-biasing results in the

continuous variables.

— please insert Figure 1 about here —

Factor variable distributions. Again, the biased sample showed a large KL-divergence ($KLD = .60$), only moderate density overlap ($A = .57$) and a very large effect size for the factor level distributions ($w = 1.21$) in the first factor f_1 . For the other factors, no bias existed, as expected from the construction of the dataset.

Again, all surveyed methods were able to improve the bias metrics on the first factor. The direct density ratio estimation method showed the largest KL-divergence (mean $KLD = .13$, lowest overlap (mean $A = .80$), and largest effect size after de-biasing (mean $w = .46$). The combined density ratios method seemed to introduce some bias to the f_2 factor. For all other methods, the de-biasing showed good improvements with respect to the factor distributions. The results can be obtained from Table 2.

— please insert Table 2 about here —

Correlation differences. In the population sample, the correlation between v_1 and v_2 was $\rho = .67$. A good de-biasing strategy should be able to achieve a similar correlation and a low effect size for the differences between the population correlations and the sample correlations.

The best results with respect to the correlation between v_1 and v_2 were obtained by the single nearest-neighbour and synthetic kNN matching methods, who showed correlations of $r = .66$ and $r = .69$, respectively, with effect sizes of virtually zero, as seen in Table 3. The combined density ratios and direct density estimation methods increased the differences between the correlation coefficients. The single variable and sequential importance sampling methods did not alter the differences between the correlation coefficients.

— please insert Table 3 about here —

For the uncorrelated variables, the single variable sampling, combined density ratios, and single nearest-neighbour matching methods introduced small distortions in the correlation between v_1 and v_4 .

For the correlation between v_3 and v_4 with $\rho = .44$ in the original sample, most methods showed good results, with the exception of the synthetic kNN matching method, which introduced a moderate difference between the correlation coefficients with $r = .66$ and $q = -.32$.

Discussion

The results from the analysis in the artificial dataset showed that all methods were able to reduce the large bias towards high values of v_1 in the biased sample. However, some methods performed better than others. The single variable importance sampling, sequential importance sampling and single nearest-neighbour matching methods showed good results in all evaluations, and did not introduce significant bias towards other variables or relationships. While the synthetic kNN matching was able to perform very well with respect to the factor biases, it introduced bias towards the correlation between v_3 and v_4 and can thus be considered unstable. The combined matching and direct density ratio estimation methods performed worst overall.

The single variable importance sampling performed surprisingly well. We attribute this to the fact that the bias was introduced only with respect to this specific variable. The results might be different if the bias was introduced towards a combination of variables. However, if there is knowledge about the nature of the bias, single variable matching is a straightforward way to de-bias a subsample and should thus be considered for simple models.

The sequential importance sampling method showed good results in all evaluations. Thus, if the bias was introduced with respect to multiple variables, sequential importance sampling is good candidate for a de-biasing strategy.

Another matching method that worked well was single nearest-neighbour matching. The downside of this method is that it needs much computation time and resources, which scales with the size of the population dataset. Since for each case of the population dataset a matching sample is collected, it might make sense to downsample the population sample in cases where it is very large.

Although synthetic k-nearest-neighbour matching performed well in most evaluations, it did not perform better than simple single nearest-neighbour matching. The main advantage of kNN matching might lie in machine learning applications, where synthetic data prevents overfitting towards specific values and single samples. More research on this method is needed should this specific use-case arise in the future.

Overall, we propose sequential importance sampling and single nearest-neighbour matching as promising de-biasing strategies for our purposes of examining factor distributions and correlation patterns in personality research.

Study 2: Matching music preferences from the Million Song Dataset

In the second study, we aimed to utilize our proposed method of examining large datasets on psychological characteristics on a real-world dataset. The Million Song Dataset (*MSD*) (Bertin-Mahieux et al., 2011) is a dataset containing metadata on over one million songs. The MSD has been enhanced with several additional datasets, like the *Taste Profile subset*, which includes music playback events for over one million users. Recent research examined the structure of musical preferences within these listening events and found a three-component structure comprising the dimensions *Arousal*, *Valence*, and *Depth* (Fricke et al., 2019). This three-component structure has been established in previous research (Greenberg et al., 2016; Fricke & Herzberg, 2017; Fricke, Greenberg, Rentfrow, & Herzberg, 2018) and can thus be considered stable across methods and datasets.

The Million Song Dataset Taste Profile subset data is very sparse: Other than a random unique user id and the playback events, no further information, such as age, gender, or country of origin are included in the dataset, not to speak of personality variables. But because of the large amount of users included in the dataset, it can be thought of as a population. Specifically, through the Million Song Dataset we have information on the distribution of musical preferences in a large sample of online music streaming users. But to draw conclusions about the relationships between music preferences and personality within this population, we have to survey a new sample,

which might be biased towards specific personality characteristics, demographics, or taste profiles.

In this study, we surveyed a new sample of online music streaming users and collected their demographic information, personality characteristics, and music playback statistics. We then used the methods reviewed in the first study to de-bias the dataset, i.e. to generate a sample that matched the population in terms of the distribution of the music preference dimensions. After de-biasing, we examined selected hypotheses concerning the relationship between music preferences and personality that have been found in previous research in a reasonably large dataset (see e.g. Fricke & Herzberg, 2017).

Hypotheses

Following previous research (e.g. Fricke & Herzberg, 2017), we examined the following hypotheses which have been pre-registered on the OSF platform¹:

- Older participants prefer low-arousal music
- Older participants prefer high-depth music
- Women prefer high arousal music less than men
- Women prefer high valence music more than men
- Women prefer deep music more than men
- Openness correlates with Depth
- Conscientiousness correlates negatively with Arousal

Method

This survey examined users of the Spotify music streaming service. The Spotify service gives access to an API that enables collection of playback statistics with full user control of the data that is to be shared. Combined with the vast user base of the service, Spotify is an ideal platform for the collection of real-world playback statistics.

¹<https://osf.io/wg4x5>

Participants. We collected data from a custom survey tool hosted on the lead author's website. The survey has been distributed to music communities, research communities, and has been posted to participant recruitment sites. In three months, 418 users participated, of which 345 provided their music playback statistics. Of those 345 participants, 20 have been excluded because they rated their most listened songs as not reflective of their general music taste.

Participants were on average 22.6 years old ($SD = 5.00$). 178 (54.8 %) were from the US, 29 (8.9 %) from Canada, 21 (6.5 %) from the UK, and the rest was distributed to further 38 countries. 171 (52.6 %) participants were female, 144 (44.3 %) male, 3 (.9 %) other, and the rest did not answer this question.

Measures. The survey platform collected the year of birth, sex, and country of residence from the participants. Participants then completed the STOMP-R (Rentfrow & Gosling, 2003), an instrument for the assessment of musical style preferences. The STOMP-R reports self-stated music genre preferences for the dimensions Mellow, Unpretentious, Sophisticated, Intense, and Contemporary. Afterwards, participants completed the Mini IPIP (Donnellan, Oswald, Baird, & Lucas, 2006) as a measure for the Big Five personality dimensions. The Mini IPIP reports scores on four of the five Big Five dimensions (Conscientiousness, Extraversion, Agreeableness, and Neuroticism), and on Imagination, a facet of Openness.

Participants were then asked to connect the survey tool to their Spotify accounts for one-time access to their listening data via the Spotify API. Data collected included account information (account type (free or premium), account country, number of followers), playlist content for the playlists "Your Top Songs 2018" and "Tastebreaker", the 50 most recently listened songs, the 50 most listened songs (divided by the Spotify API in short term, medium term, and long term), and the 50 most listened artists (again divided by term). All data collected via the Spotify API was fully disclosed before the connection was made and could be reviewed afterwards.

On the next page, participants were shown their five most listened to songs and artists (long term), as well as the five most recently listened songs, as provided by the

Spotify API. They were asked to rate how much these songs and artists reflected their listening preferences on a 1 to 7 Likert scale. This way we could examine the validity of these variables as a proxy measurement for music preferences. We only included participants who gave an average rating of 4 to their five most played songs of all time.

After data collection, we collected the top 10 most played songs (long term) of all participants, and downloaded 30 second samples for each song from the Spotify service. Such samples were not available for all songs. We ended up with 14,434 unique songs. We then proceeded to analyze these songs with the ESSENTIA music analysis software (Bogdanov et al., 2013), following previous research (Fricke et al., 2018, 2019). For each participant we then calculated the average of the music features for the top 10 most played songs. On average, 8.7 songs ($SD = 1.6$) of their top songs were available for each user. Lastly, we used the component loadings from Fricke et al. (2019) to calculate the music preference for the Arousal, Valence, and Depth music feature preference dimensions established in previous research (the *AVD* model, see e.g. Greenberg et al., 2016; Fricke et al., 2018). These dimensions capture the preference for musical features, which gives a valid estimate of musical preferences that is untainted by interpretations or knowledge about musical genres.

De-biasing. Following our results from Study 1, we employed two de-biasing strategies: Sequential importance sampling, and single nearest-neighbour matching. The population distributions for the Arousal, Valence, and Depth dimensions were retrieved from the Million Song Dataset Taste Profile subset, following the analyses of Fricke et al. (2019). For the sequential importance sampling, we sampled 1000 cases for the first two iterations and 325 cases for the last iteration to keep the same total number of cases while providing enough stability in the sampling process. For the single-nearest neighbour matching, we used a random subset of 10,000 users of the population. Since this sample has been drawn independently, its characteristics are virtually the same as those of the full dataset.

Again, we used bootstrapping with $n = 100$ for the sequential importance sampling to get average parameter estimates. Demographic statistics are presented with

standard error estimates for the mean instead of standard deviations.

Results

De-biasing. The biased sample showed large mean differences (e.g. $d = -1.31$ for Arousal), a large KL-divergence of $KLD = .97$ and an overlap of $A = .52$, as seen in Table 4. Both de-biasing strategies were able to reduce the mean differences. With the sequential importance sampling, the mean difference in the Arousal dimension remained moderate (mean $d = -.48$). The single nearest-neighbour matching was able to reduce the highest difference in the Depth domain to $d = .24$.

— please insert Table 4 about here —

Sequential importance sampling reduced the KL-divergence and increased the density function overlap area in all dimensions. However, the single nearest-neighbour matching worsened these metrics: The KL-divergences increased ($KLD \geq 1.05$ for all three variables) and the overlap area decreased for the Valence and Depth dimensions ($A \leq .59$ for all three variables). In this dataset, it seems that the single nearest-neighbour matching was unable to remove the bias from the biased sample. Specifically, while the mean differences decreased (which is expected due to the idea of matching close cases), the distribution has been altered so that it does not fit the population distribution. Visually, we could confirm this in Figure 2: Taking the Depth dimension as an example, we see that the sequential importance sampling is able to achieve a good fit to the distribution function, while the single nearest-neighbour approach leads to a jagged multimodal distribution that does not follow a normal distribution. Thus, we decided to proceed our analyses with the de-biased sample from the sequential importance sampling.

— please insert Figure 2 about here —

Descriptive statistics. The de-biased datasets comprised 58.4 % female ($SE = 5.6\%$) and 39.3 % male ($SE = 5.5\%$) participants, indicating a shift towards more females (biased sample 52.6 % and 44.3 %). 61.0 % ($SE = 5.5\%$) were from the US, 9.9 % ($SE = 3.3\%$) from Canada, and 5.6 % ($SE = 3.2\%$) from the UK, suggesting

that the de-biasing resulted in a shift towards English-speaking countries. The mean age was 22.2 years ($SE = 0.5$), which is only slightly younger than in the biased sample (22.6 years).

Relationship between music preferences and personality characteristics.

We examined the relationship between music preferences, demographic variables, and personality characteristics in the de-biased sample.

The hypothesized correlations between age and preference for low-arousal and high-depth music were small and not statistically significant. We found significant correlations of age with Mellow (mean $r = .16$), Sophisticated (mean $r = .21$), and Intense (mean $r = .19$, all $p < .05$). However, since our age range is very narrow (80 % of the participants were between 18 and 28 years old), the validity of these findings is uncertain.

T-tests revealed that men scored higher on all three music preference dimensions. On the Arousal and Valence dimensions, the effect was small (mean $d = .13$ and mean $d = .11$ with $p < .001$), and on the Depth dimension moderate (mean $d = .37$ with $p < .001$). Only the Arousal difference was in accordance with our hypotheses.

Correlations between personality characteristics and music feature preference showed relationships between Arousal and Neuroticism (mean $r = -.25$), Valence and Conscientiousness (mean $r = .19$), Valence and Neuroticism ($r = -.17$), and Valence and Imagination ($r = -.20$, all $p < .05$), as seen in Table 5. The two hypothesized correlations between Openness and Depth and Conscientiousness and Arousal were not significant. However, we found established links between self-reported music genre preference and personality, for instance between Sophisticated and Openness ($r = .35$), and Contemporary and Extraversion ($r = .40$).

— Please insert Table 5 about here —

Discussion

In this study, we examined how a biased subsample of the large Million Song Dataset population can be de-biased with respect to continuous distributions in order to

generate insights about the large population. We showed that sequential importance sampling was able to generate a sample whose distributions of the music preference dimension Arousal, Valence, and Depth resembled that of the population. With low effect sizes for mean differences, low KL-divergences, and a large density function overlap, we consider the de-biasing successful.

The results on the relationship between demographic variables, personality characteristics, and music preferences only confirm one of our hypotheses, oppose two, and did not show any significant effect for the last four. However, it should be noted that results in previous research were not always aligned. For instance, Fricke and Herzberg (2017) also couldn't confirm the correlation of Conscientiousness with Arousal found by Greenberg et al. (2016). For music feature preference and the AVD model, only those two research projects reported on the relationships with personality variables at all. The method of assessment of music feature preferences through the analysis of music playback statistics has never been examined in relation to personality before. Thus, differing results are not necessarily a sign that results of previous or the present study contradict each other, but could be due to methodological differences in the assessment of musical preferences.

An interesting result was that the de-biasing of the subsample lead to a shift towards English-speaking countries, even more than before the de-biasing. Even if we cannot say that the users of the Million Song Dataset are similarly distributed (our large bias towards these countries can't be corrected without a sufficiently diverse sample), we can assume that most users of the MSD showed preference for western-style music which is commonly successful in these countries.

General discussion

The purpose of this research project was to examine how we can gain insights on large behavioral datasets that did not collect data on any other metrics. We showed that collecting a subsample of the dataset for which we assess the variables of interest is a viable way to gain information on the dataset. We showed that we can de-bias the

sample such that biases like a self-selection bias disappear. The de-biased sample then gives a realistic estimate of the properties of the large (population) sample.

We then showed that this method can be used on a real-world dataset to examine characteristics of a large sample of online music streaming service users. Although we could not confirm hypotheses on correlations derived from previous research, we showed that our de-biasing method leads to a good approximation of the population data and can be used to infer insights on the composition of the sample.

A drawback lies in the surveyed methods. Broken down, we only examined two methods, importance sampling and k nearest-neighbour matching, although with different variants. There are probably more methods that could be adjusted to fit our task, and future research should look into adapting or developing other methods for removing bias from smaller samples with respect to the distributions in a population.

A limitation for our second study lies in the fact that the relationships between personality characteristics and music feature preference assessed with the computational analysis of music playback statistics have not been reported in previous research. It is thus hard to see if our de-biasing method lead to the same results as would be found in other studies. Future research could either try to report on these relationships in further samples, or review the validity of this approach of music feature preference assessment. Currently, only the three-component structure has been confirmed (Fricke et al., 2019), and the equivalence at least for the Depth dimension to the original description of Greenberg et al. (2016) can be questioned.

Currently, we only report statistics and relationships within the de-biased sample, and assume that these hold for the population sample. In further studies, classifiers and regressions could be used to create labels and synthetic values for the population sample from the de-biased sample. That way we would fully augment the population dataset with synthetic data, and could then examine if these data can be successfully used for tasks such as predictions and other machine learning tasks. Such experiments can be made with artificial datasets, and different approaches can then be compared, just like in this report.

We hope that future research builds on our findings in the examination of behavioral datasets for personality research. For instance, with existing large datasets on movie ratings, it would be interesting to see if generalizable results on the relationship between movie preferences and personality can be found.

While we propose de-biasing of samples with respect to target variables whose distributions in the population are known as a good approach to enhance the validity and generalizability of research results, this approach is not always feasible. Sometimes, the collected sample could be too small, or too much biased towards specific variables. In those cases, trying to correct for the bias might result in a odd dataset comprising few real cases and many duplicates. In those cases we propose to just report on the differences in the variable distributions and not to try to correct the bias by importance sampling or kNN matching.

The examination of large behavioral datasets is a valuable, yet challenging task for personality research. Our research found the collection and de-biasing of subsamples to be a viable approach to gain insights on these datasets. Sequential importance sampling showed the best de-biasing results in our two experiments. We hope that these results help researchers with the evaluation of large datasets and the creation of samples that resemble real world populations with respect to arbitrary target variables.

References

- Baumgartner, J. M. (2015). *Complete public Reddit comments corpus*. Web page. Retrieved from https://archive.org/details/2015_reddit_comments_corpus (Accessed: 2018-11-22)
- Bertin-Mahieux, T., Ellis, D. P., Whitman, B., & Lamere, P. (2011). The million song dataset. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*.
- Bogdanov, D., Wack, N., Gómez, E., Gulati, S., Herrera, P., Mayor, O., ... Serra, X. (2013). ESSENTIA: An audio analysis library for music information retrieval. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)* (pp. 493–498).
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16, 321–357.
- de Vries, R., Gosling, S., & Potter, J. (2011). Income inequality and personality: Are less equal us states less agreeable? *Social science & medicine*, 72(12), 1978–1985.
- Donnellan, M. B., Oswald, F. L., Baird, B. M., & Lucas, R. E. (2006). The mini-ipp scales: tiny-yet-effective measures of the big five factors of personality. *Psychological assessment*, 18(2), 192.
- Facebook Inc. (2018). *Company Info / Facebook Newsroom*. Web page. Retrieved from <https://newsroom.fb.com/company-info/> (Accessed: 2018-11-22)
- Fan, W., & Yan, Z. (2010). Factors affecting response rates of the web survey: A systematic review. *Computers in human behavior*, 26(2), 132–139.
- Fricke, K. R., Greenberg, D. M., Rentfrow, P. J., & Herzberg, P. Y. (2018). Computer-based music feature analysis mirrors human perception and can be used to measure music preference. *Journal of Research in Personality*, 75, 94–102.
- Fricke, K. R., Greenberg, D. M., Rentfrow, P. J., & Herzberg, P. Y. (2019). Measuring musical preferences from listening behavior: Data from 1 million people and 200,000 songs. *Psychology of Music*. Accepted manuscript.

- Fricke, K. R., & Herzberg, P. Y. (2017). Personality and self-reported preference for music attributes. *Journal of Research in Personality*, 68, 114–123.
- Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and analytics. *International Journal of Information Management*, 35(2), 137–144.
- Gerlach, M., Farb, B., Revelle, W., & Amaral, L. A. N. (2018). A robust data-driven approach identifies four personality types across four large data sets. *Nature Human Behaviour*, 2(10), 735.
- Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., & Gough, H. G. (2006). The international personality item pool and the future of public-domain personality measures. *Journal of Research in personality*, 40(1), 84–96.
- Greenberg, D. M., Kosinski, M., Stillwell, D. J., Monteiro, B. L., Levitin, D. J., & Rentfrow, P. J. (2016). The song is you. *Social Psychological and Personality Science*, 7, 597–605.
- Harper, F. M., & Konstan, J. A. (2016). The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4), 19.
- He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *2008 ieee international joint conference on neural networks (ieee world congress on computational intelligence)* (pp. 1322–1328).
- Johnson, J. A. (2014). Measuring thirty facets of the five factor model with a 120-item public domain inventory: Development of the ipip-neo-120. *Journal of Research in Personality*, 51, 78–89.
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 201218772.
- Makiyama, K. (2019). densratio: Density ratio estimation [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=densratio> (R package version 0.1.1)

- Marcus, B., & Schütz, A. (2005). Who are the people reluctant to participate in research? personality correlates of four different types of nonresponse as inferred from self-and observer ratings. *Journal of personality*, 73(4), 959–984.
- Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J., ... Seligman, M. E. (2015). Automatic personality assessment through social media language. *Journal of personality and social psychology*, 108(6), 934.
- Perrin, A. (2018). *Americans are changing their relationship with facebook*. Web page. Retrieved from <http://www.pewresearch.org/fact-tank/2018/09/05/americans-are-changing-their-relationship-with-facebook/> (Accessed: 2018-11-22)
- Peters, B. (2018). *We analyzed 43 million facebook posts from the top 20,000 brands (new research)*. Web page. Retrieved from <https://blog.bufferapp.com/facebook-marketing-strategy> (Accessed: 2018-11-22)
- Rentfrow, P. J., & Gosling, S. D. (2003). The do re mi's of everyday life: the structure and personality correlates of music preferences. *Journal of Personality and Social Psychology*, 84, 1236–1254.
- Rentfrow, P. J., Jokela, M., & Lamb, M. E. (2015). Regional personality differences in great britain. *PLoS One*, 10(3), e0122245.
- Rogelberg, S. G., Conway, J. M., Sederburg, M. E., Spitzmüller, C., Aziz, S., & Knight, W. E. (2003). Profiling active and passive nonrespondents to an organizational survey. *Journal of Applied Psychology*, 88(6), 1104.
- Sheather, S. J., & Jones, M. C. (1991). A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(3), 683–690.
- Srivastava, S., John, O. P., Gosling, S. D., & Potter, J. (2003). Development of personality in early and middle adulthood: Set like plaster or persistent change? *Journal of personality and social psychology*, 84(5), 1041.
- Stillwell, D. J., & Kosinski, M. (2012). myPersonality project: Example of successful

- utilization of online social networks for large-scale social research. In *Proceedings of the 1st ACM Workshop on Mobile Systems for Computational Social Science (MobiSys)*.
- Sugiyama, M., Nakajima, S., Kashima, H., Buenau, P. V., & Kawanabe, M. (2008). Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in neural information processing systems* (pp. 1433–1440).
- Torgo, L., Ribeiro, R. P., Pfahringer, B., & Branco, P. (2013). Smote for regression. In *Portuguese conference on artificial intelligence* (pp. 378–389).
- Vigliensoni, G., & Fujinaga, I. (2018). The Music Listening Histories Dataset. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)* (pp. 96–102).
- Youyou, W., Kosinski, M., & Stillwell, D. (2015). Computer-based personality judgments are more accurate than those made by humans. *Proceedings of the National Academy of Sciences*, 112, 1036–1040.

Table 1

De-biasing results of the artificial dataset with respect to the continuous variables.

Method	v_1			v_2			v_3			v_4		
	KLD	A	d	KLD	A	d	KLD	A	d	KLD	A	d
Biased sample	.81	.54	-1.26	.34	.69	-.82	.01	.97	-.06	.00	.98	-.02
Single variable	.09 (.04)	.91 (.02)	-.15 (.03)	.08 (.02)	.90 (.02)	-.18 (.03)	.04 (.01)	.90 (.02)	-.07 (.03)	.02 (.01)	.93 (.01)	-.11 (.03)
Sequential sampling	.16 (.05)	.87 (.02)	-.19 (.05)	.05 (.02)	.94 (.02)	-.09 (.05)	.01 (.00)	.97 (.01)	-.00 (.03)	.05 (.02)	.91 (.03)	-.07 (.05)
Combined ratios	.52 (.06)	.74 (.02)	.25 (.03)	.45 (.05)	.72 (.02)	.56 (.03)	.19 (.03)	.80 (.02)	.05 (.03)	.21 (.04)	.77 (.02)	-.09 (.03)
Density ratio	.30 (.03)	.76 (.01)	-.52 (.03)	.12 (.02)	.87 (.01)	-.18 (.03)	.06 (.01)	.88 (.01)	.04 (.03)	.01 (.00)	.96 (.01)	.01 (.03)
Single NN	.61	.81	-.12	.25	.83	-.03	.10	.85	-.01	.15	.80	-.03
Synthetic kNN	.40	.83	-.25	.17	.93	-.05	.03	.92	-.02	.46	.62	-.01

Note. KLD = Kullback-Leibler-Divergence (mean of both directions), A = Area of overlap between the density functions, d = Cohen's d. SEs are given in parantheses for methods that rely on random sampling. Bootstrap n = 1000.

Table 2

De-biasing results of the artificial dataset with respect to the factor variables.

Method	f_1			f_2			f_3		
	KLD	A	w	KLD	A	w	KLD	A	w
Biased sample	.60	.57	1.21	.00	.97	.09	.00	.97	.07
Single variable	.02 (.01)	.92 (.01)	.20 (.03)	.03 (.01)	.89 (.02)	.25 (.03)	.03 (.01)	.90 (.01)	.26 (.03)
Sequential sampling	.05 (.02)	.89 (.02)	.28 (.06)	.03 (.01)	.91 (.02)	.22 (.04)	.03 (.01)	.90 (.02)	.24 (.06)
Combined ratios	.07 (.01)	.86 (.01)	.33 (.03)	.08 (.01)	.85 (.01)	.38 (.03)	.05 (.01)	.87 (.02)	.30 (.03)
Density ratio	.13 (.02)	.80 (.01)	.46 (.03)	.03 (.01)	.92 (.01)	.22 (.02)	.01 (.00)	.96 (.01)	.10 (.03)
Single NN	.03	.90	.23	.03	.90	.22	.02	.91	.24
Synthetic kNN	.03	.89	.25	.02	.94	.17	.01	.94	.14

Note. KLD = Kullback-Leibler-Divergence (mean of both directions), A = Area of overlap between the density functions, w = Cohen's w. SEs are given in parantheses for methods that rely on random sampling. Bootstrap n = 1000.

Table 3

De-biasing results of the artificial dataset with respect to the correlations between continuous variables.

Method	v_1, v_2		v_1, v_3		v_1, v_4		v_2, v_3		v_2, v_4		v_3, v_4	
	r	q	r	q	r	q	r	q	r	q	r	q
Original	.67		.00		.00		.00		-.01		.44	
Biased sample	.59	.13	.02	-.02	.01	-.02	.00	.00	-.00	-.01	.42	.02
Single variable	.59 (.02)	.14 (.03)	-.06 (.03)	.06 (.03)	-.18 (.03)	.17 (.03)	.01 (.03)	-.01 (.03)	-.04 (.03)	.03 (.03)	.38 (.03)	.07 (.03)
Sequential sampling	.58 (.03)	.14 (.05)	-.06 (.05)	.06 (.05)	-.14 (.05)	.14 (.05)	-.01 (.06)	.01 (.06)	-.03 (.06)	.02 (.06)	.37 (.05)	.08 (.05)
Combined ratios	.51 (.02)	.25 (.03)	.03 (.03)	-.03 (.03)	-.18 (.03)	.18 (.03)	.17 (.03)	-.17 (.03)	.06 (.03)	-.07 (.03)	.33 (.03)	.13 (.03)
Density ratio	.52 (.02)	.24 (.03)	-.05 (.03)	.05 (.03)	-.05 (.03)	.04 (.03)	-.02 (.03)	.03 (.03)	-.01 (.03)	.00 (.03)	.32 (.03)	.14 (.03)
Single NN	.66	.02	-.03	.03	-.11	.11	.01	-.01	-.01	-.00	.40	.05
Synthetic kNN	.69	-.02	-.04	.04	-.06	.05	-.03	.03	-.01	.00	.66	-.32

Note. r = correlation coefficient, q = Cohen's q. SEs are given in parantheses for methods that rely on random sampling. Bootstrap n = 1000.

Table 4

De-biasing results in the Spotify users sample.

Method	Arousal			Valence			Depth		
	KLD	A	d	KLD	A	d	KLD	A	d
Biased sample	.97	.52	-1.31	.46	.67	.86	.39	.69	.77
Sequential matching	.46 (.13)	.72 (.04)	-.49 (.09)	.23 (.10)	.83 (.05)	-.12 (.09)	.15 (.11)	.89 (.04)	.15 (.07)
Single NN	1.29	.59	-.08	1.05	.54	.03	1.18	.52	.24

Note. KLD = Kullback-Leibler-Divergence (mean of both directions), A = Area of overlap between the density functions, d = Cohen's d. SEs are given in parantheses for methods that rely on random sampling. Bootstrap n = 1000.

Table 5

Correlation between music preference and personality variables.

	Year of birth	Extraversion	Agreeableness	Conscientiousness	Neuroticism	Openness
Arousal	.02 (.11)	.01 (.10)	.03 (.12)	.04 (.10)	-.25 (.09)	-.03 (.12)
Valence	.07 (.11)	.17 (.10)	-.01 (.10)	.19 (.10)	-.17 (.10)	-.20 (.11)
Depth	.00 (.10)	-.03 (.10)	.05 (.11)	.16 (.10)	-.09 (.10)	-.05 (.11)
Mellow	-.16 (.09)	.03 (.12)	.08 (.09)	-.14 (.11)	-.08 (.11)	.13 (.16)
Unpretentious	.02 (.09)	.22 (.11)	.02 (.10)	.02 (.10)	-.19 (.10)	-.06 (.12)
Sophisticated	-.21 (.09)	-.01 (.12)	.27 (.09)	-.12 (.11)	.02 (.11)	.35 (.09)
Intense	-.19 (.08)	-.08 (.12)	-.04 (.10)	-.06 (.10)	.25 (.10)	.23 (.12)
Contemporary	-.03 (.09)	.40 (.09)	.26 (.10)	-.19 (.11)	.02 (.10)	.16 (.12)

Note. Correlations in bold typeface are significant with $p < .05$.

Figure 1. Comparison of the density function between the population (grey) and the (de-biased) subsamples (black) on the v1 dimension.

Figure 2. Comparison of the density functions between the Million Song Dataset population sample (grey) and the (de-biased) subsamples (black) on the Depth dimension.

Appendix

Artificial dataset creation and analysis script

R-script for the creation of the artificial dataset and biased sample. The analysis scripts can be found in the supplemental material.

```
# R version 3.4.3 Kite-Eating Tree

set.seed(1000) # Seed for population and biased dataset

num.samples <- 10000 # Population size

num.biased <- 1000 # Biased sample size

num.matched <- 1000 # Matched sample size

###

# Create dataset

v1 <- rnorm(num.samples, mean=0, sd=1)
v2 <- rnorm(num.samples, mean=0, sd=1) + 0.9 * v1 # r ~ .67
v3 <- rnorm(num.samples, mean=0, sd=1)
v4 <- rnorm(num.samples, mean=0, sd=1) + 0.5 * v3 # r ~ .44

f1 <- as.factor(abs(round(runif(num.samples, 1, 4) + 2*v1, 0)))
# Correlated to v1 r = ~.81

f2 <- as.factor(abs(round(runif(num.samples, 1, 3) - 2*v3, 0)))
# Correlated to v3 r = ~-.77

f3 <- as.factor(sample(1:5, num.samples, replace=T))

data = data.frame(v1, v2, v3, v4, f1, f2, f3)

# Create biased sample

biased.idx <- sample(1:nrow(data), num.biased, replace=FALSE,
                    prob=(5 + v1)^8)

biased <- data[biased.idx,]
```