Assessing Algorithmic Biases for Musical Version Identification

Furkan Yesiler

furkan.yesiler@upf.edu

Music Technology Group, Pompeu Fabra University

Barcelona, Spain

Joan Serrà

Dolby Laboratories

Barcelona, Spain

Marius Miron

Music Technology Group, Pompeu Fabra University

Barcelona, Spain

Emilia Gómez

Joint Research Centre, European Commission

Sevilla, Spain

ABSTRACT

Version identification (VI) systems now offer accurate and scalable solutions for detecting different renditions of a musical composition, allowing the use of these systems in industrial applications and throughout the wider music ecosystem. Such use can have an important impact on various stakeholders regarding recognition and financial benefits, including how royalties are circulated for digital rights management. In this work, we take a step toward acknowledging this impact and consider VI systems as socio-technical systems rather than isolated technologies. We propose a framework for quantifying performance disparities across 5 systems and 6 relevant side attributes: gender, popularity, country, language, year, and prevalence. We also consider 3 main stakeholders for this particular information retrieval use case: the performing artists of query tracks, those of reference (original) tracks, and the composers. By categorizing the recordings in our dataset using such attributes and stakeholders, we analyze whether the considered VI systems show any implicit biases. We find signs of disparities in identification performance for most of the groups we include in our analyses. We also find that learning- and rule-based systems behave differently for some attributes, which suggests an additional dimension to consider along with accuracy and scalability when evaluating VI systems. Lastly, we share our dataset to encourage VI researchers to take these aspects into account while building new systems.

CCS CONCEPTS

• **Applied computing** → **Sound and music computing**; • **Information systems** → **Information retrieval**.

KEYWORDS

information retrieval, version identification, algorithmic bias

ACM Reference Format:

Furkan Yesiler, Marius Miron, Joan Serrà, and Emilia Gómez. 2022. Assessing Algorithmic Biases for Musical Version Identification. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM '22)*, February 21–25, 2022, Tempe, AZ, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3488560.3498397>

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

WSDM '22, February 21–25, 2022, Tempe, AZ, USA.

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9132-0/22/02.

<https://doi.org/10.1145/3488560.3498397>

1 INTRODUCTION

Version identification (VI) refers to the task of automatically detecting different audio renditions of a musical composition (e.g., cover songs). Considering that versions can change in several musical characteristics (including key, tempo, structure, and recording conditions), VI systems aim to look past such changes and focus on the information that is common between versions [20]. With improved accuracy and scalability, VI systems start being used in commercial applications, ranging from music recognition to music discovery, and such usage can have an impact on various stakeholders in the music ecosystem, especially regarding recognition and financial benefits.

Digital rights management (DRM) for musical recordings is a complex ecosystem, where multiple types of royalties circulate between a large amount of stakeholders [23]. When artists perform a version of an existing composition, they are obliged to pay “mechanical royalties” to the holders of “the composition copyright.” With the help of VI systems, detecting such cases could be automated. However, VI systems cannot match a recording to a composition directly; instead, they match an unknown recording to the known recordings of a composition. Therefore, we can identify three main parties involved in the process: the performing artists of the unknown (or query) recording, the performing artists of the known (or reference) recording, and the composers of the composition. Taking into account the impact that VI systems may have over such artists and composers, they can be considered as socio-technical systems [3, 19] rather than isolated technologies for such use cases.

Fairness and transparency is an emerging field that studies the societal implications of algorithmic systems [14, 15]. Any existing biases in such systems may lead to favor certain individuals or groups over others, which may result in “unfair” outcomes [9]. Although algorithmic decisions and their impact are closely related concepts, measuring fairness is a domain- and context-dependent process, mainly due to the fact that investigating societal impact requires a certain legal, cultural, and ethical framework. However, quantifying algorithmic biases is a useful step for mitigating potential unfair decisions. In music technology research, recent works have addressed potential issues in music recommendation from both individual and group fairness perspectives [5, 17] by studying gender imbalance [8, 22] and playlist diversity [16, 18]. In a recommendation context, a fairness study may investigate whether two potentially impacted groups are subject to the same exposure or not, assuming more exposure leads to larger financial gains. Nonetheless, in a music recognition task like VI, measuring fairness with respect

to such exposure may not be meaningful, as the main goal of VI systems is to correctly detect items that are objectively linked to each other as “versions of the same composition.” However, this does not imply that tasks like VI, which have objectively assigned labels, do not require any studies about fairness or algorithmic biases. For instance, a discrepancy in the identification performance of such VI systems with respect to some characteristics (e.g., demographics) of the involved parties (e.g., performing artists or composers) may put some musicians in a favored position (e.g., financially) compared to others. Therefore, an examination into whether VI systems inherit any form of bias is necessary to study the potential implications of such systems in the music ecosystem.

In this work, we propose a framework for investigating the performance discrepancies of VI systems across a selected set of attributes, in order to explore their impact on relevant parties. Although our main motivation is to raise awareness for developing fair VI systems, we cannot reach any conclusions on the fairness qualities of existing systems since such fairness measurements are context-dependent and we do not have access to a system that is interacting with the parties we consider. Instead, we aim to quantify the algorithmic bias by investigating the discrepancies in identification performances of 5 state-of-the-art VI systems of different characteristics on pairs of potentially impacted groups, mimicking a group fairness paradigm [17]. We categorize such groups using the 3 aforementioned parties (performing artists and composers) and 6 relevant side attributes (gender, popularity, country, language, year, and prevalence). To carry out our analyses, we develop the VI-Bias dataset, for which we collect metadata in terms of the attributes we investigate and assign them to songs of an existing dataset. We then gather performance data on pairs of groups and analyze the existence of potential discrepancies using the Kolmogorov-Smirnov test. Our results from a total of 115 experiments suggest that, on the one hand, the characteristics of a VI system may play a role in favoring a certain group over its counterpart but that, on the other hand, there also exist cases where all the VI systems favor the same or no group. After presenting the results, we also share our hypotheses on the possible reasons for the observed disparities. To facilitate future research, we share the instructions on how to obtain our dataset and the evaluation code¹.

2 METHODOLOGY

2.1 Systems

In this study, we consider a range of VI systems with different characteristics, in terms of whether they are learning- or rule-based, or whether they use melody- or chroma-based inputs. With this, we aim to better understand and discuss the causes of possible performance differences.

2.1.1 Qmax. Qmax [21] is a rule-based system that estimates similarities between pairs of tracks using an elaborate local alignment scheme, which results in a good performance but suffers from high computational cost. Nevertheless, it was the state-of-the-art system for VI for more than a decade, until recently. For our experiments, we use the implementation included in Essentia [4], with cremaPCP features [26] as input, which are specialized chroma representations

that correspond to an intermediate output of a structured chord estimation model [13].

2.1.2 MOVE. MOVE [25] is a deep learning-based system that features musically motivated design decisions to bring inductive biases to the model. It also uses cremaPCP features as input and transforms them into fixed-size embedding vectors, regardless of the input duration. By having explicit modules for achieving transposition and structure invariance, it has been proven effective for VI.

2.1.3 MICE-M and MICE-C. MICE [7] is a recent deep learning-based system that combines the design decisions of the model introduced by Doras et al. [6] and MOVE. Due to its input-agnostic design, MICE can be trained and used with any input feature suitable for VI. Hence, we consider two variants of the algorithm: MICE-M, which uses dominant melody features as input, and MICE-C, which uses cremaPCP features.

2.1.4 LF-c. LF-c [7] is a deep learning-based system designed to investigate the performance gains of combining models that use complementary input features (i.e., dominant melody and cremaPCP). It concatenates the embeddings obtained by MOVE and MICE-M and projects these into a new space, resulting in new embeddings. Doras et al. [7] show that using such an ensemble system improves identification performance by a large margin.

2.2 Attributes

For each of the attributes presented below, we categorize recordings into two groups. Considering various parties involved in VI, these groups can be created with respect to (1) the artist or the performance of the query track, (2) the artist or the performance of the reference or “original” track, and (3) the composer or the composition.

In the case of attributes that relate to persons (e.g., gender) rather than recordings (e.g., language), if the artist of a recording is a person, we assign the label for that recording as that of that person. If the artist is a band, we then collect the labels for all the current and past band members. To have clear distinctions between groups, we do not consider bands that have mixed labels (i.e., where some members belong to one group while the others belong to the other group). We also exclude from our analyses the cases where we cannot find a label for even a single member of a band.

We study differences in system performances by using binary labels (group 1, or G1, and group 2, or G2) for each attribute. If not stated otherwise, such labels are derived with respect to (1) the query recordings or their artists (-q), (2) the reference recordings or their artists (-r), and (3) the compositions or their composers (-c). We also analyze the performance differences where the query and the reference recordings or their artists belong to the same vs. different groups (-sd). In the case where they belong to different groups, we also check if the direction of the change (i.e., the reference in G1 and the query in G2 or the reference in G2 and the query in G1) has any effect on the performance (-d12). A summary of all the experiments and the considered attributes can be seen in Table 1.

Note that, for the -sd experiments, G1 denotes the cases where queries and references belong to the same group (e.g., queries and references from male artists) while G2 denotes that they belong

¹https://github.com/furkanyesiler/vi_bias

Table 1: Summary of the experiments. Numbers indicate the considered systems: Qmax (1), MOVE (2), MICE-M (3), MICE-C (4), and LF-c (5). The bold and underlined text indicates that the results obtained with the corresponding system show a significant difference between G1 and G2.

Feature	Query (-Q)	Ref. (-R)	Comp. (-C)	Same-Dif. (-SD)	Dif. G (-D12)
Gender (G-)	<u>12345</u>	<u>12345</u>	<u>12345</u>	<u>12345</u>	12345
Popularity (P-)	<u>12345</u>	<u>12345</u>	12345	<u>12345</u>	12345
Country (C-)	12345	12345	<u>12345</u>	12345	12345
Language (L-)	<u>12345</u>	<u>12345</u>	-	-	-
Year (Y-)	<u>12345</u>	<u>12345</u>	-	<u>12345</u>	-
Prevalence (V-)	<u>12345</u>	<u>12345</u>	<u>12345</u>	-	-

to different groups (e.g., queries from male artists and references from female artists). For the $-D_{12}$ experiments, G1 denotes the cases where queries and references are in the first and the second groups of the $-Q$ experiments, respectively (e.g., queries from male artists and references from female artists); and G2 denotes cases where the queries and references are in the second and the first groups of the $-Q$ experiments, respectively (e.g., queries from female artists and references from male artists).

2.2.1 Gender. The two groups we consider for gender are “male” (G1) and “female” (G2), and the data is gathered from MusicBrainz (MB) [24]. The experiments we perform for this attribute are denoted as G-Q (queries from male vs. female artists), G-R (references from male vs. female artists), G-C (compositions from male vs. female composers), G-SD (queries and references from the same vs. different gender groups), and G-D12 (queries from male and references from female artists vs. queries from female and references from male artists).

2.2.2 Popularity. For this attribute, we categorize the recordings as “popular” (G1) and “not so popular” (G2). We consider an artist (whether a person or a band) as popular if they ever had a number-one selling single in the top-10 music markets found in the IFPI’s Global Music Report 2021 [2] (excluding South Korea and China). By including European and Far-eastern music markets, we aim to mitigate bias toward artists from the United States. In the case of composers, we consider them popular if a popular artist (as defined above) ever played a version of any of their compositions. The number-one selling singles data is collected from Wikipedia, and the data for all the compositions of every composer are collected from MB. Although there are many ways to define “popularity,” our decision was based on the fact that number-one selling single data is the only data that can be found consistently (as opposed to, for example, top-10 singles) for all the aforementioned music markets on Wikipedia, which in itself suggests a distinction between the number-one selling artists and the rest. The experiments we perform for this attribute are denoted as P-Q, P-R, P-C, P-SD, and P-D12.

2.2.3 Country. For studying the effect of the country, we divide recordings by the ones whose artists or composers are from the

Table 2: Sample sizes per group per experiment.

Exp.	G1	G2	Exp.	G1	G2
G-Q	8,067	3,296	C-C	13,200	3,535
G-R	11,976	2,954	C-SD	424	109
G-C	16,840	1,093	C-D12	26	83
G-SD	5,943	2,040	L-Q	18,765	2,069
G-D12	603	1,437	L-R	19,487	2,172
P-Q	2,964	17,200	Y-Q	9,310	12,019
P-R	6,432	13,505	Y-R	2,579	18,517
P-C	3,171	1,255	Y-SD	7,737	11,129
P-SD	11,638	6,360	V-Q	11,025	11,342
P-D12	1,608	4,752	V-R	10,211	9,882
C-Q	1,879	638	V-C	9,024	13,388
C-R	2,967	594			

United States or the United Kingdom (G1), and the rest (G2), as the songs in our dataset consist of mostly genres originated in music traditions of those countries. In MB, there exist two annotations for the country: “area,” the place where the artist currently resides, and the “begin-area,” the place where the artist is originally from. In cases where an artist has both annotations, we use the begin-area for our analysis. The experiments we perform for this attribute are denoted as C-Q, C-R, C-C, C-SD, and C-D12.

2.2.4 Language. The two categories we consider for the language are “English” (G1) and other (G2). The language data is gathered from SecondHandSongs.com (SHS). The experiments for this attribute are denoted as L-Q and L-R. We exclude the other experiments since we cannot obtain the data for the language of compositions (L-C) and we do not have enough samples for analyzing L-SD and L-D12.

2.2.5 Year. For this attribute, we divide the recordings into two groups by whether they are from the year 2000 or after (G1) or from before (G2). Although this threshold can be set in many different ways, we use the year 2000 as the peer-to-peer networks that changed the music ecosystem drastically became popular by then. The release year data is gathered from SHS. For this attribute, the experiments we perform are denoted as Y-Q and Y-R. We also analyze the system performances in cases where the year gap between the query and the reference recordings is less than 10 years (G1) vs. greater than or equal to 10 years (G2) and denote the experiment as Y-SD.

2.2.6 Prevalence. This last attribute simply characterizes either how prevalent it is for a composition to have a version, quantified by the number of versions a composition has, or how prevalent it is for an artist to perform versions of other artists’ songs, quantified by the number of versions an artist performed. We include this attribute in our analyses to see whether songs that have been played by many artists (or artists that have played songs from many others) are in an advantageous position in terms of DRM or not. The data we use to form the categories are collected from SHS. For the artist prevalence experiments (V-Q and V-R), the two groups we consider are the artists who performed less than or equal to 25 versions (G1)

Table 3: Detailed results. For each system and experiment, we report the raw p -value of the KS test and the mean reciprocal rank of groups 1 and 2 ($\bar{\psi}$ -G1 and $\bar{\psi}$ -G2). Bold text with * denotes statistical significance after the Holm-Bonferroni correction. Overall system performance with the considered dataset is indicated in parenthesis in the top row.

Exp.	Qmax ($\bar{\psi}=0.42$)			MOVE ($\bar{\psi}=0.63$)			MICE-M ($\bar{\psi}=0.50$)			MICE-C ($\bar{\psi}=0.59$)			LF-c ($\bar{\psi}=0.75$)		
	p	$\bar{\psi}$ -G1	$\bar{\psi}$ -G2	p	$\bar{\psi}$ -G1	$\bar{\psi}$ -G2	p	$\bar{\psi}$ -G1	$\bar{\psi}$ -G2	p	$\bar{\psi}$ -G1	$\bar{\psi}$ -G2	p	$\bar{\psi}$ -G1	$\bar{\psi}$ -G2
G-Q	<0.01	0.41	0.45	<0.01*	0.63	0.70	<0.01*	0.46	0.61	<0.01*	0.58	0.65	<0.01*	0.73	0.81
G-R	0.41	0.42	0.44	<0.01*	0.63	0.69	<0.01*	0.48	0.64	<0.01*	0.59	0.65	<0.01*	0.74	0.83
G-C	0.19	0.42	0.39	<0.01	0.63	0.67	<0.01*	0.49	0.60	<0.01	0.58	0.63	<0.01*	0.74	0.80
G-SD	0.14	0.44	0.42	<0.01*	0.65	0.70	<0.01*	0.51	0.56	<0.01*	0.60	0.65	<0.01*	0.74	0.81
G-D12	0.05	0.38	0.44	1.00	0.70	0.71	0.09	0.52	0.57	1.00	0.65	0.65	1.00	0.80	0.81
P-Q	0.06	0.44	0.42	<0.01*	0.71	0.62	<0.01*	0.58	0.49	<0.01*	0.66	0.57	<0.01*	0.82	0.74
P-R	<0.01*	0.47	0.40	<0.01*	0.70	0.60	<0.01*	0.55	0.48	<0.01*	0.66	0.56	<0.01*	0.80	0.72
P-C	0.57	0.41	0.39	<0.01	0.63	0.58	0.83	0.50	0.48	0.01	0.59	0.54	0.38	0.75	0.73
P-SD	<0.01*	0.41	0.45	<0.01*	0.61	0.69	<0.01*	0.48	0.53	<0.01*	0.57	0.64	<0.01*	0.73	0.79
P-D12	0.02	0.43	0.46	1.00	0.69	0.69	0.03	0.55	0.53	0.99	0.63	0.64	1.00	0.79	0.79
C-Q	0.96	0.42	0.41	1.00	0.61	0.61	0.98	0.43	0.45	0.78	0.55	0.57	0.54	0.69	0.73
C-R	0.72	0.43	0.44	0.34	0.61	0.65	0.78	0.43	0.45	0.02	0.57	0.63	0.89	0.71	0.73
C-C	0.62	0.41	0.42	<0.01	0.63	0.65	<0.01*	0.49	0.56	<0.01	0.58	0.61	<0.01*	0.74	0.79
C-SD	1.00	0.45	0.47	0.30	0.65	0.58	0.88	0.42	0.39	0.63	0.60	0.56	0.66	0.72	0.69
C-D12	0.66	0.59	0.44	0.22	0.73	0.54	1.00	0.40	0.39	0.54	0.67	0.53	1.00	0.69	0.69
L-Q	1.00	0.42	0.41	0.07	0.64	0.66	<0.01*	0.51	0.64	0.12	0.59	0.61	<0.01*	0.75	0.82
L-R	0.98	0.42	0.41	0.06	0.64	0.66	<0.01*	0.50	0.63	0.07	0.59	0.61	<0.01*	0.75	0.82
Y-Q	<0.01*	0.46	0.39	<0.01*	0.66	0.61	<0.01*	0.52	0.48	<0.01*	0.62	0.56	<0.01*	0.77	0.73
Y-R	<0.01*	0.48	0.41	<0.01*	0.67	0.63	<0.01*	0.64	0.48	<0.01*	0.65	0.58	<0.01*	0.82	0.73
Y-SD	<0.01*	0.44	0.40	<0.01*	0.67	0.60	<0.01*	0.53	0.45	<0.01*	0.63	0.55	<0.01*	0.78	0.71
PR-Q	<0.01*	0.43	0.40	0.16	0.62	0.64	<0.01*	0.49	0.52	0.37	0.58	0.59	0.18	0.74	0.75
PR-R	0.11	0.43	0.41	<0.01*	0.62	0.65	<0.01*	0.49	0.52	0.04	0.58	0.60	<0.01	0.74	0.76
PR-C	<0.01*	0.45	0.40	<0.01*	0.65	0.61	1.00	0.50	0.50	<0.01*	0.61	0.57	<0.01*	0.76	0.74

and more than 25 versions (G2). For the composition prevalence experiment (V-c), the two groups we consider are the compositions that have less than or equal to 5 versions (G1) and more than 5 versions (G2). Thresholds 25 and 5 are chosen as they are the median values of their respective distributions.

2.3 Dataset and Evaluation

2.3.1 Dataset. We perform our analyses using a subset of the SHS4-dataset [6], for which we populated the metadata and label annotations. We publicly release this subset under the name VI-Bias, with Creative Commons BY-NC-SA 4.0 license [1]. The reference set of VI-Bias includes only one recording per composition, and we specifically chose them as the original recordings for all the compositions (as stated in SHS). Therefore, every query has exactly one correct item in the reference set. The reason for this decision is simply to imitate industrial cases where it is common to have only the original recording in the reference set. However, there can be multiple queries that have the same reference recording as the correct item. The total number of queries and references in VI-Bias is 22,428 and 15,293, respectively. All the deep learning-based models are trained using the SHS+ dataset [6]. For the training details of the considered systems, see Yesiler et al. [25] and Doras et al. [7].

2.3.2 Evaluation. We assess system performances using the reciprocal rank metric, which we denote by $\bar{\psi}$. We motivate this choice

by the fact that $\bar{\psi}$ penalizes the differences in lower ranks more than the differences in higher ranks, which we consider as a good proxy for real use cases (e.g., the difference between the correct answer being at rank 1 or 11 should have more impact than the difference of it being at rank 31 or 41). Note that, since there is only one correct item in the reference collection for each query, $\bar{\psi}$ corresponds to the same score as average precision.

For each group in our study, we collect $\bar{\psi}$ scores for all the items belonging to that group and use them to form performance distributions. To compare the obtained distributions of two groups, we use the two-sample Kolmogorov-Smirnov (KS) test [12], which has several characteristics that are desirable for our setup (e.g., it is a non-parametric exact test, and it is not sensitive to imbalanced sample sizes between groups). The null hypothesis of the KS test is that the samples are drawn from the same distribution (i.e., that the underlying distribution is the same). To reject the null hypothesis, we consider the threshold of 0.05 for the p -value; however, since we test multiple hypotheses, we apply the Holm-Bonferroni correction [11]. For computing the KS test and plotting rank distributions (see Figure 1), we collapse all the ranks above 10, to remove potential differences in the tails (i.e., high ranks that, with high probability, are not going to be considered by practitioners). For each analysis, we report sample sizes (see Table 2) and the mean $\bar{\psi}$ scores, $\bar{\psi}$, for each group.

Figure 1: Rank distributions for a selected set of experiments. The full set of 115 experiments is available in our repository.

3 RESULTS

Table 3 presents the detailed list of results for all the experiments and the considered VI systems, and Figure 1 presents rank distribution illustrations for a selected set of experiments. We include the rank distribution illustrations for all the experiments along with the detailed scores in our repository. We here highlight the main findings with respect to each considered attribute. Possible reasons for the observed disparities are discussed in Section 4.1.

3.0.1 Gender. The results for G-Q, G-R, and G-C suggest that the identification performances of the learning-based models form two distinct distributions for most systems (see Figure 1), and that performances are significantly higher for the recordings performed or composed by female artists (G2). On the contrary, for the only rule-based system we consider, the obtained p -values for G-Q, G-R, and G-C do not suggest a difference between the performance for male (G1) and female (G2) artists/composers. The results for G-SD and G-D12 show that the learning-based systems perform better when the gender of the query and the reference artists are different (G2), and that the direction of the change (female reference to male query, or male reference to female query) does not create a significant difference.

3.0.2 Popularity. Results for the popularity attribute suggest that all the systems perform better for queries and references from popular artists (G1); only for P-Q using Qmax does not result in a significant difference. However, such a difference between groups is not observed for the composers (P-C). Moreover, when the artists of the query and the reference tracks belong to different popularity groups (G2), identification performance is higher (P-SD). As with the case with gender, the direction of the change does not seem to create much of a difference (P-D12).

3.0.3 Country. Based on the results for the country attribute, we do not see any differences between performances with respect

to the country of the performing artists (C-Q and C-R). However, MICE-M and LF-c systems show such a difference with respect to the country of the composer, favoring composers from outside the US/UK (G2) group (C-c). We also observe no differences for the cases where the query and the reference recordings belong to the same (G1) vs. different (G2) groups (C-SD) and when the direction of the change is different (C-D12), but note that the sample sizes for the last two experiments are smaller than the previous ones.

3.0.4 Language. The results for the language attribute suggest a performance difference between groups for the systems that use melody-based features (i.e., MICE-M and LF-c) but not for the ones that use only chroma-based features (L-Q and L-R).

3.0.5 Year. All the considered experiments for the year category suggest the same result across all the systems: identification performance for queries and references after 2000 (G1) is higher compared to queries and references before that year (Y-Q and Y-R). Also, when the year gap between the query and the reference tracks is less than a decade (G1), the systems perform better.

3.0.6 Prevalence. We find no pattern behind the differences in the results for query recordings from artists that perform versions of other artists' compositions less (G1) or more (G2) prevalently (V-Q): while the performances of the considered groups seem similar for MOVE, MICE-C, and LF-c systems, MICE-M and Qmax favor different groups. For the experiments that compare the same quality for reference recordings (V-R), however, we see that the learning-based systems favor the cases where the reference recordings are from artists that perform versions more prevalently (G2) than the others (G1), while Qmax favors the opposite (G1 over G2). Lastly, in terms of comparing compositions with fewer (G1) or more (G2) versions (V-c), all the systems except MICE-M suggest a significant difference between the considered groups, favoring compositions with fewer versions (G1) over the others (G2).

4 DISCUSSION

Before entering into any discussion of the results, we shall first remind ourselves of the limitations of our data collection process. Firstly, the data sources we use (MB, SHS, and Wikipedia) contain human-annotated data, but we still would like to point out the possibility of human error. Secondly, since we excluded the cases where we could not find a label for a certain attribute, this may have created a potential bias on the data we use. Lastly, although we have implemented many heuristics to obtain the correct data from our sources, we have only performed random checks whether we succeeded on that, rather than checking all the obtained annotations one by one. Therefore, there may have been a small amount of metadata matching problems. However, we assume that, even if such cases exist, they should not drastically affect our results. In addition to the data collection process, the limitations of creating binary groups may also have an impact on the results. For example, creating such binary groups based on other criteria, or using non-binary partitions might have changed the outcomes of the presented analyses. Keeping these in mind, we now present our hypotheses on the causes of the observed disparities between groups and discuss the potential fairness implications of the results.

4.1 Causes of disparities

Overall, we observe that the learning-based systems work better for underrepresented groups (see, for instance, G-Q, G-R, P-Q, and L-Q). Our first hypothesis for this is that the groups with fewer samples may show less variety between versions in terms of the musical characteristics; thus, the identification performance may end up higher. A second reason for this may be the sample sizes of the groups in the training set, but we see that, for example, the gender ratio (G1 vs. G2) in the training set is 2.5, which is not very different than the one of the analysis set, which is 3.2.

When analyzing the attributes separately, we observe that the performances of the learning-based systems are higher for female artists and composers, but also the differences between $\hat{\psi}$ for G1 and G2 are higher for the systems that use melody features. We think that this may be related to the differences between male and female voices and the ability of the melody extraction algorithm to correctly estimate them. Also, looking at the experiments where the query and the reference recordings belong to the same vs. different groups for gender and popularity attributes (G-SD and P-SD), we see that the learning-based systems perform better for the cases where they belong to different groups (G2). We think that when the groups of the query and the reference change, the resulting version may appear more faithful to the original, and this may positively affect system performance.

For the popularity experiments, we see that all the systems tend to perform better for popular artists and composers (G1), though not all the experiments show a significant difference (P-Q and P-C). We hypothesize that this may be due to the input feature extraction algorithms, and the fact that they may have been optimized using tracks from such popular artists. Especially among the learning-based systems, MOVE and MICE-C show larger disparities between groups, and the chord estimation model used for cremaPCP was trained using a dataset including mainly popular tracks.

The results for the language experiments show disparities only for the systems using melody features, with recordings in languages other than English (G2) having better results than recordings in English (G1). A couple of reasons could be that, as mentioned earlier, having fewer samples may cause less variety between versions, and that the phonetic qualities of languages may affect the success of melody estimation algorithms (for example, Italian is argued to be the most suitable language for opera [10]).

In terms of the year attribute, we see that all the systems perform better for queries and references released in the year 2000 or after (G1). Possible reasons for this include the trends in the music creation process, which may affect the degree of variance between versions of a composition; the fact that G1 spans over a narrower window of years, which may naturally limit the possibility of having more different versions; and having fewer samples for G1 (as discussed above). For the year gap experiments (Y-SD), we argue that when the year gap between versions is longer (G2), it becomes more probable to see versions that are less faithful to the original track, which may cause lower system performances.

4.2 Considerations for fairness implications

Fairness evaluations are generally applied to systems that already interact with humans, where their direct impact on the considered groups is studied. In our case, however, we do not have access to a deployed VI system nor any real metrics on how the considered groups are affected (e.g., real amount of circulated royalties). Therefore, we now discuss the “potential” implications of the results presented in Section 3, rather than evaluating a real-life scenario. Although the following considerations are speculative in nature, they are based on our quantitative analyses, and we believe that not having access to a real-life system should not preclude research on such problems.

For interpreting the performance differences as fairness outcomes, we take the application of VI in DRM. In our multi-stakeholder scenario, we assume a direct connection between identification performance and the financial impact on various parties: if a query is not correctly identified, the performing artist of the query does not pay and the composer does not receive due royalties. To that extent, the reciprocal rank metric may have opposite meanings for the involved parties. The artist of the query may want a lower identification performance while the composer may desire the opposite.

To better illustrate this point, we can take a look at the gender experiments. We observe that the learning-based systems work better for female artists/composers compared to males. While this implies that female composers are likely to be rewarded more by these systems, in contrast, female artists that perform a version of an existing composition (i.e., artists of the queries) are likely to pay more royalties. Therefore, we here have a case where interpreting the fairness outcomes should be considered independently for all the involved parties, as a result of having a multi-sided structure.

Lastly, we observe that both the learning- and the rule-based systems show performance disparities for certain groups. Specifically, the learning-based systems show disparities for 54.4% of the cases while this ratio is only 30.4% for the rule-based system. While this clearly presents a disadvantage for the learning-based systems, their advantages regarding accuracy and scalability make them

irreplaceable for the industry. Therefore, to make sure such systems do not create any unfair conditions that favor certain groups, evaluating VI systems that may have a real-life impact should incorporate fairness and algorithmic bias metrics, along with metrics for measuring accuracy and scalability.

5 CONCLUSION

In this work, we have investigated the disparities in the performance of VI systems that may have an effect on fairness regarding the involved stakeholders. We have used 6 attributes to create potentially impacted groups based on the performing artists of the query, the performing artists of the reference, and the composer of the composition, which we identify as the related parties in a VI workflow. For our analyses, we have collected annotations for the attributes we consider and created VI-Bias dataset, which we publicly share. As the result of 115 experiments in total, we have seen that VI systems may indeed perform differently on certain groups. Their behavior may vary depending on whether they are learning- or rule-based, or whether they use melody- or chroma-based input features, but potentially other design choices could have an impact. After presenting the obtained results, we have proposed our hypotheses on the possible causes of such performance disparities. We leave confirming or rejecting them as future work. Lastly, we have discussed the potential implications of the obtained results on multiple stakeholders involved in VI from a fairness perspective. We encourage future VI research to incorporate fairness- and algorithmic bias-related evaluation metrics along with the existing accuracy- and scalability-related ones, to get ahold of any wrongful practices.

ACKNOWLEDGMENTS

This work is supported by the MIP-Frontiers project, the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 765068.

REFERENCES

- [1] [n.d.]. *Creative Commons - Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)*. Retrieved August 03, 2021 from <https://creativecommons.org/licenses/by-nc-sa/4.0/>
- [2] [n.d.]. *IFPI Global Music Report 2021*. Retrieved August 03, 2021 from <https://gmr2021.ifpi.org/report>
- [3] Gordon Baxter and Ian Sommerville. 2011. Socio-technical systems: from design methods to systems engineering. *Interacting with Computers* 23, 1 (2011), 4–17.
- [4] Dmitry Bogdanov, Nicolas Wack, Emilia Gómez, Sankalp Gulati, Perfecto Herrera, Oscar Mayor, Gerard Roma, Justin Salamon, José R. Zapata, and Xavier Serra. 2013. ESSENTIA: an audio analysis library for music information retrieval. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 493–498.
- [5] Alexandra Chouldechova and Aaron Roth. 2018. The frontiers of fairness in machine learning. *arXiv:1810.08810*
- [6] Guillaume Doras and Geoffroy Peeters. 2019. Cover detection using dominant melody embeddings. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 107–114.
- [7] Guillaume Doras, Furkan Yesiler, Joan Serrà, Emilia Gómez, and Geoffroy Peeters. 2020. Combining musical features for cover detection. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 279–286.
- [8] Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. Break the loop: gender imbalance in music recommenders. In *Proc. of the 2021 Conf. on Human Information Interaction and Retrieval*. 249–254.
- [9] Richard Ribón Fletcher, Audace Nakeshimana, and Olusubomi Olubeko. 2021. Addressing fairness, bias, and appropriate use of artificial intelligence and machine learning in global health. *Frontiers in Artificial Intelligence* 3 (2021), 116.
- [10] Matteo Dalle Fratte. 2018. *Italian is a language built to be sung*. Retrieved August 03, 2021 from <https://www.melonetica.com/italian-is-a-language-built-to-be-sung/>
- [11] Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70.
- [12] Frank J. Massey. 1951. The Kolmogorov-Smirnov test for goodness of fit. *J. Amer. Statist. Assoc.* 46, 253 (1951), 68–78.
- [13] Brian McFee and Juan Pablo Bello. 2017. Structured training for large-vocabulary chord recognition. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 188–194.
- [14] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.* 54, 6, Article 115 (2021), 35 pages.
- [15] Alexandra Olteanu, Jean Garcia-Gathright, Maarten de Rijke, Michael D. Ekstrand, Adam Roegiest, Aldo Lipani, Alex Beutel, Alexandra Olteanu, Ana Lucic, Ana-Andreea Stoica, Anubrata Das, Asia Biega, Bart Voorn, Claudia Hauff, Damiano Spina, David Lewis, Douglas W. Oard, Emine Yilmaz, Faegheh Hasibi, Gabriella Kazai, Graham McDonald, Hinda Haned, Iadh Ounis, Ilse van der Linden, Jean Garcia-Gathright, Joris Baan, Kamuela N. Lau, Krisztian Balog, Maarten de Rijke, Mahmoud Sayed, Maria Panteli, Mark Sanderson, Matthew Lease, Michael D. Ekstrand, Preethi Lahoti, and Toshihiro Kamishima. 2021. FACTS-IR: fairness, accountability, confidentiality, transparency, and safety in information retrieval. In *ACM SIGIR Forum*, Vol. 53. 20–43.
- [16] Lorenzo Porcaro and Emilia Gómez. 2019. 20 years of playlists: a statistical analysis on popularity and diversity. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 130–136.
- [17] Tim Rüz. 2021. Group fairness: independence revisited. In *Proc. of the ACM Conf. on Fairness, Accountability, and Transparency*. 129–137.
- [18] Kyle Robinson, Dan Brown, and Markus Schedl. 2020. User insights on diversity in music recommendation lists. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 446–453.
- [19] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proc. of the ACM Conf. on Fairness, Accountability, and Transparency*. 59–68.
- [20] Joan Serrà, Emilia Gómez, and Perfecto Herrera. 2010. Audio cover song identification and similarity: background, approaches, evaluation, and beyond. In *Advances in Music Information Retrieval*. Vol. 274. 307–332.
- [21] Joan Serrà, Xavier Serra, and Ralph G. Andrzejak. 2009. Cross recurrence quantification for cover song identification. *New Journal of Physics* 11 (2009), 093017.
- [22] Dougal Shakespeare, Lorenzo Porcaro, Emilia Gómez, and Carlos Castillo. 2020. Exploring artist gender bias in music recommendation. In *Proc. of the ImpactRS Workshop at ACM RecSys*.
- [23] Isaac Shepard. 2019. *How does the music industry work?* Retrieved August 03, 2021 from <https://www.taxi.com/transmitter/1906/how-does-the-music-industry-work/>
- [24] Aaron Swartz. 2002. Musicbrainz: a semantic web service. *IEEE Intelligent Systems* 17, 1 (2002), 76–77.
- [25] Furkan Yesiler, Joan Serrà, and Emilia Gómez. 2020. Accurate and scalable version identification using musically-motivated embeddings. In *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. 21–25.
- [26] Furkan Yesiler, Chris Tralie, Albin Andrew Correya, Diego F Silva, Philip Tovstogan, Emilia Gómez, and Xavier Serra. 2019. Da-TACOS: a dataset for cover song identification and understanding. In *Proc. of the Int. Society for Music Information Retrieval Conf. (ISMIR)*. 327–334.