

Recommending With, Not For

Co-Designing Recommender Systems for Social Good

MICHAEL D. EKSTRAND, AFSANEH RAZI, and ALEKSANDRA SARCEVIC, Drexel University, USA

MARIA SOLEDAD PERA, Delft University of Technology, NL

ROBIN BURKE, University of Colorado, USA

KATHERINE LANDAU WRIGHT, University of Oregon, USA

Recommender systems are usually designed by engineers, researchers, designers, and other members of development teams. These systems are then evaluated based on goals set by the aforementioned teams and other business units of the platforms operating the recommender systems. This design approach emphasizes the designers' vision for how the system can best serve the interests of users, providers, businesses, and other stakeholders. Although designers may be well-informed about user needs through user experience and market research, they are still the arbiters of the system's design and evaluation, with other stakeholders' interests less emphasized in user-centered design and evaluation. When extended to recommender systems for social good, this approach results in systems that reflect the social objectives as envisioned by the designers and evaluated as the designers understand them. Instead, social goals and operationalizations should be developed through participatory and democratic processes that are accountable to their stakeholders. We argue that recommender systems aimed at improving social good should be designed *by* and *with*, not just *for*, the people who will experience their benefits and harms. That is, they should be designed in collaboration with their users, creators, and other stakeholders as full co-designers, not only as user study participants.

CCS Concepts: • **Human-centered computing** → **Participatory design; HCI design and evaluation methods**; • **Information systems** → **Recommender systems**; • **General and reference** → *Evaluation*.

Additional Key Words and Phrases: co-design, co-evaluation, participatory design, recommender systems

ACM Reference Format:

Michael D. Ekstrand, Afsaneh Razi, Aleksandra Sarcevic, Maria Soledad Pera, Robin Burke, and Katherine Landau Wright. 2025. Recommending With, Not For: Co-Designing Recommender Systems for Social Good. *ACM Trans. Recomm. Syst.*, Just Accepted (August 2025), 24 pages. <https://doi.org/10.1145/3759261>

1 Introduction

Recommender systems have a substantial impact not only on the people who directly use them, but also on other stakeholders¹ and on the broader information and economic ecosystems in which they are embedded. As people use recommender systems to locate information, entertainment, goods, and services that meet their often personalized and contextual needs, the system has a

¹We use the term “stakeholder” for familiarity and consistency with the multistakeholder recommendation literature [Abdollahpouri et al. 2017; Burke 2017]. However, we note that this term is contested [Sharfstein 2016] and obscures the varying power relationships that participatory design in its full form seeks to directly confront.

Authors' Contact Information: Michael D. Ekstrand, ekstrand@acm.org; Afsaneh Razi, ar3882@drexel.edu; Aleksandra Sarcevic, as3653@drexel.edu, Drexel University, Philadelphia, PA, USA; Maria Soledad Pera, M.S.Pera@TUDelft.nl, Delft University of Technology, Delft, NL; Robin Burke, robin.burke@colorado.edu, University of Colorado, Boulder, CO, USA; Katherine Landau Wright, KLWright@uoregon.edu, University of Oregon, Eugene, OR, USA.



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

© 2025 Copyright held by the owner/author(s).

ACM 2770-6699/2025/8-ART1

<https://doi.org/10.1145/3759261>

pivotal role in providing or withholding exposure to providers and their content, influencing user perceptions of the information space and aiding or impeding access to different types of information. In practice, this means that recommender systems have the potential to promote diverse and equitable marketplaces [Mehrotra et al. 2018], influence users towards personally- and socially-beneficial activity [Knijnenburg et al. 2016; Starke et al. 2017], and advance other pro-social outcomes. However, recommender systems can also entrench established players, facilitate the spread of disinformation and propaganda, promote social division or unsustainable consumption practices, and generally catalyze societal harm [Belkin and Robertson 1976; Milano et al. 2020].

Recommender systems research and development is carried out by the engineers, researchers, designers, and other development team members, typically employed by the owner of the platforms where the systems will be deployed (collectively, the recommender *designers*). These designers are usually the primary agents in designing and evaluating recommendation experiences, and users are involved only as research subjects or occasionally through early-stage design activities such as focus groups. It is rare for users to be involved in the design of the system or in the design and interpretation of the evaluations used to validate it. Because evaluation and design are often based on observation of behavior, users have little opportunity to influence design or evaluation outside the behavioral measurements defined by the designers. Ekstrand and Willemssen [2016] critiqued this trend, calling for increased involvement of explicit user preferences, while Seaver [2018] examines recommendation and behavior through the lens of “trapping.” Current practice, however, overwhelmingly emphasizes the designers’ notions of what it means for the recommender system and its effects to be “good,” both for individual users and for society. As the recommender systems community looks beyond providing user and business benefit to projects that advance social good, the importance of incorporating diverse perspectives into research and development only increases.

Recent work has sought to qualitatively understand the various human factors involved in recommendation [Pera et al. 2024] and what users [Harambam et al. 2019; Smith et al. 2020], artists [Ferraro et al. 2021], and other stakeholders want from recommendation [Bauer et al. 2024, §4.4]. Stray et al. [2024] discussed the need for eliciting values from a wide range of perspectives and disciplines and incorporating them into the recommender design process, but the field so far has few examples of how to implement this in practice.

In this paper, we examine **who is, or can be, meaningfully empowered to determine the goals, methods, and design of recommender systems**, both with respect to their ability to meet individual users’ needs and their impact on other people, organizations, and society. We argue that recommender systems should be designed in collaboration with users, providers, and other stakeholders as full co-designers instead of merely user study participants. Our claim is that to effectively support both individual and social good through recommendation, designers must fully involve the many different stakeholders of the recommender ecosystem — including the users it will serve, the providers whose content visibility it will affect, and others impacted by the system — in the process of designing and evaluating recommender systems. We build on and extend prior arguments that explicit user goals and preferences should be elicited and considered [Ekstrand and Willemssen 2016; Lyngs et al. 2018] to propose engaging users, providers, and other stakeholders as **designers and evaluators** of recommender systems [Charisi et al. 2022]. Recommender system design can and should be a collaborative, participatory effort where designers *co-design* the system with the people it will affect, an approach we call **participatory recommendation**.

We develop this argument by first reviewing common practices, user-centric developments, and social concerns in recommender systems (§2) and providing a brief primer on the principle of co-design from human-computer interaction (§3). We then lay out our vision for co-design of

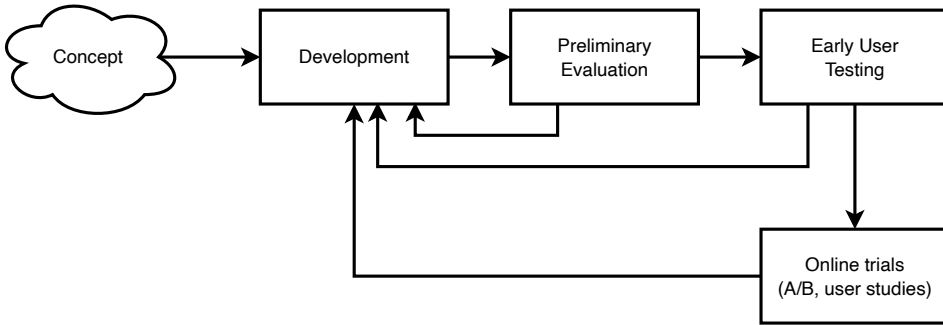


Fig. 1. Diagram of typical recommender system design and evaluation process.

recommender systems (§4), and describe the impact of this vision through several case studies across diverse domains and applications (§5).

2 Current Recommender System Design and Evaluation Practice

Current practice in recommender system design and evaluation typically starts with the designer as the source of ideas and requirements for the recommendation experience. We briefly review this practice and the associated literature as a foundation for our argument about re-negotiating the power relationships in recommender design to enable participatory recommendation.

2.1 Design Lifecycle and Participants

Recommender systems evaluation researchers have presented several perspectives on recommender system development and evaluation with slightly different terms and taxonomies [Gunawardana et al. 2022; Knijnenburg et al. 2012; Zangerle and Bauer 2022]. However, the fundamental approach remains relatively consistent, with the typical lifecycle for designing and evaluating a recommender system proceeding as follows (Fig. 1):

- (1) The designers identify an opportunity for new or improved recommendations. This opportunity may emerge from their own insight and experience, from user studies identifying gaps in the user experience, or from other business concerns such as marketing, contractual obligations, and regulatory compliance.
- (2) The designers use existing data to perform a preliminary offline evaluation of the idea, often iterating on several design options and configurations.
- (3) The designers may involve small groups of users in early-stage product testing, particularly if they have strong human-computer interaction capabilities.
- (4) The designers test the new or revised system online with users, using techniques such as behavioral A/B testing with user activity and business metrics (e.g., clicks, engagement time, or purchases) and perceptual testing with user studies incorporating surveys and explicit user feedback.

This process is iterative, with the results of one design or testing exercise informing the next step or iteration. The specifics of this process can differ among platforms, research teams or projects, with the level of human involvement increasing or decreasing at various stages. For example, Harambam et al. [2019] involved users in a focus group study to understand user goals and desires for recommendation explanation capabilities. The key point related to our argument in this paper is that the dominant approach involves users solely or primarily as *evaluators* (sometimes explicitly,

through user studies [Knijnenburg et al. 2012], but often implicitly through A/B trials [Kohavi et al. 2020]) of designs originating from the designers. The types of feedback users can provide are typically also circumscribed by the evaluation mechanisms provided by the designers. Ekstrand and Willemsen [2016] critiqued this predominant approach, especially in its implicit, behavioral form, arguing that greater explicit user involvement across the design and evaluation lifecycle will enable more useful and pro-social recommendation technology.

A significant limitation of engaging with users only as recipients and evaluators of the designers' ideas is that it restricts the set of *values* [Stray et al. 2024] that are to be incorporated into the recommendation experience: the design team or platform's values take priority, supplemented with their view of other stakeholder values. When considering recommender systems for social good, some goals may be easily discoverable, such as energy savings [Starke et al. 2017] and sustainability [Patro et al. 2020; Tomkins et al. 2018]. Still, design and evaluation need mechanisms to incorporate a much broader set of perspectives on what social good actually is and how recommendations can contribute to that good to develop and deploy systems that are beneficial for humanity and the natural, communal, and socio-technical ecosystems it inhabits.

2.2 Multistakeholder Recommendation

Users are not the only people with a stake in recommender systems and their behavior and impact. *Multistakeholder recommendation* [Abdollahpouri et al. 2017; Burke 2017; Sonboli et al. 2022] takes into account the needs and desires of a broader set of stakeholders, who directly or indirectly stand to benefit from (or be harmed by) the recommendation process. Common stakeholders include *providers* (the artists, authors, and other creators of the items being recommended, and who benefit from their exposure to potentially interested users), *subjects* (the people discussed by an item, such as communities affected by the events described by a news article), *platforms* or *systems*, *developers*, and other groups. Burke and Sylvester [2024] take this further, arguing for a **relational** approach to understanding recommender systems that emphasizes the relationships between many different groups affected by the system, including their relationships to each other (e.g., the relationship between readers and authors) affected but not mediated by the system.

The last decade has seen increased attention to the concerns of various stakeholders within the traditional evaluation paradigm, such as assessing *provider exposure* [Biega et al. 2018; Diaz et al. 2020; Raj and Ekstrand 2022; Singh and Joachims 2018] and examining the distributions of effects between and within different stakeholder groups [Ekstrand et al. 2024b]. However, deep engagement with these groups that sees their values and goals reflected in the design and evaluation process is still in its infancy [Bauer et al. 2024, §4.4].

2.3 Design and Evaluation for Specific Users and Contexts

In the pursuit of adaptation and personalization, and with the advances of technology, recommender systems — particularly those aiming to promote social good — are now more than ever focused on the “human” and their values and desires, and not just on their role as a “user.” This expansion of human concerns is in addition to considerations inherent to particular recommendation contexts [Pera et al. 2024]. This shift, however, presents new challenges and demands, expanding the scope of more traditional design and evaluation practices. Instead of solely designing generic solutions, optimizing algorithms, and evaluating in terms of average accuracy, there is now a need to consider human factors (e.g., personality, emotional and cognitive states) and specific context requirements that influence every stage of user and item modeling, system design, evaluation, and ultimately deployment.

Research and industry have produced a variety of design and evaluation practices that best fit the needs of specific user groups or contexts. For example, some recommendation strategies are

specifically created for users with different cognitive abilities, at various developmental stages, or for use in settings such as education or healthcare, which differ from the more typical focus on e-commerce [Banskota and Ng 2020; Gómez et al. 2021; Mauro et al. 2020; Pera and Ng 2014a; Tran et al. 2021]. Similarly, platforms like YouTube² and Spotify³ have customized their recommendations to better serve users like children [Papadamou et al. 2020]. In these cases, evaluation strategies have also been adapted, incorporating assessment criteria that best fit these unique users and contexts [Gómez Gutiérrez et al. 2021]. However, a common denominator in these efforts has been the lack of direct involvement from stakeholders who can provide an “insider perspective,” as they are the ones who are best positioned to express their needs or who have roles and perspectives specific to the particular recommendation context.

2.4 Social Goods and Harms

In addition to effectiveness or utility for a range of stakeholders, designers have long considered the various types of social impacts — both good and harmful — that recommender systems may have.

Some social benefits are closely aligned with the historical core goal of recommender systems to match users with products, information, artistic works, or other creations. For example, a system that effectively matches music listeners with artists whose work matches their nuanced tastes may provide exposure to a broader set of similar artists than a system that over-emphasizes popularity or demographics [Fleder and Hosanagar 2009; Ungruh et al. 2024]. These types of benefits must still be measured and deliberately pursued, as the system may appear to be performing well on some metrics while failing to provide the benefits desired by various stakeholders [Ekstrand et al. 2024b]. This process is made easier because it is an extension, rather than a realignment, of the objectives of recommender systems, i.e., these benefits are in line with recommendation effectiveness as long as the effectiveness is equitably distributed. To date, several lines of work have explored fairness [Deldjoo et al. 2023; Ekstrand et al. 2022; Wang et al. 2023; Zehlike et al. 2022], marketplace equity [Mehrotra et al. 2018], information access equity [Venkatasubramanian et al. 2021], and other social impact and social welfare dimensions of recommendation.

In other work, social impacts are studied as additional or even attenuating objectives on recommendation, ensuring that a system advances social objectives while meeting the needs of users and other stakeholders. Examples include research on harm-aware recommendation [Tommasel et al. 2021] and on ensuring that a wide range of users can use the system [Mauro et al. 2022; Milton et al. 2019; Ng and Pera 2018; Pera and Ng 2014a; Santos et al. 2014; Schedl and Bauer 2019].

Another type of social recommender work directly includes beneficial social impacts as general objectives to guide design and objective functions for model optimization and evaluation. Recommender systems in this category are designed to promote social objectives such as energy savings [Starke et al. 2017], sustainability [Patro et al. 2020; Tomkins et al. 2018], elder safety and quality of care [Gutiérrez et al. 2022], and access to capital [Burke et al. 2022, 2020; Smith et al. 2023].

While there has been important progress, significant work remains to ensure that recommender systems promote healthy, sustainable societies. Directly related to our present argument, there is also an urgent need to consider who is involved in identifying social objectives and designing and evaluating recommender systems against them.

²<https://www.youtube.com/myfamily/>

³<https://support.spotify.com/uk/article/spotify-kids/>

3 Participatory Design and Co-Design

Participatory design (PD) is both a *methodological approach* and a *political stance* that centers on shared decision making and the redistribution of design power. Originating in Scandinavian labor movements to address workplace inequities, PD was designed to incorporate marginalized voices in the design process. This section introduces PD and co-design, including their roots, their application to systems that incorporate significant artificial intelligence (AI) or machine learning (ML) components, and the need for recommender designers to not only adopt PD as a method but also to mindfully consider the politics of the design and design process.

3.1 History and Principles of Participatory Design and Co-Design

PD can be best described as a set of analytic and constructive commitments that are either political or methodological, depending on the PD version being used [Randall et al. 2007]. Because PD began as part of the Scandinavian workplace democracy movement, the two foundational commitments have always been (1) a broad concern with the politics of design and (2) user participation. Early Scandinavian PD projects focused on supporting and strengthening trade unions to provide workers with an opportunity to influence their work situation and use of technology in the workplace [Bjerknes and Bratteteig 1995; Gregory 2003; Kensing and Blomberg 1998]. Later projects were also based on the values and ideas from the trade union projects, but focused on the use of computers within organizational contexts. Over time, the Scandinavian system development tradition has adapted to the shifts in organizational development (e.g., from the employee to the customer), the changing role of unions, and new technological advances, while still contributing to democracy in evolving workplaces [Bjerknes and Bratteteig 1995].

In human-computer interaction (HCI) research, PD has been described as a method for building “third spaces,” or “hybrid spaces,” that are neutral and unfamiliar to both participants and researchers [Muller and Druin 2012; Muller and Kuhn 1993]. The most common setting for PD research is the participatory workshop. The workshop forms and facilitates a hybrid space by allowing participants and researchers to communicate and reach mutual understanding. To engage users during the workshops, researchers can use a range of PD techniques, methods, and practices developed over the years, including ethnographic methods, storytelling, semi-structured conferences, low-tech prototyping, mock-ups, video and storyboard prototyping, and envisioning future solutions [Muller et al. 1993]. These tools and techniques can help mediate interactions among the participants, building on their specialized knowledge [Brandt et al. 2012; Ehn 2017]. By offering participants both agency and opportunities to contribute to system design, PD approaches ensure that users hold equal status with designers throughout a democratic design process [Kensing and Blomberg 1998]. The HCI community has adopted PD for its ability to engage diverse stakeholders — especially individuals from marginalized and underrepresented groups [Haimson et al. 2020; Yip et al. 2017] — in shaping future design agendas and influencing the development of technologies that will have a direct and immediate impact on people’s lives. To date, HCI researchers have used the PD approaches for designing systems and user experiences in various application domains, including e-government and public organization [Pilemalm 2018], immigration [Duarte et al. 2018], healthcare [Kusunoki et al. 2015; Miller et al. 2016] and emergency medical services [Kristensen et al. 2006], and smart cities [Gooch et al. 2018].

In its early stages and through the mid-1990s, the concept of power was central to PD. Power was defined as a mechanism that regulates decision-making during the design process [Bratteteig and Wagner 2012]. Making decisions was essentially seen as exercising power. In the traditional PD process, power is grounded in different resources and distributed across stakeholders. For example, a project leader (i.e., a participatory design researcher) with a strong, long-term research

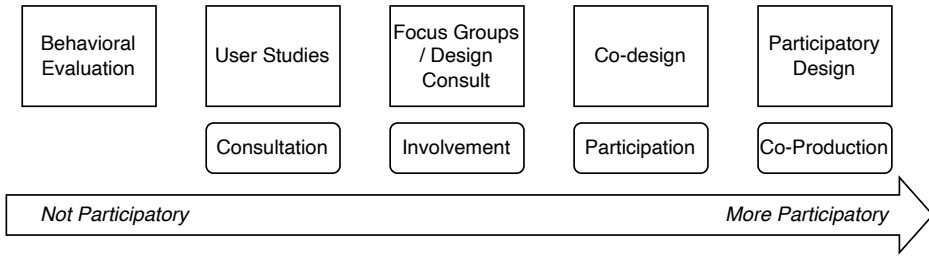


Fig. 2. Spectrum of the level of participation in various research and design paradigms.

vision and with knowledge of users and their work practices, may initiate decisions about the values and concepts, the composition of the research team, and resources allocations. The designer team is responsible for the technical design decisions, while domain experts make decisions about domain-specific issues. The domain experts also negotiate about design implementations with the outside world using their influence and connections. Finally, lay users (i.e., workshop participants) exercise their power through interacting with the prototypes, influencing the researcher’s vision and designer’s views of what they had designed [Bratteteig and Wagner 2012].

The question of power is often overlooked in contemporary discussions of PD [Bannon et al. 2018]. The concept of “user empowerment” that calls for people to be more in control when using technology [Johannsen and Kensing 2005] is distinct from user empowerment as a design framework because it mostly concerns the power of using technologies that are already designed and deployed. In the context of user empowerment, Seberger et al. [2021] distinguish between “power to” and “power over.” The “power to” concept has been defined as an individual’s capacity to act [Arendt 2019]. The power to do something within a specific context does not necessarily change the context in which the power is granted or achieved. In contrast, the concept of “power over” is based on a more traditional sociological understanding of power as the ability to compel others to act according to one’s wishes. The power over reflects inherent power imbalances and constrains the context within which “power to” can be exercised by individuals [Seberger et al. 2021]. Participatory design seeks to share “power over” not just the technology itself, but the process through which it is designed and evaluated.

The political implications of balancing power between stakeholders throughout the PD process are complicated. Several approaches to balancing power structures in a participatory design process have been proposed. One approach is to collaboratively develop power-balanced processes with the community, fostering trust between designers and social actors and allowing the community to guide decisions and share power [Tomasini Giannini and Mulder 2022]. This approach necessitates a reinterpretation of the designer’s role, promoting the de-learning of traditional practices and fostering dialogue. A power-balanced PD process views the community as an equal partner rather than merely an end-user. Another approach encourages reflection and action, allowing periodic evaluations of power dynamics to prevent biases and oppressive attitudes in the collaborative effort [Freire 1978]. As a third option, a power relation triangle has also been proposed, which includes participants, designers, and external conditions, with decision-making positioned at the core of the interactions for understanding power imbalances [Volkman et al. 2023]. In this approach, external factors may impact power relations and shifts in initiative can alter the power dynamics.

Co-design is a related participatory approach to designing technologies and systems. Although participatory design and co-design share many commonalities and are often used interchangeably, they represent two different approaches to designing systems based on the level and nature of user

participation. The participatory power of users in design research can be seen as a spectrum, starting with *consultation*, where users are informally asked for opinions or preferences, to *involvement* or more substantive opportunities for users to contribute to decision making, to *participation* or allowing users to guide the design process and even define the agenda, to *co-production*, where participants have an equal say in deciding the project goals and outcomes [Delgado et al. 2023; Sanders and Stappers 2008] (Figure 2). User participation in co-design is typically at the *participation* level, whereas user participation in PD usually reaches the *co-production* level because PD also involves the political component within the larger socio-technical system being designed.

3.2 Co-Designing AI-Based Systems

In recent years, researchers have explored the application of PD in designing AI-based technologies as “participatory AI” [Delgado et al. 2023]. Because emerging technologies such as AI can negatively affect people’s lives through algorithmic bias, opaque decision making, and the spread of misinformation [Demartini et al. 2020], applying participatory approaches to designing AI-based systems is critical.

“Participatory ML” [Kulynych et al. 2020] has emerged as a co-design approach for considering the main components of ML systems from a human-centered perspective. Critical considerations across different stages of ML design and development include whether (1) the datasets are ecologically valid for the targeted population; (2) the algorithms are grounded in human theory, understanding, and knowledge; (3) the evaluation metrics are including both quantitative and qualitative metrics; and (4) the artifacts and outcomes are safely deployed in real-world settings based on human values [Razi et al. 2021]. PD approaches have been used to address each of these critical aspects of ML design, including dataset building [Pushkarna et al. 2022], design and validation of ML algorithms [Theodorou et al. 2021], and understanding of algorithmic accountability [of Queer in AI et al. 2023] and equity [Katell et al. 2020].

Adopting PD for designing AI-based systems across different application domains raises several methodological questions [Zytke et al. 2022]:

- (1) Who will conduct participatory AI design research?
- (2) Which aspects of AI can or should be co-designed?
- (3) What will be the process of user participation in AI co-design?
- (4) What artifacts will be used in the process?
- (5) How will the collaboration between co-designers and researchers with different domain expertise be facilitated?
- (6) What are the persistent challenges to participatory AI design?

Recent AI co-design efforts have been criticized for minimal engagement activities with users through a single design workshop or focus group aimed at obtaining feedback on designer-generated artifacts [Zytke et al. 2022]. This superficial level of engagement, only reaching the involvement level of participation (Figure 2), diminishes the fundamental objective of PD: to foster active and sustained involvement of user and influence the final system design through a highly interactive and iterative process [Zytke et al. 2022]. Ongoing engagement and sharing of power is also a key distinction between fully participatory AI and human-centered AI, as articulated by [Shneiderman 2022], which focuses on consulting with users through traditional HCI techniques rather than fully involving them as co-producers of AI. Other critiques of practices in participatory ML include equitable work with participants [Corbett et al. 2023; Delgado et al. 2023; Feffer et al. 2023], “participation-washing” [Sloane et al. 2022], and co-option of participatory work [Birhane et al. 2022].

A recent interview study of “participation brokers” who lead participatory ML projects highlighted the need for equitably balancing the value generated through participatory ML with educating users and advocating for participant rights [Cooper and Zafiroglu 2024]. Alleviating frustration from indirect stakeholders by converting qualitative stories into main ML processes is also important in the participatory context [Cooper and Zafiroglu 2024]. Several efforts have been evaluated that use generative AI to replace human participation, thereby reducing the cost of research and development and increasing the diversity of collected data [Dillion et al. 2023]. However, replacing research participants with AI-generated data conflicts with the essential values of representation, inclusion, and understanding of the human [Agnew et al. 2024; Harding et al. 2024].

PD and co-design offer more than methods for creating systems — they represent a commitment to equitable design practices that center diverse voices and share decision-making power. While AI/ML-intensive systems present unique challenges, integrating participatory approaches requires not only methodological shifts but also a fundamental rethinking of design politics. Adopting such commitments in recommender system design can lead to more inclusive, accountable, and impactful innovations that serve diverse user bases and advance democratically-developed social good.

4 Participatory Recommendation

Our central argument (outlined in §1) is that recommender systems — particularly those aimed at social good — should be designed *with*, not just *for*, their users, providers, society at large, and various specific people and groups impacted by the system. Under this approach of *participatory recommendation*, designers and other stakeholders work together in a *participatory design team* (PD team) to design and evaluate a recommender system. Participatory methods are important for all types of recommendation goals, not just those aimed at social benefit. However, the explicit role of values in determining and pursuing social good heightens the necessity of participatory methods because they provide a means for ensuring the values incorporated into the system come from the community in which it will operate.

A commitment to participatory recommendation will pervade the entire recommender system design and evaluation process. The objectives for the recommender system, the computational and interactional mechanisms used to meet those objectives, and the evaluation of whether the system achieves its objectives should all be developed in collaboration with the system’s beneficiaries. This participatory collaboration is necessary not only for the kinds of individual benefit or utility objectives frequently considered in recommender system design but also for objectives related to social impact and welfare. PD team collaborations should be genuinely power-sharing relationships in which the people who will use and be impacted by the system are not only consultants or subjects in user studies, but have meaningful control over the final product and its initial and ongoing evaluation. Practically, participatory recommendation extends the arguments by Ekstrand and Willemsen [2016] to avoid purely behaviorist approaches to designing and evaluating recommender systems, providing a realization of the argument made by Burke and Sylvester [2024] that the recommender systems community must attend to the relationships between various parties in a recommendation interaction.

Human-centered and user-centered approaches to AI generally [Shneiderman 2022] and recommender systems specifically [Konstan and Terveen 2021], grounded in traditional human-computer interaction methods of user-centered design and evaluation through user studies, provide a valuable starting point for participatory recommendation. However, realizing a participatory vision of full co-production of recommender systems, or even participatory co-design, requires more. First, more stakeholders should be involved. In addition to including users, participatory recommendation design needs to involve producers and other people affected by the recommender system’s

operation. Second, traditional human-centered methods retain power, particularly design and evaluation authority, with the designers, whereas participatory recommendation shares that power between the designers and other stakeholders. Finally, while “control” is important in both visions, human-centered AI is primarily concerned with designing systems that allow user control over their operation, while participatory recommendation and participatory AI envision ensuring users and other stakeholders have power in the design and evaluation process itself.

Participatory recommendation is not an *easy* agenda to pursue (technically, organizationally, or politically), but it is necessary to truly achieve the goals of recommendation for social good, and in turn, recommendation for general human benefit. Ursula Franklin [2004] distinguishes between *relational* technologies built on reciprocity and mutual respect and unidirectional *broadcast* technologies. Too often, recommender systems operate as broadcast technologies, but we posit that participatory approaches are key to unlocking recommender systems as a relational technology. Relational understandings of recommendation are not new, as they were articulated in the field’s early days by Hill et al. [1995] and more recently highlighted by Burke and Sylvester [2024]. In an era of techlash [Atkinson et al. 2019], we also hypothesize that people will be more accepting and trusting of recommender systems if they or their communities had a meaningful say in creating them [Ekstrand and Willemsen 2016].

Participatory recommendation will likely need to start with small-scale systems, aimed at solving information access and discovery problems for small, well-defined use cases. Lessons from designing smaller systems will inform the methods needed to design or re-design global systems. Global designs will likely arise as amalgamations of approaches from many smaller, more localized design efforts because of the challenges in representing a global population of users, sellers, subjects, or others in a manageable participatory design setting. A robust marketplace of small- to medium-scale information ecosystems and e-commerce platforms may also be socially beneficial for other reasons [Rajendra-Nicoluccis et al. 2023].

Effective design, development, and evaluation will iterate between the various stages, in close collaboration with people affected in different ways and with different degrees of power. This collaboration is not unidirectional, but a true collaborative effort with shared power. The recommender system designers should also not be reduced to stenographers of designs produced by others, as they bring significant expertise on technical possibilities and the computational and systemic consequences of potential designs. In a participatory approach, designers are heavily involved but do not serve as the final arbiters of the design. Instead, they contribute their expertise to a conversation and collaboration with people who will use and be impacted by the design. Implementing participatory recommendation requires careful consideration of the power of different people involved in the processes. Access to data and computational resources are also a dimension along which designers and the platform employing them may have crucial power not held by other participants in the design process.

We next describe how different elements of the recommender systems development lifecycle can be reenvisioned as a participatory process. Our discussion here is forward-looking, as much of the work on adapting and applying participatory design and co-design, and developing the technical and methodological capabilities to support these methods in information access and AI systems has yet to be done.

4.1 Conception and Problem Definition

Recommender system design begins with conceiving the initial application and defining the problem that the system would address. A participatory approach will start by defining the problem with people who will use or be affected by the system, without pre-supposing that a recommender

system is necessary or appropriate. This approach applies to greenfield platforms or applications and to new or enhanced recommendation capabilities in existing systems.

Participatory recommendation starts from a place of humility and begins with a page from information science, information retrieval, and HCI: understanding what users and other stakeholders *need*. Music recommender systems can serve as a simple example. Users may need to know when their favorite artists have released new music. Depending on the platform dynamics, artists may desire to have their newer music surfaced if current listening and recommendation patterns favor their older hits. The platform operator needs their business to generate enough profit to be sustainable. Developers may have a range of needs, from continued employment to an intrinsic interest in supporting music discovery. Although all of these needs could point to a recommendation experience that helps locate new music, participatory methods will help elicit these needs and resulting designs directly from the platform stakeholders.

Some of the participatory engagement at this stage of the recommendation design can be done at the *consultation* level of participation (Figure 2), such as working with focus groups of users or artists. Participatory methods begin to differentiate from this consultation as more people are involved in the later stages, and as the recommender system designers commit to collaborate with the community participants.

Another key part of defining the problem is defining the stakeholder groups and recruiting participants for the design process. Stakeholder identification will often start with the designers' ideas of what perspectives are important for a particular recommender system. Designers should engage in that planning, as the goals of participatory collaboration with shared power suggest inviting the early participants to recommend additional people or groups that are not yet accounted for. In such designer-originated participatory recommendation projects, perhaps designed to make existing human-centered research and development more participatory, participant recruitment and compensation will most likely happen through existing organizational channels or through building such mechanisms where they do not yet exist. In some cases, however, the recommender system concept may originate from a stakeholder group themselves, such as an artist collective wanting a recommender system to help clients or patrons locate and select from the work of the different artists. Such ground-up, community-originated efforts, where they exist, are a clear realization of the full co-production level of participatory design, and already come with at least one set of stakeholders and participants.

In both problem definition and the remaining phases of participatory recommender systems research, different participants may have conflicting goals or design suggestions. They may also propose ideas that expand beyond the capability of current systems or with negative social or system consequences. Navigating these conflicts through collaboration is vital to ensuring a truly participatory design and evaluation process, although the precise details will depend on the context and application. Rau et al. [2013] provide one approach to resolving conflicting design goals through identifying hierarchies of goals that may be useful in participatory recommendation, as long as that hierarchy itself is an outcome of the participatory collaboration rather than imposed by the designers.

4.2 Design

Participatory design and co-design make a clearer difference as we move from problem definition to the actual design of the recommender system and its surrounding application and user experience. Once the PD team has collaboratively agreed on the need for a particular recommender system, the actual design proceeds in continued collaboration.

Co-designing a recommender system will involve meaningful input and decision-making authority from multiple stakeholders on all aspects of the design, including:

- Location of the recommendations within the user workflow.
- Criteria for item inclusion and prioritization.
- User interface elements (display, layout, feedback or control mechanisms).
- Data used (or not used) as a basis for recommendation.
- Permission structures around opting in or out of the recommendations (for both users and providers).
- Key principles of evaluation and objective functions.

We discuss the last point above — evaluation and objectives — in more detail in the next section, but as recommender model behavior is driven in large part by the designs of its objective functions and the evaluations by which it is assessed and optimized, these metrics and analyses should be collaboratively developed by the PD team. Schellingerhout et al. [2023] have made some early progress in co-designing aspects of recommender systems, finding that different stakeholders in a system may have very different needs from its outputs (in their study, explanations).

The degree of agency and ability the different members of the PD team have to shape the final outcomes determines where on the participatory spectrum a particular recommender system design effort will fall (Figure 2). Co-design focuses on collaboration during the design phase, involving stakeholders (e.g., users or communities) to contribute their perspectives and ideas and provide feedback, reaching the “participation” level. Full co-production includes multiple stakeholders as equal partners across all stages of a project. Both of these levels represent an improvement in participation over standard recommender system design practice.

4.3 Evaluation

Evaluation is a key aspect of recommendation, particularly its algorithmic and modeling components. As noted in §4.2, it is deeply intertwined with other design efforts: designing an evaluation and related objective functions (for model training, hyperparameter optimization, design evaluation, etc.) is a part of recommender system design.

Substantial research is needed to determine how to effectively co-design — and “co-evaluate” — recommender systems. Metrics and evaluation setups are complex, mathematically nuanced, and statistical (and sometimes directly stochastic). Therefore, they can be difficult even for experienced recommender system designers to interpret. Participatory recommendation during the evaluation phases will robustly involve people with various perspectives on the system and widely varying levels of expertise in other topics, but rarely in recommender systems. Facilitating this involvement will likely require advances in both participatory protocols and computational evaluation methods. The full PD team can and should be involved in many different ways, including:

- Qualitatively evaluating the system based on their own experience.
- Designing and validating evaluation metrics.
- Reviewing and selecting evaluation data and data preparation steps (e.g., integration, splitting, setting up simulated tasks).
- Designing online evaluation (e.g., what metrics in an online field trial would indicate system success).
- Reviewing evaluation results to assess whether the system meets the problem and design parameters (as discussed in §4.1 and §4.2, respectively).

Some aspects of the system can be directly and qualitatively assessed by design participants, particularly user interface and directly user-observable aspects (e.g., inclusion or exclusion of appropriate items, or perceived diversity). Many teams already do this internally by reviewing their own recommendations. Expanding this initial user involvement to a broader set of representatives will improve the evaluation. Other aspects of recommender system behavior and impact are harder

to directly observe, either because their effects are behind the scenes or they are statistical (and therefore cannot be assessed from individual recommendation results or interactions). Fairness objectives are one class that often fall into this category: looking at a single recommendation list, or even many recommendations for a single user, is not sufficient to assess many reasonable fairness criteria, such as equal opportunity for content providers [Biega et al. 2018; Diaz et al. 2020].

Fully participatory co-production of a recommender system arises when the user and other stakeholder authority is equal to that of the recommender system designers in deciding whether the evaluation shows the system is fit for purpose and should be deployed (or not). Smith et al. [2024] incorporated recommender system practitioners in the process of co-designing recommender system evaluation, but much more work is needed to incorporate additional stakeholder participants and account for power differentials and stakeholder conflict.

4.4 Monitoring and Review

Recommender system design and evaluation is not a one-and-done process, but rather an iterative one in which the system is continually reviewed and improved, and is monitored for changes in effectiveness or behavior as new content, users, interactions, and other data arrive.

Building on the iterative nature of participatory design, participatory recommendation should incorporate participatory input and ongoing co-design efforts in at least three ways:

- (1) Designing the ongoing monitoring and evaluation procedures to be carried out by the platform. What needs to be measured and monitored to ensure the system's ongoing fidelity to the agreed-upon problem definition and design?
- (2) Regular review of system behavior and emerging impacts, not just by the designers but by users.
- (3) Conception and design of improvements, enhancements, or removals of the design based on ongoing experience and community input (e.g., what users hear from other members of their communities).

Ongoing review and revision by the full PD team will ensure that the recommender system continues to meet the needs of the many different people with an interest in its operation and behavior.

5 Participatory Approaches in Practice

To demonstrate our argument, we discuss how participatory methods could impact the design and evaluation of recommender systems in several application contexts.

5.1 Fair Multi-sided Marketplaces

Many platforms employing recommender systems are some form of *multi-sided marketplace*, connecting people to each other through the products, media, or opportunities being recommended. In these settings, people have different interests and needs with regard to the recommender system's usefulness, fairness, and other properties, depending on their role in any particular transaction or interaction [Sonboli et al. 2022]. A recent line of work has proposed methods for human-centered operationalization of recommendations. This approach acknowledges that fairness is a contested concept that will have different meanings to different individuals and begins with exploratory semi-structured interviews eliciting how fairness issues arise in everyday experiences of users. One example is an interview study with employees of the crowd-funded microlending platform Kiva Microloans [Smith et al. 2023]. The findings showed a wide range of fairness logics articulated by the employees, who, in their organizational roles, had greater or lesser interaction with different stakeholders served by Kiva: lenders, borrowers, partner organizations, corporate sponsors, and

others. Although using a human-centered approach, this study did not incorporate actual end users (recommender consumers) of Kiva, whose concepts of fairness were examined in a previous study [Sonboli et al. 2021]. The same approach was pursued in an interview study with journalists that looked into the way journalistic values are (or are not) being realized in news recommender systems [Alaqabawy et al. 2023].

While interviews and subsequent analysis help recommender system designers develop general categories and logics of fairness, further steps are needed to fully operationalize and implement the ideas identified through interviews. It is up to system designers to articulate the concepts users express in computational terms. This process of operationalization is uncertain because crucial nuances, obvious but implicit for practitioners, might be lost in the (technical) translation. Once system designers have developed mathematical definitions of proposed fairness metrics, these can be brought back to stakeholders for discussion in focus groups. Smith et al. [2024] illustrate this two-step focus group process in the context of provider-oriented fairness for recommendations in social media and dating sites. In addition to examining designers' fairness measures, Smith et al. encouraged providers to formalize their own fairness notions.

In participatory recommendation, participation in the assessment of fairness continues through the design and implementation of transparency mechanisms. Although considerable research exists on recommendation explanation, it is heavily focused on the explanation of individual recommendation results [Tintarev and Masthoff 2015] and not on the types of insights into overall system behavior required by a participatory design effort. In addition, explanation and transparency for stakeholders other than recommendation consumers, such as providers of recommended items, are rarely studied in recommender research. Varasteh et al. [2024] discuss some of the challenges that would be involved in developing provider-side transparency mechanisms. Still, such tools would be essential for ongoing monitoring by providers throughout a recommender system's lifecycle. Research on operationalizing fairness objectives in collaboration with non-developer stakeholders, or with multiple stakeholder groups in a single effort, is also limited.

5.2 Social Media and Online Safety

Recommender systems are influential in shaping user experiences in *social media* and influencing what content users see and interact with. However, the mechanisms that make these systems engaging and personalized also raise critical concerns about their impact on user **online safety** and their potential to amplify harmful content. Recommender ranking algorithms, often prioritizing engagement over safety, can inadvertently expose users to harmful, inappropriate, or distressing content. This exposure can range from explicit content like violence and hate speech to subtler forms of harm like misinformation, stereotypes, and content that promotes unhealthy behaviors [Ungruh and Pera 2024]. For more vulnerable users such as children, the exposure to harmful content can lead to "harmful pathways," where users encounter increasingly extreme content over time [Ungruh and Pera 2024]. For instance, for individuals in recovery from conditions like eating disorders, encountering pro-eating disorder content through recommendations can be particularly distressing and can even lead to relapse. The challenge is that these types of recommender systems often struggle to discern nuanced contexts and may associate pro-eating disorder content with recovery content, leading to harmful recommendations [Golbeck 2025].

As in other domains, social media recommender systems can exacerbate existing social and economic inequalities, exposure to misinformation, hate speech, and harmful stereotypes. The opacity of many systems, particularly those using complex AI models, makes it difficult for users to understand how these systems work and why they are being recommended certain content. This lack of transparency limits users' ability to control their online experiences, challenge potentially

biased or harmful recommendations, and hold platforms accountable for the outcomes of their algorithms [Smith et al. 2022].

Participatory approaches will involve a wider range of stakeholders to collaboratively understand and address the potential harms of recommender systems in social media. Participatory design can be used to develop explanations for why certain content is being recommended and can empower users to make informed choices and identify potential biases or harms. Using co-design strategies, users could ensure more control over their recommendation settings, allowing them to customize content preferences, filter out specific topics, or opt out of certain types of recommendations [Smith et al. 2022]. Participatory and co-design approaches involving various stakeholders can help identify and define potential harms that may not be apparent to developers or researchers. Stakeholders can provide valuable input on developing metrics that accurately reflect the values they care about. In addition, the lived experiences and insights of users can inform the development of interventions and mitigation strategies that are both practical and responsive to their needs and the needs of their communities.

5.3 Education

The ideas of participatory recommendation are not limited to recommender systems, but are applicable to a range of information access technologies, including search engines and natural language tools for information tasks. Several of the authors have been working on developing recommender systems and search tools for *educational settings*, particularly to support K–12 teachers in finding reading materials for their students [Ekstrand et al. 2020; Jason Hall 2017; Murgia et al. 2019; Pera and Ng 2014b].

One specific example has as an educational goal supporting children’s literacy development and proficiency by engaging them with diverse *informational texts* in their early education [Ekstrand et al. 2020]. The majority of texts read by adults are informational [Atkinson et al. 2009], and comprehension of these non-narrative, nonfiction texts is key to academic success [Duke and Bennett-Armistead 2003]. The degree of text “authenticity” — whether the author wrote the text purely for pedagogical purposes (e.g., a textbook) or for a wider, real-world audience (e.g., a newspaper) — also significantly impacts learning. For example, Purcell-Gates et al. [2007] found that elementary children who engaged in authentic science literacy tasks were better able to comprehend and produce informational texts in later grades.

While research provides guidance on best practices on using informational texts in the classroom [e.g., Bradley and Donovan 2010; Maloch 2008], it fails to address the challenge of finding these readings. Ekstrand et al. [2020] found that existing search technology does not suffice when teachers are looking for authentic informational texts for classroom use. In the reported interviews, teachers frequently discussed that they could not find what they were looking for on Google, and that typical search engines often led them to materials blocked by a paywall. This work also highlighted several specific problems that can serve as a starting point for technical innovation — current search and recommender systems are not designed to support teachers in identifying resources that are relevant, appropriate, and accessible to their students. In addition, the results showed teacher interest in a recommender system that would (1) allow them to locate a variety of texts that match student interests, (2) filter for readability, and (3) support their ability to locate free resources that match curriculum needs.

Fully realizing this kind of information access system in a manner that *empowers* teachers and integrates with their work and teaching contexts demands a participatory approach in which teachers are full collaborators — if not leaders — in its design and evaluation. Because teachers are experts in both the curriculum and pedagogy of their classrooms, they are best positioned to accurately identify what will best serve their instructional needs, both in terms of the recommended

texts and the technology. The educational technology market is replete with technologies that have been developed and “thrown over the wall,” designed to serve the interests of administrators and technologists instead of students. Many are built on surveillance-capitalist technology paradigms [Hillman and Esquivel 2022; Stockman and Nottingham 2022], bypassing teachers by connecting directly with students (and harvesting data from them). Re-envisioning the information access and educational technology design process around participatory design with teachers at the center is a promising approach to delivering technology that will integrate well with classroom activities and meet real, on-the-ground needs faced by teachers and students. It will also result in technology that is better adapted to the particularities of different teaching contexts, easing adoption and improving impact. Teachers’ involvement in design should not be limited to visuals, layouts, and high-level features, but also to designing the machine learning objectives and, ultimately, the metrics and experimental procedures that will be used to design and iterate on the technical advances needed to bring teachers’ design ideas to life.

6 The Vision of Participatory Recommendation

We envision a future where people are meaningfully involved in designing and evaluating the recommender systems that affect them. The participatory agenda we have described is not limited to recommender systems, but also applies to other information access systems, such as search engines and informational AI systems.

The vision of participatory recommendation includes two major components:

- **Methodological changes** to incorporate participatory methods into recommender system design, evaluation, and deployment.
- **Political commitments** about how power is allocated, whose opinions on recommender systems matter, and who is empowered to meaningfully affect both individual systems and the broader landscape of recommendation, retrieval, and personalization.

These methodological and political moves are particularly important in the pursuit of recommender systems for social good because such efforts necessarily entail judgments about what constitutes a “social good.” To build recommender systems that truly empower communities and societies to pursue good, it is necessary to use meaningful participatory approaches in which designers relinquish power in determining what is ultimately built. Traditional design paradigms risk building systems that paternalistically impose visions of social good defined by organizations (or individuals) that do not share the values of those most directly impacted by the system or of well-meaning systems harming the communities they are intended to support.

We do not diminish or devalue technical work aimed at social goals. Directly pursuing social goals through technical research and development, in addition to establishing and evaluating technical prerequisites that can be applied to positive social outcomes, are valuable and necessary. We hope to see such work grounded and contextualized in a robust understanding of social good, including how that notion of social good was derived. We also hope to see any claims of actual attainment of social good assessed by the people who will be affected — positively or negatively — by the proposed good.

Embracing the challenge of participatory recommendation requires extensive research on several important points, including those listed in §3.2 for participatory AI in general. Key research needs include:

- Tools and methods, validated across multiple applications, to support participatory design and co-design of recommender systems, information retrieval, and similar technologies. Existing methods are highly applicable and recent efforts such as participatory AI playbooks

[Chang and Dudy 2022; Pushkarna et al. 2022] will be useful for recommender systems; new or revised tools may be necessary to support effective co-design.

- Evaluation techniques, metrics, and reporting concepts, along with participatory protocols using them, that enable non-experts to meaningfully assess a recommender system’s performance or behavior.
- No-code or low-code recommender system design tools that allow communities to adjust recommendations to their particular needs.
- Adaptable, low-data recommendation techniques that can be deployed and customized without the need to collect extensive interaction data by means typically only available to large platforms.
- Accessible, participatory auditing methods and tools to enable people to rigorously interrogate and understand the behavior of the recommender systems they already use.

Recommender systems and other information access technologies have significant capability both to advance social good and to catalyze harm. Deliberate design to effectively meet user needs, provide opportunities for providers, and advance social welfare objectives is vital to ensuring the technology benefits the people it affects.

Although our present argument emphasizes the need, value, and practical guidance for designing recommender systems with a focus on social good, we also recognize that these systems do not operate in isolation and do not serve a single objective. Multiple stakeholders contribute to and influence the requirements of recommender systems, which must ultimately be balanced and prioritized [Abdollahpouri et al. 2020; Sürer et al. 2018]. This balancing effort is challenging because there is no one-size-fits-all solution — the criteria, objectives, and priorities of different participants’ concerns vary depending on the specific use case.

One of the most immediate conflicts in participatory recommendation, or even recommendation for social goals through traditional design and research methods, is the inherent tension between profit-driven incentives and the pursuit of broader social good. Many recommender systems are integrated within commercial platforms that primarily aim to maximize sales, engagement, retention, and ultimately revenue [Jannach and Bauer 2020]. These financial incentives can sometimes conflict with ethical considerations or with the power-sharing objectives of participatory design. For instance, engagement-driven algorithms may inadvertently prioritize sensationalist or polarizing content because it generates more clicks and user interaction, even though such content can have negative societal implications. Addressing this tension, along with those arising from multiple stakeholder considerations, is an ongoing discussion. In this paper and our prior work [Burke et al. 2025; Ekstrand et al. 2024a] we are attempting to advance that vital conversation about the importance of balancing diverse needs and perspectives.

Bringing the participatory turn in AI [Delgado et al. 2023] to the project of building useful and socially-beneficial recommender systems will give more people the opportunity to directly contribute to these design efforts, and hold such efforts accountable to the people and communities they are intended to benefit.

Information access systems will best serve social good when people who experience their consequences are empowered to participate in their design and evaluation. We invite you to join us in making participatory recommendation a reality.

Acknowledgments

Robin Burke was supported by the National Science Foundation under grant IIS-2107577.

References

- Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder Recommendation: Survey and Research Directions. *User Modeling and User-Adapted Interaction* 30, 1 (March 2020), 127–158. doi:10.1007/s11257-019-09256-1
- Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Recommender Systems as Multistakeholder Environments. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*. ACM, 347–348. doi:10.1145/3079628.3079657
- William Agnew, A. Stevie Bergman, Jennifer Chien, Mark Díaz, Seliem El-Sayed, Jaylen Pittman, Shakir Mohamed, and Kevin R. McKee. 2024. The Illusion of Artificial Inclusion. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, New York, NY, USA, 1–12. doi:10.1145/3613904.3642703
- Nihal Alaqabawy, Aishwarya Satwani, Amy Volda, and Robin Burke. 2023. “It’s Just a Robot That Looks at Numbers”: Restoring Journalistic Voice in News Recommendation. *Proceedings of the International Workshop on News Recommendation and Analytics* (Sept. 2023). <https://ceur-ws.org/Vol-3561/paper1.pdf>
- Hannah Arendt. 2019. *The Human Condition*. University of Chicago Press.
- Robert D. Atkinson, Doug Brake, Daniel Castro, Colin Cunliff, Joe Kennedy, Michael McLaughlin, Alan McQuinn, and Joshua New. 2019. *A Policymaker’s Guide to the “Techlash” —What It Is and Why It’s a Threat to Growth and Progress*. Technical Report. Information Technology & Innovation Foundation. <https://itif.org/publications/2019/10/28/policymakers-guide-techlash/>
- Terry S. Atkinson, Melissa N. Matusевич, and Lisa Huber. 2009. Making Science Trade Book Choices for Elementary Classrooms. *The Reading Teacher* 62, 6 (2009), 484–497. doi:10.1598/RT.62.6.3
- Liam Bannon, Jeffrey Bardzell, and Susanne Bødker. 2018. Reimagining Participatory Design. *Interactions* 26, 1 (Dec. 2018), 26–32. doi:10.1145/3292015
- 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*. ACM, New York, NY, USA, 14–22. doi:10.1145/3340631.3394867
- Christine Bauer, Alan Said, and Eva Zangerle. 2024. Evaluation Perspectives of Recommender Systems: Driving Research and Education (Dagstuhl Seminar 24211). *Dagstuhl Reports* 14, 5 (2024), 58–172. doi:10.4230/DagRep.14.5.58
- Nicholas J Belkin and Stephen E Robertson. 1976. Some Ethical and Political Implications of Theoretical Research in Information Science. In *Proceedings of the ASIS Annual Meeting*. <https://www.researchgate.net/publication/255563562>
- Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. ACM, 405–414. doi:10.1145/3209978.3210063
- Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the People? Opportunities and Challenges for Participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '22)*. ACM, New York, NY, USA, 1–8. doi:10.1145/3551624.3555290
- Gro Bjørknes and Tone Bratteteig. 1995. User Participation and Democracy: A Discussion of Scandinavian Research on System Development. *Scandinavian Journal of Information Systems* 7, 1 (1995), 73–98. <https://aisel.aisnet.org/sjis/vol7/iss1/1>
- Linda Golson Bradley and Carol A Donovan. 2010. Information Book Read-Alouds as Models for Second-Grade Authors. *The Reading Teacher* 64, 4 (2010), 246–260. doi:10.1598/RT.64.4.3
- Eva Brandt, Thomas Binder, and Elizabeth B-N Sanders. 2012. Tools and Techniques: Ways to Engage Telling, Making and Enacting. In *Routledge International Handbook of Participatory Design*. Routledge, 145–181.
- Tone Bratteteig and Ina Wagner. 2012. Disentangling Power and Decision-Making in Participatory Design. In *Proceedings of the 12th Participatory Design Conference: Research Papers - Volume 1 (PDC '12)*. ACM, New York, NY, USA, 41–50. doi:10.1145/2347635.2347642
- Robin Burke. 2017. Multisided Fairness for Recommendation. (July 2017). <http://arxiv.org/abs/1707.00093>
- Robin Burke, Gediminas Adomavicius, Toine Bogers, Tommaso Di Noia, Dominik Kowald, Julia Neidhardt, Özlem Özgöbek, Maria Soledad Pera, Nava Tintarev, and Jürgen Ziegler. 2025. De-Centering the (Traditional) User: Multistakeholder Evaluation of Recommender Systems. arXiv:2501.05170 [cs] doi:10.48550/arXiv.2501.05170
- Robin Burke, Pradeep Ragothaman, Nicholas Mattei, Brian Kimmig, Amy Volda, Nasim Sonboli, Anushka Kathait, and Melissa Fabros. 2022. A Performance-Preserving Fairness Intervention for Adaptive Microfinance Recommendation. In *Proceedings of the KDD Workshop on Online and Adapting Recommender Systems (OARS)*. <https://par.nsf.gov/biblio/10351344>
- Robin Burke and Morgan Sylvester. 2024. Post-Userist Recommender Systems: A Manifesto. arXiv:2410.11870 doi:10.48550/arXiv.2410.11870
- Robin Burke, Amy Volda, Nicholas Mattei, and Nasim Sonboli. 2020. Algorithmic Fairness, Institutional Logics, and Social Choice. In *Harvard CRCS Workshop on AI for Social Good at the 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*. https://projects.iq.harvard.edu/files/crcs/files/ai4sg_2020_paper_29.pdf

- Michael Alan Chang and Shiran Dudy. 2022. Expansive Participatory AI: Supporting Dreaming within Inequitable Institutions. arXiv:2211.12434 doi:10.48550/arXiv.2211.12434
- Vicky Charisi, Stéphane Chaudron, Rosanna Di Gioia, Riina Vuorikari, Marina Escobar-Planas, Ignacio Sanchez, and Emilia Gomez. 2022. *Artificial Intelligence and the Rights of the Child: Towards an Integrated Agenda for Research and Policy*. Publications Office of the European Union. doi:10.2760/012329
- Ned Cooper and Alexandra Zafiroglu. 2024. From Fitting Participation to Forging Relationships: The Art of Participatory ML. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, New York, NY, USA, 1–9. doi:10.1145/3613904.3642775
- Eric Corbett, Emily Denton, and Sheena Erete. 2023. Power and Public Participation in AI. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '23)*. ACM, New York, NY, USA, 1–13. doi:10.1145/3617694.3623228
- Yashar Deldjoo, Dietmar Jannach, Alejandro Bellogin, Alessandro Difonzo, and Dario Zanzonelli. 2023. Fairness in Recommender Systems: Research Landscape and Future Directions. *User Modeling and User-Adapted Interaction* (April 2023). doi:10.1007/s11257-023-09364-z
- Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '23)*. ACM, New York, NY, USA, 1–23. doi:10.1145/3617694.3623261
- Gianluca Demartini, Stefano Mizzaro, and Damiano Spina. 2020. Human-in-the-Loop Artificial Intelligence for Fighting Online Misinformation: Challenges and Opportunities. *The Bulletin of the Technical Committee on Data Engineering* 43, 3 (Sept. 2020), 65–74. <http://sites.computer.org/debull/A20sept/p65.pdf>
- Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating Stochastic Rankings with Expected Exposure. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*. ACM. doi:10.1145/3340531.3411962
- Danica Dillion, Niket Tandon, Yuling Gu, and Kurt Gray. 2023. Can AI Language Models Replace Human Participants? *Trends in Cognitive Sciences* 27, 7 (July 2023), 597–600. doi:10.1016/j.tics.2023.04.008
- Ana Maria Bustamante Duarte, Nina Brendel, Auriol Degbelo, and Christian Kray. 2018. Participatory Design and Participatory Research: An HCI Case Study with Young Forced Migrants. *ACM Transactions on Computer-Human Interaction* 25, 1 (Feb. 2018), 3:1–3:39. doi:10.1145/3145472
- Nell K. Duke and Susan V. Bennett-Armistead. 2003. *Reading & Writing Informational Text in the Primary Grades*. Scholastic Teaching Resources.
- Pelle Ehn. 2017. Scandinavian Design: On Participation and Skill. In *Participatory Design*. CRC Press, 41–77.
- Michael D. Ekstrand, Lex Beattie, Maria Soledad Pera, and Henriette Cramer. 2024a. Not Just Algorithms: Strategically Addressing Consumer Impacts in Information Retrieval. In *Proceedings of the 46th European Conference on Information Retrieval (Lecture Notes in Computer Science, Vol. 14611)*. Springer. doi:10.1007/978-3-031-56066-8_25
- Michael D. Ekstrand, Ben Carterette, and Fernando Diaz. 2024b. Distributionally-Informed Recommender System Evaluation. *ACM Transactions on Recommender Systems* 2, 1 (March 2024), 6:1–27. doi:10.1145/3613455
- 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. doi:10.1561/15000000079
- Michael D Ekstrand and Martijn C Willemsen. 2016. Behaviorism Is Not Enough: Better Recommendations through Listening to Users. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 221–224. doi:10.1145/2959100.2959179
- Michael D Ekstrand, Katherine Landau Wright, and Maria Soledad Pera. 2020. Enhancing Classroom Instruction with Online News. *Aslib Journal of Information Management* 72, 5 (Jan. 2020), 725–744. doi:10.1108/AJIM-11-2019-0309
- Michael Feffer, Michael Skirpan, Zachary Lipton, and Hoda Heidari. 2023. From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research. In *Proceedings of the 2023 AAIL/ACM Conference on AI, Ethics, and Society (AIIES '23)*. ACM, New York, NY, USA, 38–48. doi:10.1145/3600211.3604661
- Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. What Is Fair? Exploring the Artists' Perspective on the Fairness of Music Streaming Platforms. In *Proceedings of the 18th IFIP International Conference on Human-Computer Interaction (INTERACT 2021)*. <http://arxiv.org/abs/2106.02415>
- Daniel M Fleder and Kartik Hosanagar. 2009. Blockbuster Culture's next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science* 55, 5 (May 2009), 697–712. doi:10.1287/mnsc.1080.0974
- Ursula M Franklin. 2004. *The Real World of Technology* (revised edition ed.). House of Anansi Press, Toronto, Ont.; Berkeley, CA. <http://www.cbc.ca/radio/ideas/the-1989-cbc-massey-lectures-the-real-world-of-technology-1.2946845>
- Paulo Freire. 1978. Pedagogy of the Oppressed*. In *Toward a Sociology of Education*, John Beck, Chris Jenks, Nell Keddie, and Michael F. D. Young (Eds.). Routledge.
- Jennifer Golbeck. 2025. Recommender System-Induced Eating Disorder Relapse: Harmful Content and the Challenges of Responsible Recommendation. *ACM Trans. Intell. Syst. Technol.* 16, 1 (Jan. 2025), 15:1–15:13. doi:10.1145/3675404

- Elizabeth Gómez, Carlos Shui Zhang, Ludovico Boratto, Maria Salamó, and Mirko Marras. 2021. The Winner Takes It All: Geographic Imbalance and Provider (Un)Fairness in Educational Recommender Systems. In *SIGIR '21*. ACM, New York, NY, USA, 1808–1812. doi:10.1145/3404835.3463235
- Emilia Gómez Gutiérrez, Vicky Charisi, and Stephane Chaudron. 2021. Evaluating Recommender Systems with and for Children: Towards a Multi-Perspective Framework. In *Proceedings of the Perspectives on the Evaluation of Recommender Systems Workshop 2021*, Vol. 2955. CEUR Workshop Proceedings. <https://ceur-ws.org/Vol-2955/paper2.pdf>
- Daniel Gooch, Matthew Barker, Lorraine Hudson, Ryan Kelly, Gerd Kortuem, Janet Van Der Linden, Marian Petre, Rebecca Brown, Anna Klis-Davies, Hannah Forbes, Jessica Mackinnon, Robbie Macpherson, and Clare Walton. 2018. Amplifying Quiet Voices: Challenges and Opportunities for Participatory Design at an Urban Scale. *ACM Transactions on Computer-Human Interaction* 25, 1 (Jan. 2018), 2:1–2:34. doi:10.1145/3139398
- Judith Gregory. 2003. Scandinavian Approaches to Participatory Design. *International Journal of Engineering Education* 19, 1 (2003), 62–74. <https://www.ijee.ie/articles/Vol19-1/IJEE1353.pdf>
- 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, New York, NY, 547–601. doi:10.1007/978-1-0716-2197-4_15
- Francisco Gutiérrez, Nyi Nyi Htun, Vero Vanden Abeele, Robin De Croon, and Katrien Verbert. 2022. Explaining Call Recommendations in Nursing Homes: A User-Centered Design Approach for Interacting with Knowledge-Based Health Decision Support Systems. In *Proceedings of the 27th International Conference on Intelligent User Interfaces (IUI '22)*. ACM, New York, NY, USA, 162–172. doi:10.1145/3490099.3511158
- Oliver L. Haimson, Dyke Gorrell, Denny L. Starks, and Zu Weinger. 2020. Designing Trans Technology: Defining Challenges and Envisioning Community-Centered Solutions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, New York, NY, USA, 1–13. doi:10.1145/3313831.3376669
- Jaron Harambam, Dimitrios Bountouridis, Mykola Makhortykh, and Joris van Hoboken. 2019. Designing for the Better by Taking Users into Account: A Qualitative Evaluation of User Control Mechanisms in (News) Recommender Systems. In *Proceedings of the 13th ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 69–77. doi:10.1145/3298689.3347014
- Jacqueline Harding, William D'Alessandro, N. G. Laskowski, and Robert Long. 2024. AI Language Models Cannot Replace Human Research Participants. *AI & SOCIETY* 39, 5 (Oct. 2024), 2603–2605. doi:10.1007/s00146-023-01725-x
- William Hill, Larry Stead, Mark Rosenstein, and George Furnas. 1995. Recommending and Evaluating Choices in a Virtual Community of Use. In *CHI '95: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 194–201. doi:10.1145/223904.223929
- Velislava Hillman and Molly Esquivel. 2022. How Surveillance Capitalism Is Spreading to Schools. *The Progressive Magazine* (June 2022). <https://progressive.org/api/content/6ac9c2dc-e8fa-11ec-a11d-12274efc5439/>
- Dietmar Jannach and Christine Bauer. 2020. Escaping the McNamara Fallacy: Toward More Impactful Recommender Systems Research. *AI Magazine* 41, 4 (Dec. 2020), 79–95. doi:10.1609/aimag.v41i4.5312
- Maria Soledad Pera Jason Hall. 2017. UBR: A Book Search – Recommender Hybrid. In *Proceedings of the Thirtieth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2017)*. AAAI. <https://aaai.org/papers/665-flairs-2017-15422/>
- Nis Johanssen and Finn Kensing. 2005. Empowerment Reconsidered. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility (CC '05)*. ACM, New York, NY, USA, 203–206. doi:10.1145/1094562.1094599
- Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guetler, Aaron Tam, Corinne Bintz, Daniella Raz, and P M Krafft. 2020. Toward Situated Interventions for Algorithmic Equity: Lessons from the Field. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA, 45–55. doi:10.1145/3351095.3372874
- Finn Kensing and Jeanette Blomberg. 1998. Participatory Design: Issues and Concerns. *Computer Supported Cooperative Work (CSCW)* 7, 3 (Sept. 1998), 167–185. doi:10.1023/A:1008689307411
- Bart P Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender Systems for Self-Actualization. In *RecSys '16*. ACM, New York, NY, USA, 11–14. doi:10.1145/2959100.2959189
- Bart P Knijnenburg, Martijn C Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell. 2012. Explaining the User Experience of Recommender Systems. *User Modeling and User-Adapted Interaction* 22, 4 (Oct. 2012), 441–504. doi:10.1007/s11257-011-9118-4
- Ron Kohavi, Diane Tang, and Ya Xu. 2020. *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press, Cambridge, United Kingdom. <https://doi.org/10.1017/9781108653985>
- Joseph Konstan and Loren Terveen. 2021. Human-Centered Recommender Systems: Origins, Advances, Challenges, and Opportunities. *AI Magazine* 42, 3 (Nov. 2021), 31–42. doi:10.1609/aimag.v42i3.18142
- Margit Kristensen, Morten Kyng, and Leysia Palen. 2006. Participatory Design in Emergency Medical Service: Designing for Future Practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 161–170. doi:10.1145/1124772.1124798

- Bogdan Kulynych, David Madras, Smitha Milli, Inioluwa Deborah Raji, Angela Zhou, and Richard Zemel. 2020. Participatory Approaches to Machine Learning. In *International Conference on Machine Learning Workshops*, Vol. 7.
- Diana Kusunoki, Aleksandra Sarcevic, Zhan Zhang, and Maria Yala. 2015. Sketching Awareness: A Participatory Study to Elicit Designs for Supporting Ad Hoc Emergency Medical Teamwork. *Computer Supported Cooperative Work (CSCW)* 24, 1 (Feb. 2015), 1–38. doi:10.1007/s10606-014-9210-5
- Ulrik Lyngs, Reuben Binns, Max Van Kleek, and Nigel Shadbolt. 2018. "So, Tell Me What Users Want, What They Really, Really Want!". In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, alt04:1–alt04:10. doi:10.1145/3170427.3188397
- Beth Maloch. 2008. Beyond Exposure: The Uses of Informational Texts in a Second Grade Classroom. *Research in the Teaching of English* 42, 3 (2008), 315–362. jstor:40171756 <https://www.jstor.org/stable/40171756>
- 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*. ACM, New York, NY, USA, 163–172. doi:10.1145/3340631.3394845
- Noemi Mauro, Liliana Ardissono, and Federica Cena. 2022. Supporting People with Autism Spectrum Disorders in the Exploration of PoIs: An Inclusive Recommender System. *Commun. ACM* 65, 2 (Jan. 2022), 101–109. doi:10.1145/3505267
- 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 & Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. ACM, 2243–2251. doi:10.1145/3269206.3272027
- Silvia Milano, Mariarosaria Taddeo, and Luciano Floridi. 2020. Recommender Systems and Their Ethical Challenges. *AI & Society* 35, 4 (Dec. 2020), 957–967. doi:10.1007/s00146-020-00950-y
- Andrew D. Miller, Sonali R. Mishra, Logan Kendall, Shefali Haldar, Ari H. Pollack, and Wanda Pratt. 2016. Partners in Care: Design Considerations for Caregivers and Patients during a Hospital Stay. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 756–769. doi:10.1145/2818048.2819983
- 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>
- Michael J Muller and Allison Druin. 2012. Participatory Design: The Third Space in Human-Computer Interaction. In *Human Computer Interaction Handbook*. CRC Press, 1125–1153. doi:10.1201/B11963-57
- Michael J. Muller and Sarah Kuhn. 1993. Participatory Design. *Commun. ACM* 36, 6 (June 1993), 24–28. doi:10.1145/153571.255960
- Michael J. Muller, Daniel M. Wildman, and Ellen A. White. 1993. Taxonomy of PD Practices: A Brief Practitioner's Guide. *Commun. ACM* 36, 6 (1993), 26. doi:10.1145/153571.255960
- 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>
- 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*. ACM, New York, NY, USA, 209–213. doi:10.1145/3240323.3240405
- Organizers of Queer in AI, Anaelia Ovalle, Arjun Subramonian, Ashwin Singh, Claas Voelcker, Danica J. Sutherland, Davide Locatelli, Eva Breznik, Filip Klubicka, Hang Yuan, Hetvi J. Huan Zhang, Jaidev Shriram, Kruno Lehman, Luca Soldaini, Maarten Sap, Marc Peter Deisenroth, Maria Leonor Pacheco, Maria Ryskina, Martin Mundt, Milind Agarwal, Nyx Mclean, Pan Xu, A Pranav, Raj Korpan, Ruchira Ray, Sarah Mathew, Sarthak Arora, St John, Tanvi Anand, Vishakha Agrawal, William Agnew, Yanan Long, Zijie J. Wang, Zeerak Talat, Avijit Ghosh, Nathaniel Dennler, Michael Noseworthy, Sharvani Jha, Emi Baylor, Aditya Joshi, Natalia Y. Bilenko, Andrew Mcnamara, Raphael Gontijo-Lopes, Alex Markham, Evyn Dong, Jackie Kay, Manu Saraswat, Nikhil Vytla, and Luke Stark. 2023. Queer in AI: A Case Study in Community-Led Participatory AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. ACM, New York, NY, USA, 1882–1895. doi:10.1145/3593013.3594134
- Kostantinos Papadamou, Antonis Papasavva, Savvas Zannettou, Jeremy Blackburn, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Michael Sirivianos. 2020. Disturbed YouTube for Kids: Characterizing and Detecting Inappropriate Videos Targeting Young Children. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. AAAI, 522–533. doi:10.1609/icwsm.v14i1.7320
- Gourab K Patro, Abhijnan Chakraborty, Ashmi Banerjee, and Niloy Ganguly. 2020. Towards Safety and Sustainability: Designing Local Recommendations for Post-Pandemic World. In *Proceedings of the Fourteenth ACM Conference on*

- Recommender Systems*. ACM, 358–367. doi:10.1145/3383313.3412251
- Maria Soledad Pera, Federica Cena, Monica Landoni, Cataldo Musto, and Alain D. Starke. 2024. Human Factors in User Modeling for Intelligent Systems. In *A Human-Centered Perspective of Intelligent Personalized Environments and Systems*, Bruce Ferwerda, Mark Graus, Panagiotis Germanakos, and Marko Tkalčić (Eds.). Springer Nature Switzerland, Cham, 3–42. doi:10.1007/978-3-031-55109-3_1
- Maria Soledad Pera and Yiu-Kai Ng. 2014a. Automating Readers’ Advisory to Make Book Recommendations for K-12 Readers. In *Proceedings of the 8th ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 9–16. doi:10.1145/2645710.2645721
- M S Pera and Y K Ng. 2014b. How Can We Help Our K-12 Teachers?: Using a Recommender to Make Personalized Book Suggestions. In *2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, Vol. 2. 335–342. doi:10.1109/WI-IAT.2014.116
- Sofie Pilemalm. 2018. Participatory Design in Emerging Civic Engagement Initiatives in the New Public Sector: Applying PD Concepts in Resource-Scarce Organizations. *ACM Transactions on Computer-Human Interaction* 25, 1 (Jan. 2018), 5:1–5:26. doi:10.1145/3152420
- Victoria Purcell-Gates, Nell K Duke, and Joseph A Martineau. 2007. Learning to Read and Write Genre-Specific Text: Roles of Authentic Experience and Explicit Teaching. *Reading Research Quarterly* 42, 1 (2007). doi:10.1598/RRQ.42.1.1
- Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 1776–1826. doi:10.1145/3531146.3533231
- Amifa Raj and Michael D Ekstrand. 2022. Measuring Fairness in Ranked Results: An Analytical and Empirical Comparison. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 726–736. doi:10.1145/3477495.3532018
- Chand Rajendra-Nicoluccis, Michael Sugarman, and Ethan Zuckerman. 2023. *The Three-Legged Stool: A Manifesto for a Smaller, Denser Internet*. Technical Report. Initiative for Digital Public Infrastructure. <https://publicinfrastructure.org/2023/03/29/the-three-legged-stool/>
- David Randall, Richard Harper, and Mark Rouncefield. 2007. *Fieldwork for Design: Theory and Practice*. Springer Science & Business Media.
- Martina A. Rau, Vincent Aleven, Nikol Rummel, and Stacie Rohrbach. 2013. Why Interactive Learning Environments Can Have It All: Resolving Design Conflicts between Competing Goals. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’13)*. ACM, New York, NY, USA, 109–118. doi:10.1145/2470654.2470670
- Afsaneh Razi, Seunghyun Kim, Ashwaq Alsoubai, Gianluca Stringhini, Thamar Solorio, Munmun De Choudhury, and Pamela J. Wisniewski. 2021. A Human-Centered Systematic Literature Review of the Computational Approaches for Online Sexual Risk Detection. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 465:1–465:38. doi:10.1145/3479609
- Elizabeth B.-N. Sanders and Pieter Jan Stappers. 2008. Co-Creation and the New Landscapes of Design. *CoDesign* 4, 1 (March 2008), 5–18. doi:10.1080/15710880701875068
- Olga C Santos, Jesus G Boticario, and Diana Pérez-Marín. 2014. Extending Web-Based Educational Systems with Personalised Support through User Centred Designed Recommendations along the e-Learning Life Cycle. *Science of Computer Programming* 88 (2014), 92–109. doi:10.1016/j.scico.2013.12.004
- Markus Schedl and Christine Bauer. 2019. Online Music Listening Culture of Kids and Adolescents: Listening Analysis and Music Recommendation Tailored to the Young. (Dec. 2019). <http://arxiv.org/abs/1912.11564>
- Roan Schellingerhout, Francesco Barile, and Nava Tintarev. 2023. A Co-Design Study for Multi-Stakeholder Job Recommender System Explanations. In *Explainable Artificial Intelligence*, Luca Longo (Ed.). Springer Nature Switzerland, Cham, 597–620. doi:10.1007/978-3-031-44067-0_30
- Nick Seaver. 2018. Captivating Algorithms: Recommender Systems as Traps. *Journal of Material Culture* 24, 4 (Dec. 2018). doi:10.1177/1359183518820366
- John S. Seberger, Marissel Llavoré, Nicholas Nye Wyant, Irina Shklovski, and Sameer Patil. 2021. Empowering Resignation: There’s an App for That. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI ’21)*. ACM, New York, NY, USA, 1–18. doi:10.1145/3411764.3445293
- Joshua M. Sharfstein. 2016. Banishing “Stakeholders”. *The Milbank Quarterly* 94, 3 (Sept. 2016), 476–479. doi:10.1111/1468-0009.12208
- Ben Shneiderman. 2022. *Human-Centered AI*. Oxford University Press, Oxford, United Kingdom.
- Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, New York, NY, USA, 2219–2228. doi:10.1145/3219819.3220088
- Mona Sloane, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2022. Participation Is Not a Design Fix for Machine Learning. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*

- (EAAAMO '22). ACM, New York, NY, USA, 1–6. doi:10.1145/3551624.3555285
- Jessie J. Smith, Anas Buhayh, Anushka Kathait, Pradeep Ragothaman, Nicholas Mattei, Robin Burke, and Amy Volda. 2023. The Many Faces of Fairness: Exploring the Institutional Logics of Multistakeholder Microlending Recommendation. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. ACM, New York, NY, USA, 1652–1663. doi:10.1145/3593013.3594106
- Jessie J. Smith, Lucia Jayne, and Robin Burke. 2022. Recommender Systems and Algorithmic Hate. In *Proceedings of the 16th ACM Conference on Recommender Systems (RecSys '22)*. ACM, New York, NY, USA, 592–597. doi:10.1145/3523227.3551480
- Jessie J. Smith, Aishwarya Satwani, Robin Burke, and Casey Fiesler. 2024. Recommend Me? Designing Fairness Metrics with Providers. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. ACM, New York, NY, USA, 2389–2399. doi:10.1145/3630106.3659044
- Jessie J. Smith, Nasim Sonbolil, Casey Fiesler, and Robin Burke. 2020. Exploring User Opinions of Fairness in Recommender Systems. In *Fair & Responsible AI Workshop @ CHI 2020*. <https://fair-ai.owlstown.net/publications/1459>
- Nasim Sonboli, Robin Burke, Michael Ekstrand, and Rishabh Mehrotra. 2022. The Multisided Complexity of Fairness in Recommender Systems. *AI Magazine* 43, 2 (June 2022), 164–176. doi:10.1002/aaai.12054
- Nasim Sonboli, Jessie J. Smith, Florencia Cabral Berenfus, Robin Burke, and Casey Fiesler. 2021. Fairness and Transparency in Recommendation: The Users' Perspective. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*. ACM, New York, NY, USA, 274–279. doi:10.1145/3450613.3456835
- Alain Starke, Martijn Willemsen, and Chris Snijders. 2017. Effective User Interface Designs to Increase Energy-Efficient Behavior in a Rasch-based Energy Recommender System. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 65–73. doi:10.1145/3109859.3109902
- Caroline Stockman and Emma Nottingham. 2022. Surveillance Capitalism in Schools: What's the Problem? *Digital Culture and Education* 14, 1 (Jan. 2022), 1–15. <https://www.digitalcultureandeducation.com/volume-14-1/stockman-2022>
- Jonathan Stray, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Chloe Bakalar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, Connie Moon Sehat, Sara Johansen, Lianne Kerlin, David Vickrey, Spandana Singh, Sanne Vrijenhoek, Amy Zhang, Mckane Andrus, Natali Helberger, Polina Proutskova, Tanushree Mitra, and Nina Vasan. 2024. Building Human Values into Recommender Systems: An Interdisciplinary Synthesis. *ACM Transactions on Recommender Systems* 2, 3 (June 2024), 20:1–57. doi:10.1145/3632297
- Özge Sürer, Robin Burke, and Edward C Malthouse. 2018. Multistakeholder Recommendation with Provider Constraints. In *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 54–62. doi:10.1145/3240323.3240350
- Lida Theodorou, Daniela Massiceti, Luisa Zintgraf, Simone Stumpf, Cecily Morrison, Edward Cutrell, Matthew Tobias Harris, and Katja Hofmann. 2021. Disability-First Dataset Creation: Lessons from Constructing a Dataset for Teachable Object Recognition with Blind and Low Vision Data Collectors. In *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21)*. ACM, New York, NY, USA, 1–12. doi:10.1145/3441852.3471225
- Nava Tintarev and Judith Masthoff. 2015. Explaining Recommendations: Design and Evaluation. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, Boston, MA, 353–382. doi:10.1007/978-1-4899-7637-6_10
- Fabiana Tomasini Giannini and Ingrid Mulder. 2022. Towards a Power-Balanced Participatory Design Process. In *Proceedings of the Participatory Design Conference 2022 - Volume 2 (PDC '22, Vol. 2)*. ACM, New York, NY, USA, 111–117. doi:10.1145/3537797.3537819
- Sabina Tomkins, Steven Isley, Ben London, and Lise Getoor. 2018. Sustainability at Scale: Towards Bridging the Intention-Behavior Gap with Sustainable Recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, 214–218. doi:10.1145/3240323.3240411
- Antonela Tommasel, Daniela Godoy, and Arkaitz Zubiaga. 2021. Second Workshop on Online Misinformation- and Harm-Aware Recommender Systems: Preface. In *Proceedings of the Second Workshop on Online Misinformation- and Harm-Aware Recommender Systems (CEUR Workshop Proceedings, Vol. 3012)*, Antonela Tommasel, Daniela Godoy, and Arkaitz Zubiaga (Eds.). CEUR, Virtual Event, Amsterdam, i–iv. <https://ceur-ws.org/Vol-3012/#OHARS2021-preface>
- Thi Ngoc Trang Tran, Alexander Felfernig, Christoph Trattner, and Andreas Holzinger. 2021. Recommender Systems in the Healthcare Domain: State-of-the-Art and Research Issues. *Journal of Intelligent Information Systems* 57, 1 (Aug. 2021), 171–201. doi:10.1007/s10844-020-00633-6
- Robin Ungruh, Karlijn Dinissen, Anja Volk, Maria Soledad Pera, and Hanna Hauptmann. 2024. Putting Popularity Bias Mitigation to the Test: A User-Centric Evaluation in Music Recommenders. In *Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24)*. ACM, New York, NY, USA, 169–178. doi:10.1145/3640457.3688102
- Robin Ungruh and Maria Soledad Pera. 2024. Ah, That's the Great Puzzle: On the Quest of a Holistic Understanding of the Harms of Recommender Systems on Children. In *Designing for Children's Digital Well-being: A Research, Policy and Practice Agenda (DCDW '24)*, Co-Located with ACM IDC 2024. arXiv. arXiv:2405.02050 [cs] doi:10.48550/arXiv.2405.02050

- Meysam Varasteh, Elizabeth McKinnie, Amanda Aird, Daniel Acuña, and Robin Burke. 2024. Comparative Explanations for Recommendation: Research Directions. In *IntrRS'24: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems*.
- Suresh Venkatasubramanian, Carlos Scheidegger, Sorelle Friedler, and Aaron Clauset. 2021. Fairness in Networks: Social Capital, Information Access, and Interventions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. ACM, New York, NY, USA, 4078–4079. doi:10.1145/3447548.3470821
- Torben Volkman, Markus Dresel, and Nicole Jochems. 2023. Balancing Power Relations in Participatory Design: The Importance of Initiative and External Factors. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. ACM, New York, NY, USA, 1–6. doi:10.1145/3544549.3585864
- Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Transactions on Information Systems* 41, 3 (Feb. 2023), 52:1–52:43. doi:10.1145/3547333
- Jason C. Yip, Kiley Sobel, Caroline Pitt, Kung Jin Lee, Sijin Chen, Kari Nasu, and Laura R. Pina. 2017. Examining Adult-Child Interactions in Intergenerational Participatory Design. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 5742–5754. doi:10.1145/3025453.3025787
- Eva Zangerle and Christine Bauer. 2022. Evaluating Recommender Systems: Survey and Framework. *Comput. Surveys* 55, 8 (Dec. 2022), 170:1–170:38. doi:10.1145/3556536
- Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2022. Fairness in Ranking, Part I: Score-Based Ranking. *Comput. Surveys* (April 2022). doi:10.1145/3533379
- Douglas Zytco, Pamela J. Wisniewski, Shion Guha, Eric P. S. Baumer, and Min Kyung Lee. 2022. Participatory Design of AI Systems: Opportunities and Challenges across Diverse Users, Relationships, and Application Domains. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. ACM, New York, NY, USA, 1–4. doi:10.1145/3491101.3516506