

Measuring Commonality in Recommendation of Cultural Content: Recommender Systems to Enhance Cultural Citizenship

Andres Ferraro
andresferraro@acm.org
McGill University
Montréal, Canada

Fernando Diaz
Canadian CIFAR AI Chair
Google
Montréal, QC, Canada
diazf@acm.org

Gustavo Ferreira
gustavo.ferreira@mila.quebec
McGill University
Montréal, Canada

Georgina Born
University College London
London, United Kingdom
g.born@ucl.ac.uk

ABSTRACT

Recommender systems have become the dominant means of curating cultural content, significantly influencing the nature of individual cultural experience. While the majority of academic and industrial research on recommender systems optimizes for personalized user experience, this paradigm does not capture the ways that recommender systems impact cultural experience in the aggregate, across populations of users. Although existing novelty, diversity, and fairness studies probe how recommender systems relate to the broader social role of cultural content, they do not adequately center culture as a core concept and challenge. In this work, we introduce commonality as a new measure of recommender systems that reflects the degree to which recommendations familiarize a given user population with specified categories of cultural content. Our proposed commonality metric responds to a set of arguments developed through an interdisciplinary dialogue between researchers in computer science and the social sciences and humanities. With reference to principles underpinning non-profit, public service media (PSM) systems in democratic societies, we identify universality of address and content diversity in the service of strengthening cultural citizenship as particularly relevant goals for recommender systems delivering cultural content. Taking diversity in movie recommendation as a case study in enhancing pluralistic cultural experience, we empirically compare the performance of recommendation algorithms using commonality and existing utility, diversity, novelty, and fairness metrics. Our results demonstrate that commonality captures a property of system behavior complementary to existing metrics and suggest the need for alternative, non-personalized interventions in recommender systems oriented to strengthening cultural citizenship across populations of users. In this way, commonality contributes to a growing

body of scholarship developing ‘public good’ rationales for digital media and machine learning systems.

CCS CONCEPTS

• **Information systems**; • **Computing methodologies** → *Artificial intelligence*;

KEYWORDS

recommender systems, evaluation, cultural content, movies, public service media

ACM Reference Format:

Andres Ferraro, Gustavo Ferreira, Fernando Diaz, and Georgina Born. 2022. Measuring Commonality in Recommendation of Cultural Content: Recommender Systems to Enhance Cultural Citizenship. In *Sixteenth ACM Conference on Recommender Systems (RecSys '22)*, September 18–23, 2022, Seattle, WA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3523227.3551476>

1 INTRODUCTION

Online platforms that host cultural content (e.g. music, movies, and books) often use recommender systems to suggest and distribute items from their catalogs. These algorithmic recommendations are often designed to optimize personalization objectives such as precision or click-through rate [16] and, as such, are aligned with user-level metrics like retention and subscriptions.

Academic and industry research on recommender systems has converged on personalization as a paradigm, and while metrics linked to personalization are productive, they do not capture the wider shaping effects of recommender systems in aggregate nor measure the effects of recommender systems across a population of users. Research demonstrates that recommendations have cumulative effects in shaping the wider cultures and societies within which they are being used [1]. That said, these effects and how they inform recommender system evaluation and design remains relatively unexplored [6].

We advocate for the design of recommender systems delivering cultural content to be motivated not only by personalization—or associated commercial interests—but also by appropriate normative principles. For guidance we turn to the normative principles



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

RecSys '22, September 18–23, 2022, Seattle, WA, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9278-5/22/09.
<https://doi.org/10.1145/3523227.3551476>

underpinning non-profit, public service media (PSM) systems. We identify universality of address and content diversity in the service of strengthening cultural citizenship as particularly relevant for recommender systems delivering cultural content. If personalization attempts to maximize individual user satisfaction with a platform, then these PSM principles aim more to enhance the commonality of diverse cultural experiences across a population, building cultural citizenship. In this way, we contribute to a growing body of scholarship developing ‘public good’ rationales for digital media and machine learning systems [2, 4, 5, 26, 28, 39, 41]

In this light, in this paper we propose and discuss a new evaluation metric that measures the degrees to which a system familiarizes a given user population with a specified category or group of categories of cultural content. Our proposed commonality metric responds to a set of arguments developed through an interdisciplinary dialogue between researchers in computer science and the social sciences and humanities.

As a case study in the application of these normative principles and their goal of enhancing cultural citizenship across a population, we consider movie recommendations. We empirically compare the performance of more than twenty recommendation algorithms using the proposed commonality metric with existing utility, diversity, novelty, and fairness metrics. We analyze how our commonality metric complements previous metrics and discuss how it is aligned with the intended purpose. Our aim with this work is to measure commonality in the consumption of diverse categories of movies that are generally under-represented by existing recommender systems. To date, criticisms of recommender systems and machine learning systems for their capacity to reproduce forms of bias and discrimination have been met by interventions designed to redress such biases at the level of individual users. However, it is also important to identify means of counteracting biases and enhancing diversity as common experiences across a population of users – by developing recommender systems that promote diversity through counteracting racism, sexism, and the neglect of non-Western content as common experiences. Only in this way can the wider cultural changes called for by anti-racist and feminist critics as well as those sympathetic to these criticisms from the RecSys community be delivered.

2 BACKGROUND

Recommender systems have become the dominant means of curating cultural content in the digital era. Curation – or the selection and promotion of content to be distributed to consumers – is, however, a historical constant: consumers have always encountered the cultural content they wish to consume via some type of curation. Today, when cultural curation supplied by recommender systems is multiplied across billions of recommendations presented to users by online platforms, they significantly influence the nature of individual cultural experiences [38]. But this influence is also magnified and multiplied across time and across populations, regions, and cultures. In the short term, recommender systems influence individual cultural consumption and taste. In the medium and long terms, by employing data on consumer behavior and influencing consumer choices, they can shape cultural literacies as well as population-wide trends in consumption and taste [6]. As a

result, unlike previous means of distribution, there is a high degree of automatized intervention in the way people and communities encounter and experience cultural content. Despite their personalized address, recommender systems therefore have cumulative effects in shaping the wider cultures and societies within which they are being used. These effects have been relatively unexplored as a focus of research in the RecSys community.

Academic and industrial research on recommender systems has converged on personalization as a paradigm. While metrics linked to personalization and ‘user relevance’ are productive (e.g., NDCG, precision, recall), they do not capture the wider, aggregate shaping effects of recommender systems on patterns of cultural consumption, taste and literacy as described above. That is, they do not measure the effects of recommender systems’ use across populations of users.

Although concerned with and sensitive to the broader social role of cultural content, existing diversity and fairness evaluations of recommender systems do not adequately support the rich set of goals system designers might have. Typically, diversity metrics are limited to the goal of capturing the variety of content offered *within* a recommendation list; they may consider categorizations of the content, distances in a latent space, or simply how many different items are recommended [1, 17, 20, 34]. Aligned with the goals of personalization, the formulation of these diversity metrics sometimes optionally considers the relevance of the content for users, assuming that what a user consumed in the past indicates what they are still interested in, so recommendations should be limited to such categories. While related novelty metrics measure the newness of items or categories of recommendations, they are still individualized and agnostic about *what type* of content is new to the user. The work on fairness addresses specific topics of increasing biases and under-representation of particular groups [11]. *Provider fairness* metrics typically consider how many different groups of content providers appear in recommendations and assume a given distribution that it is desired to match. *Consumer fairness* considers disparate treatments of the system to different groups of consumers. Recent research proposed more general *multi-stakeholder fairness*, acknowledging the impact recommender systems have for the different groups of individuals [8, 24, 35].

We suggest that it is timely for the design of recommender systems delivering cultural content to be motivated not only by individualized interests but by appropriate normative principles oriented to furthering the democratic development of contemporary societies. By normative we refer to principles considered to provide models of morally, ethically and/or politically right or just action or behavior in the interests of democratic societies as well as individuals. In this way we contribute to growing scholarship developing ‘public good’ rationales for digital media and ML systems [2, 4, 5, 26, 28, 39, 41], advocating ‘a computational politics wedded to emancipation and human flourishing’ [36]. For guidance we turn to the principles underpinning non-profit, public service media (PSM) systems, proposing that ‘a public service rationale is as pertinent as ever in the digital era’ [2]. The normative ideas underpinning public service media developed over the last century in the context of democratic states committed to enhancing democratic and representative channels of communication.

A substantial body of research in media and political theory has identified the normative principles informing PSM systems, among them universality, citizenship, and diversity [3–5, 7, 31, 33]. We consider this triad – universality (or commonality) of address, citizenship, and diversity of content – as particularly relevant for recommender systems delivering cultural content, since together they answer calls for digital media systems to enhance cultural citizenship [3]. The concept of cultural citizenship has become foundational for democratic political theories in the last two decades. It responds to recognition of the challenges posed by globalization, migration, the growing heterogeneity of the populations of nation states, and the intensification of identity politics among subaltern and marginalized groups [25, 32]. It draws attention to a ‘new domain of cultural rights [involving] the right to symbolic presence, dignifying representation’, and ‘the maintenance and propagation of distinct cultural identities’ [30]. In this light, PSM – and other democratic distribution and curation media – should promote cultural citizenship by curating and disseminating a plurality of cultural content stimulating intercultural dialogue and ‘acceptance of, and respect for, cultural diversity’. In this way PSM and other democratic media can act both as a force ‘for social cohesion and integration’ and as a forum for pluralistic cultural experience among groups and communities coexisting in democratic societies [19]. Both universality or commonality—the provision of common cultural experiences—and diversity of content are therefore essential to the strengthening of cultural citizenship: ‘mutual cultural recognition and the expansion of cultural referents... are dynamics essential to the well-being of pluralist societies. But this does not obviate the need also for integration—for the provision of common [cultural] experience and the fostering of common identities’ [3]. Scholarship on these matters emphasizes that implementing principles like universality (commonality), citizenship, and diversity require ‘alternative success metrics... focused on PSM’s impact on democracy and the public sphere’ which address users ‘as citizens and not just... as consumers’ [39]. Such metrics will enable democratic media to adapt to the present by advancing ‘cultural citizenship and the needs of the digital society’ [19].

As a case study we take diversity in movie recommendation, with the aim of measuring commonality in the consumption of diverse categories of movies that are as yet under-represented in existing recommender systems. Current recommender systems are criticized for a tendency to reproduce or exacerbate wider forms of cultural and social discrimination [13, 29]. Previous works show that movie recommender systems may generate feedback loops between recommendation and consumption. As a result, they amplify popularity bias, promote narrow options in terms of cultural content, and homogenize users’ identity profiles, with a stronger effect on gender minorities [22]. In movie and music domains, commercial popular content may be privileged for users in dominant groups according to age and gender [12]. From the provider side in the music domain, female and non-binary artists can be under-represented in recommendations, reflecting industry gender biases and reducing the diversity of content being recommended which may affect users’ future streams [14, 15]. These provider biases shown by under-representation of cultural content with respect to gender, race, class, and region correspond to ‘core-periphery dynamics and geographical inequality’ in the cultural industries [9, 37, 42].

Recommender systems often mirror these inequalities, promoting Western-centric popular cultural content in the English language, released by major producers [37, 43]. Generally, criticisms of recommender systems and machine learning systems for reproducing such forms of bias and discrimination have been met by personalized recommender systems interventions aimed at redressing bias at the level of individual users. In light of our discussion of culture, we contend that it is also important to identify means of counteracting bias and enhancing the diversity of content offered across a population of users. The proposed commonality metric achieves this by measuring common experiences of diversity at the aggregate level. Assuming a democratic media environment, the metric provides a way of assessing whether recommender systems are contributing to the strengthening of cultural citizenship by systematically promoting diversity within a given type of cultural content (here, movies). At the same time, it has the potential to assist in counteracting racism, sexism, and the neglect of non-Western and non-mainstream content across a user population (via commonality). In this way, the metric is a means of measuring the extent of the kinds of wider cultural changes called for by anti-racist and feminist critics as well as by those sympathetic to criticisms of existing recommender systems.

3 MEASURING COMMONALITY

Recall that we are interested in measuring the extent to which users will gain a shared familiarity with a set of promoted categories. The promoted categories, we suggest, will be identified and curated by experts in a relevant field (here, movies) seeking to promote a plurality of cultural content in the service of strengthening cultural citizenship. We contrast this with Mehrotra et al.’s purely statistical method for selecting under-represented categories [23], which may surface under-represented content misaligned with our twin goals of boosting a diversity of cultural experience across a user population, and at the same time enhancing their common experience of that diversity. An expert may opt to promote, for example, movies by female directors as well as those produced for non-Western markets.

The *commonality* of a system captures the probability that every user simultaneously gains familiarity with the editorially-selected categories. Let \mathcal{U} be the set of system users where $m = |\mathcal{U}|$ and \mathcal{D} the catalog of items where $n = |\mathcal{D}|$. Given a user $u \in \mathcal{U}$, a system produces a ranking π_u of items in \mathcal{D} ; we will use π to refer to the set of all rankings. We are interested in measuring the degree to which the system supports the normative value of commonality. Given a set of editorially-selected categories \mathcal{G} , we can compute the commonality of a system with respect to a single category $g \in \mathcal{G}$ as the probability that every user has become familiar with g under the system’s ranking,

$$\begin{aligned} C_g(\pi) &= p(F_{1,g}, \dots, F_{m,g} | \pi) \\ &= \prod_{u \in \mathcal{U}} p(F_{u,g} | \pi_u) \end{aligned} \quad (1)$$

where $F_{u,g}$ is a binary random variable indicating the familiarity that user u has with category g .

In order to estimate $p(F_{u,g} | \pi_u)$, the familiarity of a user with a category after the recommender intervention, we adopt a standard

browsing model from existing evaluation metrics. Specifically, we are interested in the comprehensiveness of a user’s exposure to items from g after browsing the ranking,

$$p(F_{u,g}|\pi_u) = \sum_{i=1}^n p(i)R(\pi_u, i, g) \quad (2)$$

where $p(i)$ is the probability that the user stops at rank i and $R(\pi, k, g)$ is the recall of items with category g in π at rank cut-off k . For our experiments, we adopt an exponential discount, $p(i) = (1 - \gamma)\gamma^{i-1}$ based on the rank-biased precision browsing model [27]. For a set of categories, we measure commonality using the mean commonality, $C_G(\pi) = \mathbb{E}_{g \in \mathcal{G}} [C_g(\pi)]$.

4 EXPERIMENTS

The goal of these experiments is to empirically compare commonality with existing metrics when ranking different recommender systems. Using movies as a case study, we focus on analyzing how commonality complements previous metrics in the consumption of diverse categories of movies that are generally under-represented by existing recommender systems. We then discuss how commonality is aligned with the intended purpose of promoting a shared plurality of cultural content in the service of strengthening cultural citizenship.

Data: We use the movielens-1m dataset, which contains 1,000,209 ratings of 3,900 movies from 6,040 users from the movielens platform. Using a separate dataset¹, we augmented the movielens movies with metadata including country of production, gender of the director, original language, and keywords collected from the movie’s description. For this dataset, we used rankings from multiple recommendation systems prepared by Valcarce et al. [40]. Following the method described by the authors we converted to binary relevance labels considering ratings of 4 and 5 as relevant.

Categories: We selected categories of movies that are typically under-represented by movie recommender systems. Specifically, we consider female directors (under-representation by gender); independent film (under-representation by industry sector); and several sources of non-Western film (under-representation by geographical and linguistic inequality). We use categorical gender data, acknowledging the limitations of this framing [18]. For geographic categories, we use the country of production for the following regions of the world: South America, Central America, North Africa, South Africa, West Africa, Mid Africa, Southeast Asia, South Asia, Western Asia, Central Asia and East Asia. We consider, broadly, non-English language movies as a separate category. And, finally, we use keywords to create categories with selected movies whose categories contain “independent films”, “LGBT”, and “transgender”. We manually checked whether these keywords can be trusted to represent specific identities.

Baseline Metrics: We compare commonality against three classes of metrics: utility metrics, diversity metrics, and fairness metrics [10]. For utility, we consider normalized discounted cumulative gain (NDCG) and reciprocal rank (RR). For diversity, we consider α -NDCG and intent-aware expected reciprocal rank (ERR-IA). For fairness, we use ranking-based statistical parity (RSP) and ranking-based equal-opportunity (REO) [44].

¹<https://www.themoviedb.org/>

Table 1: Correlation between commonality with existing metrics. Kendall’s τ between rankings of runs. ** indicates $p < 0.05$.

NDCG	RR	α -NDCG	ERR-IA	REO	RSP
-0.4476**	-0.5048**	-0.5333**	-0.5143**	0.2316	0.1053

5 RESULTS

We were first interested in understanding the redundancy between commonality and existing metrics. To this end, we measured Kendall’s τ between system rankings according to commonality with system rankings according to our baseline metrics. We present the results of this analysis in Table 1. As expected, the correlation between commonality and utility metrics (NDCG, RR) is negative and significant, indicating that we would tradeoff commonality and utility (personalization) if selecting from existing systems. Moreover, this suggests that, for existing systems, if we select the higher utility systems during model development, it will compromise our goals of commonality and cultural citizenship. We observed a similar negative correlation with existing diversity metrics, largely due to the utility-orientation of the metrics. Finally, we did not find evidence of correlation between commonality and fairness metrics. Collectively, these results suggest that commonality measures properties of a set of recommendations not present in existing metrics.

In order to explore the relationship between commonality and utility, we plotted the NDCG and commonality values for our twenty runs in Figure 1a. The run with the highest NDCG, Bayesian personalized ranking matrix factorization (BPRMF), also exhibited a very low average commonality across categories, relative to other runs; we contrast this with the run with the lowest NDCG, random, that exhibited comparatively much higher average commonality across categories. The run with the highest average commonality, SVD, on the other hand, generated lower-utility rankings.

In Figure 1b, we further disaggregate the metric for three runs: BPRMF, random, and popularity. We observed that popularity-based ranking offers the highest commonality for the “female” and “north africa” categories while random ranking offers a the best commonality for all categories except “south america”. In general, personalization-based recommendation had low commonality across all categories.

6 CONCLUSIONS AND FUTURE WORK

In this work, motivated by defining metrics for recommendation of cultural content, we developed a method to measure alignment with principles of cultural citizenship that we adapted from the PSM literature.

Our proposed commonality metric emphasizes shared familiarity, by which is intended the simultaneous exposure of users to content from selected categories. This definition, captured by the joint probability of familiarity events, is worth exploring for both theoretical and pragmatic reasons. Theoretically, our approach to shared exposure is very conservative, penalizing the metric if even a single user has a low probability of exposure. That said, this conservativeness attenuates systematic under-exposure to certain categories. Pragmatically, since the joint probability is often quite small for our evaluated systems, we can encounter numerical stability issues.




Figure 1: Behavior of commonality. 1a Scatterplot of mean NDCG versus mean commonality across categories. Each point is a single recommender system run. 1b Per-category commonality values for recommendations based on random, popularity, and BPRMF models. Horizontal lines indicate the mean commonality across categories.

This is especially salient if we average commonality across categories, where a single category may have an exponentially larger value and dominate the mean. However, this may indicate gross and systemic under-performance of existing systems in terms of commonality.

In addition to commonality, we introduce a relatively simple model of familiarity based on recall. We believe there is opportunity to develop alternative models of familiarity that consider a user’s previous experience with the category or other contextual information. However, the design of a familiarity model should be aligned with the concept of shared experience, meaning that, even if a user has engaged with content from a category in the past, *re-exposing* them may promote commonality at the risk of over-satiating users with niche interests, a topic of recent research [21].

Our results demonstrate that existing high-utility recommendation algorithms under-perform in terms of commonality. We believe that exploring the space of commonality-informed recommendation can produce algorithms that perform substantially better in terms of commonality while maintaining high utility.

In summary, our preliminary results indicate a new class of evaluation metrics specifically aimed at measuring the effects of recommendation systems in terms of common exposure to diverse items within a given area of cultural content. Our future work includes assessing important evaluative properties of commonality such as generalizability and robustness. We are also interested in the joint optimization of both personalization and commonality as a new class of algorithm. Future extensions to our commonality metric may also consider the cumulative effects over time of recommender systems attuned to enhancing common experience of diverse content and thereby cultural citizenship.

ACKNOWLEDGMENTS

This research project forms part of the research program ‘Music and Artificial Intelligence: Building Critical Interdisciplinary Studies’ (MusAI), directed by PI Prof. Georgina Born, which has received

funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101019164). This project was supported in part by the Canada CIFAR AI Chairs program.

REFERENCES

- [1] 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*. Association for Computing Machinery, New York, NY, USA, 2155–2165. <https://doi.org/10.1145/3366423.3380281>
- [2] Mark Andrejevic. 2013. Public service media utilities: Rethinking search engines and social networking as public goods. *Media International Australia* 146, 1 (2013), 123–132.
- [3] Georgina Born. 2006. Digitising Democracy. In *What can be done?: making the media and politics better*, John Lloyd and Jean Seaton (Eds.). Oxford: Blackwell, 102–23.
- [4] Georgina Born. 2018. Principles of Public Service for the 21st Century. In *A Future for Public Service Television*, Des Freedman and Vana Goblot (Eds.). [The MIT Press]. <https://doi.org/10.7551/mitpress/9781906897710.003.0015>
- [5] Georgina Born. 2018. Taking the Principles of Public Service Media into the Digital Ecology. In *A Future for Public Service Television*, Des Freedman and Vana Goblot (Eds.). [The MIT Press], 181–190. <https://doi.org/10.7551/mitpress/9781906897710.003.0025>
- [6] Georgina Born, Jeremy Morris, Fernando Diaz, and Ashton Anderson. 2021. Artificial Intelligence, Music Recommendation, and the Curation of Culture. *Schwartz Reisman Institute for Technology and Society White Paper* (2021).
- [7] Georgina Born and Tony Prosser. 2001. Culture and Consumerism: Citizenship, Public Service Broadcasting and the BBC’s Fair Trading Obligations. *The Modern Law Review* 64, 5 (2001), 657–687. <http://www.jstor.org/stable/1097275>
- [8] Robin D Burke, Himan Abdollahpouri, Bamshad Mobasher, and Trinadh Gupta. 2016. Towards multi-stakeholder utility evaluation of recommender systems. *UMAP (Extended Proceedings)* 750 (2016).
- [9] Minerva Campos-Rabadán. 2020. Inequalities within the international film arena. A framework for studying Latin American film festivals. *Comunicación y Medios* 29, 42 (2020), 70–82. <https://doi.org/10.5354/0719-1529.2020.57224>
- [10] Rocio Cañameres, Pablo Castells, and Alistair Moffat. 2020. Offline evaluation options for recommender systems. *Information Retrieval Journal* (2020).
- [11] Michael D. Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. 2022. Fairness in Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, New York, NY, 679–707. https://doi.org/10.1007/978-1-0716-2197-4_18
- [12] Michael D Ekstrand, Mucun Tian, Ion Madrazo Azpiazu, Jennifer D Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All the cool kids, how do they fit in?: Popularity and demographic biases in recommender evaluation and effectiveness. In *Conference on fairness, accountability and transparency*. PMLR, 172–186.

- [13] Michael D. Ekstrand, Mucun Tian, Mohammed R. Imran Kazi, Hoda Mehrpouyan, and Daniel Klüver. 2018. Exploring Author Gender in Book Rating and Recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) (*RecSys '18*). Association for Computing Machinery, New York, NY, USA, 242–250. <https://doi.org/10.1145/3240323.3240373>
- [14] Avriel Epps-Darling, Henriette Cramer, and Romain Takeo Bouyer. 2020. Artist gender representation in music streaming.. In *ISMIR*. 248–254.
- [15] Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. Break the loop: Gender imbalance in music recommenders. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. 249–254.
- [16] Asele Gunawardana, Guy Shani, and Sivan Yogev. 2022. Evaluating Recommender Systems. Springer US, New York, NY, 547–601.
- [17] Jungkyu Han and Hayato Yamana. 2017. A Survey on Recommendation Methods Beyond Accuracy. *IEICE Trans. Inf. Syst.* 100-D (2017), 2931–2944.
- [18] Anna Lauren Hoffmann. 2019. Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society* 22, 7 (2019), 900–915.
- [19] Karol Jakubowicz. 2007. Public Service Broadcasting in the 21st century. *From public service broadcasting to public service media. RIPE@ 2007* (2007), 29–49.
- [20] Yucheng Jin, Nava Tintarev, and Katrien Verbert. 2018. Effects of Individual Traits on Diversity-Aware Music Recommender User Interfaces. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization* (Singapore, Singapore) (*UMAP '18*). Association for Computing Machinery, New York, NY, USA, 291–299. <https://doi.org/10.1145/3209219.3209225>
- [21] Liu Leqi, Fatma Kilinc-Karzan, Zachary C. Lipton, and Alan L. Montgomery. 2021. Rebounding Bandits for Modeling Satiation Effects.
- [22] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 2145–2148.
- [23] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness and Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM conference on Information and knowledge management*.
- [24] Silvia Milano, Mariarosaria Taddeo, and Luciano Floridi. 2021. Ethical aspects of multi-stakeholder recommendation systems. *The information society* 37, 1 (2021), 35–45.
- [25] Toby Miller. 2001. Introducing... cultural citizenship. *Social Text* 19, 4 (2001), 1–5.
- [26] Hallvard Moe. 2008. Dissemination and dialogue in the public sphere: a case for public service media online. *Media, Culture & Society* 30, 3 (2008), 319–336.
- [27] Alistair Moffat and Justin Zobel. 2008. Rank-biased Precision for Measurement of Retrieval Effectiveness. *ACM Trans. Inf. Syst.* 27, 1 (Dec. 2008), 2:1–2:27.
- [28] Graham Murdock. 2005. Building the digital commons. *Cultural dilemmas in public service broadcasting* (2005), 213–31.
- [29] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kicman. 2019. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *Frontiers in Big Data* 2 (2019). <https://doi.org/10.3389/fdata.2019.00013>
- [30] Jan Pakulski. 1997. Cultural citizenship. *Citizenship studies* 1, 1 (1997), 73–86.
- [31] Marc Raboy. 1996. *Public broadcasting for the 21st century*. Vol. 17. Indiana University Press.
- [32] Renato Rosaldo. 1994. Cultural Citizenship in San Jose, California. *PoLAR* 17 (1994), 57.
- [33] Paddy Scannell. 1989. Public service broadcasting and modern public life. *Media, Culture & Society* 11, 2 (1989), 135–166.
- [34] Markus Schedl and David Hauger. 2015. Tailoring Music Recommendations to Users by Considering Diversity, Mainstreamness, and Novelty. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Santiago, Chile) (*SIGIR '15*). Association for Computing Machinery, New York, NY, USA, 947–950. <https://doi.org/10.1145/2766462.2767763>
- [35] Nasim Sonboli, Robin Burke, Michael Ekstrand, and Rishabh Mehrotra. 2022. The multisided complexity of fairness in recommender systems. *AI Magazine* 43, 2 (2022), 164–176. <https://doi.org/10.1002/aaai.12054>
- [36] Luke Stark. 2018. Algorithmic psychometrics and the scalable subject. *Social Studies of Science* 48, 2 (2018), 204–231.
- [37] Tamas Tofalvy and Júlia Koltai. 2021. “Splendid Isolation”: The reproduction of music industry inequalities in Spotify’s recommendation system. *new media & society* (2021), 14614448211022161.
- [38] Matus Tomlein, Branislav Pecher, Jakub Simko, Ivan Srba, Robert Moro, Elena Stefancova, Michal Kompan, Andrea Hrcokova, Juraj Podrouzek, and Maria Bielikova. 2021. An Audit of Misinformation Filter Bubbles on YouTube: Bubble Bursting and Recent Behavior Changes. In *Fifteenth ACM Conference on Recommender Systems*. 1–11.
- [39] Klaus Unterberger and Christian Fuchs. 2021. *The public service media and public service internet manifesto*. University of Westminster Press.
- [40] Daniel Valcarce, Alejandro Bellogín, Javier Parapar, and Pablo Castells. 2018. On the Robustness and Discriminative Power of Information Retrieval Metrics for Top-N Recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) (*RecSys '18*). Association for Computing Machinery, New York, NY, USA, 260–268. <https://doi.org/10.1145/3240323.3240347>
- [41] Hilde Van den Bulck and Hallvard Moe. 2018. Public service media, universality and personalisation through algorithms: mapping strategies and exploring dilemmas. *Media, Culture & Society* 40, 6 (2018), 875–892.
- [42] Deb Verhoeven, Bronwyn Coate, and Vejune Zemaityte. 2019. Re-distributing gender in the global film industry: Beyond# MeToo and# MeThree. *Media Industries Journal* 6, 1 (2019).
- [43] Michael Matthias Voit and Heiko Paulheim. 2021. Bias in Knowledge Graphs—an Empirical Study with Movie Recommendation and Different Language Editions of DBpedia. *arXiv preprint arXiv:2105.00674* (2021).
- [44] Ziwei Zhu, Jianling Wang, and James Caverlee. 2020. Measuring and mitigating item under-recommendation bias in personalized ranking systems. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*. 449–458.