



Follow the guides: disentangling human and algorithmic curation in online music consumption

Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail,
Camille Roth

► To cite this version:

Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, Camille Roth. Follow the guides: disentangling human and algorithmic curation in online music consumption. RecSys'21: Fifteenth ACM Conference on Recommender Systems, Association for Computing Machinery (ACM), pp.380-389, 2021, RecSys '21: Fifteenth ACM Conference on Recommender Systems, 978-1-4503-8458-2. 10.1145/3460231.3474269 . hal-03467626

HAL Id: hal-03467626

<https://hal.science/hal-03467626v1>

Submitted on 6 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Follow the guides: disentangling human and algorithmic curation in online music consumption

QUENTIN VILLERMET, Computational Social Science Team, Centre Marc Bloch, Germany

JÉRÉMIE POIROUX, Computational Social Science Team, Centre Marc Bloch, Germany and CNRS, France

MANUEL MOUSSALLAM, Deezer Research, France

THOMAS LOUAIL, CNRS, Géographie-cités, France

CAMILLE ROTH*, Computational Social Science Team, Centre Marc Bloch, Germany and CNRS, France

The role of recommendation systems in the diversity of content consumption on platforms is a much-debated issue. The quantitative state of the art often overlooks the existence of individual attitudes toward guidance, and eventually of different categories of users in this regard. Focusing on the case of music streaming, we analyze the complete listening history of about 9k users over one year and demonstrate that there is no blanket answer to the intertwinement of recommendation use and consumption diversity: it depends on users. First we compute for each user the relative importance of different access modes within their listening history, introducing a trichotomy distinguishing so-called ‘organic’ use from algorithmic and editorial guidance. We thereby identify four categories of users. We then focus on two scales related to content diversity, both in terms of dispersion – how much users consume the same content repeatedly – and popularity – how popular is the content they consume. We show that the two types of recommendation offered by music platforms – algorithmic and editorial – may drive the consumption of more or less diverse content in opposite directions, depending also strongly on the type of users. Finally, we compare users’ streaming histories with the music programming of a selection of popular French radio stations during the same period. While radio programs are usually more tilted toward repetition than users’ listening histories, they often program more songs from less popular artists. On the whole, our results highlight the nontrivial effects of platform-mediated recommendation on consumption, and lead us to speak of ‘filter niches’ rather than ‘filter bubbles’. They hint at further ramifications for the study and design of recommendation systems.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; • **Information systems** → **Multimedia streaming**; **Recommender systems**.

ACM Reference Format:

Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, and Camille Roth. 2021. Follow the guides: disentangling human and algorithmic curation in online music consumption. In *Fifteenth ACM Conference on Recommender Systems (RecSys '21)*, September 27–October 1, 2021, Amsterdam, Netherlands. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3460231.3474269>

1 INTRODUCTION

The contribution of algorithmic guidance to user preferences and behavior is at the heart of an active debate backed by an increasing number of empirical studies. This literature typically relies on some notion and measure of diversity, and on some dichotomy between user activity with and without recommendation. The latter, diversely denoted as “organic”, “autonomous”, “self-selected” behavior, generally serves as a reference point of what users would do had they not been

*Corresponding author: roth@cmb.hu-berlin.de

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
Manuscript submitted to ACM

advised or even diverted by algorithms. Nevertheless, many studies focused on user practices have emphasized how user navigation may both passively and actively rely on recommendation devices. The concept of “navigational queries” [12] first described how search engines may be used to access content that one is already aware of, and where, for instance, autocompletion suggestions merely play a limited role, essentially by reading the user’s mind and accelerating their navigation [28]. Beyond this simple old case, a growing literature points at the diversity of user practices and expectations toward platform-based recommendation; in the context of music streaming platforms, see [9, 24, 46]. Recommendation most likely does not apply similarly on all users. In other words, its influence should be studied by differentiating user classes who distinctly rely on, and plausibly respond to, recommendation.

This issue should furthermore be framed within a broader context where guidance may stem both from algorithmic devices and from human recommendations. The latter is arguably not new: friends and relatives, mainstream media and salesclerks have traditionally played a prime role in suggesting and shaping consumption preferences. In the music realm, the composition of radio playlists has typically long been decided by humans. The flourishing development of human-mediated, or “editorial”, guidance [5, 10] may however sometimes appear as a blind spot in the appraisal of platform-based recommendation. Building upon this recent trend, we extend the usual algorithmic/organic dichotomy by considering editorial guidance as a specific, distinct mode of recommendation. Contemporary music streaming platforms epitomize this trichotomy for they feature the three types of access to content: *algorithmic* (songs suggested by a recommendation engine, for instance in autoplay mode), *editorial* (songs curated by the platform staff, for instance as “featured playlists”) and *organic* (self-selected songs, either from personal playlists and favorite albums, or directly sought through the search bar). Note that in all three cases, it is obviously the user who chooses the mode of access: what matters in this trichotomy is that song titles are proposed by an algorithm, the editorial staff, or the user themselves.

The case of music streaming further makes it empirically feasible to connect the role of platform-based recommendation with traditional content curation. Specifically, we are able to contrast listening practices of users who rely more or less on some types of platform recommendation affordances with the mass (and somewhat traditional) guidance proposed by music radios, as an offline or “off-platform” reference point. Our study is thus based on two datasets: first, anonymized yet comprehensive listening histories for several thousands of users active on a major music streaming platform over a certain period of time. Second, playlists for a selection of mainstream radios in France over the same period.

On the whole, we aim to concretely examine (i) whether there exists distinct behavioral patterns with respect to the three modes of access to content and (ii) whether some type of guidance tends to pull listeners towards exploring certain types of content. We characterize the exploratory nature of listening histories with two distinct notions, dispersion and artist popularity, which correspond to two dimensions of diversity — i.e., dispersion at the user level (skewed popularity of content within a given user’s listening portfolio: redundancy or not) and dispersion at the macro level (skewed popularity of artists over all users: mainstream or not).

The main contributions are as follows:

- we hypothesize the existence of distinct classes of user attitudes towards recommendation before examining its possible influence on content consumption patterns;
- we further differentiate between two types of platform-side recommendation, stemming either fully from algorithmic suggestions or primarily from human curation;
- we establish a bridge between online platform recommendation and one of the traditional types of music guidance from the non-platform world, exemplified by radio music playlists.

2 RELATED WORK

Diversity and recommendation systems. The role of recommendation devices in fostering diversity is assuredly at the heart of a fast-growing literature focused on a myriad of platform types, where music streaming is only one application case among many. Contrarily to popular assumptions about so-called “filter bubbles”, the emerging empirical picture suggests that recommendation algorithms generally seem to increase diversity and serendipity [1, 4, 19, 29, 35, 36], even though recent results on specific platforms such as Spotify or YouTube tend to suggest otherwise [3, 37], while explicit personalization or “self-selection” also appear to induce algorithmic reinforcement and confinement, for instance regarding news consumption [14, 51]. Most of this literature works at the aggregate level without distinguishing populations of users who may differently use or respond to algorithmic guidance. Several studies nonetheless expressly differentiate users who are eager for recommendation [31], diversity [30], or exploration [17, 21]. This hints at the existence of different user behaviors and expectations towards recommendation [24] prior to it influencing users.

Use of music streaming platforms. Music streaming platforms feature an array of recommendation uses. They may be employed as “intimate experts” [24] who are able to counsel users in the most personalized manner and contribute to the self-formation of taste. They also provide editorial guidance which induces a hybrid type of curation, distinct from pure algorithmic recommendation [10]. The co-existence of algorithmic and editorial guidance is not unique to music streaming: their differential effect has already been studied on news platforms, where editorial recommendations appear to outperform algorithms in terms of diversity or in terms of concentration [5]. This issue connects with a relatively older literature appraising the role of online retail platforms in the consumption of content of varying popularity, especially at the top and bottom of the so-called “long tail” [15, 18]. Curation on music platforms also opens specific opportunities with respect to the access to niche content [6]. In this regard, quantitative studies show that platform use in general is associated with listening to less mainstream and less redundant content [13], even though diversity seems to plateau for the most active users [34]. However, the proper contribution of recommendation in these aggregate studies remains difficult to assess. Two recent studies differentiate organic from guided navigation and appear to yield conflicting results, whereby the former [3] or contrarily the latter [9] is shown to be correlated with increased consumption diversity. Notwithstanding, here again, little is known on the existence of underlying types of users (or even classes of users in the sociological sense, as in [46]) who make use of the platform affordances in different ways — and, in turn, how the same affordances distinctly affect these user populations in their exploration of different regions of the musical space.

Music recommendation goals. The intertwinement between expectations, uses and influences is not foreign to the more technical issue of the design of recommendation. The introduction of recommending devices is generally posterior to the creation of music streaming platforms, which were often initially operating as pure digital libraries (a notable exception being Pandora). While recommendation may first have been construed as an “attentional trap” i.e., to help retaining the user longer on the platform [40], the question of what it aims to optimize and how it intends to improve user experience provides further context to our contribution. The work of [11] offers a comprehensive overview of the current techniques for automated playlist creation, and their evaluation. It emphasizes in particular the importance of taking into account song popularity (e.g., to tackle the cold-start problem [7]), while radio playlists are said to typically “contain popular tracks in a relatively homogeneous style”. Models that generate and rely on categories of user tastes have also been proposed [45, 50] and seem to go much further than the above-mentioned literature in framing algorithmic recommendation in terms of implicit user classes. Regardless, the conception of music recommendation

systems already makes significant use of so-called “beyond accuracy” measures [27] by attempting at integrating diversity-related metrics [variously denoted as coverage, serendipity, novelty, see 23, 39, 49]. How these optimization principles may be translated, integrated and combined into platform-wide practices largely remains an open and active question.

3 DATASET

We obtain song listening histories from Deezer, a French private music streaming platform which started operations in 2007. Deezer currently has about 14m monthly active users worldwide, about half of whom are paid subscribers as of January 2019. Our data set describes twelve months of anonymized activity over 2019 for 8,639 randomly chosen paying subscribers based in France who were registered to the service before December 31st, 2018 and active in January 2019. We limit ourselves to users with a paid subscription, even though a different listening behavior of users of the free service might be assumed from previous work [48]. We discard events lasting less than 30s, assuming these are so-called “skips”. The dataset ultimately totals about 51m timestamped streaming events (plays) describing which user listened to which song by which artist for how long i.e., 5.9k plays per user on average (or around 16 per day).

The data also indicates which product feature users accessed songs through. Note that the platform may be indifferently used from a dedicated desktop app, a mobile app, or directly from a browser: the interface generally provides the same functions irrespective of the chosen device. On Deezer, users may directly look for music titles, albums or artists using a search bar. They can tag songs, artists and albums as favorites and build their own playlists. We denote these modes of access as “*organic*” for they entirely rely on user choices, whereby users do look for a specific item in their or the platform’s library. Users can also navigate to a home page where they find tailored recommendations which are either assembled and labeled by human editors at Deezer (such as recommended playlists variously called “10s electronic”, “Rock & Chill”, etc.) or algorithmically curated (such as the so-called “flow”, which is a personalized automatic mix). We denote the former as “*editorial*”, since the content is curated and recommended by human editors, and the latter as “*algorithmic*”, for it entirely relies on an interaction between Deezer algorithms, platform-wide data and users’ listening histories. In some cases, editorial playlists are algorithmically selected to be presented to users. We classify these as “*editorial*” since content selection remains primarily the choice of human editors. There is admittedly some porosity in the catalogs related to each of these three modes: for one, a user may include an algorithmically- or editorially-discovered song into one of their playlists; similarly, algorithmic recommendation may feature songs a user already listened to organically. Transfers between various access modes have been analyzed in [42]. We focus on the properties of song listening portfolios associated to such and such type of user and affordance. Combining this analysis with transfers would make it possible to characterize adoption and influence dynamics in the mid- to long-term but remains beyond the scope of the present paper and is left for future research. Whenever a song is listened to through some access mode, we consider that it is because the user decided to rely on that access mode, and thus describes listening patterns associated to that mode, irrespective of their more distant origin.

We also collect comprehensive playlist histories for a selection of 39 mainstream radios in France over the same period. This includes the top 15 French national musical stations in terms of measured audience during 2019, along with a relatively arbitrary selection of more specialized stations and webradios. For each radio, broadcast logs were collected and songs matched against Deezer’s whole catalog.




Fig. 1. **Left:** Use profiles and classes, where each dot on the ternary plot corresponds to a user of barycentric coordinates (p_a, p_e, p_o) , and each color refers to one of the four classes **a** (blue), **e** (red), **o** (green), **o+** (yellow). **Right:** Box and violin plots of activity (number of plays P , log-scale) for each user class.

4 USER PRACTICES

4.1 Modes of access and user behavior classes

The quantitative literature does not often pay attention to the possible existence of distinct classes of users when it comes to their behaviors and listening habits on the platform, classes for which the function and effect of recommendation may differ significantly [17]. Many studies report aggregate effects averaged over binary categories of users, for instance depending on a heavy vs. limited use of recommendation [13, 31] or categorical variables such as gender [41] or age [3]. We assume (and confirm) that a pre-grouping of users depending on broad use classes may reveal distinct sensitivities to recommendation. In other words, we contend that users who, say, rely more on editorialized playlists may respond differently to algorithmic recommendation than users who are mainly organic.

Let us denote P the number of plays in a user’s listening history, of which proportions p_a , p_e and $p_o = 1 - (p_a + p_e)$ have respectively been accessed algorithmically, editorially and organically. Their access mode profiles may thus be described by a triplet (p_a, p_e, p_o) defining barycentric coordinates in a ternary space, as shown in Figure 1-left. Even though a significant portion of users visibly rely to a large extent on the organic mode, we nonetheless observe a great deal of heterogeneity which hints at distinct use behaviors. We further define user groups by performing a simple Voronoi partitioning of profiles using a k -means algorithm. Choosing $k = 4$ explains around 80% of the variance with limited improvement for $k > 4$. Other clustering methods have been tried but generally yield a single cluster, most likely because data density increases roughly monotonously in the direction of the “organic” vertex, thereby preventing the formation of marked boundaries. As such, this partitioning only aims at defining bins as areas of the ternary space, rather than well-separated clusters per se. As a result, our four clusters should be understood as a simple approach to create an independent variable of use profiles among and between which we observe marked behavioral heterogeneity. It is similar, albeit not equivalent, to defining, say, quartiles or deciles of income, which notoriously obey a hybrid log-normal/Pareto distribution with no clear-cut boundaries.

In the following we thus consider four user bins that we denote as “**a**” (rather “algorithmic”, 989 users), “**e**” (rather “editorial”, 655 users), “**o**” (rather “organic”, 1614 users) and “**o+**” (“very organic”, 5381). To avoid any ambiguity with

<i>user class</i>	<i>population</i>		<i>access mode share</i>		
	<i>#users</i>	<i>(% dataset)</i>	<i>organic</i>	<i>algorithmic</i>	<i>editorial</i>
a	989	(11%)	36%	58%	6%
e	655	(8%)	55%	7%	38%
o	1614	(19%)	70%	23%	7%
o+	5381	(62%)	94%	2%	4%

Table 1. Population of each user class and proportion of plays by access mode for users chosen as class centroids (rows sum to 100%).

the various other practical bins that we define and use later on, we will now denote these user bins as “*user classes*”. On the whole, organic classes (**o** and **o+**) comprise 7062 users i.e., roughly 80% of the dataset, for whom at least half of all plays have been accessed autonomously by users. Table 1 gives the shares of plays in each of the three modes for the centroids of these four classes. The number of plays P is also a proxy of the user’s activity on the platform. Its distribution spans several orders of magnitude, in a manner similar for all classes, as can be seen in Figure 1-right — in other words, there are weakly and strongly active users in all classes. We note that average activity is nonetheless slightly smaller for “editorial” users.

4.2 Two dimensions of diversity

We rely on these four classes to appraise the role of each access mode on music consumption, in terms of *where* each mode brings users to, and *how much*. We characterize the portfolio of songs the users listen to by focusing on two fundamental diversity measures:

- (1) dispersion, denoting the lack of redundancy in the listening history;
- (2) artist popularity, denoting the tilt toward songs by more popular artists.

We explain below (4.2.2) why we chose not to use genre meta-data to appraise diversity.

4.2.1 Dispersion and activity. Dispersion may simply be computed as S/P , where S is the number of unique songs among the P plays. It thus equals 1 when each song is played exactly once, and goes to 0 as plays of the same songs are repeated. In other words, low dispersion implies high redundancy of a listening history. S/P is also called “exploratory ratio” in [26] and further alludes to the trade-off between exploration and exploitation. Whichever the preferred interpretation, this ratio relates to a *functional* understanding of the diversity of a user’s listening behavior. While the data display large fluctuations, we fit for each user class a linear model to highlight the general tendency. This reveals an inverse relationship between dispersion and activity, for all user types (see Figure 2-left): a user’s catalog tends to saturate and its dispersion decreases with listening time (put differently, redundancy i.e., exploitation of the catalog increases). For a given activity level, dispersion is nevertheless lower for very organic users, while it is comparable for moderately organic and for recommendation-intensive users; i.e., as soon as some form of recommendation is involved. This indicates a general inclination for redundancy with users who rely most on organic access modes. Focusing on recommendation-intensive users, dispersion is slightly lower for editorial users. However, their overall lower activity, which in turn corresponds to higher dispersion, leads to levels of dispersion comparable with those of algorithmic users, as shown on Figure 2-right.

We refine this picture by breaking down dispersion values per access mode. Figure 3 reveals that, for all user classes, the major access mode generally exhibits lower dispersion values — all the more so when activity is higher for a given access mode, as shown by the fact that the darkest bars (i.e., highest activity) are generally to the left of these histograms




Fig. 2. Dispersion profiles for each user class. *Left*: scatterplot (each dot is a user) and binned averages (solid lines) as a function of activity (P , log-scaled). *Right*: boxplots per user class.

(i.e., lowest dispersion). Put simply, for algorithmic users (**a**), dispersion is lower for algorithmic recommendations and higher for editorial ones. For editorial users (**e**), dispersion is lower for editorial recommendations, and so on: the same applies to organic (**o**) and even more so for very organic users (**o+**) who exhibit high dispersion in both algorithmic and editorial access.

This suggests a phenomenon of specialization in each class of users, who tend to prioritize exploitation in their preferred mode of access, and favor exploration in the others, in terms of more dispersed plays. We shall here keep in mind that very organic users exhibit the lowest dispersion – using the platform heavily as a *digital library* connotes more exploitation of the catalog than using it as a *digital librarian*. In other words, even a moderate use of some form of recommendation is generally associated with a higher level of exploration.

4.2.2 Artist popularity and access modes. Artist popularity, by contrast, describes music consumption in terms of content created by popular or less popular artists. We denote it as v . It relates to a *semantic* understanding of diversity. The literature traditionally proposes two main approaches to describe such diversity. On the one hand, some studies appraise diversity in spatial terms by hypothesizing an underlying “cultural” or music similarity space, and by embedding songs [2, 3, 47], individuals [25, 38] or both [22] into a graph or a low-dimensional space where proximity roughly corresponds to semantic similarity (e.g., similar musical genres or tastes). Such spaces are also used by platform algorithms to recommend songs not too far from one another (e.g., avoiding recommending manouche jazz after a heavy metal song).

On the other hand, some studies characterize consumption diversity through content popularity, most notably through the share of top artists [9, 13]. We follow here this latter approach for three main reasons.

First, designing an embedding space requires choosing between many possible similarity metrics and construction approaches, each based on strong hypotheses on the semantic encoding. For instance, distances based on the co-occurrences of content tags (e.g. taxonomies [20] or folksonomies [32]), will typically suffer from definition ambiguities and socio-cultural inconsistencies [8, 16] and more prosaically from annotation noise and biases. Distances based on acoustic similarities [43] or usage [22] are typically tailored for neighborhood consistency, but will not generally bear any meaning for, say, the distance between metal and jazz songs, compared with the distance between jazz




Fig. 3. Breakdown of dispersion values for each user class and for each access mode. Histograms are further binned by deciles of increasing S/P values (from 0 to 1 from left to right) and indicate how many users of each class (**a**, **e**, **o** and **o+**) exhibit which dispersion value for a certain access mode (algorithmic, editorial or organic). Average activity values for users of each decile bar are further indicated by a grayscale, where darkest shades correspond to highest P values.

and classical pieces. In addition, many users, particularly omnivores [33, 44] may show interest for more than one genre. Navigation profiles may thus exhibit several distinct centroids in possibly distant regions. In turn, this potential “musical polycentrism” might blur the meaning of notions solely based on the geometric extent of listening profiles in multi-dimensional spaces. Conversely, *artist popularity* remains a mono-dimensional notion, whose mean and deviation are simple yet likely robust indicators of the position and span of a user’s musical content consumption with respect to the whole field.

Second, this will ease the comparison with radio playlists and thus offline editorial recommendation. The radios of our dataset indeed address a quite diverse collection of music genres, some being very generalist (e.g., ‘FIP’), some much more specialized (e.g., ‘TSF Jazz’). To discuss specialization and eclectism across genres and radios, we believe that artist popularity may act as a better sort of *lingua franca* than, say, extents in a vector space loosely encoding cultural similarities.

Third, artist popularity directly connects with an older debate on whether the almost infinite catalogs of online platforms foster consumption of more mainstream or more niche content, or both, and by which types of users [15].

We first define artist popularity by computing how many times their songs have been played in the whole dataset as a proxy of their popularity on the platform. This is referred as “artist playcounts” in [7], also connected with the mainstreamness of artists. We then distinguish four bins of popularity such that each bin gathers a similar numbers of plays i.e., artists which, taken together, represent the same total amount of plays ($\frac{1}{4} \sum_{\text{users}} P$). This ensures that a play chosen in a uniformly random way from the listening history has the same likelihood of belonging to any of the four bins. As a result, the first bin v_1 gathers the top 73 artists, the fourth bin v_4 the bottom 166,869 artists.

Table 2 gathers the proportion of access modes for each popularity bin. In general, songs are principally listened to organically, irrespective of artist popularity — around 80% of all plays on average, which is assuredly expected given the prevalence of organic access modes (figure 1). We however see that organic access is non-monotonous, in that the

bin	# artists	<i>access mode</i>			
		algorithmic	editorial	organic	
v_1	73	9%	8%	83%	100%
v_2	319	16%	8%	76%	100%
v_3	1462	18%	5%	77%	100%
v_4	166869	15%	5%	80%	100%
all	164955	14%	7%	79%	100%

Table 2. Proportion of access modes for each artist popularity bin (preferred bins for each access mode are marked in bold).

highest shares are found in both the most and least popular artist bins [consistent with 18]. A different picture emerges for guided access modes. Algorithmic access puts the highest share on intermediate bins, and by far the lowest share on the highest popularity bin. Editorial access, by contrast, is monotonously more frequent for higher popularity bins. In other words, content by popular artists seems to be more often proposed by editorial rather than algorithmic picks. To further exhibit this effect, we compute the average popularity bin for each access mode and find 2.68 for algorithmic plays, 2.49 for organic ones, and 2.23 for editorial ones (the average popularity bin for all plays is 2.50, by construction).

Finally, we observe that dispersion and artist popularity are associated. Part of it is likely mechanical: since the most popular bin features much less artists as well as less songs (18,425 songs for v_1 vs. 1,161,257 songs for v_4), it also induces a lower likelihood of dispersion, all other things being equal. Notwithstanding, we essentially observe that users who consume content by less popular artists exhibit more dispersed listening histories; see Figure 4-left. We refine this picture by computing the dispersion of the listening history of each user restricted to songs of a given popularity bin, shown in the right panel of Figure 4, which also decreases for content made by popular artists.

To summarize, average dispersion increases with activity and decreases with artist popularity. Since very organic users have a lower dispersion, this further hints at a positive relationship between the use of recommendation and dispersion.

4.3 User types and access mode biases

We may now appraise the interplay between user types, access modes and diversity from the user perspective. Let us start with artist popularity. For each user, we compute the proportion of plays that fall in each popularity bin and divide it by its expected value i.e., we use a null hypothesis where all plays would uniformly fall in equal amounts into each bin (typically proportional to $P/4$ by construction of popularity bins). These ratios yield a divergence profile for each user in terms of their appetency for some bins over others. We then average these profiles over all users of each class and plot in Figure 5 (top panel) the mean over- or under-consumption per popularity bin.

We find that algorithmic users (**a**) rather under-consume content by popular artists while editorial users (**e**) exhibit an almost opposite profile by over-consuming such content. Interestingly, **o** and **o+** users markedly differ in their appetency for popular artist content: whereas very organic users tend to slightly over-listen to songs from both the most and the least popular artists [again, consistent with consumer behavior on retail platforms 18], organic users exhibit a monotonous profile favoring less popular content in a strictly increasing manner; in this regard, they resemble more algorithmic users.

We refine this analysis by further plotting over- and under-consumption with respect to access modes. More precisely, given a user type, we perform the normalization locally for each access mode: the ratio between actual and expected




Fig. 4. Artist popularity and dispersion. *Left*: dispersion of listening histories as a function of the average artist popularity profile of each user $\bar{v}$. Each dot is a user, the solid line represents the best linear fit as a guide to the underlying trend. *Right*: boxplot of dispersions restricted to content of a certain popularity bin.




Fig. 5. Relative consumption of content from each popularity bin, average log-ratio with respect to a uniformly random baseline for each bin (0 corresponds to no deviation, the x-axis is ordered from v_1 to v_4 i.e., for musical content from more to less popular artists). *Top*: average over all plays. *Bottom*: breakdown by access mode.

quantities is computed using respectively $p_a P$, $p_e P$ and $p_o P$ instead of P . The results are shown on the lower panel of figure 5. For instance, the curve that corresponds to, say, editorial access of **o+** users, indicates to what extent their editorial plays fall into each popularity bins.

We generally confirm that algorithmic access tends to correspond to content from less popular artists, even for very organic users (**o+**) who normally over-consume content from popular artists, and also for editorial users (**e**) who normally consume even more content from popular artists, yet to a lesser extent. By contrast, on the whole, and for all

user types, editorial access pulls towards popular artists. Organic access, finally, tends to somewhat mimic average behavior (and expectedly all the more so for the most organic users), but in a less marked manner, suggesting that a good part of the deviation in trends observed on average values stems from recommendation-based access modes, be it editorial or algorithmic.

More broadly, these observations contribute to shed light on a typical chicken-and-egg problem i.e., whether access modes (and thus platform features and affordances) influence user behavior, or user behavior influences how access modes are being used, and for what. In effect, if we focus on organic access as a reference of what users autonomously look for on the platform, we see that user classes exhibit markedly different appetencies, for instance with respect to content from least popular artists. We also know that user classes roughly entail specialized ways of using the platform, whereby dispersion is lower for the main corresponding access mode (as per Figure 3), which indicates that the major access mode leans more toward exploitation than exploration. Further, the use of recommendation, algorithmic or not, appears to tilt consumption towards more exploration, all other things being equal. This is further emphasized by the fact that editorial users exhibit a higher dispersion on average than very organic users, even if they generally consume more popular content which is otherwise related to lower dispersion in aggregate. If we now shift to the use of recommendation-based access modes, we see that they tend to exhibit a general pattern, irrespective of user types: editorial favors popular, algorithmic leans toward least popular, in very broad and rough terms. In this regard, Figure 5 contributes to disentangle how both types of effects yield a general consumption profile — i.e. *both* due to user types and access modes. It is legitimate to suggest here that we observe and, at least in part, deconstruct the intertwining of both (1) underlying user types, corresponding to different ways of using the platform, and (2) overlaying access modes, corresponding to the different ways in which the platform may affect user consumption.

5 THE DIVERSITY OF HUMAN-ASSISTED GUIDANCE

Notwithstanding the preferences and expectations of users on the platform, the picture that emerges is that of a split between two types of recommendation in terms of fostering diversity: irrespective of dispersion trends, editors rather than algorithms appear to sustain the consumption of songs from popular artists (and vice versa for least popular artists). To put this observation in perspective, we need to rely on an external reference. We contend that radio programs constitute a relevant instance in this regard. Radio-based playlists can be construed as one of the closest offline equivalent of editorial playlists on Deezer — and, from a human computing viewpoint, one of the oldest large-scale music recommendation systems.

In practice, it would be difficult to access detailed radio listening histories for a number of people, let alone for the subset of users we considered here. To circumvent this issue, we adjust the way we carry out computations on users to make both sources as comparable as possible. On the one hand, we use artist popularity values from the Deezer data set as a general proxy. Around 83% of songs played on radios are matched with the user data set. These songs inherit their respective artists' popularity computed from the Deezer data, and we ignore unmatched songs. By construction, this likely induces an overestimation of popular content, and, in turn, of the popularity bin for radios that play less popular artists.

On the other hand, we define hourly listening sessions for both users and radios. This additional focus on hours aims at taking into account the fact that radio playlist may be heavily shaped by the existence of programs broadcast at specific moments of the day, or night. Concretely, for each radio and each user we consider their time-ordered sequence of plays over the year, and we setup counters to keep track of the number of plays $P(h)$ that occurred between hour h




Fig. 6. Dispersion and artist popularity for a selection of radios and the four user classes. Boxplots are computed on hourly values. *Left*: dispersion ranked by decreasing average values, colored by average popularity bins. *Right*: dually, popularity ranked by decreasing average bin number, colored by average dispersion.

and hour $h + 1$ during the entire period, along with the number of new songs $S(h)$ that were played during this hour, that is, songs that had never been played previously by this radio (resp. by this user). This way for each radio and user we calculate twenty-four hourly P , S , S/P and $\bar{v}$ values. We then averaged these hourly values over users that belong to the same class, and compare the four user classes previously identified with radio stations, according to the range of their hourly dispersion S/P and hourly average popularity $\bar{v}$ values.

Results are gathered on figure 6. We select a representative sample of 15 radios out of 39 for clarity purposes. On the left panel, we compute dispersion boxplots both for radios and user types. All items, radios or user types, are ordered from top to bottom by decreasing value of average dispersion. We further color boxplots according to the average popularity bin. Table 3 additionally gathers the detailed breakdown of plays in each popularity bin, for all user types and a few selected broadcasters.

We observe that most user types exhibit more dispersion than most radios, indicating that radio programming is tilted toward the exploitation of a more limited catalog. More precisely, there appears to be an inflection point around “TSF Jazz” and “RFM” which roughly splits the set of items between a larger half with a significantly low dispersion (below 0.10) and a smaller half with much higher dispersions (generally above 0.25). In this picture, user classes exhibit among the highest values. Remarkably, two radios are above all user classes — France Musique and FIP, which are public-funded, predominantly musical and rather eclectic broadcasters. Besides, a mild correlation between popularity and dispersion is visible: the diversity of the catalog in terms of distinct played items is roughly linked to its diversity in terms of playing artists from the less popular bins.

A dual computation consists in representing boxplots of observed mainstreamness averaged over hours, while coloring them with the average dispersion — see the right panel of Figure 6. While the correlation between dispersion and popularity is not immediately visible, the Spearman rank correlation between S/P and the proportion of songs in v_4 is equal to 0.827, confirming the similar inverse ordering of dispersion and popularity: radios with more dispersed

	v_1	v_2	v_3	v_4	$\langle S/P \rangle$
a	17%	27%	29%	27%	0.39
e	29%	30%	22%	19%	0.41
o	22%	25%	26%	27%	0.38
o+	29%	23%	23%	25%	0.28
Radio Meuh	0%	2%	16%	82%	0.23
FIP	2%	7%	20%	71%	0.60
RFM	7%	27%	36%	30%	0.10
Fun Radio	18%	61%	13%	8%	0.04
NRJ	41%	23%	23%	7%	0.03

Table 3. Average breakdown of content played into each popularity bin, for all user types and a selection of radios, and average dispersion.

programs also play more songs from less popular artists. There are notable exceptions such as “Radio Meuh”, whose lower bound on popularity is the highest of all radios —and it is indeed considered informally as a very eclectic broadcaster— while having a relatively moderate dispersion (as seen in Table 3). The new interesting take-away of this second plot lies in the positions of user types. They are now much closer to the median value of this selection, even though organic and algorithmic users appear to be a little bit above it. Put differently, while online music listening practices seem to foster functional diversity (dispersion), in terms of semantic diversity (popularity) they seem, on average, to be neither significantly above nor significantly below the mass of the radios we focus on.

6 CONCLUDING REMARKS

Our analysis of recommendation, in the broad sense, on a music streaming platform started by assuming the existence of distinct attitudes towards platform affordances. This produced a typology of four user classes who markedly differ in their modes of access to music and, more importantly, in the way these modes are associated with distinct effects on the diversity of consumed songs. We showed more broadly that there is no blanket answer to the question of the influence of recommendation devices: rather, we contend that it primarily depends on users. This observation may have important ramifications in the appraisal of the effect of recommender systems on consumption diversity — we suggest to speak of filter niches, rather than bubbles.

Moreover, we showed that the framing of recommendation on platforms and, thus, their algorithmic governance, would benefit from paying specific attention to the various types of guidance. Human-mediated curation is associated with distinct effects on the exploration of content and fostering diversity, and with distinct user classes. Both types of guidance seem to often have opposite effects. They may also fulfill distinct roles and thus have to obey to distinct design principles. We further introduced a bridge between platform-based recommendation and traditional offline curation, exemplified by radio playlists, which we see as a fruitful point of comparison to frame how platforms may or may not affect user preferences and access to content. If user classes were radios (and assuming it is valid to speak of an average user for each user class), they would certainly not appear to be more repetitive than most radio stations, while they would exhibit a wide range of serendipity. Interestingly, while algorithmic access modes seem to generally be associated with an avoidance of popular content, it is what we call organic users who appear to focus most on the least popular content — yet, these users precisely exhibit a relatively balanced diet of platform affordances, generally combining their autonomous navigation with both sorts of recommendation. In other words, we may hypothesize that the most “expert” users best exploit the platform capabilities to explore the long tail. Editorial access, on the other hand, and even more

so for editorial users, appears to fulfill a role traditionally ascribed to mainstream radios — putting forward a higher proportion of popular artists.

A step further from this study would consist in analyzing temporal aspects of the partitioning. Indeed, the chronology of access modes to a song probably carries valuable information. For instance, a song may first be discovered through an editorial recommendation, then later accessed organically and symmetrically, organically accessed content may be subsequently picked by the algorithm later on. More generally, user trajectories and transitions between classes should be the focus of future research that would shed light on the acclimation of users to platform features. In particular, the causal relationship that may exist between the presence or absence of some content in editorial and algorithmic recommendations and the evolution of their platform-wide popularity would also be a fruitful field of investigation. With respect to the ubiquitous debate on whether platform recommendation fosters diversity or not, a mixed picture emerges depending on which types of recommendation one is talking about, as well as which type of audience — whereby radios provide an insightful reference. On the whole, this comparison may shed light on the relative position of platforms and their contribution to cultural diversity with respect to traditional cultural prescribers and, again, how their algorithmic principles should be governed. Understanding user expectations not only in terms of genres and tastes, but in terms of their use and perhaps understanding of the technical features of a platform, may help refine the design of recommender systems *ex ante*, rather than optimizing their results *ex post*. The corresponding takeaway for practitioners is to be attentive to the differences between access modes which may fulfill distinct roles for listeners e.g., active exploration vs. background music listening. Longitudinal analyses of the evolution of user tastes, specifically in terms of learning and of long-term effects of distinct platform affordances, as well as qual-quant analyses based on user surveys and interviews would be most helpful in further disentangling the chicken from the egg.

ACKNOWLEDGMENTS

This paper has been partially realized in the framework of the “RECORDS” grant (ANR-2019-CE38-0013) funded by the ANR (French National Agency of Research). We are grateful to Dougal Shakespeare for useful comments.

REFERENCES

- [1] Luca Maria Aiello and Nicola Barbieri. 2017. Evolution of Ego-networks in Social Media with Link Recommendations. In *Proc. 10th ACM Intl. Conf. on Web Search and Data Mining (WSDM '17)*. ACM, New York, NY, 111–120.
- [2] Massimo Airoidi, Davide Beraldo, and Alessandro Gandini. 2016. Follow the algorithm: An exploratory investigation of music on YouTube. *Poetics* 57 (2016), 1–13.
- [3] Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of The Web Conference 2020*. ACM, New York, 2155–2165.
- [4] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [5] Jack Bandy and Nicholas Diakopoulos. 2020. Auditing news curation systems: A case study examining algorithmic and editorial logic in Apple News. In *Proceedings of the 14th ICWSM International AAAI Conference on Web and Social Media*. AAAI Press, Palo Alto, 36–47.
- [6] Emilia Barna. 2017. “The perfect guide in a crowded musical landscape:” Online music platforms and curatorship. *First Monday* 22, 4 (2017).
- [7] Christine Bauer and Markus Schedl. 2019. Global and country-specific mainstreamness measures: Definitions, analysis, and usage for improving personalized music recommendation systems. *PLoS one* 14, 6 (June 2019), e0217389.
- [8] David Beer. 2013. Genre, boundary drawing and the classificatory imagination. *Cultural Sociology* 7, 2 (2013), 145–160.
- [9] Jean-Samuel Beuscart, Samuel Coavoux, and Sisley Maillard. 2019. Music recommendation algorithms and listener autonomy – The listening patterns of a panel of music streaming users. *Revue de la Recherche Musicale* 213, 1 (2019), 17–47.
- [10] Tiziano Bonini and Alessandro Gandini. 2019. “First Week Is Editorial, Second Week Is Algorithmic”: Platform Gatekeepers and the Platformization of Music Curation. *Social Media + Society* 5, 4 (2019), 1–1.
- [11] Geoffroy Bonnin and Dietmar Jannach. 2014. Automated generation of music playlists: Survey and experiments. *ACM Computing Surveys (CSUR)* 47, 2 (2014), 1–35.

- [12] Andrei Broder. 2002. A taxonomy of web search. *ACM SIGIR Forum* 36, 2 (2002), 3–10.
- [13] Hannes Datta, George Knox, and Bart J Bronnenberg. 2018. Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science* 37, 1 (2018), 5–21.
- [14] Ivan Dylko, Igor Dolgov, William Hoffman, Nicholas Eckhart, Maria Molina, and Omar Aaziz. 2017. The dark side of technology: An experimental investigation of the influence of customizability technology on online political selective exposure. *Computers in Human Behavior* 73 (2017), 181 – 190.
- [15] Anita Elberse. 2008. Should you invest in the long tail? *Harvard Business Review* 86, 7/8 (2008), 88–96.
- [16] Elena V Epure, Guillaume Salha, Manuel Moussallam, and Romain Hennequin. 2020. Modeling the Music Genre Perception across Language-Bound Cultures. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 4765–4779.
- [17] Jean Garcia-Gathright, Brian St. Thomas, Christine Hosey, Zahra Nazari, and Fernando Diaz. 2018. Understanding and Evaluating User Satisfaction with Music Discovery. In *Proc. 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 55–64.
- [18] Sharad Goel, Andrei Broder, Evgeniy Gabrilovich, and Bo Pang. 2010. Anatomy of the long tail: ordinary people with extraordinary tastes. In *Proc. WSDM'10 ACM 3rd Intl Conf on Web Search and Data Mining*. ACM, New York, NY, 201–210.
- [19] Mario Haim, Andreas Graefe, and Hans-Bernd Brosius. 2018. Burst of the Filter Bubble? Effects of personalization on the diversity of Google News. *Digital Journalism* 6, 3 (2018), 330–343.
- [20] Romain Hennequin, Jimena Royo-Letelier, and Manuel Moussallam. 2018. Audio based disambiguation of music genre tags. In *Proceedings of the 19th ISMIR Conference*. International Society of Music Information Retrieval, 645–652.
- [21] Christine Hosey, Lara Vujović, Brian St. Thomas, Jean Garcia-Gathright, and Jennifer Thom. 2019. *Just Give Me What I Want: How People Use and Evaluate Music Search*. Association for Computing Machinery, New York, NY, USA, 1–12.
- [22] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. 2017. Collaborative metric learning. In *Proceedings of the 26th international conference on world wide web*. 193–201.
- [23] Marius Kaminskas and Derek Bridge. 2016. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 7, 1 (2016), 1–42.
- [24] Nedim Karakayali, Burc Kostem, and Idil Galip. 2018. Recommendation systems as technologies of the self: Algorithmic control and the formation of music taste. *Theory, Culture & Society* 35, 2 (2018), 3–24.
- [25] Renaud Lambiotte and Marcel Ausloos. 2005. Uncovering collective listening habits and music genres in bipartite networks. *Physical Review E* 72, 6 (2005), 066107.
- [26] Thomas Louail and Marc Barthelemy. 2017. Headphones on the wire Statistical patterns of music listening practices. *arXiv preprint 1704.05815* (2017).
- [27] Sean M McNee, John Riedl, and Joseph A Konstan. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 extended abstracts on Human factors in computing systems*. 1097–1101.
- [28] Bhaskar Mitra, Milad Shokouhi, Filip Radlinski, and Katja Hofmann. 2014. On User Interactions with Query Auto-Completion. In *Proc. ACM SIGIR 37th Intl. Conf. on Research & development in information retrieval*. ACM, 1055–1058.
- [29] Judith Möller, Damian Trilling, Natali Helberger, and Bram van Es. 2018. Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. *Information, Communication & Society* 21, 7 (2018), 959–977.
- [30] Sean A. Munson and Paul Resnick. 2010. Presenting Diverse Political Opinions: How and How Much. In *Proc. CHI 2010: Expressing and Understanding Opinions in Social Media, April 10–15, 2010, Atlanta, GA, USA*. 1455–1466.
- [31] Tien T Nguyen, Pik-Mai Hui, F Maxwell Harper, Loren Terveen, and Joseph A Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*. ACM, 677–686.
- [32] Minsu Park, Ingmar Weber, Mor Naaman, and Sarah Vieweg. 2015. Understanding musical diversity via online social media. In *Proceedings of the 9th ICWSM International AAAI Conference on Web and Social Media*. AAAI Press, Palo Alto, 308–317.
- [33] Richard A Peterson and Roger M Kern. 1996. Changing highbrow taste: From snob to omnivore. *American sociological review* 61, 5 (1996), 900–907.
- [34] Rémy Poulain and Fabien Tarissan. 2020. Investigating the lack of diversity in user behavior: The case of musical content on online platforms. *Information processing & management* 57, 2 (2020), 102169.
- [35] Cornelius Puschmann. 2019. Beyond the bubble: Assessing the diversity of political search results. *Digital Journalism* 7, 6 (2019), 824–843.
- [36] Camille Roth. 2019. Algorithmic Distortion of Informational Landscapes. *Intellectica* 70, 1 (2019), 97–118.
- [37] Camille Roth, Antoine Mazières, and Telmo Menezes. 2020. Tubes and bubbles topological confinement of YouTube recommendations. *PloS one* 15, 4 (2020), e0231703.
- [38] Mike Savage and Modesto Gayo. 2011. Unravelling the omnivore: A field analysis of contemporary musical taste in the United Kingdom. *Poetics* 39, 5 (2011), 337–357.
- [39] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. 2018. Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval* 7, 2 (2018), 95–116.
- [40] Nick Seaver. 2019. Captivating algorithms: Recommender systems as traps. *Journal of Material Culture* 24, 4 (2019), 421–436.

- [41] Dougal Shakespeare, Lorenzo Porcaro, Emilia Gómez, and Carlos Castillo. 2020. Exploring Artist Gender Bias in Music Recommendation. In *Proc. 2nd Workshop on the Impact of Recommender Systems (ImpactRS 2020) at 14th ACM Conf. on Recommender Systems (RecSys 2020)*. <https://arxiv.org/abs/2009.01715>.
- [42] Dougal Shakespeare and Camille Roth. 2021. Tracing Affordance and Item Adoption on Music Streaming Platforms. In *Proceedings of the 22nd ISMIR Conference*. International Society of Music Information Retrieval.
- [43] Aaron van den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Advances in Neural Information Processing Systems*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.), Vol. 26. Curran Associates, Inc.
- [44] Koen Van Eijck. 2001. Social differentiation in musical taste patterns. *Social forces* 79, 3 (2001), 1163–1185.
- [45] Gabriel Viglienconi and Ichiro Fujinaga. 2016. Automatic Music Recommendation Systems: Do Demographic, Profiling, and Contextual Features Improve Their Performance?.. In *Proceedings of the 17th ISMIR Conference*. International Society of Music Information Retrieval, 94–100.
- [46] Jack Webster. 2020. Taste in the platform age: Music streaming services and new forms of class distinction. *Information, Communication & Society* 23, 13 (2020), 1909–1924.
- [47] Kris West and Paul Lamere. 2006. A model-based approach to constructing music similarity functions. *EURASIP Journal on Advances in Signal Processing* 2007, 1 (2006), 024602.
- [48] Nils Wlömert and Dominik Papies. 2016. On-demand streaming services and music industry revenues—Insights from Spotify’s market entry. *International Journal of Research in Marketing* 33, 2 (2016), 314–327.
- [49] Yuan Cao Zhang, Diarmuid Ó Séaghdha, Daniele Quercia, and Tamas Jambor. 2012. Auralist: introducing serendipity into music recommendation. In *Proceedings of the fifth ACM international conference on Web search and data mining*. 13–22.
- [50] Elena Zheleva, John Guiver, Eduarda Mendes Rodrigues, and Nataša Milić-Frayling. 2010. Statistical models of music-listening sessions in social media. In *Proceedings of the 19th international conference on World wide web*. 1019–1028.
- [51] F. J. Zuiderveen Borgesius, D. Trilling, J. Möller, B. Bodó, C. H. de Vreese, and N. Helberger. 2016. Should we worry about filter bubbles? *Internet Policy Review* 5, 1 (2016).