Unboxing the Algorithm

Designing an Understandable Algorithmic Experience in Music Recommender Systems

Anna Marie Schröder [1] and Maliheh Ghajargar [1, 2]

[1] *Malmö University, Nordenskiöldsgatan 1, Malmö, Sweden*
[2] *Malmö University, Internet of Things and People (IOTAP), Nordenskiöldsgatan 1, Malmö, Sweden*

Abstract[1]

After decades of the existence of algorithms in everyday use technologies, users have developed an algorithmic awareness, but they still lack the confidence to grasp them. This study explores how understandability as a principle drawn from sociology, design, and computing can enhance the algorithmic experience in music recommendation systems. The preliminary results of this Research-Through-Design showed that users had limited mental models so far but had a curiosity to learn. Further, it confirmed that explanations as a dialogue could improve the algorithmic experience in music recommendation systems. Users could comprehend recommendations the best when they were easy to access and understand, directly related to user behavior, and when they allowed the user to correct the algorithm. To conclude, our study reconfirms that designing experiences that help users to understand the algorithmic workings will make authentic recommendations from intelligent systems more applicable in the long run.

Keywords

Human-centered computing, Interaction design, Empirical studies in interaction design, algorithmic experience, music recommendation systems, transparency, machine learning, explainable AI

1. Introduction

Nowadays, music streaming supplies users with endless music choices through on-demand services. Algorithms often assist digital music explorations and choices. More precisely, they recommend content. Based on user data, recommender systems continuously learn during the music listening experience. System users might notice the existence of machine learning algorithms (ML) in plenty of touchpoints in digital services. However, they are still black-boxed for users, which can lead to confusion when experiencing the algorithm [2].

In 2003 Rogers described in Diffusion of Innovations that the diffusion and improvement of innovative technologies like music recommendation systems (MRSs) could only progress if users accept and understand the innovation. Coming from another field, product designer Dieter Rams captured ten principles for good design, which include making “a product understandable” [8, 35]. Music streaming services face central design challenges related to understandability and try hard to find “the sweet spot of calibrated [user] trust” and explore users’ “mental models” to manage expectations [20]. Listeners tend to distrust algorithms in music recommendation platforms when they do not understand how they work [7]. Even more challenging is that users notice unpleasant recommendations more than enjoyable ones [2, 3]. Algorithmic hick-ups in ML services could dissatisfy users, who might pick another

[1] Find a video presentation at https://vimeo.com/598749992

Perspectives on the Evaluation of Recommender Systems Workshop (PERSPECTIVES 2021), September 25th, 2021, co-located with the 15th ACM Conference on Recommender Systems, Amsterdam, The Netherlands
EMAIL: schroederanna@outlook.com (A. 1); maliheh.ghajargar@mau.se (A. 2)
ORCID: 0000-0001-5008-7392 (A. 1); 0000-0003-1852-3937 (A. 2)

music streaming service from the highly competitive market to find a more satisfying algorithmic experience (AX) [26]. ML and intelligence embodied in AI systems have been considered as a design material and as such it requires designers to understand the matter without needing to have technical ML knowledge [9, 18, 30, 36].

The ML field should start to prioritize AX over prediction accuracy [11]. This paper explores the dialogue between music recommendation systems and their users to reveal (1) user understandability of the recommender system, and (2) frustrating user experiences, to (3) propose solutions and offer discussions points around user understandability and AX of MRSs. We pursued the research questions: How can understandability as a design principle improve the algorithmic experience in music recommendation systems? What mental models exist, and how can expectations be met to improve the UX?

The ongoing research confirms the value of a good user experience and the self-perceived confidence of users of intelligent recommendation systems[3] [18]. Specifically, this paper contributes knowledge to Explainable AI and Algorithmic Experience design applied to music recommender systems.

2. Background

User Experience design (UX) is a broad human-centered design field, aiming to improve an existing experience or to design for a specific new experience in interaction design processes [12]. As our interactive systems have started to become more intelligent and to embody ML algorithms, questions regarding how that affects the user experience or how we can design a better UX considering the opportunities that ML can offer, have been raised [15]. Alvarado & Waern's studies on human-ML algorithms interactions resulted in the conceptual Algorithmic Experience (AX) framework. They explored how social media users experience algorithmic automation and how specific knowledge affects a user's interaction with the algorithm [2]. When it comes to algorithm awareness and AX, digital users divide into user types varying from "the unaware", "the uncertain", "the affirmative" over "the neutral", "the skeptic", to "the critical" [17].

2.1. Explainable AI

As part of human-centered AI, Explainable AI (XAI) explores how users gain more trust and a better experience with intelligent systems through receiving explanations of its "inner-workings" [14, 27]. "Explanations are post-hoc descriptions of how a system came to a given conclusion or behavior" [28]. They should "cover both the 'know how' and 'know what' of a system", which connects to the knowledge types by Rogers as elaborated in section 2.3 [27, 29]. Still, ML systems cannot offer true decision transparency due to the autonomous learning process complexity [27, 28]. It might be "post-hoc interpretations" instead, that "explain an output without reference to the inner workings of the system" to answer how a system came to a decision [27]. With the rising system complexity comes a discrepancy between "a stronger need for explainable AI" but "a greater gap in achieving it" [19]. Explanations should be meaningful according to the user's "complexity capacity" and the use situation [19].

Explainable recommendations and "explanatory debugging" describe when a "system explains the reasons for its predictions to its end-user, who explains corrections back to the system" [24, 34]. Tintarev & Masthoff also described how to design explanations in

[3] This study uses the terms algorithm, machine learning, music recommender, music recommendation system, intelligent system, and AI-infused systems synonymously to cover the broad application instead of distinctly differentiating between them, if not stated differently.

recommender systems striving for transparency, scrutability, trust, effectiveness, persuasiveness, efficiency, and user satisfaction. They also point out how explanations serve as means for evaluation of the whole recommender system and play into users' acceptance of the system [33]. Among other explanation types, iterative explanation dialogues might build and correct users' mental models [24]. Current research links explainable recommendations to address "the information asymmetry in data collection and use" within music recommenders [6].

2.2. Understandability

It is always assumed that a well-designed AI is predictable, and its user interface is understandable to contribute to better usability and UX [4]. Creating explanations of technicalities that make sense to the users, computer scientists, designers, and copywriters work hand in hand through a multidisciplinary lens [19].

The term is used with different meanings — sociologist Rogers embedded understandability within the Diffusion-of-Innovation theory, following which users claim awareness-knowledge, how-to-knowledge, and principles-knowledge about an innovation. These include knowledge of the existence of an invention or technology, further about its application, and lastly, understanding how it is functioning [29]. Product designer Rams presented understandability as one of ten principles for good design as "clarify[ing] the product" and "making the product talk" [16]. Jakob Nielsen and Ben Shneiderman created heuristics and UX rules that align with user expectations, e.g., to make the system speak users' language [1]. Shneiderman recommends "informative feedback", creating a dialogue that permits "easy reversal of actions" [31]. Understandability is further defined as "the extent to which a person with relevant technical background can accurately and confidently" explain a system [5]. "Unmanaged complexity" is declared as "the primary enemy of understandability", which connects to Rogers' adoption factor of complexity [5, 29].

2.3. Music Recommendation Systems (MRSs)

MRSs present items from music databases to users based on specific item attributes and guide users through large music libraries [32]. When designing MRS interfaces, transparent communication of technical mechanics needs to be weighed against the user's mental capacity to comprehend such complex systems [20]. Many MRSs apply three main ML models to customize music recommendations: Firstly, collaborative filtering results in "predictions about the user's preferences based on similar user preferences" [21]. Further, the algorithm scans publicly available information using Natural Language Processing (NLP) to analyze blog posts and online discussions. Lastly, MRSs' algorithms compare tracks in their music library by BPM, style, and rhythm through audio modeling to identify similar songs [21].

Currently, several MRSs are found to violate against examined guidelines for interactive ML [4]. From a study by Amershi et al., five principles for designing human-AI interactions include the notions of explainability and understandability, which became relevant for this study and supported the design and research processes.

3. Methodology

The point of departure in our research was the connection we found between Rogers' diffusion-of-innovation theory, Moore's technology marketing approach *Crossing the Chasm*

[25, 29], and the User Experience of AI (e.g. Algorithmic Experience) fields. We observed that what links all these theories together is positioning the user understandability as a crucial factor in innovation and in designing any interactive systems. Especially music recommenders moved into the center of our research attention, as their algorithmic recommendations are often appreciated and praised as surprisingly effective from a user perspective, while they were found to violate some of the guidelines for human-AI interaction [4].

This contrast grounded how this study came together. After first investigating the literature around the topic of understandability and explainability, subsequently, a short user study was conducted following a research-through-design process [37]. After designing the digital prototype of the user interface, we conducted user testing sessions with six end-users.

3.1. Focus User Sessions

We recruited participants based on their availability through digital tools and their level of engagement with ML music recommendation features of music streaming services, for semi-structured interviews. Therefore, mostly young adults affine to digital music joined the participant pool. In the selection of people, we also considered reaching as much diversity in gender, nationality, and platform use as possible.

In individual 60-minute sessions, the researchers invited five participants (two male and three female 23-38 YO) to reveal personal moments of discontent with MRSs: "Often the mood is not quite matched. When it doesn't match at all, I notice that very consciously. Then I wonder, 'Huh, where does that come from?'"; another participant voiced, "sometimes it doesn't work out at all. I can't really explain why." Most users were sure that algorithms work based on their listening behavior and liked songs. Three of five users communicated to skip recommended content they do not enjoy: "I skip until I find something that I like", expecting the algorithm to consider that. Four users said they rarely use explicit feedback features: "I never liked a song."

Along with their shared personal listening experiences, the researchers introduced participants to ML principles applied in MRSs as described in 2.4. None of the participants were surprised about the use of algorithms in these ways. One participant added: "Our generation is used to algorithms, so I can guess how they work", another person said: "I never read any scientific explanation. It's just a gut feeling". As stated before, all participants had a rough idea how and where algorithms come into play for their personal music discovery, still one interesting thought was: "I think that's just information that you have to come into contact with to further think about. Very few people even know [about music algorithms], and even fewer people question how [they] actually work". In general, people voiced a tendency to act passively and accept recommendations as a surprise. Statements here were: "I like the unexpected", "I like to leave the control to the machine and let myself be surprised" and "I feel 5/10 in control, because I want to be surprised". All users also tended to trust the algorithm with picking the best recommendations; one even said: "I am never annoyed, because I know why I get certain recommendations, it's my 'fault'". To the question how they would react if they would continuously receive unliked recommendations, several answers went into the direction of "when [bad recommendations] happen often, I would stop using the service".

3.2. Prototyping And User Testing

The digital prototype consisted of a purely visual interactive interface that mimicked a personal recommendation feature. Four focus users from the exploration phase and two additional users (five female and one male 23-28 YO) explored interactive wireframe prototypes compared different interactions to retrieve explanations, as well as different versions of explanations of algorithmic principles. While some sessions could take part in person, most of them happened virtually (figure 1).

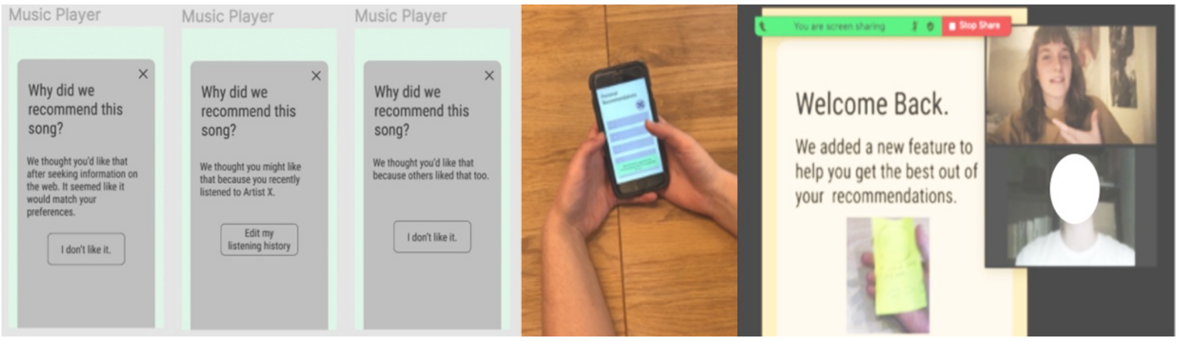

Figure 1: Prototype and User Testing Sessions (in-person and virtual)

With quick iterations between sessions, we evaluated how much users appreciated each option afterward.

4. Analysis and Results

The conversations with users confirmed the value of explanatory recommendations [34]. Users' mental capacity allowed them to understand underlying ML principles only under certain conditions. A general need to promote awareness-knowledge came out of statements like: “I don’t know why I get recommended what I get recommended.” People tended to act passively and accept recommendations as a surprise. Statements here were: “I like to leave the control to the machine and let myself be surprised.”

Any explanation should directly refer to the user’s actions and offer to take control in each case. Therefore, all explanations offer users to edit their listening history to alter their data points in hindsight. In the end, retrieving understandable explanations through press-and-hold turned out as the most feasible interaction. Participants appreciated claiming explanations by interacting with the song directly instead of dialing through a menu.

All testers reacted positively to the understandable dialogue with the prototype: “Cool to see that the system pays attention to my behavior.” Users found it more realistic to retrospectively correct their listening than to pay attention not to “ruin the algorithm” beforehand, especially in social listening situations. Users appreciated all explanations but found collaborative filtering and audio modeling easier to grasp than NLP: “I feel better when explanations refer to me. Then it seems less random and more personalized.” Additional thoughts spanned from “I really like that I get explanations because I wouldn’t search for them on my own”, over “It would be only fair to know what data is used for my recommendations.”

We found that users are missing an understanding of MRSs and have limited mental models. Most of them bring an intuitive algorithm awareness due to growing up as digital natives. Still, it is mostly “just a gut feeling.” “When it doesn’t match at all, I notice that very consciously.” “As soon as they clash with my preferences, I don’t bother to take a look at the rest.” These statements about bad recommendations demonstrated the need to improve AX through understandable explanations at this point of social adoption of algorithmic music recommendations. Listeners tended to quickly leave one streaming service after interaction breakdowns with the algorithm: “I would stop using the service.” Participants were further overwhelmed by the complexity of algorithmic explanations. Therefore, systems should introduce easily digestible information bites about their inner workings.

The mentioned algorithm intuition aligns with Rogers' awareness- knowledge. Pure principles-knowledge overloaded the capacity of most participants' mental MRS models. One participant said: "Our generation is used to algorithms, so I can guess how they work", another person said: "I never read any scientific explanation. It's just a gut feeling." Informative explanations should only address principles-knowledge if the user receives how-to-knowledge as well. Explanations directly referencing their behavior were easiest to understand.

"Information should come with control, otherwise it's bluff." - Agreeing with Tintarev & Masthoff's aim of scrutability and Shneiderman's principles, the study showed that allowing users to take control to adjust single data points (e.g., songs in their history) made them more open to receiving and explanative information after all [31, 33]. This affordance aligns with Dudley & Kristensson's solutions that IML should promote rich interactions and engage the user [10].

Post-hoc explanations improved the AX in MRSs when they were easy to access and digest, directly related to user behavior, and gave control to correct the algorithm, as sketched out in the finalized interactive prototype.

4.1. #SKIPSMATTER

Participants expected their music algorithms to notice whenever they skip a song and take this behavior into account for following recommendations. Hence, the MRS would inform as soon as it counted three skips for one song, which gives the listeners a level of control to correct the system and bring the track back at another time (see figure 2).

Figure 2: Sketch, example wireframes, and final design for #SkipsMatter

4.2. #UNDERSTANDWHY

On a press-and-hold interaction with a music item in a list or directly from the player-view, an MRS shows the user an explanation of why it recommended that song. The user can then directly remove the track from their list or even edit their listening history to correct single data points from the past (see figure 3).

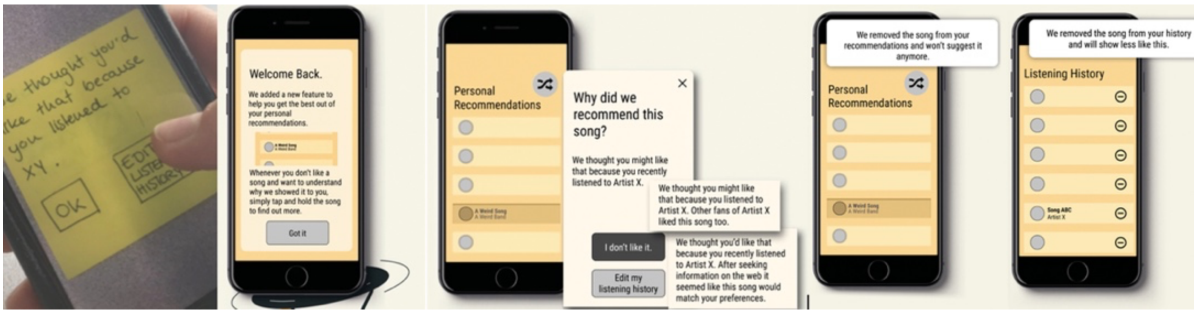

Figure 3: Design outcome for interface elements and interaction flows for #UnderstandWhy

Figure 4 shows explanations of the three main ML models as described in 2.4 that the prototype system provides after certain user interactions.

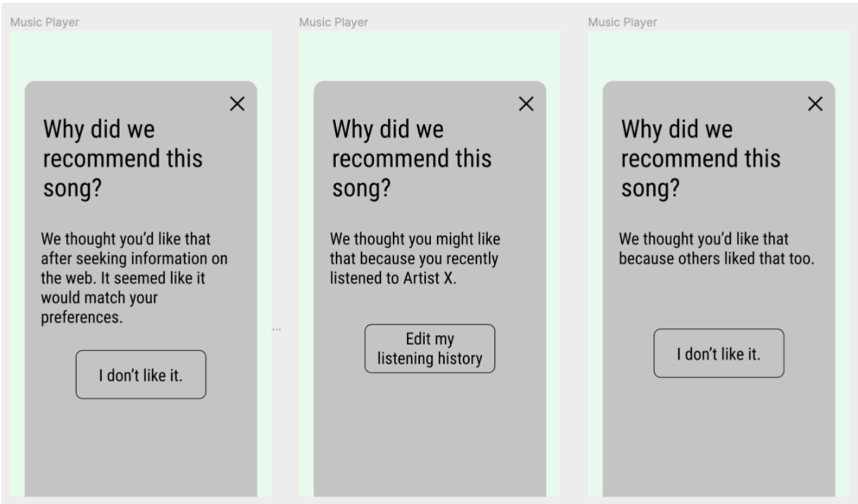

Figure 4: Wireframes for Explanation Presentations (NLP, Audio Modeling, Collaborative Filtering

5. Discussion

5.1. Answering Conditional Curiosity

The study explored how to simplify systems for non-experts to understand complex ML. Users were interested to receive information about their music algorithms that directly referred to their behavior and meant a benefit or more control, they were curious to learn about certain explanations. This mental capacity or conditional curiosity should be accepted by leaving it to the users to decide if and how much they want to learn.

Users valued the element of surprise during algorithmic music discovery. An attempt to unbox the algorithm, therefore, decreases user satisfaction as it removes any lasting 'magical' aspect of how MRSs seem as intelligent as they do. However, bad recommendations might be part of a good AX in MRSs, while recommendations that feel *too good to be true* might lead to algorithm aversion. User testing showed how quickly data safety concerns can raise and presenting explanations could reduce that mistrust.

5.2. Iterative Process to Evaluate the Recommender System

The iterative and interactive nature of design methodologies such as testing the user experience with sketchy wireframe prototypes certainly helps to bring new perspectives in evaluating existing RSs. But thanks to its iterative and sketchy fashion, it creates an open environment for the users to express themselves and be empowered, hence bringing more user-centered insights that can be of help to describe and evaluate the system along with theoretical frameworks. (e.g. Knijnenburg et al. [23]. We especially observed that a design-driven approach that goes beyond usability studies is beneficial to designing algorithmic experiences and is more powerful to capture the nuances of the experience.

5.3. Algorithmic Experience to Bridge the Chasm

Users showed an intuitive algorithmic understanding. This might seem advantageous at first, but algorithm knowledge and corresponding expectations should be corrected and extended instead. Algorithm awareness and algorithm control can support an honest AX. This "should include [...] building [...] algorithmic literacy and critical capacity of the users" [22]. With Moore's chasm between early adopters and the early majority in mind, truly accessible and understandable explanations could avoid a frustrated early majority – for which an honest human-system dialogue might bring confidence and awareness.

5.4. The Empowered User of the Future

Interaction and interface designers for MRSs should aim to increase algorithmic awareness in the system outcome and empower the user through system transparency to control MRSs. Transparency considerations push businesses' boundaries that rely on data-driven business deals running in the background. How much longer can recommendations be authentic that still button-up behind algorithmic black-box walls?

AX in music exploration might not change the world, but findings from this ongoing research apply to other contexts with automated decision-making (e.g. medical treatments, political courses). This can be a future progression for users to handle algorithmic decision-making. What algorithmic social intercourse are we training ourselves into, and is that ethically sustainable? Part of the purpose of designing more graspable and understandable recommendation systems should undeniably be about how algorithms influence decisions. Designing tangible interaction can be helpful to make the system graspable, and expose its shortcomings and open the system for critique, and that does not need to be attained by compromising the main functionalities of the recommender system [13, 14]. Critically differentiating between human-made versus machine-made recommendations should persist, even if it only comes to what song to listen to next.

6. Conclusions And Future Work

This ongoing study explored how to improve the AX in MRSs by encouraging system understanding and graspability. We were able to confirm current theories applied to intelligent music recommendation and concretized them into some design requirements. Two design openings evolved, which we iterated in user testing: #UnderstandWhy and #SkipsMatter. From the prototypes, users positively adopted knowledge about the algorithmic inner workings when the information (1) was easy to access and digest, (2) directly related to user behavior, and (3) provide the opportunity to correct the algorithm.

Our study suggests that designers and developers of AI systems need to make users aware of algorithms and make corresponding knowledge accessible to empower them.

Lastly, designing more understandable music streaming apps might be one step towards building an ethical and social discourse around algorithms that soon escape their magic black-box.

This work was carried out by prototyping and testing a traditional mobile UI, within a controlled environment. Therefore, we believe the participants' feedback was highly influenced by the nature of the experiment, while not considering other possible ways that users might interact with a music recommender system. There is certainly future work to be done regarding a comparison of interaction modalities (e.g., existing interface interactions, Tangible Interaction, AR/VR), and the ecology of listening to music artifacts (e.g., headphones, speakers, etc.) depending on the context of use. This study needs to be expanded to understand how users would experience the explainability of the system in situations where direct interaction with the UI is not possible or desirable, for instance when users listen to music while running or cooking.

Another interesting area to explore is to consider the relationships between the music recommender system and other artifacts that are connected to help to design a coherent experience of listening to music, while also promoting the explainability of the system.

7. References

[1] 10 usability heuristics for user interface design: 2020. *https://www.nngroup.com/articles/ten-usability-heuristics/*. Accessed: 2021-04-16.

[2] Alvarado, O. and Waern, A. 2018. Towards algorithmic experience: initial efforts for social media contexts. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC Canada, Apr. 2018), 1–12.

[3] Amershi, S., Cakmak, M., Knox, W.B. and Kulesza, T. 2014. Power to the people: the role of humans in interactive machine learning. *AI Magazine*. 35, 4 (Dec. 2014), 105–120. DOI:https://doi.org/10.1609/aimag.v35i4.2513.

[4] Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R. and Horvitz, E. 2019. Guidelines for human-ai interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow Scotland Uk, May 2019), 1–13.

[5] Boeuf, J., Kern, C. and Reese, J. 2020. Design for Understandability. *Building secure and reliable systems best practices for designing, implementing, and maintaining systems*. O'Reilly Media.

[6] Born, G., Morris, J., Diaz, F. and Anderson, A. 2021. *Artificial Intelligence, Music Recommendation, and the Curation of Culture*. University of Toronto; Schwartz Reisman Institute for Technology and Society; CIFAR.

[7] Deezer | a case ux case study: 2020. *https://medium.com/@gasparslayfordkatie/deezer-a-case-ux-case-study-4721b408c23c*. Accessed: 2021-04-13.

[8] Dieter rams: 10 timeless commandments for good design: 2020. *https://www.interaction-design.org/literature/article/dieter-rams-10-timeless-commandments-for-good-design*. Accessed: 2021-04-12.

[9] Dove, G., Halskov, K., Forlizzi, J. and Zimmerman, J. 2017. Ux design innovation: challenges for working with machine learning as a design material. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver Colorado USA, May 2017), 278–288.

[10] Dudley, J.J. and Kristensson, P.O. 2018. A review of user interface design for interactive machine learning. *ACM Transactions on Interactive Intelligent Systems*. 8, 2 (Jul. 2018), 1–37. DOI:https://doi.org/10.1145/3185517.

[11] Five ways to make AI a greater force for good in 2021: 2021. *https://www.technologyreview.com/2021/01/08/1015907/ai-force-for-good-in-2021/*. Accessed: 2021-04-01.

[12] Forlizzi, J. and Battarbee, K. 2004. Understanding experience in interactive systems. *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques* (New York, NY, USA, Aug. 2004), 261–268.

[13] Ghajargar, M. and Bardzell, J. 2021. Synthesis of Forms: Integrating Practical and Reflective Qualities in Design. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery. 1–12.

[14] Ghajargar, M., Bardzell, J., Renner, A.S., Krogh, P.G., Höök, K., Cuartielles, D., Boer, L. and Wiberg, M. 2021. From ”explainable ai” to ”graspable ai”. *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Salzburg Austria, Feb. 2021), 1–4.

[15] Ghajargar, M., Persson, J., Bardzell, J., Holmberg, L. and Tegen, A. 2020. The UX of Interactive Machine Learning. *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (New York, NY, USA, Oct. 2020), 1–3.

[16] Good design: *https://www.vitsoe.com/us/about/good-design*. Accessed: 2021-04-16.

[17] Gran, A.-B., Booth, P. and Bucher, T. 2020. To be or not to be algorithm aware: a question of a new digital divide? *Information, Communication & Society*. (Mar. 2020), 1–18. DOI:https://doi.org/10.1080/1369118X.2020.1736124.

[18] Hamilton, K., Karahalios, K., Sandvig, C. and Eslami, M. 2014. A path to understanding the effects of algorithm awareness. *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (Toronto, Ontario, Canada, Apr. 2014), 631–642.

[19] Hind, M. 2019. Explaining explainable AI. *XRDS: Crossroads, The ACM Magazine for Students*. 25, 3 (Apr. 2019), 16–19. DOI:https://doi.org/10.1145/3313096.

[20] How can design thinking build trust in the age of machine learning? 2019. *https://spotify.design/article/how-can-design-thinking-build-trust-in-the-age-of-machine-learning*. Accessed: 2021-04-13.

[21] How spotify uses machine learning models? | analytics steps: *https://www.analyticssteps.com/blogs/how-spotify-uses-machine-learning-models*. Accessed: 2021-04-12.

[22] Klumbyte, G., Lücking, P. and Draude, C. 2020. Reframing ax with critical design: the potentials and limits of algorithmic experience as a critical design concept. *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn Estonia, Oct. 2020), 1–12.

[23] Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H. and Newell, C. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*. 22, 4–5 (Oct. 2012), 441–504. DOI:https://doi.org/10.1007/s11257-011-9118-4.

[24] Kulesza, T., Burnett, M., Wong, W.-K. and Stumpf, S. 2015. Principles of explanatory debugging to personalize interactive machine learning. *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA, Mar. 2015), 126–137.

[25] Moore, G. 2014. *Crossing the Chasm: Marketing and Selling Disruptive Products to Mainstream Customers*. Harper Business.

[26] Music Subscriber Market Shares Q1 2020: 2020. *https://www.midiaresearch.com/blog/music-subscriber-market-shares-q1-2020*. Accessed: 2021-04-20.

[27] Preece, A. 2018. Asking ‘Why’ in AI: Explainability of intelligent systems - perspectives and challenges. *Intelligent Systems in Accounting, Finance and Management*. 25, 2 (Apr. 2018), 63–72. DOI:https://doi.org/10.1002/isaf.1422.
[28] Riedl, M.O. 2019. Human-centered artificial intelligence and machine learning. *Human Behavior and Emerging Technologies*. 1, 1 (Jan. 2019), 33–36. DOI:https://doi.org/10.1002/hbe2.117.
[29] Rogers, E.M. 2003. *Diffusion of innovations*. Free Press.
[30] Rozendaal, M.C., Ghajargar, M., Pasman, G. and Wiberg, M. 2018. Giving Form to Smart Objects: Exploring Intelligence as an Interaction Design Material. *New Directions in Third Wave Human-Computer Interaction: Volume 1 - Technologies*. M. Filimowicz and V. Tzankova, eds. Springer International Publishing. 25–42.
[31] Shneiderman’s eight golden rules will help you design better interfaces: 2020. *https://www.interaction-design.org/literature/article/shneiderman-s-eight-golden-rules-will-help-you-design-better-interfaces*. Accessed: 2021-04-16.
[32] Stray, J., Adler, S. and Hadfield-Menell, D. 2020. What are you optimizing for? Aligning Recommender Systems with Human Values. *Participatory Approaches to Machine Learning* (Jul. 2020).
[33] Tintarev, N. and Masthoff, J. 2011. Designing and evaluating explanations for recommender systems. *Recommender Systems Handbook*. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer US. 479–510.
[34] Tsukuda, K. and Goto, M. 2020. Explainable recommendation for repeat consumption. *Fourteenth ACM Conference on Recommender Systems* (Virtual Event Brazil, Sep. 2020), 462–467.
[35] What is “good” design? A quick look at dieter rams’ ten principles.: *https://designmuseum.org/discover-design/all-stories/what-is-good-design-a-quick-look-at-dieter-rams-ten-principles*. Accessed: 2021-04-12.
[36] Wirfs-Brock, J., Mennicken, S. and Thom, J. 2020. Giving voice to silent data: designing with personal music listening history. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu HI USA, Apr. 2020), 1–11.
[37] Zimmerman, J., Forlizzi, J. and Evenson, S. 2007. Research through design as a method for interaction design research in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose California USA, Apr. 2007), 493–502.