Matching Consumer Fairness Objectives & Strategies for RecSys*

Michael D. Ekstrand†
ekstrand@acm.org
People and Information Research Team
Boise State University
Boise, Idaho, USA

Maria Soledad Pera†‡
M.S.Pera@TUDelft.nl
Web Information Systems Group
Delft University of Technology
Delft, The Netherlands

ABSTRACT

The last several years have brought a growing body of work on ensuring that recommender systems are in some sense *consumer-fair* — that is, they provide comparable quality of service, accuracy of representation, and other effects to their users. However, there are many different strategies to make systems more fair and a range of intervention points. In this position paper, we build on ongoing work to highlight the need for researchers and practitioners to attend to the details of their application, users, and the fairness objective they aim to achieve, and adopt interventions that are appropriate to the situation. We argue that consumer fairness should be a creative endeavor flowing from the particularities of the specific problem to be solved.

1 PATHS TO CONSUMER FAIRNESS

Fair recommendation is a complex and multi-sided problem [20, 49], with a significant focus on providing a fair experience to one or both of two main stakeholders: producers (who provide the items or services to be suggested) and users (who consume the provided recommendations) [13]. We are particularly interested in the latter group, for whom recommender systems (**RS**) have to offer appealing items while considering that "the best items for one user may be different than those for another" [13]. *Consumer fairness* [12] is the aspect of fairness concerned with ensuring that the users (or "consumers") of a RS are treated fairly in the quantitative and/or qualitative aspects of their experience. The relevant literature considers several ideas of what it means to be "fair" to consumers, along with different techniques to measure or attain such fairness; one particularly common goal is to ensure that certain users or groups of users do not receive a systematically lower-quality or less-useful experience than others [22, 35, 41].

This interest mirrors a line of work on specific user audiences. Ekstrand et al. [22] show recommender performance can differ between users of different genders and ages. Explorations of children's media use [37, 50] reveal that preferred traits in songs and books vary from childhood to early adulthood, indirectly urging RS work to treat "children" not as a monolithic entity, but as individuals to better serve them. Researchers have suggested going beyond traditional popularity- or collaborative-filtering algorithms that would inevitably prioritize the majority of the consumers (i.e., adults) to explicitly consider factors like the readability levels (comprehension), familiarity with concepts covered in the classroom (learning), and explainability (engagement and improve task performance), if suggestions are to be suitable–and therefore apt for consumption [38, 39, 44, 45, 53]. Literature bringing awareness to autism [8, 34, 43] emphasizes that RS should account for user-specific sensory aversions or skill limitations of recommended items are to be compatible with what these users require, and hence useful.

Several concerns from the broader RS literature can also be regarded as forms of consumer fairness. Examples include macro-averaging evaluation metrics by user [24, 52] to assess the experience of all users instead of emphasizing highly-active users [19, 20] and providing good results to new users [20].

The works presented thus far share the common goal of providing effective, often personalized, experiences to all their users. They do so through a variety of definitions, methods, and points of intervention (where the RS is changed to advance the goal). Ekstrand et al. [20, §5] have cataloged many of the existing strategies and noted some challenges in matching a strategy to specific fairness objectives. Expanding on that argument, our proposition in this paper is that **researchers and practitioners need to select interventions that are appropriate to the specific fairness goal(s) and particularities of an application context**. More importantly, we hope to see a robust discussion between researchers, practitioners, and stakeholder representatives from different disciplinary perspectives to understand how best to promote RS that are "good" — in multiple relevant ways — for everyone who uses them.

2 TYPES OF FAIRNESS OBJECTIVES

Numerous fairness objectives have been studied under the banner of consumer fairness. Perhaps the most well-known is *equity of utility*: ensuring that a RS (or other information access system) provides comparable quality of service to all users or groups of users [e.g. 22, 28, 31, 35, 50], typically measured by online or offline effective measures such as nDCG or click-through rate. A related objective is *equity of usability*: ensuring that people can actually use the system, either in addition to or independent of considering equity in the utility of results [6, 27, 34, 43, 47]. Accessibility is a clear concern here, as a system that does not work with screen readers, for example, cannot be used as easily by visually-impaired users [4, 15, 23, 30, 36, 54]. Other works focus on attending to *specific information needs* that a particular group of users may have that are not effectively met by systems more attuned to needs common among the majority of the population [5, 7, 47].

Looking past the effectiveness and usability of a RS, some consumer fairness work has looked at issues of *fair representation* or *representational harms* [16], in terms of either the RS's internal representation of the user (e.g. avoiding user embeddings that may lead to stereotyped recommendations [11]) or the recommended items themselves. One example of this last concern is the objective of *recommendation independence* [29]: this goal is satisfied if the

*Presented at FAccTRec 2022.
†Both authors contributed equally to this paper.
‡Work begun while at Boise State University.

probability of a particular item being recommended to the user is independent of their gender or other protected status.

Effective consumer fairness must begin by identifying an objective to pursue or problem to solve, as the choice of operationalization and intervention (§3) follows from the objective [21].

3 INTERVENTION STRATEGIES

Prior work has proposed various strategies to advance one or more fairness objectives (§2). Here, we mention some salient ones, grouped by the stage of the RS at which they intervene.

Design interventions. Designing RS to adapt to users in the quest for consumer fairness involves the use of multiple interfaces that are matched to users' needs. For example, Deldjoo et al. [17] proposed a child-oriented TV/movie recommendation interface for in-home set-top boxes that incorporated tangible interaction: the child could hold up a toy truck to get recommendations for shows about trucks. Another common alternative is to detect the particular group a user belongs to and adapt RS behavior and/or interface to the corresponding group. Practical applications of this strategy include, upon identification of the grade or skill of the target user, modifying the types of queries that are recommended Madrazo Azpiazu et al. [33], showcasing different multi-modal cues to point users towards suitable spelling suggestions [18], or adapting presented choices to enable knowledge acquisition [47].

Algorithmic interventions. Modifying recommendation algorithms is common. This means, for instance, including the inter-user equity objective into the loss function [26, 55, 56], sometimes through a regularization term [29, 58]. Reranking [31] can also reduce gaps in utility by post-processing recommendations from an existing model. These can be applied to many objectives beyond equity of utility. Penalizing dependence between recommended items and user attributes [29] is another alternative.

Adversarial learning methods can also help reduce unfairness. Beutel et al. [11] use a discriminator to learn user embeddings that are not predictive of sensitive attributes such as race or gender; more broadly, fair representation learning [32, 59] can be applied to consumer fairness [57]. There are also many other algorithmic strategies considered as well, such as changing neighborhoods [14].

Data interventions. Some strategies manipulate the RS's input data to improve fairness, e.g. by injecting fake user profiles [46] or removing spam reviews [48].

Process interventions. Improvements to engineering and quality assurance processes can be useful for providing consumer fairness. Regular auditing for violations of fairness objectives [25], through disaggregated evaluations [9, 35] or other means, identify problems and help detect regressions on past fairness improvements.

The engineering process is another place to improve a system's fairness. Little has been little published on this, but studies that reveal *why* a fairness problem occurs may enable engineers and model owners to identify and prioritize software improvements that will address the problem, even if they are not directly fairness interventions. For example, if a music recommender performs poorly for users from a particular region due to lower-quality song metadata, investing in that data could improve equity of utility.

Marketplace interventions. Consumer fairness can also call for the development of new RS targeted at under-served groups. This can be done either by new entrants to the market or existing firms seeking to shore up their market position. Consider popular sites like Goodreads and Amazon: the segment of their user base producing the most interactions, and hence driving recommendation algorithms, are adults. In turn, the resulting experience may not suit children. Some startups are trying to fill this perceived gap by creating new sites specifically for children; examples include ABC Mouse [1], BiblioNasium [2], or Pickatale [3]. As for examples of sites aiming to expand their target audiences, we find Netflix offering recommendations specific to children and families [42] or Spotify, which now offers Spotify Kids [51] as an alternative to better support children.

4 MATCHING OBJECTIVES AND STRATEGIES

Our central proposition in this paper is that the choice of *where* in the RS and its sociotechnical context to intervene, and *how* to intervene at that point, needs to be well-matched to the specific fairness objective and details of the application, domain, and users.

Some pairings of strategies and outcomes are better-matched than others. E.g., auditing differences in utility [9, 22, 35] can identify unfair utility and provide an empirical starting point for many potential strategies, including design and process interventions, but not every intervention strategy is likely a good fit for this objective. Ekstrand et al. [20] note that useful recommendations in most domains are not a subtractable good [10] (users do not compete with each other for good recommendations). The inequity itself is not the problem, but rather a symptom of the system not providing some of its users with good recommendations. Training to minimize differences in utility [e.g. 26, 31, 40] can ensure equity, but at the risk of placing users in competition with each other, sometimes with significant majority-group utility loss [31]. Positive-sum rather than zero-sum utility aggregates avoids the competition problem [55], as do interventions that seek to directly address the causes of under-serving a segment of the user base [20].

A better-matched pairing involving algorithmic intervention is Beutel et al. [11]'s use of adversarial learning to remove unwanted correlations between user embeddings and user group membership in hopes of producing less stereotypical recommendations.

A single objective may have significantly more complexity and nuance than is accounted for by simple strategies. For example, what counts as a good recommendation may differ between groups [27] and contexts. Here, pursuing an objective such as equity of utility should consider whether metrics accurately measure utility across the varied constituencies and contexts in an RS's usage.

We invite the broad community of people concerned with ensuring fair access to information through RS and related information access systems to think carefully and interdisciplinarily about the specific problems to be solved and select appropriate, not just convenient or familiar, interventions. Further research is needed to understand how to implement the various interventions in §3 (and more not listed) most effectively, and to more thoroughly decompose the problem space of consumer fairness. Such research will identify when different interventions may or may not be appropriate, and provide evidence-based guidance for future practice.

ACKNOWLEDGMENTS

Partially supported by National Science Foundation grant 17-51278.

REFERENCES

[1] 2022. ABC Mouse. https://www.abcmouse.com/.
[2] 2022. BiblioNasium. https://www.biblionasium.com/.
[3] 2022. Pickatale. https://pickatale.com/.
[4] Vasyl Andrunyk, Volodymyr Pasichnyk, Tetiana Shestakevych, and Natalya Antonyuk. 2019. Modeling the Recommender System for the Synthesis of Information and Technology Complexes for the Education of Students with Autism. In *Proceedings of IEEE 14th International Conference on Computer Sciences and Information Technologies (CSIT)*, Vol. 3. 183–186. https://doi.org/10.1109/STC-CSIT.2019.8929776
[5] Vasyl Andrunyk, Tetiana Shestakevych, Volodymyr Pasichnyk, and Nataliia Kunanets. 2020. Information Technologies for Teaching Children with ASD. In *Advances in Computer Science for Engineering and Education II*. Springer International Publishing, 523–533. https://doi.org/10.1007/978-3-030-16621-2_49
[6] Oghenemaro Anuyah, Michael Green, Ashlee Milton, and Maria Soledad Pera. 2019. The Need for a Comprehensive Strategy to Evaluate Search Engine Performance in the Classroom. In *KidRec '19: Workshop in International and Interdisciplinary Perspectives on Children & Recommender and Information Retrieval Systems, Co-located with ACM IDC*. https://kidrec.github.io/papers/KidRec_2019_paper_1.pdf
[7] Ion Madrazo Azpiazu, Nevena Dragovic, Maria Soledad Pera, and Jerry Alan Fails. 2017. Online searching and learning: YUM and other search tools for children and teachers. *Information Retrieval Journal* 20, 5 (2017), 524–545. https://doi.org/10.1007/s10791-017-9310-1
[8] Alisha Banskota and Yiu-Kai Ng. 2020. Recommending Video Games to Adults with Autism Spectrum Disorder for Social-Skill Enhancement. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*. ACM, 14–22. https://doi.org/10.1145/3340631.3394867
[9] Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Krones, Meredith Ringel Morris, Jennifer Wortman Vaughan, W Duncan Wadsworth, and Hanna Wallach. 2021. Designing Disaggregated Evaluations of AI Systems: Choices, Considerations, and Tradeoffs. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*. ACM, 368–378. https://doi.org/10.1145/3461702.3462610
[10] C Dustin Becker and Elinor Ostrom. 1995. Human Ecology and Resource Sustainability: The Importance of Institutional Diversity. *Annual Review of Ecology and Systematics* 26, 1 (1995), 113–133. https://doi.org/10.1146/annurev.es.26.110195.000553
[11] Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. 2017. Data Decisions and Theoretical Implications when Adversarially Learning Fair Representations. (2017). http://arxiv.org/abs/1707.00075
[12] Ludovico Boratto, Gianni Fenu, Mirko Marras, and Giacomo Medda. 2022. Consumer Fairness in Recommender Systems: Contextualizing Definitions and Mitigations. In *Lecture Notes in Computer Science*. Springer International Publishing, 552–566. https://doi.org/10.1007/978-3-030-99736-6_37
[13] Robin Burke. 2017. Multisided Fairness for Recommendation. (2017). http://arxiv.org/abs/1707.00093
[14] Robin Burke, Nasim Sonboli, and Aldo Ordonez-Gauger. 2018. Balanced Neighborhoods for Multi-sided Fairness in Recommendation. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A Friedler and Christo Wilson (Eds.). PMLR, 202–214. http://proceedings.mlr.press/v81/burke18a.html
[15] Wei Chen, Li-Jun Zhang, Can Wang, Chun Chen, and Jia-Jun Bu. 2008. Pervasive Web News Recommendation for Visually Impaired People. In *Proceedings of IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT '08)*. 119–122. https://doi.org/10.1109/WIIAT.2008.43
[16] Kate Crawford. 2017. The Trouble with Bias. Neural Information Processing Systems 2017. https://youtu.be/fMym_BKWQzk
[17] Yashar Deldjoo, Cristina Frà, Massimo Valla, Antonio Paladini, Davide Anghileri, Mustafa Anil Tuncel, Franca Garzotto, and Paolo Cremonesi. 2017. Enhancing Children's Experience with Recommendation Systems. In *KidRec 2017*. https://yasdel.github.io/files/KidRec17_deldjoo.pdf
[18] Brody Downs, Maria Soledad Pera, Katherine Landau Wright, Casey Kennington, and Jerry Alan Fails. 2021. KidSpell: Making a difference in spellchecking for children. *International Journal of Child-Computer Interaction* 32 (2021), 100373. https://doi.org/10.1016/j.ijcci.2021.100373
[19] Michael Ekstrand, John Riedl, and Joseph A Konstan. 2010. Collaborative Filtering Recommender Systems. *Foundations and Trends® in Human-Computer Interaction* 4, 2 (2010), 81–173. https://doi.org/10.1561/1100000009
[20] Michael D Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. 2022. Fairness in Information Access Systems. *Foundations and Trends® in Information Retrieval* 16, 1-2 (2022), 1–177. https://doi.org/10.1561/1500000079
[21] 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, 679–707. https://doi.org/10.1007/978-1-0716-2197-4_18
[22] 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 *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A Friedler and Christo Wilson (Eds.). PMLR, 172–186. https://proceedings.mlr.press/v81/ekstrand18b.html
[23] Adam Fourney, Meredith Ringel Morris, Abdullah Ali, and Laura Vonessen. 2018. Assessing the Readability of Web Search Results for Searchers with Dyslexia. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. ACM, 1069–1072. https://doi.org/10.1145/3209978.3210072
[24] Asela Gunawardana, Guy Shani, and Sivan Yogev. 2022. Evaluating Recommender Systems. In *Recommender Systems Handbook* (third ed.), Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, 547–601. https://doi.org/10.1007/978-1-0716-2197-4_15
[25] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, III, Miro Dudik, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do Industry Practitioners Need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19, Paper 600)*. ACM, 1–16. https://doi.org/10.1145/3290605.3300830
[26] Wen Huang, Kevin Labille, Xintao Wu, Dongwon Lee, and Neil Heffernan. 2020. Achieving User-Side Fairness in Contextual Bandits. (Oct. 2020). http://arxiv.org/abs/2010.12102
[27] Theo Huibers, Monica Landoni, Maria Soledad Pera, Jerry Alan Fails, Emiliana Murgia, and Natalia Kucirkova. 2021. What Does Good Look Like? Report on the 3rd International and Interdisciplinary Perspectives on Children & Recommender and Information Retrieval Systems (KidRec) at IDC 2019. *SIGIR Forum* 53, 2 (2021), 76–81. https://doi.org/10.1145/3458553.3458561
[28] Theo Huibers and Thijs Westerveld. 2019. Relevance and utility in an educational search environment. In *KidRec '19: Workshop in International and Interdisciplinary Perspectives on Children & Recommender and Information Retrieval Systems, Co-located with ACM IDC*. https://kidrec.github.io/papers/KidRec_2019_paper_4.pdf
[29] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2018. Recommendation Independence. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A Friedler and Christo Wilson (Eds.). PMLR, 187–201. http://proceedings.mlr.press/v81/kamishima18a.html
[30] Elisa Kreiss, Cynthia Bennett, Shayan Hooshmand, Eric Zelikman, Meredith Ringel Morris, and Christopher Potts. 2022. Context Matters for Image Descriptions for Accessibility: Challenges for Referenceless Evaluation Metrics. (2022). http://arxiv.org/abs/2205.10646
[31] Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2021. User-oriented Fairness in Recommendation. In *Proceedings of the Web Conference (WWW '21)*. ACM, 624–632. https://doi.org/10.1145/3442381.3449866
[32] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. 2018. Learning Adversarially Fair and Transferable Representations. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 3384–3393. https://proceedings.mlr.press/v80/madras18a.html
[33] Ion Madrazo Azpiazu, Nevena Dragovic, Oghenemaro Anuyah, and Maria Soledad Pera. 2018. Looking for the Movie Seven or Sven from the Movie Frozen? A Multi-perspective Strategy for Recommending Queries for Children. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval (CHIIR '18)*. ACM, 92–101. https://doi.org/10.1145/3176349.3176379
[34] Noemi Mauro, Liliana Ardissono, and Federica Cena. 2020. Personalized Recommendation of PoIs to People with Autism. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*. ACM, 163–172. https://doi.org/10.1145/3340631.3394845
[35] Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. 2017. Auditing Search Engines for Differential Satisfaction Across Demographics. In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW '17 Companion)*. International World Wide Web Conferences Steering Committee, 626–633. https://doi.org/10.1145/3041021.3054197
[36] Ashlee Milton, Garrett Allen, and Maria Soledad Pera. 2021. To Infinity and Beyond! Accessibility is the Future for Kids' Search Engines. In *Proceedings of the IR for Children 2000-2020: Where Are We Now? Workshop co-located with SIGIR 2021 (IR4C '21)*. http://arxiv.org/abs/2106.07813
[37] Ashlee Milton, Levesson Batista, Garrett Allen, Siqi Gao, Yiu-Kai D Ng, and Maria Soledad Pera. 2020. "Don't Judge a Book by its Cover": Exploring Book Traits Children Favor. In *14th ACM Conference on Recommender Systems (RecSys '20)*. ACM, 669–674. https://doi.org/10.1145/3383313.3418490
[38] Ashlee Milton, Emiliana Murgia, Monica Landoni, Theo Huibers, and Maria Soledad Pera. 2019. Here, There, and Everywhere: Building a Scaffolding for Children's Learning through Recommendations. In *Proceedings of the 1st*

Workshop on the Impact of Recommender Systems co-located with 13th ACM Conference on Recommender Systems, Vol. 2462. CEUR-WS. http://ceur-ws.org/Vol-2462/short2.pdf

[39] Emiliana Murgia, Monica Landoni, Theo Huibers, Jerry Alan Fails, and Maria Soledad Pera. 2019. The Seven Layers of Complexity of Recommender Systems for Children in Educational Contexts. In *Proceedings of the Workshop on Recommendation in Complex Scenarios co-located with 13th ACM Conference on Recommender Systems*, Vol. 2449. CEUR-WS. http://ceur-ws.org/Vol-2449/paper1.pdf

[40] Mohammadmehdi Naghiaei, Hossein A Rahmani, and Yashar Deldjoo. 2022. CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. ACM, 770–779. https://doi.org/10.1145/3477495.3531959

[41] Nicola Neophytou, Bhaskar Mitra, and Catherine Stinson. 2022. Revisiting Popularity and Demographic Biases in Recommender Evaluation and Effectiveness. In *Proceedings of Advances in Information Retrieval: 44th European Conference on IR Research (ECIR '22)*. Springer-Verlag, 641–654. https://doi.org/10.1007/978-3-030-99736-6_43

[42] Netflix. 2022. Children & Family Movies. https://www.netflix.com/browse/genre/783.

[43] Yiu-Kai Ng and Maria Soledad Pera. 2018. Recommending Social-interactive Games for Adults with Autism Spectrum Disorders (ASD). In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. ACM, 209–213. https://doi.org/10.1145/3240323.3240405

[44] Maria Soledad Pera, Emiliana Murgia, Monica Landoni, and Theo Huibers. 2019. With a Little Help from My Friends: Use of Recommendations at School. In *Proceedings of ACM RecSys 2019 Late-breaking Results*. http://ceur-ws.org/Vol-2431/paper13.pdf

[45] Maria Soledad Pera and Yiu-Kai Ng. 2014. Automating Readers' Advisory to Make Book Recommendations for K-12 Readers. In *Proceedings of the 8th ACM Conference on Recommender Systems (RecSys '14)*. ACM, 9–16. https://doi.org/10.1145/2645710.2645721

[46] Bashir Rastegarpanah, Krishna P Gummadi, and Mark Crovella. 2019. Fighting Fire with Fire: Using Antidote Data to Improve Polarization and Fairness of Recommender Systems. In *Proceedings of the 12th ACM International Conference on Web Search and Data Mining (WSDM '19)*. ACM, 231–239. https://doi.org/10.1145/3289600.3291002

[47] Meagan Rothschild, Takeshi Horiuchi, and Marie Maxey. 2019. Evaluating "Just Right" in EdTech Recommendation. In *KidRec '19: Workshop in International and Interdisciplinary Perspectives on Children & Recommender and Information Retrieval Systems, Co-located with ACM IDC*. https://kidrec.github.io/papers/KidRec_2019_paper_6.pdf

[48] Anu Shrestha, Francesca Spezzano, and Maria Soledad Pera. 2021. An Empirical Analysis of Collaborative Recommender Systems Robustness to Shilling Attacks. In *Proceedings of the Second Workshop on Online Misinformation- and Harm-Aware Recommender Systems co-located with 15th ACM Conference on Recommender Systems (OHARS '21)*. CEUR-WS, 45–57. http://ceur-ws.org/Vol-3012/OHARS2021-paper4.pdf

[49] 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

[50] Lawrence Spear, Ashlee Milton, Garrett Allen, Amifa Raj, Michael Green, Michael D Ekstrand, and Maria Soledad Pera. 2021. Baby Shark to Barracuda: Analyzing Children's Music Listening Behavior. In *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys 2021 Late-Breaking Results)*. ACM Press. https://doi.org/10.1145/1122445.1122456

[51] Spotify. 2022. Spotify Kids. https://www.spotify.com/us/kids/.

[52] Jean Tague-Sutcliffe. 1992. The Pragmatics of Information Retrieval Experimentation, Revisited. *Information Processing & Management* 28, 4 (1992), 467–490. https://doi.org/10.1016/0306-4573(92)90005-K

[53] Konstantinos Tsiakas, Emilia Barakova, Javed Vassilis Khan, and Panos Markopoulos. 2020. BrainHood: Towards an Explainable Recommendation System for Self-regulated Cognitive Training in Children. In *Proceedings of the 13th ACM International Conference on PErvasive Technologies Related to Assistive Environments (PETRA '20, Article 73)*. ACM, 1–6. https://doi.org/10.1145/3389189.3398004

[54] Alexandra Vtyurina, Adam Fourney, Meredith Ringel Morris, Leah Findlater, and Ryen W White. 2019. Bridging Screen Readers and Voice Assistants for Enhanced Eyes-Free Web Search. In *The World Wide Web Conference (WWW '19)*. ACM, 3590–3594. https://doi.org/10.1145/3308558.3314136

[55] Lequn Wang and Thorsten Joachims. 2021. User Fairness, Item Fairness, and Diversity for Rankings in Two-Sided Markets. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval (ICTIR '21)*. ACM, 23–41. https://doi.org/10.1145/3471158.3472260

[56] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint Multisided Exposure Fairness for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. ACM, 703–714. https://doi.org/10.1145/3477495.3532007

[57] Le Wu, Lei Chen, Pengyang Shao, Richang Hong, Xiting Wang, and Meng Wang. 2021. Learning fair representations for recommendation: A graph-based perspective. In *Proceedings of the Web Conference 2021 (WebConf '21)*. ACM, New York, NY, USA. https://doi.org/10.1145/3442381.3450015

[58] Sirui Yao and Bert Huang. 2017. Beyond Parity: Fairness Objectives for Collaborative Filtering. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS '20)*, I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett (Eds.). Curran Associates, Inc., 2925–2934. http://papers.nips.cc/paper/6885-beyond-parity-fairness-objectives-for-collaborative-filtering.pdf

[59] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 28)*, Sanjoy Dasgupta and David McAllester (Eds.). PMLR, 325–333. https://proceedings.mlr.press/v28/zemel13.html