

MIT Open Access Articles

*Building Human Values into Recommender Systems: An
Interdisciplinary Synthesis*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Stray, Jonathan, Halevy, Alon, Assar, Parisa, Hadfield-Menell, Dylan, Boutilier, Craig et al. "Building Human Values into Recommender Systems: An Interdisciplinary Synthesis." ACM Transactions on Recommender Systems.

Published Version: <http://dx.doi.org/10.1145/3632297>

Publisher: ACM

Permanent Link: <https://hdl.handle.net/1721.1/153135>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Building Human Values into Recommender Systems: An Interdisciplinary Synthesis

Building Human Values into Recommender Systems

JONATHAN STRAY

UC Berkeley Center for Human-compatible AI, USA

ALON HALEVY

Meta AI, USA

PARISA ASSAR

Adobe Inc. USA

DYLAN HADFIELD-MENELL

Department of Electrical Engineering and Computer Science, MIT, USA

CRAIG BOUTILIER

Google Research, USA

AMAR ASHAR

Spotify Inc., USA

CHLOE BAKALAR

Meta Inc., USA

LEX BEATTIE

Spotify Inc., USA

MICHAEL EKSTRAND

Department of Information Science, Drexel University, USA

CLAIRE LEIBOWICZ

Partnership on AI, USA

CONNIE MOON

SEHAT Hacks/Hackers, USA

SARA JOHANSEN

Department of Psychiatry and Behavioral Sciences, Stanford University, USA

LIANNE KERLIN

BBC, UK

DAVID VICKREY

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored.

Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2770-6699/2023/1-ART1 \$15.00

<http://dx.doi.org/10.1145/3632297>

Meta Inc., USA

SPANDANA SINGH

Meta Inc., USA

SANNE VRIJENHOEK

Institute for Information Law, University of Amsterdam, Netherlands

AMY ZHANG

Allen School of Computer Science & Engineering, University of Washington, USA

MCKANE ANDRUS

Partnership on AI, USA

NATALI HELBERGER

Institute for Information Law, University of Amsterdam, Netherlands

POLINA PROUTSKOVA

BBC, UK

TANUSHREE MITRA

Information, School, University of Washington, USA

NINA VASAN

Department of Psychiatry and Behavioral Sciences, Stanford University, USA

Recommender systems are the algorithms which select, filter, and personalize content across many of the world's largest platforms and apps. As such, their positive and negative effects on individuals and on societies have been extensively theorized and studied. Our overarching question is how to ensure that recommender systems enact the values of the individuals and societies that they serve. Addressing this question in a principled fashion requires technical knowledge of recommender design and operation, and also critically depends on insights from diverse fields including social science, ethics, economics, psychology, policy and law. This paper is a multidisciplinary effort to synthesize theory and practice from different perspectives, with the goal of providing a shared language, articulating current design approaches, and identifying open problems. We collect a set of values that seem most relevant to recommender systems operating across different domains, then examine them from the perspectives of current industry practice, measurement, product design, and policy approaches. Important open problems include multi-stakeholder processes for defining values and resolving trade-offs, better values-driven measurements, recommender controls that people use, non-behavioral algorithmic feedback, optimization for long-term outcomes, causal inference of recommender effects, academic-industry research collaborations, and interdisciplinary policy-making.

CCS CONCEPTS •Information systems~Information retrieval~Retrieval tasks and goals~Recommender systems
•Human-centered computing~Collaborative and social computing~Collaborative and social computing theory, concepts and paradigms •Human-centered computing~Human computer interaction (HCI)~HCI design and

evaluation methods •Social and professional topics~Computing / technology policy~Government technology policy~Governmental regulations

Additional Keywords and Phrases: recommender systems, value-sensitive design, technology policy

1 INTRODUCTION

Recommender systems are the algorithms which select, filter, and personalize content across social media [195], news aggregators [79], music and video streaming services [155,389], online shopping [210], online ad targeting [387], and other systems. As such, their positive and negative effects on individuals and societies have been extensively theorized and studied. In the context of social media recommendation there has been mixed evidence regarding both positive and negative effects on adolescent well-being [266], polarization [6,16] and news consumption [118]. In news recommender systems, diversity of opinion and content sourcing is a major concern [33,150]. Recommender systems used to promote job openings may be discriminatory if they do not consider the balance of distribution across legally protected user attributes [196]. Product recommendations used for online shopping could shift large-scale behavioral patterns with significant economic, environmental, or social effects [137]. Even systems designed purely for entertainment, such as film and music streaming services, must consider the fair allocation of attention to the artists who create content [222]. One overarching question across all of these contexts is, *how can we make recommender systems enact the values of the individuals and societies that they serve?*

This paper is an interdisciplinary, multi-stakeholder effort to define a common language and review the current state of the art for addressing problems related to designing and operating recommender systems in support of a wide array of human values. We proceed by proposing a set of relevant values based on a synthesis of previous work, and refined through a cross-sector expert workshop. Based on this set of values, we then review current industrial practice and emerging design techniques for building value-aligned recommender systems. This is a wide-ranging exercise in value-sensitive design [122], which we undertake by synthesizing knowledge from a variety of perspectives including computer science, ethics, economics, psychology, sociology, journalism, philosophy, and law. Although any inquiry into values brings up deep theoretical questions, our orientation is fundamentally practical: we want to know what can be done today, or in the near future, by those who build and operate contemporary large scale recommenders.

We use the term “recommender systems” to focus on the core problem of personalized content selection across many domains. Recommender systems often operate without an explicit user query, though the user may also ask for more tailored recommendations (e.g., “politics podcast”). This contrasts with search functionality which requires an explicit query and where results tend to be much less personalized [74,190,199]. Social media is a major application of recommender systems, but we note that the two are not synonymous, and the effects of social media depend on many other design choices including content moderation and other ways of finding content. We focus primarily on recommenders for social media, news content, and entertainment streaming services. However, many of the issues and approaches we highlight are broadly applicable to other important categories of recommender systems, including online shopping, targeted advertising, recruitment, healthcare, and education.

There are several widely used frames for discussing the normative implications of AI systems, most of which apply to our narrower context of recommender systems. “Alignment” is concerned with ensuring that AI systems enact the intentions of its designers and of its users despite the impossibility of specifying the correct action in all possible cases [141] and provides us with the language of “value alignment.” “Fairness” or “bias” is primarily concerned with the distribution of benefits and harms between people or groups [64,99]. “Integrity” refers to identifying and moderating content that violates platform policies for a variety of reasons, including obscenity, copyright, criminal activity, and misinformation [142]. “Well-being” is an umbrella term for a wide variety of sociological measures used across sectors including government, public health, and research, which are starting to be applied to AI systems [291]. There is a body of work on the “ethics” of recommender systems [226]. The potential effects of AI systems have also been analyzed within a “human rights” framework [89]. There is substantial overlap between these categories, each of which encompasses a range of more specific concerns such as misinformation [114], polarization [321] or addiction [7].

Rather than trying to reconcile these diverse frameworks, we take all of these to be concerned with *values*. We draw on the field of value-sensitive design [122] to define values as “what a person or group of people consider important in life” [40:1]. Building real systems requires both deep technical practice and grounding in the realities of diverse human lives, including an understanding of the actual effects of deployed recommenders. Further, values and our understanding of them are constantly evolving, so the whole exercise depends on moral and philosophical reflection. Therefore, building human values into recommender systems requires a mix of conceptual, empirical, and technical work [121].

Our focus is on existing or near-future techniques and current industrial practice, because we want to ensure that this work is applicable to the developers of large commercial recommender systems. This paper contributes to previous work by answering three main research questions:

- What human values are most relevant to the design and operation of recommender systems? We answer this by reviewing previous compilations, then gathering structured feedback from a specific set of stakeholders: experts in academia, civil society, and industry.
- What organizational processes might be used to encode values into production recommender systems? We answer this by proposing an idealized version of current industrial practice, drawing on public reports and the authors’ own experience at several platforms.
- What specific software designs might support the values identified above? We answer this with a synthesis of existing and emerging approaches to item ranking algorithms and recommender user interfaces.

The paper is structured as follows. Section 2 reports our work identifying values and summarizes the state of research and practice around each, including the empirical effects of existing recommender systems. Section 3 summarizes current industry practice, including the design of contemporary recommenders and a widely used framework for implementing values in commercial recommender systems. Section 4 takes up the challenges of measuring values in practice, including issues around operationalization and data sources. Section 5 synthesizes previous published work in recommender design that is relevant to implementing these values. Section 6 discusses policy

approaches that might incentivize the implementation of these values. Section 7 concludes by listing major open problems.

2 VALUES FOR RECOMMENDER SYSTEMS

Our first research question is, what values are applicable to recommender systems? We did not aim to create a comprehensive list of all relevant values (which is probably not possible). Instead, our goal was to ground the many discussions happening in industry, academia, and policy in a reasonably broad set of values so progress can be made.

We began by searching for previous compilations of values, issues, and risks relevant to either recommenders in particular or AI systems in general. We found four sources that were primarily compilations: a survey of AI ethics policy documents [116], a large user research project undertaken by the BBC [183], a compilation of AI-relevant well-being metrics [161], and an IEEE process for ethical system design [162]. We extracted and combined these lists to produce an initial set of values.

One of the questions that arises when assessing relevant values is “valuable to whom?” Recommender systems are fundamentally multi-stakeholder in nature, as they have effects on consumers, creators, platforms, and non-users including society in general [1]. There are also fairness considerations that apply across different subgroups [99] and cultural considerations that apply in different parts of the world. Therefore, we extended our initial list through a multi-stakeholder approach.

In an effort to include values relevant to a wide variety of stakeholders, we gathered input through the Partnership on AI (PAI), a global non-profit partnership of roughly 100 academic, civil society, industry, and media organizations working towards positive outcomes from AI for people and society. We put out a broad invitation to people from these organizations to attend a discussion of values in recommender systems, and approximately 40 people responded. At the subsequent workshop, these participants discussed our initial list of values in small groups to further develop their conception of how values apply in recommender systems.

Of the 29 people who identified themselves in a pre-workshop survey, 14 were from civil society, 8 were from industry, and 7 were from academic institutions [202]. Almost all workshop participants were based in the US. Anticipating this bias, we included a set of global “techno-moral virtues” [350:15,120] as well as the feminist ethic of care [309] and philosophical traditions in Africa that emphasize the relationships and bonds among people [243]. We asked each participant to provide further citations to any literature they considered relevant. We combined our initial list derived from the compilations above, all values discussed in transcripts of the workshop deliberations, the values mentioned in all participant-submitted references, and anything mentioned in a brief post-workshop survey. The authors then refined this long list by merging and editing entries through several rounds of reviews to produce our final list.

This approach of literature review followed by expert deliberation complements attempts to articulate the values that apply to recommender systems through user research. We note that one of our initial compilation sources [183] is a large user survey, and that our list of values includes those found in qualitative recommender user studies such as [68]. We also note that there is much overlap with the UN Sustainable Development Goals [341], in particular good health and well-being, quality education, decent work and

economic growth, industry and innovation, reduced inequalities, climate action, and peace, justice, and strong institutions.

The goal of this exercise was not to arrive at a set of global, universally applicable values, but to list some of the key values that have been identified to be particularly relevant to various stakeholders in various recommender contexts. The main challenge in devising our list was coming up with a set of definitions at an appropriate level of granularity that together cover a wide set of overlapping and often vague concepts. We aimed for a level of abstraction between overly general statements (e.g., “do good”) and specific formulations (e.g., particular metrics). Our list is presented in table 1. In Appendix A we expand this table to include citations for definitions, example indicators or metrics for each value, and example design changes which could promote that value.

Table 1: Our list of values relevant to recommender systems. See Appendix A for detailed definitions, metrics, and related design approaches.

Theme	Value
Usefulness	Usefulness
	Control
	Agency, Autonomy, Efficacy
Well-being	Well-being
	Connection
	Physical health
	Mental Health
	Community, Belonging
	Recognition, Acknowledgment
	Self-expression, authenticity
	Care, Compassion, Empathy

Theme	Value
	Self-actualization, Personal Growth
	Inspiration, Awe
	Entertainment
Legal and Human Rights	Privacy
	Freedom of Expression
	Liberty
	Fairness, Equality, Equity
	Accessibility, Inclusiveness
	Transparency & Explainability
	Accountability
Public Discourse	Accuracy (Factuality)
	Diversity
	Knowledge, Informativeness
	Civic Engagement
	Tolerance, Constructive Discourse
Safety	Safety, Security
Societal values	Progress
	Labor

Theme	Value
	Tradition, History
	Environmental Sustainability
	Duty

By nature, there are tensions between values, and those tensions lead to many of the difficulties in operationalizing them in recommender systems. For example, the value of *free expression* is in tension with the value of *safety* because if we allow users to say anything they want on social media, others may feel threatened. As another example, *privacy* can be in tension with *usefulness*. *Privacy* suggests that a platform should not try to infer whether a user might have a particular disease, even though early information and intervention might be helpful to them.

There are tradeoffs not just between values, but between different people and groups of people. For example, if we exclusively focus on the *well-being* of individuals, society may suffer [254]. Other tensions arise because recommenders must simultaneously serve several types of stakeholders including users, content producers, platforms, and non-users, as studied in the field of multi-stakeholder recommendation (Abdollahpouri et al., 2020). Values can also operate on varying time scales. Giving users entertaining content may satisfy their short term needs, but providing more informative content may have longer term benefits. Having informed users may also be a societal value, and individual values often operate on a shorter timescale than societal values.

Many of the values in our list are closer in nature to instrumental goals for an AI system, by which higher-order values are achieved. For example, the bioethics principles of respect for *autonomy*, *beneficence*, and *justice* are often considered primary, without resorting to other values for justification [56]. By contrast, other values seem to derive much of their importance because they contribute to these principles, including *privacy*, *agency*, *control*, *transparency*, *accessibility/inclusiveness*, and *accountability*.

Rather than discussing each value individually, we approach them through a number of themes that characterize many discussions on these topics: usefulness, well-being, legal and human rights, public discourse, safety, and societal values.

2.1 Usefulness

Platforms use recommender systems because they believe them to be useful to users, content creators, and themselves. The most straightforward distinction between recommendation and search is that a recommender can suggest items without an explicit query, which is valuable in a variety of contexts. For example, news items cannot be selected through user queries alone, because the user is unaware of new events. While the value of presenting a previously unknown post, article, person, movie, song or ad varies, all of these can lead to positive and novel outcomes for people. Modern recommender systems have their roots in collaborative filtering systems in the 1990s, and the need for intelligent

filtering has only increased since then as the pool of available information has exploded. We call this value “*usefulness*” to distinguish it from “utility” which has a more technical meaning from economics (discussed in section 5.3.1).

Usefulness is closely related to *control*, which we take to mean that users should be able to select which content they are seeing, the type and degree of personalization (if any), and understand the processes that determine what they see. In large part this requires that platforms provide features to support such choices (e.g., playlists on a music streaming service or topic selection on a news service) though there are also important questions about community governance [201,385]. On social networks, interfaces for describing which posts to see are particularly complex given the breadth of available content. Note that feeling in control and actually having control are different. Placebo controls may increase user satisfaction without offering any actual improvement [348] while users may not perceive even large effects of functional controls [207]. *Agency* is a similar value, but we use it to refer to control over other elements of the users’ life, not just the recommender. For example, a recommender could assist a user with education [81] or direct them to job opportunities.

There are complex tensions between *agency* and *control* and other values, as users might make choices that harm themselves or others. This grounds out in concrete ethical questions such as: if it is possible to infer that someone has an eating disorder [377], and if there is research indicating that viewing dieting videos leads to bad outcomes for such a person, is it reasonable or even obligatory to thwart their expressed intention to see such material?

Recommender systems can also be useful to both content creators and platforms. While recommenders are a commercially important technology, algorithmic content optimization does not translate directly into revenue in many contexts [169]. For example, subscription services must maximize user retention, while current recommender designs struggle with long-term outcomes. Direct optimization for revenue, or more precisely profit, is best developed in the context of online shopping [80,83,168,210]. Even for ad-driven services, content personalization and ad selection are often handled by different recommenders which optimize against different objectives [333]. Other organizations which operate recommenders may not intend to make money by doing so, such as news publishers. While not all recommender systems need to be profitable, they all must be useful to the operator by some measure.

2.2 Well-being

Well-being is fundamental to human experience, and recommender systems have the potential to affect users’ *well-being* in many ways. While the phrase has specific implications around positive subjective experience, it is also widely used in the policy community as an umbrella framework which encompasses other values [106,135,250] and has been explored as an important end goal for AI systems in general [161]. *Well-being* as a subjective experience is one of the values in our list, but many other values intersect with it strongly and are often included in well-being frameworks.

Well-being is a complex concept, and there is little consensus on how to define it [88]. Objective measures such as employment, lack of crime, and economic prosperity were historically used as proxies for *well-being*. More recently there has been increasing focus on a more holistic understanding of *well-being* based on both subjective and objective measures [86,153,191]. These subjective *well-being* measures account

for people’s cognitive and affective evaluations of their lives by asking subjects to rate how much they agree with statements like, “The conditions of my life are excellent” [87].

The values of *connection, community and belonging, recognition and acknowledgement, self expression, care, compassion, and empathy* all relate to the concept of *well-being*. Many different types of content can contribute to increased *well-being*, such as through education, motivation, or personal relationships. *Entertainment* can also contribute to *well-being*, especially as a short-term emotional experience, and hedonism is often considered a basic drive [293].

The effect of recommender systems on *well-being* has been most widely studied in the context of social media use. There is conflicting evidence for both positive and negative effects, in part because *well-being* in this context is variously defined and often represented by other factors, from self-esteem and life satisfaction to *civic engagement*, social capital, and user satisfaction. Also, most studies have not been designed to separate the effects of algorithmic content selection from other aspects of social media such as user creation and sharing.

In one study of over 2,000 college students, social media use was associated with improvement in various facets of *psychological well-being* such as overall life satisfaction, *civic engagement*, and social trust [349]. Some investigators have suggested that the amount of time spent online is less important than the quality of that time; active use may promote well-being, whereas passive use and emotional connection to use may have a negative impact [28,123,354]. The literature also points toward negative health effects related to social media use. One longitudinal study compared social media use with *mental health, physical health*, and found a decrease of 5%-8% of a standard deviation in self-reported *mental health* [296]. In another study, deactivating social media accounts for four weeks resulted in increased time in offline interactions and improved subjective *well-being* [6]. Upward social comparison has been proposed as one potential link between social media use and *mental health* disorders such as depression and anxiety [225,367].

Some recommender-based products may encourage addictive tendencies. Allcott et al. [7] found that abstaining from social media use for a time or allowing people to set future screen-time limits produced a decrease in subsequent use, suggesting that social media use may result in habit formation and self-control issues. These effects may not be due to personalized recommendations *per se*, as broadcast television, a non-personalized medium, has also been found to be addictive in this sense [120].

There is classic work on the advantages of both strong and weak social ties [13]. These benefits are less well studied in the context of recommender systems, but a few studies of social media are worth noting. Social media use has been associated with increased *social connection* and social capital in online and offline social networks, with particular benefit for users experiencing low self-esteem and low life-satisfaction [101,315,326]. Larger social networks may be associated with greater perceived social support, reduced stress, and improved *well-being* [238]. Social media use has been shown to increase intergroup contact and reduce prejudice when offline social-network diversity is low [16]. Recommenders can also contribute to *civic engagement*. A large-scale experiment showed that messaging on social media can increase voter turnout on the order of 1% [38]. While the possibility that recommendations connect people to violent or

extremist groups has been widely discussed (and we review the evidence below) the converse of this concern is the possibility of connecting people to constructive communities or social movements.

We see long term *well-being* as a major open problem for recommender systems, with a number of challenging sub-problems: defining *well-being* in contextually appropriate ways, measuring user *well-being* in production, and algorithmically optimizing for relatively long-term outcomes (months to years).

2.3 Legal and Human Rights

Some of the values in our list can be considered rights, in the sense of obligations to various parties. Previous work has explored the legal and human rights potentially impacted by recommender systems [226] and by AI in general [89]

Privacy has at least three different interpretations in the recommender context. The first sense refers to information the system knows, infers or estimates about the user. The second refers to the possibility of direct revelation of user information to other users or third-parties. There is also the possibility of indirect (often noisy) revelation of user information by the actions of the recommender, e.g. a recommendation made to one user may allow them to infer something about the items viewed by another user. This form of information revelation is the subject of a large body of research in the area of differential privacy [69,96,166]. *Privacy* trades off against other values such as *usefulness*, *transparency*, and *fairness*. The question of demographic inference which might help the user is well discussed in the context of algorithmic fairness [11] but analogous concerns arise with many types of user knowledge. Granting users *control* over whether and how they want platforms inferring and using their personal data to recommend useful content might be one way to reconcile *privacy* and ethical use of personal data.

Like *privacy*, *transparency* and *explainability* are broad concepts. *Explainability* might mean that a recommender system gives reasons why a certain piece of content is being shown to a user [386]. *Transparency* could require platform disclosure of various types of data including metrics characterizing different types of content or aggregated user activity. *Transparency* can be important for building users' trust [76] and is often discussed as a policy tool to promote understanding and *accountability*. Explanations and disclosures can come in multiple forms depending on their intended audience, such as users, researchers, governance bodies or auditors, and meaningful transparency needs to be informed by concrete individual and societal information needs [93,303]. From a technical perspective, explanations can be tricky to generate because many recommender techniques (including deep learning) do not lend themselves to easy explanations [366].

Fairness, *equity* and *equality* are closely tied to human rights (Universal Declaration of Human Rights, UN). In everyday situations where recommender systems are employed, multiple parties have some interest in the outcome [2]. Fairness has a variety of meanings in recommender systems, including considering the disparate impacts of recommendations across user classes [203]. Content providers want to be fairly treated in terms of the exposure and benefit they receive from the system [85,174]; users want to receive good quality of service, and do not want to be under-served relative to other users [100]; other stakeholders, such as the system operator, content creators, and society broadly each have ideas of what it may mean for recommendation to be "fair" [85,98]. This has deep ties to various conceptions of *equity*, in particular equity

of attention [37:201,302]. Recommenders may also create externalities that affect people who do not use the platform, such as by directing many new people to a formerly obscure place, or encouraging the consumption of products with environmental consequences [279]. The field of multi-stakeholder recommendation has emerged to tackle these dynamics [1].

Given all the parties involved through the platform or via externalities, *fairness* is a multifaceted problem where different stakeholders have different objectives and needs from the system. Often the desires of different stakeholders are in conflict; not everyone can have exactly what they want, though it might still be possible to give something worthwhile to everyone.

2.4 Public Discourse

Some of the most intense recent discussion around recommender systems has centered around how they affect public discourse. *Accuracy* of information and *diversity* of content are two prime examples.

Access to accurate and factual information supports human decision making and understanding across myriad domains, ranging from healthcare to politics to economics [154,295]. While falsehoods and misleading content have threatened truth for centuries [14,253], online user-generated media may have amplified the threats to sensemaking and decision making. For example, some evidence suggests fabricated news articles (as identified by third-party fact checkers) spread significantly faster through Twitter than genuine news articles, especially articles about politics [356]. It is not clear to what extent this is an effect of recommender algorithms, as opposed to non-algorithmic social media creation and sharing dynamics [156,236]. Similarly, COVID-19 misinfo was common on social media globally, especially early in the pandemic [8,48] and there is evidence that suggests misinformation contributed to the spread of the virus via reduced vaccination rates [208,264]. This suggests that personalized recommendations may have had a causal effect on disease spread, though this has not been directly studied.

Many authors have argued that designers of news recommenders have editorial responsibility similar to news editors [115,148,244,311]. Getting the right information to the right people at the right time is a key normative concern that goes well beyond ensuring accuracy, a value that we call “informativeness.” We can look to the tradition of public-service journalism to inform recommender design [115,311] but personalized news delivery is a new technology and specific editorial theories are still developing. One approach would be to try to deliver a news item if the user previously expressed an interest in a particular topic, if it reports on events that affect their life, or if there is an opportunity for the user to help others [319]. These sorts of ideas have yet to be effectively translated into algorithmic terms.

Diversity is a value that can be relevant to consumers, content creators, and society in general. In industrial settings *diversity* has mostly been studied because increased *diversity* typically results in greater user satisfaction and user engagement, at least up to a point [155,194,222,369]. Many recommenders use diversification algorithms for the practical task of ensuring that users are not continuously shown the same type of item in their recommendations, often implemented as a re-ranking pass [132,390]. Other parties also have an interest in diversity. A streaming music or television platform needs to ensure that the long tail of less popular artists or producers have enough exposure to make it worthwhile for them to stay on the platform [222], and increased diversity may contribute to equity of attention [37,302].

The experience of a lack of diversity in personalized content consumption has been described as a “filter bubble” or “echo chamber.” However, this language is somewhat vague and has been used to describe a wide variety of phenomena including self-selected consumption behavior, homophily in social networks, and algorithmic feedback effects [50]. We consider such possibilities more specifically throughout the rest of this paper.

Diversity has been most specifically studied in the context of news recommendations, where it might serve a variety of democratic goals including consumer choice, civic participation, pluralist tolerance, or challenging the status quo [33,148,357]. The meaning of *diversity* can differ between news organizations, depending on their editorial missions and the balancing of other values that matter to the organization (such as personal relevance, engagement, or time spent). In practice *diversity* is often measured using item-similarity metrics [194] but such formulations do not capture the complexity that the social sciences have brought to the debate about media diversity [206].

One concern is that a lack of diversity in personalized news recommendations could prevent users from being exposed to contrary perspectives. Simulations have suggested that optimization to increase user engagement could create feedback loops that drive users into narrower selections of content [57,175,177,192,284]. However, news personalization algorithms do not seem to produce a less diverse selection than human editors [231] and the news provided to different users is quite similar [139,244]. Social media users consume a more diverse range of news sources than non-users [118] but correspondingly also consume more news from partisan sources [117].

A related concern is that recommender systems might be causing large-scale polarization of ideology (issue polarization) or attitudes (affective polarization) [321,323]. A recent review found an overall positive correlation between “digital media” use and polarization [209]. The causal evidence is more mixed. In the U.S., polarization began increasing decades before social media [43] while several other developed countries have similar internet usage but do not show increasing polarization [44]. Paying people to stop using social media for several weeks produced small declines in issue polarization measures in a study of American users [6] but *increases* in a measure of ethnic polarization in Bosnia-Herzegovina [16] which was hypothesized to result from reduced inter-group contact. While there is little evidence that filter bubbles or a lack of diversity are driving polarization, other socio-technical processes can increase polarization including partisan sorting [339]. Optimizing for engagement (see section 3.1) can prioritize divisive content, and there are now several lines of evidence that this mechanism is exacerbating polarization [323].

2.5 Safety

Safety includes the idea that people should not be bullied, attacked, or dehumanized and should not be exposed to disturbing content. For social media, safety has thus far mostly been considered in the context of content moderation, and many platforms have developed their content moderation policies based on international human rights principles and frameworks and in consultation with third party experts [90,108,260]. As a notable example, a series of human rights abuses occurred in Myanmar, characterized by hate speech and disinformation against the minority Rohingya by the military on social media [51].

Recommender systems should not promote violence. Note that “polarization,” a mass hardening of political divisions, is conceptually distinct from “radicalization,” where a small number of individuals violate mainstream norms and may resort to violence [316]. There are a number of documented cases of far-right and terrorist radicalization where online recommendations were involved [26,237,280,362]. However, these reports mention many other factors including chat rooms, personal relationships, user-directed searches, and life circumstances. More systematic studies have looked for recommender feedback effects that move users toward radicalization en masse [110,200,236,275]. These studies generally show that recommenders alter content mix in the direction of engagement, but have produced poor evidence on the radicalizing potential of recommender systems because of insufficiently powerful experimental designs, as we will discuss below. In general, causal understanding of user trajectories through recommender systems remains a major challenge.

2.6 Societal Values

There are additional values on our list that have more of a societal flavor such as *progress, labor, duty, environmental sustainability, duty, and tradition and history*. These are important values that arose in our research, though the degree to which they are considered important varies considerably from one culture to another. In our workshop and the literature we reviewed these values did not come up as particularly relevant to online platforms, but that does not mean that they cannot or should not be promoted in certain contexts.

3 THE STATE OF PRACTICE

Our second research question is, how might these values be operationalized within commercial recommender systems today? Because our focus is industrial applicability, rather than attempting to design an ideal process from first principles we document and build on current practice. While the academic literature on recommenders is vast, there is far less documentation of actual commercial operations.

In Section 3.1 we describe a general architecture that captures common design patterns in commercial recommender systems. This provides necessary context for the rest of the paper. In Section 3.2 we propose an organizational process for modifying a recommender system to support a specific value, within the context of a large company. These sections draw on public material, plus the authors’ own experience at several large platforms, to triangulate the current state of practice.

3.1 How Recommenders Work

Recommender systems are often described as “black boxes” [307] but most are constructed using similar principles. This section presents a greatly simplified, but illustrative recommender design. While it doesn’t represent the details of any particular system, many real systems share its features. Our discussion also leads to a technical definition of *engagement* in terms of behavioral data. Most recommenders are built to prioritize some form of engagement at their core, though they also consider many other types of signals.

Recommendations start with a pool of content items. These items may be produced entirely by the recommender operator, as with a news organization’s recommendations; or they may be curated from multiple sources, such as a music recommender which gets content from publishers; or items may be entirely user generated and posted without prior review, as on social media. These three categories have been called

“closed,” “curated,” and “open” recommender systems [72]. Before recommendations are generated, a moderation process identifies and removes items which violate platform policies. Platform moderation is a complex process which involves many human and machine steps from policy making to enforcement to appeals [142,380] but here we are concerned only with how moderation defines which items are available for recommendation.

Individual recommendations are generated using data about item content, user attributes, and context, and result in an output stream or set of recommended items, sometimes called a *feed* or *slate*. The context may include a wide variety of features such as the video the user is currently watching, the time of day, or a search query like “politics podcast” [25,36]. User attributes may be derived from personal information the user has provided, any explicit user feedback or control settings, and any implicit feedback contained in the history of past interactions with the system.

The recommendation process usually begins with the selection of *candidate items*. Candidate generation algorithms are tuned to retrieve an overbroad sample, but are very efficient at filtering a corpus of (potentially billions of) items available for recommendation down to a small set which might be a good fit for the user and context, typically ranging in size from a few hundred to a few thousand. These candidates are then *ranked*, that is, each is assigned a *relevance score*, which typically reflects a prediction of user engagement with the candidate item. In modern systems ranking may involve dozens or hundreds of “signals” which summarize aspects of the content, the user, the context, and how all of these interact. The top scoring items are then selected as the user’s recommendations. Many systems then re-rank the remaining items, this time comparing them to each other, rather than evaluating each individually, to achieve goals such as diversity of item topic or source [131,390].

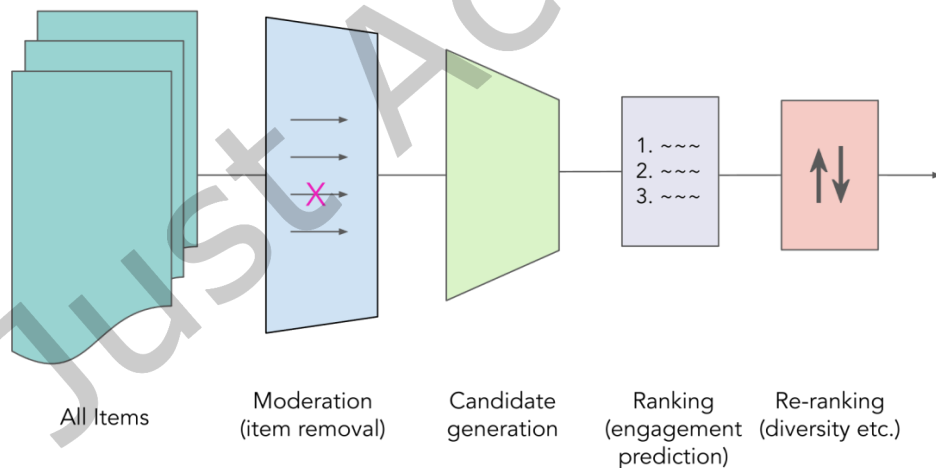


Figure 1: An illustration of how many modern recommenders work, adapted from (Google, 2020; Lada et al., 2021). The global set of items is first moderated to remove content that violates platform policies. The remaining steps happen when a user is served recommendations: candidate generation selects a wide set of items that could be relevant, ranking scores each one, and re-ranking ensures feed-level properties like diversity.

Item ranking constitutes the core of personalization. The final item score is typically a weighted combination or some other function of (i) the predicted probability of a number of different types of user responses [109,227,310,389] plus (ii) a wide variety of scoring signals that range from source credibility [261] to playlist diversity [222] to whether an item tends to be inspiring [163]. Ranking items by the probability of desired or targeted user reactions (e.g., sharing, dwell time) is informally known as optimizing for *engagement*. This may or may not optimize for *value* to various stakeholders, which is why many non-engagement signals are also used.

The word “engagement” has been used across many fields including media and technology to suggest that users are repeatedly interacting with a product, as evidenced by a wide variety of metrics. Here we propose a more specific definition, compatible with recommender design practice. We take engagement to be *a set of user behaviors, generated in the normal course of interaction with the platform, which are thought to correlate with value to the user, the platform, or other stakeholders*. This definition builds on previous work [167,379]. It is multi-stakeholder in nature, and reflects the fact that engagement signals are chosen to be indicators of value, but aren’t going to be fully aligned with specific values in all cases. It also suggests there are some signals of value that can only be derived from non-ordinary or non-behavioral data, as we discuss below.

3.2 Implementing a Product Change

In this section, we answer our second research question: what organizational processes might be used to encode values into production recommender systems? Our goal is to provide practical advice for people who are working on large recommender systems within industry. Because of this we focus on documenting the ways in which values are actually built into recommenders today. Unfortunately, there is little public documentation of the actual processes by which values are engineered into large recommender systems, so we draw heavily on the authors’ collective experience working on a variety of platform-scale recommenders.

This is a different approach than, for example, attempting to derive an ideal organizational process from first principles. Nor does it explain what stakeholders outside of the organizations building recommenders should do (but see our discussion of policy approaches in section 6). However, this approach has the advantage of codifying a process which we know is actually possible within industry today.

We start by assuming that a choice has been made to prioritize a particular value, or to adjust the status-quo tradeoff in favor of that value. We do not give a process for deciding how to prioritize values, taking into consideration the needs of all stakeholders – this is a major open problem which we articulate in section 7. Here, we focus on what can be done within an organization once the choice has been made.

We illustrate the process by which a specific human value might be incorporated into a recommender system’s design using the example value of *diversity*. Consider the development or refinement of a news recommendation platform in which the designers are concerned with users having the opportunity to develop an awareness and understanding of multiple political

perspectives—this has been called the *deliberative perspective* on recommender diversity [150]. One way that a recommender system might contribute to this value is by increasing the diversity of recommended items. This strategy rests on two key assumptions: there is currently a lack of diversity in the items users see, and showing them more diverse items will increase tolerance. Both of these assumptions are complex and available evidence is mixed [249,321]. Nonetheless, this is a non-trivial example of how metric-driven recommender engineering might proceed.

Once the decision to increase content diversity has been made, the implementation might proceed using the steps outlined below.

3.2.1 Research unintended consequences

Generally, such an effort begins with the study of potential mechanisms to incorporate this value. In the case of diversity, designers may research the history of attempts to create inter-group tolerance through diversity of exposure, including the ways in which this has failed [258,263]. User studies may be used to test the exposure diversity hypothesis with users of the platform, assessing the impact of seeing news items from diverse perspectives on people’s knowledge and attitudes. In particular, it will be important to test for backfire effects where people reject diverse information [18,95] or other unintended consequences. There must also be a reasonable expectation that the changes to the recommendation system won’t adversely affect other values that the system embodies. In concrete terms, this often means that any proposed change should not decrease other metrics (see below) more than a specified amount.

3.2.2 Develop a metric

As discussed in Section 4, metrics are central to any attempt to build values into recommender systems. A variety of diversity metrics have been proposed, both for news content [75,357] and for recommender systems more generally [60,194]. Choosing or developing a metric requires settling on a definition of diversity suitable for that particular system. It may be that multiple metrics are needed to capture different aspects of diversity. While product teams typically choose metrics, there has been experimentation involving external stakeholders in order to increase the legitimacy of such decisions [33,201,320].

Here we focus on a conception of “*productive*” diversity, where people disagree in ways that are ultimately constructive [321]. Given this concept, developing a method to measure the diversity of articles on the platform may be broken into several phases, such as:

- Developing a description of diversity for use by human raters (e.g., “Does this set of articles include constructive contributions from multiple perspectives?”).
- Creating a training data set of positive and negative examples using human-rater evaluations.
- Using this training data to develop a heuristic or machine-learned model that can predict whether a list of (recommended) articles adequately reflects “productive diversity.”

3.2.3 Consider different product designs

Before implementing a specific product change, different product-based approaches to increasing diversity will be explored. For example, the designer could change the user interface, perhaps by showing related articles from other viewpoints below each item [325], or change the ranking algorithm to try to nudge people to consume more articles that represent this type of diversity [216]. For the sake of example, we assume below that the latter change has been selected.

3.2.4 Implement the product change

Implementing this specific product refinement requires incorporating the diversity prediction model into the item ranking procedure. Whereas many current recommender systems score each item independently, diversity is a property of sets of items. One challenge is that scoring entire lists is both more complicated and more expensive than scoring individual items [160,369] which may necessitate the development of more efficient algorithms for ranking such item sets.

3.2.5 Evaluate and advocate

Once implemented, the new product feature will typically be tested using offline data [55,128] followed by A/B tests with a small group of users. If these tests show positive results, the new diversity prediction model will be deployed in production, and monitored to see whether the target diversity metric improves (this may involve gradual ramping up of the deployment to larger numbers of users, holdback tests, etc.). There may also be side effects—for example, while increasing diversity often increases engagement [194,369], this is not always the case. Suppose that as a by-product of the model deployment there is a drop in engagement, declines in other values-relevant outcomes such as quality or safety metrics, or some cost imposed on other stakeholders (e.g., content producers). In practice, this requires negotiation among internal stakeholders to decide if the increase in diversity is significant enough to justify a drop in other metrics.

3.2.6 Establish guardrails to prevent reversion

Once deployed, product teams often establish a review process so that subsequent product changes, within the originating team or elsewhere, don't indirectly revert the diversity improvement. This might include numeric "guardrails" that specify the maximum allowable decrease in diversity induced by other product or algorithmic changes.

3.2.7 Monitor outputs and outcomes

In practice, product teams will continually monitor the diversity of recommended items to detect operational failures (e.g., bugs or system failures), or (induced or exogenous) changes in the user distribution and user behavior. A survey which asks users or human raters to detect "productive" diversity may also be regularly employed to detect model drift and produce updated training data. Finally, determining whether the ultimate goal is being achieved, or is still worthy of being achieved, requires ongoing evaluation. This could involve survey methods to assess metrics such as affective polarization, and ethnographic research to understand what it means for users to be encouraged to be "more tolerant" in this way.

4 VALUE MEASUREMENT

In order for recommender systems to incorporate human values, there need to be methods for measuring how well a system is adhering to, promoting, or facilitating these values over time. This section describes how to go from a value to a metric, the issues involved in designing good metrics for values, and the data that is available on which to build such metrics. [Figure 2](#) shows the steps in operationalizing a value, including the process of selecting metrics, the interactions a user has, and the ultimate outcomes.

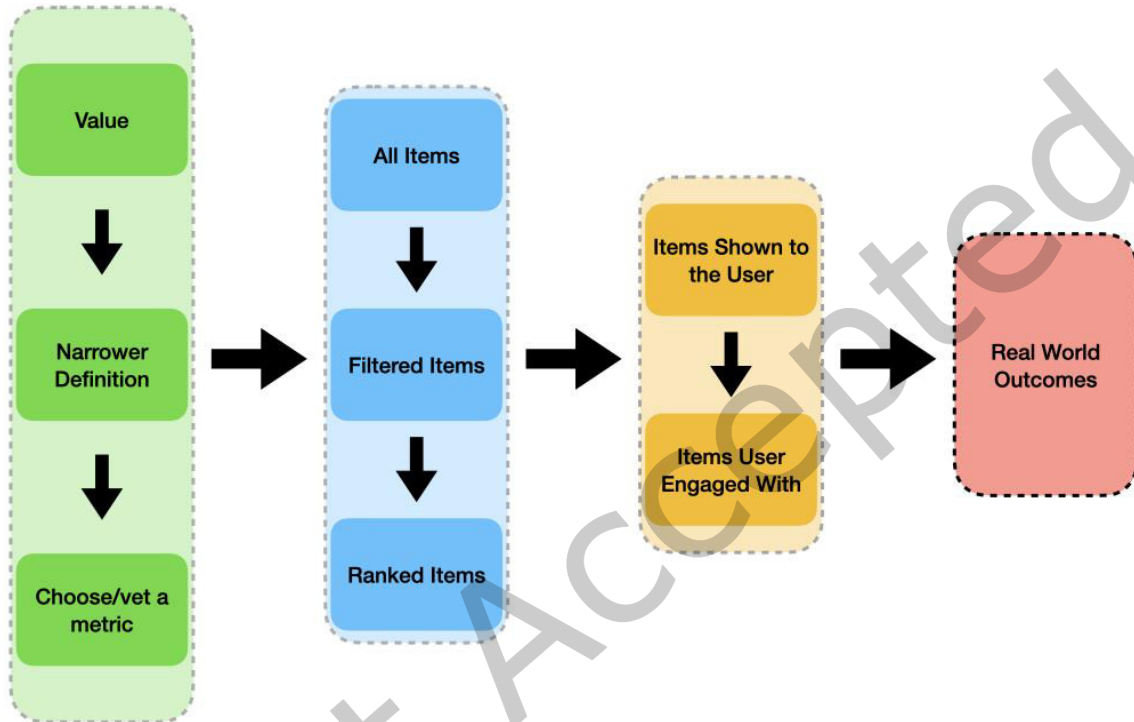


Figure 2: An illustration of the relationship between values, metrics, recommenders, items, and outcomes. The selection of a metric leads to a selection of items, some of which the user will engage with. Over time, the items a user engages with shape outcomes of interest.

4.1 From Value to Metric

Defining a metric is part of the process of specifying what, exactly, a particular human value means in the context of a real system. This process is called *operationalizing* a value. The people involved in defining metrics have considerable influence over the ultimate function of a recommender, which is why multi-stakeholder involvement in recommender metric selection may be important [201,320].

To illustrate the gap between a value and its measurement, consider the value of “safety,” and in particular, protecting users from hate speech. A precise definition of hate speech is not only hard to articulate, but under constant debate and evolution, and every choice has some set of undesirable side effects [84,134,214,214,259]. Well-being is an even more complicated example. As discussed, well-being has many

components (e.g., physical health, health of personal relationships, having purpose in life, etc.) and can be considered in the short term (e.g. entertaining content) or the longer-term (e.g. learning useful skills or fostering relationships). Making news recommendations that serve the public interest likewise requires defining concrete metrics that reflect some assessment of that interest [115,357]. At some point, a real recommender must commit to specific operationalizations of broad concepts, with the resulting tradeoffs between competing values and stakeholders.

4.2 Characteristics of a Good Value Measurement

A good measurement has a number of desirable properties, including validity and reliability [230], fairness, and legitimacy. Quinn et al. [269] summarize the situation as:

The evaluation of any measurement is generally based on its reliability (can it be repeated?) and validity (is it right?). Embedded within the complex notion of validity are interpretation (what does it mean?) and application (does it “work?”). [269:216]

Many social science theories involve quantities that are not directly observable and hence must be inferred, often with considerable uncertainty, which makes any measurement instrument implicitly a model of some purported underlying reality [164]. Construct validity is the convergence of a successful theoretical idea with a measurement that effectively captures it [318]. Jacobs and Wallach [165] propose several other types of validity in the context of measurements within computational systems, including face validity (is the metric plausible?), content validity (does the metric capture all the relevant parts of the concept of interest?), convergent validity (does this metric agree with other accepted metrics?) and discriminant validity (does this capture something different than other metrics?). Reliability is typically evaluated in terms of agreement between multiple measurements (test-retest reliability), between different human judges (inter-rater reliability) and, for surveys, between different ways of asking a question (inter-item reliability) [164].

It is rare that a particular measurement fully captures what we mean by some value in a particular context, meaning that most metrics are in fact proxies for what we actually care about. Even a good metric will change meaning when it is known to be used to make consequential decisions, e.g. student test scores must be interpreted differently when they are used to decide academic progression, because instructors will begin “teaching to the test.” This effect is sometimes known as *Goodhart’s law*, but there are a variety of different causal structures which can produce feedback processes that widen the distance between a metric and the underlying value [213,235]. In the technical community this has been most discussed in the context of the general problem of algorithmic optimization and the difficulty of objective specification [78,141]. Hence, as part of being precise about the definition of a human value, it is important to identify gaps in what is measured and monitor them over time. Human values — that is, what is considered important — also tend to change over time. Qualitative user research plays a critical ongoing role in designing and evaluating metrics.

Because modeling assumptions are required to connect a measurement to the underlying value it purports to reflect, measurement itself has fairness implications. The signals from measurements can vary across demographics, even after controlling for differences in user intent and task [221]. For example, if older users read more slowly than younger users, then a metric based on dwell time will be an over-optimistic measurement for older users regardless of their level of satisfaction. Thus, interpreting metrics at face value may systematically disadvantage and misrepresent certain demographics and user groups. Organizations

deploying recommendation systems often rely on *internal auditing* methods as a way to measure how the overall system performs across different demographics and other user attributes [113,272]. External audits may also be required by regulation, as discussed below.

It is not enough to have an accurate and fair measurement of a human value. Stakeholders in the recommender system need to be able to accept the process and outcome of the measurement as legitimate. This means that two metrics that are operationally identical (e.g. they both correlate strongly with a desired outcome) may not be interchangeable. Transparency around how the measurement is carried out may help build trust in a metric. For example, some platforms periodically release public transparency reports which include various metrics [308]. Another way to increase legitimacy is to establish accountability regarding a measurement, e.g., through independent, external audits of the measurement [329]. More ambitiously, a metric could be created or chosen through a participatory process [320]. For instance, the measurement could aggregate the opinions of a panel of users, as in the “digital juries” [112] and “citizens assembly” concepts for making platform decisions [255]. In one case, representatives of various stakeholders participated in the construction of a recommender that matched supermarket excess food donations to volunteer drivers and local food banks, using an elicitation process to define a ranking function [201].

4.3 Data Sources for Measuring Values

Our news recommender diversity example involved a complex, multi-step process for defining an algorithmic measure, ultimately drawing on human labeling to define the value of diversity. Broadly speaking, there are three main data sources for value measurements on a recommender platform.

The first source of data is the behavioral signals that users generate during normal usage, e.g. articles a user clicks on, songs a user plays, comments, emoji reactions, re-sharing of content, ads a user clicks on, purchases made, time spent on the platform and on specific items, and so on. These sorts of implicit signals of value are often called “engagement,” but we note that some behavior signals are explicitly designed to give feedback to the algorithm, such as swipe left/right or thumbs up/down. Implicit and explicit behavioral feedback both provide distinct and useful information to a recommender system [172].

The second source for signals of value is answers to survey questions that are posed to a fraction of the user base, typically a very small fraction. Surveys can ask very targeted questions, e.g., Facebook has asked users whether individual posts were “worth their time” [107] while YouTube uses user satisfaction surveys that ask users to evaluate previous recommendations [131,389]. Survey results can be used to monitor real-world outcomes, evaluate A/B tests, and ultimately recommend items that a user is predicted to respond positively to when asked about their experience.

The third source for value measurements is human annotation. This data is often produced by paid raters, though human ratings also come from users flagging or reporting items. Vast amounts of human training data are routinely used to create models for identifying particular kinds of positive or negative content for the purpose of content moderation and ranking. Platforms that feature professionally-made content can also ask creators to provide metadata for a variety of ranking purposes. For instance, the Swedish public broadcaster asks editors to rate each story in terms of importance, public service value, and lifespan [217]. While survey data is limited by how often users can be asked to fill out a form, annotation data is costly to create and

therefore available in limited quantities, and some types of labeling work may contribute to mental health issues for raters [15]. Human annotation can also be noisy, inconsistent and biased depending on how rater pools are selected, and have limited, unbalanced coverage, while the meaning or usage of labels can change over time. Research in the area of human computation [198] attempts to address such issues [111].

In contrast to surveys and annotation, the benefit of behavioral signals is that they are plentiful—many orders of magnitude more data is available. However, the behavior of users on the platform typically correlates with but does not perfectly capture any particular type of value, and is only a proxy for what different stakeholders actually care about. Moreover, behavioral signals have been shown to be sensitive to a variety of factors, such as the user and their demographics [58,146], the context in which the user is interacting with the system, and the recency of interactions. For example dwell time, a behavioral signal that has been used as a proxy for user satisfaction [379], can vary significantly depending on whether the item is the first one clicked in a list of results or not [39]. As discussed above, behavior does not represent underlying preferences for a variety of reasons including cognitive biases, information asymmetry, coercion, lowered expectations, and so on [9,32,54]. Additionally, the collection of vast amounts of user data poses significant privacy concerns.

One important approach that tries to combine the benefits of behavioral and survey signals is to learn a model that predicts survey results given user behavior, that is, it tries to predict how a particular user would answer a survey question if shown a particular set of items. This method extrapolates limited survey data to all users, and is common in industry [131,320] but has not been much discussed. We view such prediction of survey results as an emerging method for aligning recommender systems with human values, though it must be understood that predictions are just proxies and must be continually monitored for divergence from ground-truth survey responses.

In principle, surveys can elicit complex judgements of value. In addition to the challenges of complicated, multi-component values (e.g., well-being) the context and wording of a survey can significantly affect the results [317,364]. Social desirability bias, the tendency of respondents to answer in ways that others would view favorably, adds a further complication [124,239]. Another issue arises when users from different demographic groups tend to answer questions differently [102]. Since online surveys often involve casual participants, “seriousness” checks and data denoising can be very important [17].

5 DESIGN APPROACHES

We are now ready to discuss the major approaches that recommender system developers and researchers have used to promote the values discussed above, and in turn, to identify some of the challenges and open problems which require the development of new technologies. The structure of this section reflects the issues presented above. Section 5.1 discusses the core of the recommender system, namely the techniques used to select and order items shown to the user, and long-term user paths through these items. Section 5.2 discusses the affordances and controls made available to users and other stakeholders through UI and UX design. Finally, Section 5.3 considers fairness and multi-stakeholder perspectives, and describes techniques used by system designers to help optimize tradeoffs within real systems.

5.1 Item Selection

At the heart of any recommendation system are the models and algorithms used to generate one or more recommended items. Generally, the core models are trained to predict a user’s behavioral response to a candidate recommendation (e.g., watch, like, reply, purchase) using properties of the user, the candidate items, and the context. Items are then scored by combining predicted outcomes in some way [227,389]. The nature of the properties used to make predictions, the predicted responses themselves, and the way these predictions are combined into a final score all play a role in the values that a recommender embodies.

5.1.1 Item ranking signals

The properties of an item being considered for recommendation—whether a social media post, news article or musical track—play a key role in determining whether it is of interest to the user, and therefore whether showing it to a user can serve one of their values.

The “topic” of an item is an important signal. A variety of techniques have been used for text analysis in recommender systems, including latent semantic indexing in Google News [79], Latent Dirichlet allocation [173] and transformers. Image and video analysis are also used for topic assessment [82], as is audio analysis [381]. Modern recommender topic taxonomies can encompass tens of thousands, or even millions, of distinct categories and sub-categories. One of the challenges is that these classifiers are typically accurate for the popular elements of the taxonomy of topics, but often do not perform well for posts on less popular topics.

Topics do not directly correspond to specific user needs or to values. A post about “football” can be about organizing a football viewing party (possibly contributing to the value of connection) or just reporting the latest scores or player injuries (the value of knowledge). To support the value of empathy and care it would be useful to identify posts where the poster could benefit from support of their network (e.g., after losing a loved one, needing advice on a certain matter, or announcing an important event in their life). An open challenge is inferring the way a particular user will *relate* to an item, rather than analyzing properties of the item alone. There is nascent work on predicting the experience a user may have when viewing content, e.g., whether they are entertained, angry, or inspired [67,163]. These sorts of inferences could also enable better targeted or more persuasive advertising, which raises policy concerns [304].

Content analysis is also heavily used by social media platforms in order to determine whether a post violates community policies or is of low quality in some way, and should therefore be removed or demoted. The predominant method for determining such violations uses ML models trained on content labels provided by paid raters. These models increasingly use multimodal techniques, simultaneously considering text, speech/audio, images and video. This is often crucial to understanding the intent of a post, e.g., the caption “love the way you smell today” means something different when superimposed over the image of a rose vs. a skunk [184].

Increasingly, items are evaluated not just in terms of their content but their context, including properties of the poster and audience, previous user behavior, reactions from other users, and so on [142]. Incorporating embeddings of sequences of user interactions has been shown to increase accuracy in misinformation and

hate speech classification on Facebook [246] while “toxic” conversations on Twitter are better predicted by including network structure features [290].

These models are generally trained on labels or annotations generated by human raters. This process itself is potentially subject to error and particular forms of bias, both in the instructions given to evaluators and their assessment biases [287]. Also, many of these categories are both complex and politically contested [46,214] and raters often disagree, which complicates the evaluation of model accuracy [134]. The promotion or demotion of particular topics or types of language, and the errors and biases in this process, have major implications for freedom of expression, which implies a deep connection between these technical methods and broader policy considerations [91,181,380].

5.1.2 User Trajectories

We use “trajectory” to refer both to the sequence of items a user saw over time and the reactions those items evoked. A user’s past trajectory often provides information that can be used to make better predictions about their future preferences and behaviors. At the same time, we want to ensure that the user’s future trajectory supports the values they care about. Moreover, trajectories are the basic unit for studying some problems users experience on platforms, as they are central to discussions about potential long term effects on, say, well-being and polarization. For example, the “filter bubble” critique is essentially a statement about the typical course of user trajectories.

The majority of recommenders deployed in practice are “myopic” in the sense that they make predictions of a user’s immediate response to the next slate of items presented, and rank items based on these predictions. However, many recommenders use past sequences of user and system behavior in sophisticated ways.

Advances in deep learning have made possible the practical deployment of *sequence models*, such as recurrent neural networks or transformers. For example, Beutel et al. [36] describe a recurrent (specifically, an LSTM) model deployed at YouTube. Recurrent approaches explicitly model a user’s “latent” or “hidden” state, that is, they include variables which represent aspects of the user’s situation or psychological state which are unobserved but have effects on what the user wants to see next. This state might encode aspects such as user satisfaction, frustration, or current topic focus, but interpretation of the hidden state embodied in such models is challenging and is tightly coupled to the engagement metrics being predicted and optimized. Inferring user state from behavior is an important challenge and a key step toward better supporting many values. For example, a recommender could detect a user’s dissatisfaction with a certain type of content, or with too little topic diversity in the recommendation stream, or even that someone was developing an eating disorder [377].

Current ranking techniques do not offer the means to (directly) optimize a user’s *future* trajectory. A promising direction is the use of *reinforcement learning (RL)* for optimizing such futures non-myopically [5,297,330]. In particular, the use of RL allows the system to consider the impact on the user of not just immediate recommendations, but of the entire sequence of recommendations it might make over an extended period, and plan that sequence accordingly.

In our news diversity example, the redesigned recommender considered diversity within each slate of recommended items independently. An RL-based recommender would be able to consider the diversity of items over days or weeks as well. In an educational setting—where a user’s understanding of a topic might best be served by an individualized sequence of content— an RL recommender resulted in faster learning and more course completions than a baseline linear sequence [24]. As such, RL offers considerable promise as a technology for individual or user-level value alignment. To date it has been used largely to optimize user engagement over the long term, but with suitable metrics and objective criteria, it could play a vital role in better aligning recommendation trajectories with user well-being by adaptively planning recommendation sequences.

As with sequence models, advances in deep RL methods have made it more practical to deploy RL in practical recommender systems [66,126,160]. That said, a number of challenges remain. First, since RL relies on sequential interaction data with real users, such models are often trained offline using data generated by past recommendations, which may induce bias because the user was interacting with a different model [66]. The second challenge involves choosing from a large action space ranging from hundreds to millions depending on the recommendation task [94,160]. A related problem is that item recommendations are often made jointly in slates or scrolling feeds, where the interactions or interference among the visible items makes interpreting and optimizing for user responses challenging [160]. Finally, adopting RL for true value-alignment requires sophisticated models of various aspects of user latent state (e.g., preferences, satisfaction, awareness/knowledge levels, fatigue/boredom, etc.) and their dynamics. These psychological and situational states are challenging to uncover from observable user-response behavior, and may require planning over extremely long “event horizons” as user adaptation to recommender changes may take 3-6 months to fully materialize [57,229].

Some work has considered trajectories that might be unhealthy or harmful. One problem concerns minimizing the likelihood that trajectories will end up with users viewing large amounts of questionable content, such as videos that might promote unhealthy behavior, false statements about COVID vaccines [61], or other content that may not explicitly violate the platform’s policies. This might occur not because users explicitly ask for such material, but because they end up in places (e.g., groups) where this material is common. In addition to downranking such content, another solution is to avoid recommending low quality groups or content sources to users.

Another strand of work concerns items that are acceptable or appropriate when considered in isolation, but could be harmful if consumed too much or by certain vulnerable people. For example, there may be nothing wrong with a diet video but perhaps someone with an eating disorder should not be presented with an endless stream of such videos; or it may be unhealthy to watch exclusively violent movies. Singh et al. [301] address this possibility by ranking user trajectories by proportion of “unhealthy” content, then using the mean proportion of this content in the top α th percentile of user trajectories as a regularization metric. They demonstrate a safe RL approach that improves both worst-case and average-case outcomes, in terms of the fraction of “unhealthy” items recommended to any one user.

The most complex concern about trajectories is the possibility that recommender systems might change user preferences in a manner that increases engagement but harms some other value. This is commonly associated with the idea of “manipulation,” meaning unwanted changes in attitude or behavior (even if unintentional).

This sort of optimization-driven shift has been frequently suggested as a mechanism driving filter bubble, echo chamber, or polarization effects, though the empirical evidence is mixed [50,139,249,321,392]. Such models posit a feedback loop where users choose particular items (as in selective exposure effects [265]) and the recommender responds to that engagement by narrowing its output to those topics, which in turn shifts user preferences further in that direction. This effect may be a particular concern for RL systems, which may learn how to make their users more predictable so as to maximize engagement [283]. A polarizing preference shift effect has been demonstrated in multiple simulations with different specifications [57,105,175,177,192,284] which suggests that it could be a robust effect.

A range of work has attempted to determine the causal effect of recommender systems on content consumption trajectories. This is particularly important if those trajectories are correlated with offline outcomes such as violence [129]. There is replicated evidence that strongly moralizing expression spreads faster than other content on social networks [47] and such moralizing seems to precede offline violence [232]. A number of researchers interested in categories such as “far right,” “conspiracy,” or “radical” have studied the network structure of recommendations between YouTube channels [110,200,275]. While this showed that more extreme channels often link to each other, these studies do not analyze user trajectories because they were conducted without any personalization. A different approach is to program bots to selectively engage with certain topics. This has generally shown that engaging with some type of content increases its frequency [360,368] but this design models users as unchanging, so it does not provide evidence on persuasive effects. A user-tracking study of “far right” content on YouTube from 2016 to 2019 found that consumption there matched broader patterns across the web, including consumption from sources without recommender systems [156]. Separating consumption due to recommenders and consumption due to users’ seeking behavior is difficult without on-platform experiments. Using a set of Twitter users randomly assigned to receive a baseline chronological feed, Huszár et al. [159] found that Twitter’s home timeline recommender reduced the consumption of more politically extreme material.

In general, inference of the causal effects of recommenders on user preferences and outcome is a major open problem. The feedback effects between algorithms and societies are at the cutting edge of social science research, in need of both interdisciplinary and cross-sector collaboration.

5.2 Controls and Feedback

Meaningfully supporting human values requires effective communication between the user and the system, beyond the standard implicit signals (clicks, shares, dwell time, etc.). We place communication affordances on a continuous spectrum between explicit and implicit. We use controls to refer to features where the user can explicitly change certain settings, and feedback to refer to situations where the user is more passive and their preferences are elicited in some way, such as by answering a survey question, providing a like/dislike, etc.

There are a variety of documented benefits to providing users with more control over their experience. Users have reported being more satisfied with their recommendations when given greater control [176] and better controls have led to increased engagement [171]. Interestingly, even the appearance of controls can achieve some of these effects. In [348], users reported higher satisfaction with the system even though the controls had random effects. Conversely, Loecherbach [207] gave users control over the diversity of items in their feed but there was no correlation between actual and perceived diversity. In general, greater recommender control usage does not always translate into better control [359].

Existing systems offer a variety of ways for a user to customize their experience, and many more control and feedback schemes have been proposed by researchers [147]. Some involve simple direct feedback on individual items, e.g., thumbs up / thumbs down, or item-level ratings [42,388]. Others involve evaluating the relevant importance of pairs of items [41,65]. Even when controls are provided, many users do not know that they exist or what they do [157,306], find them challenging to use, or simply don't see the value in engaging with them. As a result, most users do not use recommender controls and a "passive" user experience remains the default [176]. Aside from better educating users on controls, a promising approach is to design feedback mechanisms that serve multiple purposes: such as sending a social signal (which is why the user would use it) and simultaneously providing direct feedback that pertains to some value. Some examples of this exist today, such as the "Insightful" emoji on LinkedIn, the "Care" emoji on Facebook, and the proposed Respect button [324].

One of the issues that limits the use of controls is the difficulty in creating a direct, understandable link between the input a user provides and the resulting change in the system's behavior. In some cases this link is quite direct (e.g., "don't show me any sports-related content"), whereas in others the outcome of a control/feedback action is less obvious (e.g., "show me more videos like this video I just watched"). There are many points along this spectrum. For example, instructing the system to "show less about football" will be largely predictable in terms of the content it affects (though what about political statements/protests by athletes?) but users may still not understand the expected magnitude of change. To complicate things further, there is often some ambiguity about whether the change the control offers is transient (e.g., for the current session) or for the longer term. One proposed design is to highlight items that will be added and removed from the user's feed when a control is changed [292].

Organizing content into "channels" can help users to better customize their experience. Some systems include implicit ways of specifying what sort of items are to be included, such as music recommenders which can continue a human-generated playlist [382]. Algorithmic channels could also be designed around particular goals (e.g., learning to play guitar) or needs (e.g., getting support from friends, participating in lively conversations). The primary technical challenge here is building relevance measures that capture the dimensions that users care about.

In the future, communication between the user and the system will take different forms. Conversational recommender systems are an emerging area that offer a potentially more natural, engaging, and usable interface for people to express their preferences, particularly in light of recent advances in speech, natural language and dialog technologies [3]. Applications range from integrations with the approaches listed above

to entire reimaginings of the control process [71,170]. A related idea is recommender personae, where the recommender is associated with a particular personality (e.g., explorer, diplomat, expert) to set a particular context for the conversation [144]. Conversational recommenders face many of the same challenges as other control systems, but moreso: they must interpret the metaphorical, imprecise, or subjective language a user may use to convey their needs or topic preferences.

5.2.1 Meaningful explainability

The field of explainable recommender systems has grown into a vibrant research area [386] in part because users respond positively to having good justifications for why certain items are suggested. Tintarev and Masthoff [337] argue that a good explanation can increase transparency, scrutability, trust, effectiveness, efficiency, persuasiveness and satisfaction in the system. Each of these terms expresses somewhat different goals: transparency shows how the system works and can be instrumental in accountability; scrutability allows the users to tell the system it is wrong; trust increases confidence in the system; and satisfaction increases users' sense of derived benefit; effectiveness helps the user employ the system toward their own aims; persuasiveness convinces the user to change their beliefs or behavior; efficiency can increase the value of the system. Some of these values can be contradictory and may not be achieved at the same time. For example, an explanation that increases transparency of the system does not necessarily increase trust if the explanation is not understandable or reveals undesirable behavior. Therefore, what constitutes a "good" explanation very much depends on the goal, and the field is rapidly developing [145,241,248].

With the rise of increasingly complex machine learning models, it has also become increasingly difficult to give intuitive and understandable explanations of why a user received a specific recommendation. Explanations may be shown to the user in different forms (e.g., text, visuals, highlighting relevant features) and may either attempt to explain the workings of the recommendation model itself, or may be the result of training a separate model that generates post-hoc explanations from model inputs and outputs [386]. There are also recommender algorithms specifically designed to be explainable [376]. The system of Balog et al. [19], for instance, operates based on set intersections, e.g., "recommended for you because you don't like science fiction movies except those about space exploration." This explainable-by-design approach avoids the challenges of interpreting learned models [282], but generating understandable explanations from deep learning models is an active area of research that may yet prove fruitful [185,189,328].

5.3 Making Tradeoffs

In general, there are numerous tradeoffs involved when incorporating human values into recommender systems, and a variety of techniques to evaluate these tradeoffs and make choices. There are at least three categories of tradeoffs: 1) tradeoffs between the benefits and potential harms to different people, 2) tradeoffs between different types of stakeholders, such as content creators versus content consumers, and 3) tradeoffs between values, resulting in different (but not obviously worse or better) measurable outcomes induced by various recommendation strategies.

Ultimately there must be some notion of a fair or optimal tradeoff, and techniques for making "good" tradeoffs in this sense. We hope that these tradeoffs are informed by the expressed opinions of users and other stakeholders, so we first discuss the theory of social choice, which studies how to combine the

preferences of many people. We then discuss techniques for achieving various notions of fairness, such as between different types of users or between different stakeholder categories. Finally, we discuss tools designed to optimize tradeoffs when faced with the practical necessity of tuning a recommender’s ranking function.

5.3.1 Tradeoffs in theory: Social Choice

People express their values in their everyday use of recommender systems. Indeed, many recommender controls are designed specifically for this purpose, everything from upvoting and swiping left/right to reporting violating content. This creates the problem of aggregating preferences, and negotiating between competing desires of individuals, groups, and stakeholders. The framework of *social choice*, originally developed in economics [125,294], provides a foundational tool for addressing tradeoffs at all of these levels.

Abstractly, this framework assumes that each stakeholder has a *utility function* over a set of possible outcomes for them. This is motivated by the idea that someone could say which of several situations they’d prefer, that is, that each person has *preferences*. Note that this utility function can be purely “local” (a user may care only about whether they get good recommendations) or it can involve societal values (a user may care that other users also like what they like, or that vendors are treated fairly). A *social welfare function* is a voting process or a way of “adding up” or combining stakeholder preferences to produce a single number, a “societal utility” or social welfare. The aim is then to adopt a recommender policy which maximizes the expected social welfare.

This formulation can encompass virtually any criteria that express preferences over short or long-term, individual or group outcomes. There are relatively direct mathematical expressions for penalizing addictive behaviors, group-level diversity of consumed content, fairness across individuals, and so on. Conversely, any collaborative recommender system aggregates the explicit or implicit preferences of users in some mechanical way, as when signals like upvoting and watch time are combined across many users to decide how items are ranked for a different user. Social choice theory is a key bridge between the normative and the algorithmic, useful for analysis and design.

Although social choice theory is foundational, there are two major difficulties to applying it in practice. First, determining what an individual’s “preferences” actually are is quite challenging both in theory and in practice [334]. There is a long history of formal elicitation methods such as asking users to repeatedly say which of two options they prefer, or asking them to play various kinds of economic games [53,65], but formal preference elicitation can be involved and is quite demanding of the participants. Further, someone might be uninformed, coerced, addicted, have lowered expectations [9] or want something that isn’t available [119,277]. In addition, attempts to elicit preferences can lead to strategic behavior where people misrepresent their beliefs to try to induce more favorable recommender outcomes, e.g. a content creator usually has an incentive to argue that their content is relevant to as many users as possible [29,30]. While many different kinds of feedback can provide crucial signals of what people value, behavior cannot be naively interpreted as true preferences.

Second, there is no entirely bottom-up or value-neutral method of ethics. Simply specifying the outcomes over which stakeholder preferences are defined is inextricably tied to the values being considered [32]. For example, there is the choice of what will be voted upon. There is also the question of whose preferences count in what situations, e.g. community administrators may have special voting privileges, and it may be important to encode “rights” that cannot be infringed by the preferences of others. Hence, social choice approaches cannot excuse system designers from making consequential normative choices [27]. There is more to democracy than voting systems.

5.3.2 Fairness and Tradeoffs among Stakeholders

Many of the values in our list involve making trade offs between different stakeholders (such as users vs. content providers) or among members of a stakeholder group (such as different subgroups of content providers which compete for attention). Correspondingly, there are a wide variety of notions of “fairness,” which often (but not always) are framed as some sort of trade off. The extent to which these tradeoffs are inherent is an open question in the research literature, because there are cases where it is possible to improve performance for one user subgroup without decreasing performance for other users. There are multiple challenges in this area, including defining fairness, measuring it in practice, and designing algorithms for efficient recommendation. See Ekstrand et al. [98] for an overview.

Several major categories of fairness have been proposed in the context of recommender systems, roughly corresponding to the interests of different stakeholders. “C fairness” considers how well individual information consumers are served, “P fairness” is concerned with the distribution of attention between items or providers of content, while “CP fairness” considers both simultaneously, as in a rental property recommender designed to protect the rights of both minority renters and minority landlords [52,361].

Recommender systems are mostly evaluated based on average performance across all users, but different user subgroups, such as age or gender groups, might be served with differing performance or error rates. Subgroup performance disparities can happen for a variety of reasons, including differences in group size or activity that affect the amount of training data available [100,204]. There is a large body of work on mitigating group-level unfairness in classifier models, some of which can be adapted to recommender systems. For example, [34,35] use pairwise comparisons of the ranking of different items to generalize the well-known “equality of opportunity” and “equality of odds” measures, showing that it is possible to equalize prediction error rates between user groups on a large commercial platform. However, algorithmic approaches that aim to equalize effectiveness disparities between user groups may make inappropriate tradeoffs: increasing recommendation utility for one user does not necessarily require decreasing it for other users, so it is not clear that allowing a solution that may decrease utility for well-served users is appropriate, as opposed to other approaches such as feature prioritization in the engineering process (Ekstrand et al., 2021 §5.4).

Items and their producers, on the other hand, necessarily compete for user attention. This leads to the concept of “exposure fairness” which may be formulated in a variety of ways [302]. Even if predicted “relevance” scores correctly measure the value of an item to an individual user, slightly less relevant items may get disproportionately less attention simply because they appear farther down, an effect known as

“position bias.” Several algorithms have been proposed to ensure item attention is proportional to item value on average, either across a group of items or across multiple slates of recommendations [37,85,361]. Such algorithms can be considered a correction to a type of error or “technical bias” in the fair ranking typology of Zehlike *et al.* [383]. More normative definitions of item fairness are often desirable. For example, Spotify strives to give exposure to less popular artists to counteract the “superstar economics” of cultural production [222], while a “demographic parity” conception of fairness may be appropriate when qualified candidates from different groups (say, men and women) should be shown to prospective employers at the same rate [127]. A wide variety of fairness metrics concerning the exposure of items, groups of items, or producers have been proposed, though many of these are closely related [193,270,288,384].

Where it is possible to produce reasonable quantitative estimates of utility to different stakeholders, *multi-objective optimization (MOO)* can be used to balance multiple conflicting stakeholder utility and fairness objectives. One approach is to ensure that the recommender is Pareto efficient, meaning that there should be no way to modify a slate of recommendations to make it more fair without reducing utility, or to increase utility without reducing fairness [374]. Mladenov *et al.* [228] go beyond this by proposing a recommendation method that maximizes user social welfare (total user utility) by allowing small sacrifices in utility from well-served users to drive large gains for less well-served users. Reinforcement learning has also been applied to multi-sided fairness, through contextual multi-armed bandits which simultaneously optimize stakeholder utility and fairness objectives [222,223,375]. Recently, several researchers have taken a game-theoretic approach to the study of recommender systems. Ben-Porat and Tennenholtz [30] and Ben-Porat *et al.* [29] develop approaches that account for the strategic behavior of content providers while aiming to maximize user engagement. While all these methods hold promise for value alignment in complex ecosystems, they have not seen practical deployment to date.

5.3.3 Optimizing Tradeoffs

Because there is no purely “bottom up” way of making tradeoffs, system designers must ultimately choose some set of overall objectives or “ground truth” signals to serve as overall measures of value. Increasingly, AI tools are used to help make tradeoffs over the complex design space of recommender parameters, particularly the relative “weighting” of the signals that feed into item ranking functions.

Milli *et al.* [227] determined the relative value of different user actions including viewing a Tweet, sharing it, liking it, etc. by connecting these interactions to the sparse use of the “see less often” control in a causal Bayesian network. This network represents dependencies such as the fact that a user has to view a Tweet before they can share it. By taking “see less often” as a ground truth signal of negative value it was possible to infer the value of all of the other, more common interactions. This idea generalizes to more complex methods for finding multiple weights that optimize multiple metrics.

Bayesian optimization [130,274] can be used to find the weights of a ranking function that maximize relevant (perhaps long-term) metrics, and automatically run data-gathering experiments to improve those predictions [140]. This approach requires that the designer be able to assess the overall “utility” of any vector of performance metrics, which itself induces various tradeoffs. For example, is a product that has a greater number of items viewed but less total time spent and lower reported satisfaction better than the opposite? The tools of interactive optimization and utility elicitation [42,388] could

play an important role here, though these approaches have not yet found widespread use in practical recommender design.

6 POLICY APPROACHES

In the previous section we have discussed design approaches to recommender systems and human values, that is, potential product changes. This section considers policy-making, which can be a powerful lever for change. Policy-making is informed by all of the perspectives articulated above, including ethical, procedural, measurement, and technical issues. Chosen policies impose constraints on how to build a recommender system and introduce additional technical challenges.

By policy-making we mean external governance, from government, regulators, or external bodies with appropriate authority. All large platforms have internal policies as well, particularly around which types of items are eligible for recommendation, but we focus on external governance as an important interface between recommender system operators and the rest of society. Because recommender systems are used in so many different types of products, we do not offer specific policy recommendations. Instead, our goal is to discuss the major categories of policies that have been proposed, and especially to understand how these policies could be translated into the terms of metrics and algorithms. At the current time there is a large gap in terminology and understanding between the recommender technology and policy communities, which we seek to highlight and begin to address. There are also policy-relevant technology gaps: the capability to do what policy-makers ask may not yet exist, as in the “right to explanation” provisions of the GDPR [358].

We consider policy approaches that are relevant to recommender systems specifically, as opposed to social media or online platforms generally (neither of which necessarily involve recommender systems). We do not discuss content moderation policy approaches here, but direct readers to reviews such as [91,181,380].

6.1 Risk Management vs. Value Sensitive Design

One proposed policy approach would require recommender system operators to evaluate the potential risk of harm from operating their systems. This is the approach taken by the EU Digital Services Act, which requires “very large online platforms” (currently defined as those with more than 45 million users in the EU) to perform yearly assessments for three kinds of risks: the dissemination of illegal content, negative effects on fundamental rights, and “manipulation” with effects on “public health, minors, civic discourse, or actual or foreseeable effects related to electoral processes and public security.” [104] Any harms found must be mitigated through various means including “adapting content moderation or recommender systems.” The proposed Digital Services Oversight and Safety Act of 2022 (DSOSA) in the U.S. takes a similar approach, requiring platforms of certain sizes and scope to conduct assessments and mitigate any risks, including by “adapting the content moderation or recommender systems (including policies and enforcement) of the provider.

This touches on several of the values in our table, but is relatively narrow in two ways. First, a risk-based framework is concerned only with potential harms. Second, the type of risk mitigation envisioned by the DSA and the DSOSA generally happens after a system is already built. Another approach is to require consideration of important values during the design phase, as with “privacy by design” provisions [92]. Extending this to more general values, one German law requires platforms to meet certain content diversity obligations [151].

The challenge for policy-makers or regulators is to be both precise and general about how harms are to be assessed and values are to be enacted in recommender systems. In principle this could involve monitoring certain metrics, as is already done in environmental regulation and media monopoly policy. Such regulation would face all of the challenges of choosing metrics discussed above, and certainly no one set of metrics will be appropriate for all recommender systems. Even if useful metrics for harm or good can be found, there is the difficult question of what constitutes an “acceptable” value [244].

6.2 Accountability

Policy approaches to the issue of *accountability* include provisions around transparency and evaluation mechanisms such as audits. Both are instrumental in supporting other values such as agency, control, and accountability.

Transparency has been a major focal point of legislative efforts in the United States and the European Union. Over the past two years, lawmakers in the United States have introduced numerous bills that seek to compel internet platforms to provide more transparency around how they develop and deploy algorithmic systems for content curation purposes [234,303]. The key challenge from a policy perspective is in defining the goal of transparency efforts, who they are meant to serve, and what, exactly, should be disclosed. Transparency may be intended to serve users, lawmakers, researchers, journalists, etc. [93]. There are also privacy, security, and intellectual property issues that complicate disclosure [251:4,298,307]. Suggestions for what to publish have included the recommender source code, data on users (e.g. demographic data and their interactions with the system), options for users to modify recommendation “parameters”, data used for training models, key metrics, and the rationale behind product features and ranking changes [93,322,355]. Each of these has limitations.

Production recommender code is extremely large, difficult for outsiders and non-technical individuals to understand, and may not be particularly revealing without reference to the content and user data it operates on. For example, user interactions over time can be used to understand people’s trajectories through the system, as discussed above. However, user data is difficult to share because of privacy concerns. This leads to the idea of sharing aggregated data, and many platforms already do so in the context of content moderation and targeted advertising [304,308]. Then the key policy question becomes which metrics should be disclosed and how they are defined. While recommender operators will have key insights into what is relevant to measure, relying on them to select what is shared may pose a conflict of interest. Note that even aggregated user data can be used to re-identify individuals so it may need to be protected with techniques such as differential privacy [96,240].

A related suggestion is that recommender operators should share the policies used to guide ranking and recommendation efforts, including what types of low quality or harmful content a platform downranks. The policy community has pushed recommender operators to disclose the rationale behind ongoing product changes more generally, including changes to recommender algorithms and parameters [305,322]. This leads to the idea of a “change log” or “proceedings” that details what the operators were trying to do with each change (e.g., increase news quality), what data they had in front of them (e.g., fraction of items from each news source rated false by fact checkers), and what change they made (e.g., downrank certain sources by a certain amount). This is especially important as metrics or algorithms alone will not tell the full story: the

motivation and context of a change are relevant to values. Thus far, a handful of civil society organizations have pushed platforms to adopt change logs for recommender system-related policies [305].

Algorithmic auditing of AI/ML based systems has recently gained increased attention, and has mostly focussed on fairness and discrimination concerns in decision-making systems or predictive models [353]. Although fairness remains a concern in recommender systems, in principle recommenders might be audited for any of the values discussed in this paper. A recent review of 62 academic algorithmic audits identified eight audits of recommender systems [20], of which seven were concerned with “distortion” effects such as echo chambers or lack of source diversity. Several of these audits looked for, but did not find, echo chamber or filter bubble effects on Google News, the Facebook News Feed, and Apple News [20:16]. They did find that a small number of news sources make up a large percentage of the results in Google News and Apple News.

While companies can and do perform internal algorithm audits of various types [70,272] regulation could require external audits of recommender systems to mitigate these concerns [307]. Regulation could also direct audits to evaluate specific biases or harms to users or consumers [49]. The Digital Services Act requires certain platforms to undertake yearly audits to ensure they are meeting their risk assessment and mitigation obligations [104]. Similarly the Algorithmic Fairness Act in the United States would require covered entities to conduct and retain a five-year audit trail which includes information on how an algorithmic system was developed, trained, and tested [303].

The personalized nature of recommender systems complicates auditing. Consider the problem of auditing user trajectories through recommender systems. Ribeiro et al. [276] collected data on non-personalized recommendations across YouTube channels and then evaluated trajectories using random walks, while the Wall Street Journal used 100 TikTok bots pre-configured to watch only videos on particular topics [360]. Neither of these methods model real user behavior. Conversely, Hosseinmardi et al. [156] used panel web browsing data collected by Nielsen to evaluate consumption of YouTube videos across the political spectrum, and The Markup’s Citizen Browser project asks users to install a browser extension that reports what they see on Facebook and YouTube [332]. These observational studies have the advantage of involving real users, but the methods used so far cannot produce robust causal inferences about recommender effects. For example, it is currently not clear if social media contributes to depression or if depressed people spend more time on social media, or both. If these methodological issues cannot be solved, even extensive platform data sharing may not be sufficient to answer questions of interest. In that case, on-platform experiments may be the only reliable approach to auditing the effects of recommender systems, which would require extensive collaboration between recommender system operators and external researchers.

The nature of platform access remains to be defined. Industry players assert they face constraints when participating in third party audits, including concerns around privacy, security, intellectual property, competitiveness, and cost [251:4,298,307]. Many of these concerns are also risks to a third-party auditor, who must protect shared personal data yet typically do not possess the security resources of a platform. Additionally, many third-party auditing entities do not have the necessary technical skill and resources to audit recommender systems at scale [307].

6.3 Translating Between Policy and Technology

There has long been miscommunication between the builders and regulators of technology. At the present moment, many countries are drafting or passing laws that regulate recommender systems of various kinds, especially social media, news recommenders, or targeted advertising. Unfortunately, much of the policy discussion taking place uses terms that do not map easily to recommender technical affordances [303].

For example, Article 27 of the Digital Services Act requires recommender-based platforms to disclose “the main parameters used in their recommender systems” as well as any available controls to “modify or influence those main parameters.” [104] Unfortunately, it’s not clear what a “main parameter” is [149]. While this probably doesn’t refer to the millions or billions of learned neural network weights, real recommender systems involve hundreds of major interacting components which are continually configured and tuned in complex ways. The vast majority of such internally configurable settings will not be understandable or useful to users, auditors or regulators. A better formulation might arise from considering the design of recommender controls, as discussed above.

The same provision also stipulates that recommenders offer “at least one option which is not based on profiling.” Profiling in this context is defined in the GDPR and includes “personal preferences” and “interests” [103]. Thus, exclusion of profiling in this sense may exclude even Twitter’s classic reverse chronological timeline, which cannot operate without the user indicating who to follow. An alternative approach would be to require an option that personalizes based on explicit user controls only, as opposed to implicit signals such as clicks or watch time.

“Amplification” is another word which appears in several proposed laws yet is difficult to translate into operational terms. Modern platforms can provide huge distribution to an item in a short period of time via information cascades, which result from a combination of user sharing and algorithmic recommendation [23]. While “amplification” is a reasonable name for this phenomenon, many definitions of “amplified” collapse to “shown” when parsed carefully, e.g. “the promotion of certain types of extreme content at the expense of more moderate viewpoints” [368]. The selection of any type of content for display reduces the promotion of all other types because of the zero-sum nature of selection, so amplification would mean any display of “extreme” content.

Other definitions are more specific, and compare the distribution of an item to some (often implicit) baseline. One approach is to define amplification as the prevalence of some type of content in user feeds relative to the prevalence of that type among all available content. This definition can be useful in some contexts, though it is unclear why raw prevalence should be a presumptively neutral baseline. Indeed, this formulation leads to the perverse outcome that spamming content may reduce amplification (through an increase in the denominator) even though it may increase distribution. For some systems a chronological baseline may make sense. This is plausible for Twitter because it was designed around a reverse-chronological feed, and several studies compare the algorithmic and chronological options on Twitter [21,159]. However, there is no similarly natural baseline for systems such as YouTube, Google News, Netflix, Spotify, or Amazon where a chronological feed makes little sense. Further, purely chronological feeds can suffer from problems that make them unattractive baselines [178,182] including recency bias and

the prevalence of low quality content like spam. Because of these conceptual and practical issues, “amplification” is not a well-defined measure for most recommender systems [335]. In the U.S., legislation targeting amplification in any of these senses is also likely to face 1st Amendment challenges [182].

Examples like these highlight the importance of finding new ways of bridging the knowledge divide between policy makers, specialized expertise and independent research. Possible ways forward include educational programs for policy makers and engineers alike, the embedding of technology expertise (such as TechCongress which recruits technologists and embeds them in Capitol Hill offices to inform tech-focused legislation), regulatory sandboxes which allow controlled experimentation, a strong commitment to evidence-based policy making, and other initiatives to make technical expertise more easily accessible.

7 OPEN PROBLEMS

Based on this review of values in recommender systems from a variety of perspectives, we propose the following list of open problems. Each of these is both consequential and challenging, and would benefit greatly from future work.

Multi-stakeholder processes for prioritizing values. There is no widely-agreed process for including the many stakeholders of recommender systems in consequential decisions, including which values are prioritized and tracked, how they are measured, and how trade-offs are adjudicated. This work provides important prerequisites: a list of values relevant to recommender systems (section 2), techniques to measure adherence to values (section 4), and a list of indicators for specific values (Appendix A). We have not provided a process for eliciting input from stakeholders to prioritize values and resolve tradeoffs. Broad methods such as participatory design [300] provide general approaches, but more specific methods will be needed. Approaches such as multi-stakeholder metric construction [201,320] and user juries [112] may provide a way forward, but are not yet well developed in the context of recommender systems.

Better measurements. Many values are not easy to operationalize. Collections of AI-relevant metrics such as the IEEE 7010 standard [291] provide useful compendiums, but measures developed in social science and policy may not apply directly to particular recommender contexts such as news recommendation or targeted advertising, and will need to be refined.

Controls that people want to use. The controls offered to users so far have been remarkably sparse and slow to develop, despite the obvious practical, policy, and ethical advantages of better control. In part this may be a problem of under-investment, but there is also a fundamental unsolved problem with recommender controls: most are never used by more than a few percent of users [176]. Better control designs might improve this situation, such as immediate feedback on how changing a setting changes which items are recommended [292]. Still, it is likely that most users will never adjust default settings, so a better strategy may be to attempt to design controls that elicit user feedback in the normal course of use, e.g. voting systems or a “respect” button [324]. It will be important to differentiate between giving users a feeling of agency and giving them effective controls [207,348] as both are necessary.

Non-behavioral feedback. At present essentially all feedback signals to recommender systems are behavioral, e.g. engagement. There are a variety of well-known problems with inferring preferences from

behavior, including self-control or addiction issues, uninformed users, and the contextual meaning of choices [9,334]. In particular, many of the values discussed above involve outcomes that are unlikely to be identifiable from on-platform behavior alone, e.g. well-being. One promising solution is simply to ask users about their experience by repurposing survey instruments that have been developed in the social science and policy communities. Only a small fraction of users can be surveyed, but the resulting data stream can be used to build and continuously validate models that predict survey responses, which can then be used as algorithmic objectives. This is already done in industry [131,320] but there is essentially no public research on this emerging technique.

Long-term outcomes. Most recommender algorithms today are myopic in the sense of optimizing only for immediate responses. Longer-term outcomes are typically managed by product teams who monitor richer feedback and make algorithmic adjustments. While today these teams are typically optimizing for purchases, subscriptions or user retention, recommender systems could also be managed on values-relevant outcomes. However, human management will always optimize against aggregated outcomes, while algorithmic optimization can be personalized and therefore has the potential to better serve subgroups and individuals. Emerging reinforcement learning methods may make this possible [5,229] but this may also require cheap and accurate individual-level proxy metrics for the outcomes of interest.

Causal inference of human outcomes. Recommenders may have significant effects on people and society, yet determining what those effects are remains extremely challenging due to sampling and confounding effects. Individual case studies can be instructive but are difficult to generalize to platform scale, while even large observational studies struggle to answer questions like “does social media cause depression or do depressed people use social media more?” In many cases it will not be possible to say what a platform’s effect on an individual has been because there is no counterfactual, but it may be possible to measure group-level effects through on-platform experiments. The long term effects of design decisions – and the resulting outcomes for users interacting with the system – are particularly difficult to study because so many other things are happening in a user’s life.

Industry-academic research collaborations. While a variety of external algorithmic auditing methods have been developed [286] many of the most important questions can only be answered in an ecologically valid setting using private platform data or on-platform experiments [138]. Unfortunately, it is not easy for external researchers and platforms to work together due to concerns around access, privacy, security, research integrity, intellectual property, competitiveness, and cost [251:4,298,303]. External collaborations would benefit from the development of technical methods to enhance the privacy and security of shared data, and from legal or policy approaches that set the terms of engagement so as to protect all relevant interests.

Interdisciplinary Policy-making. There is a substantial gap between the policy community and the technical community (including between scholars of both). Not only is there an enormous amount of specialized knowledge unknown to each side, but they use very different terms to understand and describe the problems of recommenders. There are probably also fundamental differences in values and politics between disciplines that will have to be resolved, which is complicated by the fact that recommender policy

touches fundamental rights such as privacy and freedom of expression. It is clear that effective policy-making must be collaborative and interdisciplinary, even if it is not yet clear how to achieve this.

8 CONCLUSION

Recommender systems are a profound technology that will continue to touch many aspects of individual and social life. At their best, they serve important interests of multiple parties. Consumers might want accurate information, good music, or useful products, producers might depend on recommenders to help them find their audience or customers, and platforms need to capture some of this value to continue operating. Yet recommenders can also cause a wide variety of harmful effects, or they may miss opportunities to do good. Building human values into recommender systems raises complex and consequential challenges, including philosophical, regulatory, political, sociological, and technical issues.

This paper contributes to this interdisciplinary conversation and guides further research in several ways. We have identified some values that are relevant to recommender systems, and discussed the main issues surrounding each. We have described current industrial practice including a sketch of modern recommender design, and an illustrative process for shifting a value in a production system. We discussed the challenges of measuring adherence to a value, including the difficulty of “operationalizing” or translating a value into metrics, and the types of data sources that might supply useful information on values. We feel it is too early to attempt a general theory of the value-sensitive design of recommenders, because the field is still emerging. Instead, we articulated an extensive menu of design techniques that have been applied or could likely be applied to production systems. Finally, we surveyed developing approaches to recommender regulation, identifying a substantial gap in knowledge and orientation between the technology and policy communities.

While the intersection of content personalization with individual and societal flourishing is a huge and varied area, we hope that this synthesis provides a shared language, useful starting points, and essential research directions for building human values into recommender systems.

ACKNOWLEDGMENTS

Thank you to Daphne Keller, Joelle Pineau, and Andrea Wong for helpful feedback. Dylan Hadfield-Menell would like to thank support from a gift from Effective Giving.

REFERENCES

- [1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. 30, 1 (2020), 127–158. DOI:<https://doi.org/10.1007/s11257-019-09256-1>
- [2] Himan Abdollahpouri and Robin Burke. 2019. Multi-stakeholder recommendation and its connection to multi-sided fairness. (2019). Retrieved from <https://arxiv.org/abs/1907.13158>
- [3] Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. Towards a Human-like Open-Domain Chatbot. arXiv:2001.09977 [cs, stat] (February 2020). Retrieved November 30, 2021 from <http://arxiv.org/abs/2001.09977>
- [4] David Adkins, Bilal Alsallakh, Adeel Cheema, Narine Kokhlikyan, Emily McReynolds, Pushkar Mishra, Chavez Procope, Jeremy Sawruk, Erin Wang, and Polina Zvyagin. 2022. Method Cards for Prescriptive Machine-Learning Transparency. Retrieved April 19, 2022 from <https://conf.researchr.org/details/cain-2022/cain-2022/12/Method-Cards-for-Prescriptive-Machine-Learning-Transparency>
- [5] M. Mehdi Afsar, Trafford Crump, and Behrouz Far. 2021. Reinforcement learning based recommender systems: A survey. arXiv:2101.06286 [cs] (January 2021). Retrieved August 25, 2021 from <https://arxiv.org/abs/2101.06286v1>

- [6] Hunt Allcott, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020. The welfare effects of social media. *American Economic Review* 110, 3 (2020), 629–676. DOI:<https://doi.org/10.1257/aer.20190658>
- [7] Hunt Allcott, Matthew Gentzkow, and Lena Song. 2021. Digital Addiction. NBER. Retrieved from https://www.nber.org/system/files/working_papers/w28936/w28936.pdf
- [8] Md. Sayeed Al-Zaman. 2021. Prevalence and source analysis of COVID-19 misinformation in 138 countries. *IFLA Journal* (August 2021), 03400352211041135. DOI:<https://doi.org/10.1177/03400352211041135>
- [9] Elizabeth Anderson. 2001. Symposium on Amartya Sen’s philosophy: 2 Unstrapping the straitjacket of ‘preference’: A comment on Amartya Sen’s contributions to philosophy and economics. *Economics & Philosophy* 17, 1 (2001), 21–38. DOI:<https://doi.org/10.1017/S0266267101000128>
- [10] Ruth A. Anderson, Lowell Worthington, William T. Anderson, and Glen Jennings. 1994. The development of an autonomy scale. *Contemp Fam Ther* 16, 4 (August 1994), 329–345. DOI:<https://doi.org/10.1007/BF02196884>
- [11] McKane Andrus, Elena Spitzer, Jeffrey Brown, and Alice Xiang. 2021. “What We Can’t Measure, We Can’t Understand”: Challenges to Demographic Data Procurement in the Pursuit of Fairness. arXiv:2011.02282 [cs] (January 2021). Retrieved September 18, 2021 from <http://arxiv.org/abs/2011.02282>
- [12] Charles Angelucci and Andrea Prat. 2021. Is Journalistic Truth Dead? Measuring How Informed Voters Are about Political News. Social Science Research Network, Rochester, NY. DOI:<https://doi.org/10.2139/ssrn.3593002>
- [13] Sinan Aral. 2016. The Future of Weak Ties. *American Journal of Sociology* 121, 6 (May 2016), 1931–1939. DOI:<https://doi.org/10.1086/686293>
- [14] Hannah Arendt. 1972. *Crises of the Republic: Lying in Politics, Civil Disobedience on Violence, Thoughts on Politics, and Revolution*. Houghton Mifflin Harcourt.
- [15] Andrew Arsht and Daniel Etcovitch. 2018. The Human Cost of Online Content Moderation. *Harvard Journal of Law & Technology* (March 2018). Retrieved January 6, 2022 from <https://jolt.law.harvard.edu/digest/the-human-cost-of-online-content-moderation>
- [16] Neja Asimovic, Jonathan Nagler, Richard Bonneau, and Joshua A. Tucker. 2021. Testing the effects of Facebook usage in an ethnically polarized setting. *Proc Natl Acad Sci USA* 118, 25 (June 2021). DOI:<https://doi.org/10.1073/pnas.2022819118>
- [17] Frederik Aust, Birk Diedenhofen, Sebastian Ullrich, and Jochen Musch. 2013. Seriousness checks are useful to improve data validity in online research. *Behav Res* 45, 2 (June 2013), 527–535. DOI:<https://doi.org/10.3758/s13428-012-0265-2>
- [18] Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Marcus Mann, Jaemin Lee, Alexander Volfovsky, and Friedolin Merhout. 2018. Exposure to opposing views on social media can increase political polarization. *PNAS* 115, 37 (2018), 9216–9221. DOI:<https://doi.org/10.1073/pnas.1804840115>
- [19] Kristzian Balog, Filip Radlinski, and Shushan Arakelyan. 2019. Transparent, Scrutable and Explainable User Models for Personalized Recommendation. ACM, New York, NY, USA, 265–274. DOI:<https://doi.org/10.1145/3331184.3331211>
- [20] Jack Bandy. 2021. Problematic Machine Behavior: A Systematic Literature Review of Algorithm Audits. arXiv:2102.04256 [cs] (February 2021). Retrieved June 4, 2021 from <http://arxiv.org/abs/2102.04256>
- [21] Jack Bandy and Nicholas Diakopoulos. 2021. More Accounts, Fewer Links: How Algorithmic Curation Impacts Media Exposure in Twitter Timelines. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1 (April 2021), 1–28. DOI:<https://doi.org/10.1145/3449152>
- [22] Simon Baron-Cohen and Sally Wheelwright. 2004. The Empathy Quotient: An Investigation of Adults with Asperger Syndrome or High Functioning Autism, and Normal Sex Differences. *J Autism Dev Disord* 34, 2 (April 2004), 163–175. DOI:<https://doi.org/10.1023/B:JADD.0000022607.19833.00>
- [23] Alon Bartal, Nava Pliskin, and Oren Tsur. 2020. Local/Global contagion of viral/non-viral information: Analysis of contagion spread in online social networks. *PLoS ONE* 15, 4 (April 2020), e0230811. DOI:<https://doi.org/10.1371/journal.pone.0230811>
- [24] Jonathan Bassen, Bharathan Balaji, Michael Schaarschmidt, Candace Thille, Jay Painter, Dawn Zimmaro, Alex Games, Ethan Fast, and John C. Mitchell. 2020. Reinforcement Learning for the Adaptive Scheduling of Educational Activities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ACM, Honolulu HI USA, 1–12. DOI:<https://doi.org/10.1145/3313831.3376518>
- [25] Christine Bauer and Alexander Novotny. 2017. A consolidated view of context for intelligent systems. *AIS* 9, 4 (June 2017), 377–393. DOI:<https://doi.org/10.3233/AIS-170445>
- [26] Philip Baugut and Katharina Neumann. 2020. Online propaganda use during Islamist radicalization. *Information, Communication & Society* 23, 11 (September 2020), 1570–1592. DOI:<https://doi.org/10.1080/1369118X.2019.1594333>
- [27] Seth D. Baum. 2020. Social choice ethics in artificial intelligence. *AI Mag* 35, 1 (2020), 165–176. DOI:<https://doi.org/10.1007/s00146-017-0760-1>
- [28] Mesfin A. Bekalu, Rachel F. McCloud, and K. Viswanath. 2019. Association of Social Media Use With Social Well-Being, Positive Mental Health, and Self-Rated Health: Disentangling Routine Use From Emotional Connection to Use. *Health Educ Behav* 46, 2_suppl (December 2019), 69S–80S. DOI:<https://doi.org/10.1177/1090198119863768>
- [29] Omer Ben-Porat, Gregory Goren, Itay Rosenberg, and Moshe Tennenholtz. 2019. From Recommendation Systems to Facility Location Games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1772–1779. DOI:<https://doi.org/10.1609/aaai.v33i01.33011772>
- [30] Omer Ben-Porat and Moshe Tennenholtz. 2018. A Game-Theoretic Approach to Recommendation Systems with Strategic Content Providers. In *Advances in Neural Information Processing Systems*, Curran Associates, Inc. Retrieved January 2, 2022 from <https://proceedings.neurips.cc/paper/2018/hash/a9a1d5317a33ae8cef33961c34144f84-Abstract.html>
- [31] Larwan Berke, Sushant Kafle, and Matt Huenerfauth. 2018. Methods for Evaluation of Imperfect Captioning Tools by Deaf or Hard-of-Hearing Users at Different Reading Literacy Levels. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ACM, Montreal QC Canada, 1–12. DOI:<https://doi.org/10.1145/3173574.3173665>

- [32] B. Douglas Bernheim. 2016. The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics. 7, 1 (2016), 12–68. DOI:<https://doi.org/10.1017/bca.2016.5>
- [33] Abraham Bernstein, Claes de Vreese, Natali Helberger, Wolfgang Schulz, Katharina Zweig, Christian Baden, Michael A. Beam, Marc P. Hauer, Lucien Heitz, Pascal Jürgens, Christian Katzenbach, Benjamin Kille, Beate Klimkiewicz, Wiebke Loosen, Judith Moeller, Goran Radanovic, Guy Shani, Nava Tintarev, Suzanne Tolmeijer, Wouter van Atteveldt, Sanne Vrijenhoek, and Theresa Zueger. 2020. Diversity in News Recommendations. (2020). Retrieved from <http://arxiv.org/abs/2005.09495>
- [34] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H. Chi, and Cristos Goodrow. 2019. Fairness in recommendation ranking through pairwise comparisons. (2019), 2212–2220. DOI:<https://doi.org/10.1145/3292500.3330745>
- [35] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Allison Woodruff, Christine Luu, Pierre Kreitmann, Jonathan Bischof, and Ed H. Chi. 2019. Putting Fairness Principles into Practice: Challenges, Metrics, and Improvements. In Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, ACM, Honolulu HI USA, 453–459. DOI:<https://doi.org/10.1145/3306618.3314234>
- [36] Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, and Ed H. Chi. 2018. Latent Cross: Making Use of Context in Recurrent Recommender Systems. In Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18), Association for Computing Machinery, New York, NY, USA, 46–54. DOI:<https://doi.org/10.1145/3159652.3159727>
- [37] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. DOI:<https://doi.org/10.1145/3209978.3210063>
- [38] Robert M. Bond, Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. A 61-million-person experiment in social influence and political mobilization. 489, 7415 (2012), 295–298. DOI:<https://doi.org/10.1038/nature11421>
- [39] Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016. A context-aware time model for web search. In Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, 205–214.
- [40] Alan Borning and Michael Muller. 2012. Next Steps for Value Sensitive Design. In CHI '12, 10.
- [41] Craig Boutilier. 2013. Computational Decision Support Regret-Based Models for Optimization and Preference Elicitation. In Comparative Decision Making. Oxford University Press, New York. DOI:<https://doi.org/10.1093/acprof:oso/9780199856800.003.0041>
- [42] Craig Boutilier, Richard S. Zemel, and Benjamin Marlin. 2003. Active collaborative filtering. In Proceedings of the 19th Conference on Uncertainty in Artificial Intelligence (UAI-03), Acapulco, Mexico, 98–106.
- [43] Levi Boxell, Matthew Gentzkow, and Jesse Shapiro. 2017. Is the Internet Causing Political Polarization? Evidence from Demographics. (2017). DOI:<https://doi.org/10.3386/w23258>
- [44] Levi Boxell, Gentzkow, Matthew, and Jesse M Shapiro. 2020. Cross-Country Trends in Affective Polarization. Retrieved from <https://www.nber.org/papers/w26669>
- [45] Lia Bozarth and Ceren Budak. 2020. Toward a Better Performance Evaluation Framework for Fake News Classification. 14, Icwsm (2020), 60–71.
- [46] Lia Bozarth, Aparajita Saraf, and Ceren Budak. 2020. Higher ground? how groundtruth labeling impacts our understanding of fake news about the 2016 U.S. presidential nominees. Icwsm (2020), 48–59. DOI:<https://doi.org/10.2139/ssrn.3340173>
- [47] William J. Brady and Jay Joseph Van Bavel. 2021. Estimating the effect size of moral contagion in online networks: A pre-registered replication and meta-analysis. Open Science Framework. DOI:<https://doi.org/10.31219/osf.io/s4w2x>
- [48] J Scott Brennan, Felix M Simon, Philip N Howard, and Rasmus Kleis Nielsen. 2020. Types, Sources, and Claims of COVID-19 Misinformation. Reuters Institute for Politics, University of Oxford. Retrieved from <https://reutersinstitute.politics.ox.ac.uk/types-sources-and-claims-covid-19-misinformation>
- [49] Miles Brundage, Shahar Avin, Jasmine Wang, Haydn Belfield, Gretchen Krueger, Gillian Hadfield, Heidy Khlaaf, Jingying Yang, Helen Toner, Ruth Fong, Tegan Maharaj, Pang Wei Koh, Sara Hooker, Jade Leung, Andrew Trask, Emma Blumke, Jonathan Lebensold, Cullen O’Keefe, Mark Koren, Théo Ryffel, JB Rubinovitz, Tamay Besiroglu, Federica Carugati, Jack Clark, Peter Eckersley, Sarah de Haas, Maritza Johnson, Ben Laurie, Alex Ingerman, Igor Krawczuk, Amanda Askell, Rosario Cammarota, Andrew Lohn, David Krueger, Charlotte Stix, Peter Henderson, Logan Graham, Carina Prunkl, Bianca Martin, Elizabeth Seger, Noa Zilberman, Adrian Weller, Brian Tse, Elizabeth Barnes, Allan Dafoe, Paul Scharré, Ariel Herbert-Voss, Martijn Rasser, Shagun Sodhani, Carrick Flynn, Thomas Krendl Gilbert, Lisa Dyer, Saif Khan, Yoshua Bengio, and Markus Anderljung. 2020. Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims. (2020). Retrieved from <https://arxiv.org/abs/2004.07213v2>
- [50] Axel Bruns. 2019. It’s Not the Technology, Stupid: How the ‘Echo Chamber’ and ‘Filter Bubble’ Metaphors Have Failed Us. (2019). Retrieved from <http://snurb.info/node/2526>
- [51] BSR. 2018. Human Rights Impact Assessment: Facebook in Myanmar.
- [52] Robin Burke. 2017. Multisided Fairness for Recommendation. arXiv:1707.00093 [cs] (July 2017). Retrieved October 5, 2021 from <http://arxiv.org/abs/1707.00093>
- [53] Colin F Camerer and Ernst Fehr. 2004. Measuring social norms and preferences using experimental games: A guide for social scientists. In Foundations of Human Sociality – Experimental and Ethnographic Evidence from 15 Small-Scale Societies, h.
- [54] Colin F. Camerer, George Loewenstein, and Matthew Rabin (Eds.). 2004. Advances in Behavioral Economics. Princeton University Press.
- [55] Rocío Cañamares, Pablo Castells, and Alistair Moffat. 2020. Offline evaluation options for recommender systems. Inf Retrieval J 23, 4 (August 2020), 387–410. DOI:<https://doi.org/10.1007/s10791-020-09371-3>
- [56] Cansu Canca. 2020. Operationalizing AI ethics principles. Commun. ACM 63, 12 (November 2020), 18–21. DOI:<https://doi.org/10.1145/3430368>
- [57] Micah Carroll, Dylan Hadfield-Menell, Anca Dragan, and Stuart Russell. 2021. Estimating and Penalizing Preference Shift in Recommender

Systems. RecSys '21: Fifteenth ACM Conference on Recommender Systems (2021). DOI:<https://doi.org/10.1145/3460231.3478849>

- [58] Ben Carterette, Evangelos Kanoulas, and Emine Yilmaz. 2012. Incorporating variability in user behavior into systems based evaluation. In Proceedings of the 21st ACM international conference on Information and knowledge management, 135–144.
- [59] Giuliana Carullo, Aniello Castiglione, and Alfredo De Santis. 2014. Friendship Recommendations in Online Social Networks. In 2014 International Conference on Intelligent Networking and Collaborative Systems, IEEE, Salerno, 42–48. DOI:<https://doi.org/10.1109/INCoS.2014.114>
- [60] Pablo Castells, Neil J. Hurley, and Saul Vargas. 2015. Novelty and diversity in recommender systems. Springer US, 881–918. DOI:https://doi.org/10.1007/978-1-4899-7637-6_26
- [61] Calvin Chan, Viknesh Sounderajah, Elisabeth Daniels, Amish Acharya, Jonathan Clarke, Seema Yalamanchili, Pasha Normahani, Sheraz Markar, Hutan Ashrafian, and Ara Darzi. 2021. The Reliability and Quality of YouTube Videos as a Source of Public Health Information Regarding COVID-19 Vaccination: Cross-sectional Study. *JMIR Public Health and Surveillance* 7, 7 (July 2021), e29942. DOI:<https://doi.org/10.2196/29942>
- [62] Jakraphan Chaopreecha. 2019. Revitalization of Tradition through Social Media: A Case of the Vegetarian Festival in Phuket, Thailand. (2019), 35.
- [63] Gilad Chen, Stanley M. Gully, and Dov Eden. 2001. Validation of a New General Self-Efficacy Scale. *Organizational Research Methods* 4, 1 (January 2001), 62–83. DOI:<https://doi.org/10.1177/109442810141004>
- [64] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2020. Bias and Debias in Recommender System: A Survey and Future Directions. arXiv:2010.03240 [cs] (October 2020). Retrieved September 2, 2021 from <http://arxiv.org/abs/2010.03240>
- [65] Li Chen and Pearl Pu. 2004. Survey of Preference Elicitation Methods. (2004), 23.
- [66] Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H. Chi. 2019. Top-k off-policy correction for a reinforce recommender system. (2019), 456–464. DOI:<https://doi.org/10.1145/3289600.3290999>
- [67] Yan-Ying Chen, Tao Chen, Winston H. Hsu, Hong-Yuan Mark Liao, and Shih-Fu Chang. 2014. Predicting Viewer Affective Comments Based on Image Content in Social Media. In Proceedings of International Conference on Multimedia Retrieval, ACM, Glasgow United Kingdom, 233–240. DOI:<https://doi.org/10.1145/2578726.2578756>
- [68] Zhilong Chen, Jinghua Piao, Xiaochong Lan, Hancheng Cao, Chen Gao, Zhicong Lu, and Yong Li. 2022. Practitioners Versus Users: A Value-Sensitive Evaluation of Current Industrial Recommender System Design. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (November 2022), 1–32. DOI:<https://doi.org/10.1145/3555646>
- [69] Steve Chien, Prateek Jain, Walid Krichene, Steffen Rendle, Shuang Song, Abhradeep Thakurta, and Li Zhang. 2021. Private Alternating Least Squares: Practical Private Matrix Completion with Tighter Rates. In Proceedings of the 38th International Conference on Machine Learning, PMLR, 1877–1887. Retrieved January 7, 2022 from <https://proceedings.mlr.press/v139/chien21a.html>
- [70] Rumman Chowdhury. 2021. Sharing learnings about our image cropping algorithm. Twitter. Retrieved October 14, 2021 from https://blog.twitter.com/engineering/en_us/topics/insights/2021/sharing-learnings-about-our-image-cropping-algorithm
- [71] Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. 2016. Towards Conversational Recommender Systems. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16), Association for Computing Machinery, New York, NY, USA, 815–824. DOI:<https://doi.org/10.1145/2939672.2939746>
- [72] Jennifer Cobbe and Jatinder Singh. 2019. Regulating Recommending: Motivations, Considerations, and Principles. 10, 3 (2019). DOI:<https://doi.org/10.2139/ssrn.3371830>
- [73] Council of Europe. 2020. Guide on Article 8 of the European Convention on Human Rights. Retrieved from https://www.echr.coe.int/documents/guide_art_8_eng.pdf
- [74] Cédric Courtois, Laura Slechten, and Lennert Coenen. 2018. Challenging Google Search filter bubbles in social and political information: Disconfirming evidence from a digital methods case study. *Telematics and Informatics* 35, 7 (October 2018), 2006–2015. DOI:<https://doi.org/10.1016/j.tele.2018.07.004>
- [75] Tim Cowlshaw, Todd Burlington, David Man, Jakub Fiala, Rhiannon Barrington, and George Wright. 2018. Personalizing the Public: Personalising Linear Radio at a Public Service Broadcaster. British Broadcasting Corporation. Retrieved from <https://www.bbc.com/news/technology-3293>
- [76] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Model User-Adap Inter* 18, 5 (August 2008), 455. DOI:<https://doi.org/10.1007/s11257-008-9051-3>
- [77] Mark Csikszentmihalyi. 2020. Confucius. In *The Stanford Encyclopedia of Philosophy* (Summer 2020), Edward N. Zalta (ed.). Metaphysics Research Lab, Stanford University. Retrieved October 21, 2021 from <https://plato.stanford.edu/archives/sum2020/entries/confucius/>
- [78] Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yian Ma, Cory McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. 2020. Underspecification Presents Challenges for Credibility in Modern Machine Learning. (2020). Retrieved from <http://arxiv.org/abs/2011.03395>
- [79] Abhinandan S. Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. Google news personalization: scalable online collaborative filtering. ACM Press, New York, New York, USA, 271. DOI:<https://doi.org/10.1145/1242572.1242610>
- [80] Aparna Das, Claire Mathieu, and Daniel Ricketts. 2009. Maximizing profit using recommender systems. (2009). Retrieved from <http://arxiv.org/abs/0908.3633>
- [81] Maria-Iuliana Dascalu, Constanta-Nicoleta Bodea, Monica Nastasia Mihailescu, Elena Alice Tanase, and Patricia Ordoñez de Pablos. 2016. Educational recommender systems and their application in lifelong learning. *Behaviour & Information Technology* 35, 4 (April 2016), 290–297. DOI:<https://doi.org/10.1080/0144929X.2015.1128977>

- [82] Yashar Deldjoo, Markus Schedl, Paolo Cremonesi, and Gabriella Pasi. 2020. Recommender Systems Leveraging Multimedia Content. *ACM Comput. Surv.* 53, 5 (October 2020), 1–38. DOI:<https://doi.org/10.1145/3407190>
- [83] Joaquin Delgado, Samuel Lind, Carl Radecke, and Satish Konijeti. 2019. Simple objectives work better. Retrieved from <http://ceur-ws.org/Vol-2440/paper5.pdf>
- [84] Ángel Díaz and Laura Hecht-Felella. 2021. Double Standards in Social Media Content Moderation. Brennan Center for Justice. Retrieved from https://www.brennancenter.org/sites/default/files/2021-08/Double_Standards_Content_Moderation.pdf
- [85] 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 & Knowledge Management (CIKM '20)*, Association for Computing Machinery, New York, NY, USA, 275–284. DOI:<https://doi.org/10.1145/3340531.3411962>
- [86] Ed Diener. 2000. Subjective well-being: The science of happiness and a proposal for a national index. *American Psychologist* 55, 1 (2000), 34–43. DOI:<https://doi.org/10.1037/0003-066X.55.1.34>
- [87] Ed Diener, Robert A. Emmons, Randy J. Larsen, and Sharon Griffin. 1985. The Satisfaction With Life Scale. *Journal of Personality Assessment* 49, 1 (February 1985), 71–75. DOI:https://doi.org/10.1207/s15327752jpa4901_13
- [88] Rachel Dodge, Annette P. Daly, Jan Huyton, and Lalage D. Sanders. 2012. The challenge of defining wellbeing. *International Journal of Wellbeing* 2, 3 (August 2012). Retrieved December 2, 2021 from <https://www.internationaljournalofwellbeing.org/index.php/ijow/article/view/89>
- [89] Eileen Donahoe and Megan Macduffee Metzger. 2019. Artificial intelligence and human rights. 30, 2 (2019), 115–126. DOI:<https://doi.org/10.1353/jod.2019.0029>
- [90] Evelyn Douek. 2020. The Limits of International Law in Content Moderation. *SSRN Journal* (2020). DOI:<https://doi.org/10.2139/ssrn.3709566>
- [91] Evelyn Douek. 2021. Governing Online Speech: From “Posts-as-Trumps” to Proportionality and Probability. 121, 3 (2021). Retrieved from <https://columbialawreview.org/content/governing-online-speech-from-posts-as-trumps-to-proportionality-and-probability/>
- [92] Matjaž Drev and Boštjan Delak. 2021. Conceptual Model of Privacy by Design. *Journal of Computer Information Systems* (July 2021), 1–8. DOI:<https://doi.org/10.1080/08874417.2021.1939197>
- [93] M Z van Drunen, N Helberger, and M Bastian. 2019. Know your algorithm: what media organizations need to explain to their users about news personalization. *International Data Privacy Law* 9, 4 (November 2019), 220–235. DOI:<https://doi.org/10.1093/idpl/izp011>
- [94] Gabriel Dulac-Arnold, Richard Evans, Hado van Hasselt, Peter Sunehag, Timothy Lillicrap, Jonathan Hunt, Timothy Mann, Theophane Weber, Thomas Degris, and Ben Coppin. 2015. Deep reinforcement learning in large discrete action spaces. *arXiv preprint arXiv:1512.07679* (2015).
- [95] Kris Dunn and Shane P. Singh. 2014. Pluralistic conditioning: social tolerance and effective democracy. *Democratization* 21, 1 (January 2014), 1–28. DOI:<https://doi.org/10.1080/13510347.2012.697056>
- [96] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (August 2014), 211–407. DOI:<https://doi.org/10.1561/04000000042>
- [97] Thomas Ehrlich. 2000. *Civic responsibility and higher education*. Oryx Press, Westport, Conn.
- [98] Michael D. Ekstrand, Anabrata Das, Robin Burke, and Fernando Diaz. 2021. Fairness in Recommender Systems. In *Recommender Systems Handbook* (3rd ed.), Francesco Ricci, Lior Roach and Bracha Shapira (eds.). Springer-Verlag.
- [99] Michael D. Ekstrand, Anabrata Das, Robin Burke, and Fernando Diaz. 2022. Fairness in Information Access Systems. *Foundations and Trends in Information Retrieval* 16, 1–2 (2022), 1–177.
- [100] Michael D. Ekstrand, Mucun Tian, Ion Madrazo Azpiazu, Jennifer D. Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All The Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, PMLR, 172–186. Retrieved November 10, 2021 from <https://proceedings.mlr.press/v81/ekstrand18b.html>
- [101] Nicole B. Ellison, Charles Steinfield, and Cliff Lampe. 2007. The Benefits of Facebook “Friends:” Social Capital and College Students’ Use of Online Social Network Sites. *Journal of Computer-Mediated Communication* 12, 4 (2007), 1143–1168. DOI:<https://doi.org/10.1111/j.1083-6101.2007.00367.x>
- [102] Jacob K. Eskildsen and Kai Kristensen. 2011. The gender bias of the Net Promoter Score. In *2011 IEEE International Conference on Quality and Reliability*, 254–258. DOI:<https://doi.org/10.1109/ICQR.2011.6031720>
- [103] European Commission. 2016. General Data Protection Regulation. Retrieved November 19, 2021 from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [104] European Commission. 2022. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act). Retrieved October 10, 2023 from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32022R2065>
- [105] Charles Evans and Atoosa Kasirzadeh. 2021. User Tampering in Reinforcement Learning Recommender Systems. *arXiv:2109.04083 [cs]* (September 2021). Retrieved September 27, 2021 from <http://arxiv.org/abs/2109.04083>
- [106] Carrie Exton and Michal Shinwell. 2018. Policy use of well-being metrics: Describing Countries’ Experiences. 33, 94 (2018). DOI:<https://doi.org/10.1787/d98eb8ed-en>
- [107] Facebook. 2019. People, Publishers, the Community. Facebook Newsroom. Retrieved October 8, 2021 from <https://about.fb.com/news/2019/04/people-publishers-the-community/>
- [108] Facebook. 2021. Facebook’s Corporate Human Rights Policy. Retrieved December 8, 2021 from <https://about.fb.com/wp-content/uploads/2021/03/Facebooks-Corporate-Human-Rights-Policy.pdf>
- [109] Facebook. 2021. Our Approach to Ranking. Facebook Transparency Center. Retrieved October 28, 2021 from

<https://transparency.fb.com/features/ranking-and-content/>

- [110] Marc Faddoul, Guillaume Chaslot, and Hany Farid. 2020. A Longitudinal Analysis of YouTube’s Promotion of Conspiracy Videos. (2020), 1–8.
- [111] Boi Faltings, Radu Jurca, Pearl Pu, and Bao Duy Tran. 2014. Incentives to Counter Bias in Human Computation. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, 59–66. Retrieved January 7, 2022 from <https://ojs.aaai.org/index.php/HCOMP/article/view/13145>
- [112] Jenny Fan and Amy X. Zhang. 2020. Digital Juries: A Civics-Oriented Approach to Platform Governance. (2020), 1–14. DOI:<https://doi.org/10.1145/3313831.3376293>
- [113] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining, 259–268.
- [114] Miriam Fernandez and Alejandro Bellogín. 2020. Recommender Systems and Misinformation: The Problem or the Solution? In OHARS Workshop, 14th ACM Conference on Recommender Systems, 9.
- [115] Benjamin Fields, Rhianna Jones, and Tim Cowlshaw. 2018. The case for public service recommender algorithms. In FATREC Workshop at RecSys 2018, Vancouver. Retrieved December 5, 2021 from <https://pirt.gitlab.io/fatrec2018/program/fatrec2018-fields.pdf>
- [116] Jessica Fjeld, Nele Achten, Hannah Hillgoss, Adam Nagy, and Madhulika Srikumar. 2020. Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI. Berkman Klein Center for Internet & Society, Cambridge, MA. Retrieved June 10, 2021 from <https://dash.harvard.edu/handle/1/42160420>
- [117] Richard Fletcher, Antonis Kalogeropoulos, and Rasmus Kleis Nielsen. 2021. More diverse, more politically varied: How social media, search engines and aggregators shape news repertoires in the United Kingdom. *New Media & Society* (July 2021), 14614448211027393. DOI:<https://doi.org/10.1177/14614448211027393>
- [118] Richard Fletcher and Rasmus Kleis Nielsen. 2018. Are people incidentally exposed to news on social media? A comparative analysis. 20, 7 (2018), 2450–2468. DOI:<https://doi.org/10.1177/1461444817724170>
- [119] Rachel Freedman, Rohin Shah, and Anca Dragan. 2020. Choice Set Misspecification in Reward Inference. In ICJCAI-PRICAI 2020 Workshop on Artificial Intelligence Safety. Retrieved August 18, 2021 from <http://arxiv.org/abs/2101.07691>
- [120] Bruno S. Frey, Christine Benesch, and Aloise Stutzer. 2007. Does Watching TV Make Us Happy? 28, 3 (2007), 283–313.
- [121] Batya Friedman, David G. Hendry, and Alan Borning. 2017. A survey of value sensitive design methods. 11, 23 (2017), 63–125. DOI:<https://doi.org/10.1561/1100000015>
- [122] Batya Friedman, Peter H Kahn, and Alan Borning. 2002. Value Sensitive Design: Theory and Methods. (2002), 8.
- [123] Eline Frison and Steven Eggermont. 2020. Toward an Integrated and Differential Approach to the Relationships Between Loneliness, Different Types of Facebook Use, and Adolescents’ Depressed Mood. *Communication Research* 47, 5 (July 2020), 701–728. DOI:<https://doi.org/10.1177/0093650215617506>
- [124] Adrian Furnham. 1986. Response bias, social desirability and dissimulation. *Personality and Individual Differences* 7, 3 (January 1986), 385–400. DOI:[https://doi.org/10.1016/0191-8869\(86\)90014-0](https://doi.org/10.1016/0191-8869(86)90014-0)
- [125] Wulf Gaertner. 2009. *A Primer in Social Choice Theory: Revised Edition*. OUP Oxford.
- [126] Jason Gauci, Edoardo Conti, Yitao Liang, Kittipat Virochsiri, Yuchen He, Zachary Kaden, Vivek Narayanan, Xiaohui Ye, Zhengxing Chen, and Scott Fujimoto. 2019. Horizon: Facebook’s Open Source Applied Reinforcement Learning Platform. arXiv:1811.00260 [cs, stat] (September 2019). Retrieved November 30, 2021 from <http://arxiv.org/abs/1811.00260>
- [127] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM, Anchorage AK USA, 2221–2231. DOI:<https://doi.org/10.1145/3292500.3330691>
- [128] Alexandre Gilotte, Clément Calauzènes, Thomas Nedelec, Alexandre Abraham, and Simon Dollé. 2018. Offline A/B Testing for Recommender Systems. In Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, ACM, Marina Del Rey CA USA, 198–206. DOI:<https://doi.org/10.1145/3159652.3159687>
- [129] Global Internet Forum to Counter Terrorism. 2021. Content-Sharing Algorithms, Processes, and Positive Interventions Working Group Part 1: Content-Sharing Algorithms & Processes. Retrieved January 3, 2022 from <https://gifct.org/wp-content/uploads/2021/07/GIFCT-CAPII-2021.pdf>
- [130] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D. Sculley. 2017. Google Vizier: A Service for Black-Box Optimization. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’17), Association for Computing Machinery, New York, NY, USA, 1487–1495. DOI:<https://doi.org/10.1145/3097983.3098043>
- [131] Christos Goodrow. 2021. On YouTube’s recommendation system. YouTube blog. Retrieved November 19, 2021 from <https://blog.youtube/inside-youtube/on-youtubes-recommendation-system/>
- [132] Google. 2020. Recommendation Systems Overview. Google Developers. Retrieved January 5, 2022 from <https://developers.google.com/machine-learning/recommendation/overview/types>
- [133] Google. 2021. AI At Google: Our Principles. Retrieved November 2, 2021 from <https://ai.google/principles/>
- [134] Mitchell L Gordon, Kaitlyn Zhou, and Michael S Bernstein. 2021. The Disagreement Deconvolution : Bringing Machine Learning Performance Metrics In Line With Reality. Retrieved from <https://dl.acm.org/doi/abs/10.1145/3411764.3445423>
- [135] Carol Graham, Kate Laffan, and Sergio Pinto. 2018. Well-being in metrics and policy. *Science* 362 (6412), (2018), 287–288.
- [136] Francis Green. 2011. Unpacking the misery multiplier: How employability modifies the impacts of unemployment and job insecurity on life satisfaction and mental health. 30, 2 (2011), 265–276. DOI:<https://doi.org/10.1016/j.jhealeco.2010.12.005>

- [137] Ulrike Gretzel and Daniel R. Fesenmaier. 2006. Persuasion in Recommender Systems. *International Journal of Electronic Commerce* 11, 2 (2006), 81–100.
- [138] Nastasia Griffioen, Marieke van Rooij, Anna Lichtwarck-Aschoff, and Isabela Granic. 2020. Toward improved methods in social media research. *Technology, Mind, and Behavior* 1, 1 (2020). DOI:<https://doi.org/10.1037/tmb0000005>
- [139] Andrew Guess, Benjamin Lyons, Brendan Nyhan, and Jason Reifler. 2018. Avoiding the Echo Chamber about Echo Chambers: Why selective exposure to like-minded political news is less prevalent than you think. Retrieved from https://kf-site-production.s3.amazonaws.com/media_elements/files/000/000/133/original/Topos_KF_White-Paper_Nyhan_V1.pdf
- [140] Viral Gupta and Yunbo Ouyang. 2020. Rise of the Machines: Removing the {Human-in-the-Loop}. In 2020 USENIX Conference on Operational Machine Learning. Retrieved January 2, 2022 from <https://www.usenix.org/conference/opml20/presentation/gupta>
- [141] Dylan Hadfield-Menell and Gillian K. Hadfield. 2019. Incomplete contracting and AI alignment. In AIES '19: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, 417–422. DOI:<https://doi.org/10.1145/3306618.3314250>
- [142] Alon Halevy, Cristian Canton-Ferrer, Hao Ma, Umut Ozertem, Patrick Pantel, Marzieh Saeidi, Fabrizio Silvestri, and Ves Stoyanov. 2022. Preserving integrity in online social networks. *Commun. ACM* 65, 2 (February 2022), 92–98. DOI:<https://doi.org/10.1145/3462671>
- [143] Jaron Harambam, Dimitrios Bountouridis, Mykola Makhortych, 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, Copenhagen Denmark, 69–77. DOI:<https://doi.org/10.1145/3298689.3347014>
- [144] Jaron Harambam, Natali Helberger, and Joris Van Hoboken. 2018. Democratizing algorithmic news recommenders: How to materialize voice in a technologically saturated media ecosystem. 376, 2133 (2018). DOI:<https://doi.org/10.1098/rsta.2018.0088>
- [145] Peter Hase and Mohit Bansal. 2020. Evaluating Explainable AI: Which Algorithmic Explanations Help Users Predict Model Behavior? arXiv:2005.01831 [cs] (May 2020). Retrieved September 16, 2021 from <http://arxiv.org/abs/2005.01831>
- [146] Ahmed Hassan and Ryen W. White. 2013. Personalized models of search satisfaction. In Proceedings of the 22nd ACM international conference on Information & Knowledge Management, 2009–2018.
- [147] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56, (2016), 9–27.
- [148] Natali Helberger. 2019. On the Democratic Role of News Recommenders. 7, 8 (2019), 993–1012. DOI:<https://doi.org/10.1080/21670811.2019.1623700>
- [149] Natali Helberger, Max von Drunen, Sanne Vrijenhoek, and Judith Möller. 2021. Regulation of news recommenders in the Digital Services Act: empowering David against the Very Large Online Goliath. *Internet Policy Review* (2021). Retrieved July 23, 2021 from <https://policyreview.info/articles/news/regulation-news-recommenders-digital-services-act-empowering-david-against-very-large>
- [150] Natali Helberger, Kari Karppinen, and Lucia D’Acunto. 2018. Exposure diversity as a design principle for recommender systems. 21, 2 (2018), 191–207. DOI:<https://doi.org/10.1080/1369118X.2016.1271900>
- [151] Natali Helberger, Paddy Leerssen, and Max Van Drunen. 2019. Germany proposes Europe’s first diversity rules for social media platforms. Retrieved October 25, 2020 from <https://blogs.lse.ac.uk/mediase/2019/05/29/germany-proposes-europes-first-diversity-rules-for-social-media-platforms/>
- [152] Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *Journal of Machine Learning Research* 21, (2020).
- [153] Stephen Hicks, Lucy Tinkler, and Paul Allin. 2013. Measuring Subjective Well-Being and its Potential Role in Policy: Perspectives from the UK Office for National Statistics. *Soc Indie Res* 114, 1 (October 2013), 73–86. DOI:<https://doi.org/10.1007/s11205-013-0384-x>
- [154] Jennifer L. Hochschild and Katherine Levine Einstein. 2015. Do Facts Matter?: Information and Misinformation in American Politics. University of Oklahoma Press.
- [155] David Holtz, Ben Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer, and Sinan Aral. 2020. The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify. In Proceedings of the 21st ACM Conference on Economics and Computation (EC '20), Association for Computing Machinery, New York, NY, USA, 75–76. DOI:<https://doi.org/10.1145/3391403.3399532>
- [156] Homa Hosseinmardi, Amir Ghasemian, Aaron Clauset, David M. Rothschild, Markus Mobius, and Duncan J. Watts. 2020. Evaluating the scale, growth, and origins of right-wing echo chambers on YouTube. (2020). Retrieved from <http://arxiv.org/abs/2011.12843>
- [157] Silas Hsu, Kristen Vaccaro, Yin Yue, Aimee Rickman, and Karrie Karahalios. 2020. Awareness, Navigation, and Use of Feed Control Settings Online. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), Association for Computing Machinery, New York, NY, USA, 1–13. DOI:<https://doi.org/10.1145/3313831.3376583>
- [158] Rosalind Hursthouse and Glen Pettigrove. 2018. Virtue Ethics. In *The Stanford Encyclopedia of Philosophy* (Winter 2018), Edward N. Zalta (ed.). Metaphysics Research Lab, Stanford University. Retrieved October 21, 2021 from <https://plato.stanford.edu/archives/win2018/entries/ethics-virtue/>
- [159] Ferenc Huszár, Sofia Ira Ktena, Conor O’Brien, Luca Belli, Andrew Schlaikjer, and Moritz Hardt. 2021. Algorithmic Amplification of Politics on Twitter. Retrieved from <https://arxiv.org/abs/2110.11010>
- [160] Eugene Ie, Vihan Jain, Jing Wang, Sanmit Narvekar, Ritesh Agarwal, Rui Wu, Heng-Tze Cheng, Tushar Chandra, and Craig Boutilier. 2019. SlateQ: A Tractable Decomposition for Reinforcement Learning with Recommendation Sets. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, International Joint Conferences on Artificial Intelligence Organization, Macao, China, 2592–2599. DOI:<https://doi.org/10.24963/ijcai.2019/360>
- [161] IEEE. 2020. IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being. IEEE. DOI:<https://doi.org/10.1109/IEEESTD.2020.9084219>

- [162] IEEE. 2021. 7000-2021: IEEE Standard Model Process for Addressing Ethical Concerns during System Design. Retrieved from <https://ieeexplore.ieee.org/document/9536679/>
- [163] Oana Ignat, Y.-Lan Boureau, Jane A. Yu, and Alon Halevy. 2021. Detecting Inspiring Content on Social Media. arXiv:2109.02734 [cs] (September 2021). Retrieved October 10, 2021 from <http://arxiv.org/abs/2109.02734>
- [164] Simon Jackman. 2008. Measurement. In *The Oxford Handbook of Political Methodology*, Janet M. Box-Steffensmeier, Henry E. Brady and David Collier (eds.). Oxford University Press. DOI:<https://doi.org/10.1093/oxfordhb/9780199286546.003.0006>
- [165] Abigail Z. Jacobs and Hanna Wallach. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, ACM, Virtual Event Canada, 375–385. DOI:<https://doi.org/10.1145/3442188.3445901>
- [166] Prateek Jain, Om Dipakbhai Thakkar, and Abhradeep Thakurta. 2018. Differentially Private Matrix Completion Revisited. In *Proceedings of the 35th International Conference on Machine Learning*, PMLR, 2215–2224. Retrieved January 7, 2022 from <https://proceedings.mlr.press/v80/jain18b.html>
- [167] Dietmar Jannach and Gediminas Adomavicius. 2016. Recommendations with a purpose. (2016), 7–10. DOI:<https://doi.org/10.1145/2959100.2959186>
- [168] Dietmar Jannach and Gediminas Adomavicius. 2017. Price and Profit Awareness in Recommender Systems. Como, Italy. Retrieved from <http://arxiv.org/abs/1707.08029>
- [169] Dietmar Jannach and Michael Jugovac. 2019. Measuring the business value of recommender systems. 10, 4 (2019), 1–22. DOI:<https://doi.org/10.1145/3370082>
- [170] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A Survey on Conversational Recommender Systems. *ACM Comput. Surv.* 54, 5 (June 2021), 1–36. DOI:<https://doi.org/10.1145/3453154>
- [171] Dietmar Jannach, Sidra Naveed, and Michael Jugovac. 2017. User Control in Recommender Systems: Overview and Interaction Challenges. In *E-Commerce and Web Technologies*, Derek Bridge and Heiner Stuckenschmidt (eds.). Springer International Publishing, Cham, 21–33. DOI:https://doi.org/10.1007/978-3-319-53676-7_2
- [172] Gawesh Jawaheer, Martin Szomszor, and Patty Kostkova. 2010. Comparison of implicit and explicit feedback from an online music recommendation service. In *Proceedings of the 1st International Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec '10)*, Association for Computing Machinery, New York, NY, USA, 47–51. DOI:<https://doi.org/10.1145/1869446.1869453>
- [173] Hamed Jelodar, Yongli Wang, Chi Yuan, Xia Feng, Xiahui Jiang, Yanchao Li, and Liang Zhao. 2019. Latent Dirichlet allocation (LDA) and topic modeling: models, applications, a survey. *Multimed Tools Appl* 78, 11 (June 2019), 15169–15211. DOI:<https://doi.org/10.1007/s11042-018-6894-4>
- [174] Olivier Jeunen and Bart Goethals. 2021. Top-K Contextual Bandits with Equity of Exposure. In *Fifteenth ACM Conference on Recommender Systems*, Association for Computing Machinery, New York, NY, USA, 310–320. Retrieved November 10, 2021 from <https://doi.org/10.1145/3460231.3474248>
- [175] Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. 2019. Degenerate Feedback Loops in Recommender Systems. (February 2019). DOI:<https://doi.org/10.1145/3306618.3314288>
- [176] Yucheng Jin, Bruno Cardoso, and Katrien Verbert. 2017. How Do Different Levels of User Control Affect Cognitive Load and Acceptance of Recommendations? 8.
- [177] Dimitris Kalimeris, Smriti Bhagat, Shankar Kalyanaraman, and Udi Weinsberg. 2021. Preference Amplification in Recommender Systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, ACM, Virtual Event Singapore, 805–815. DOI:<https://doi.org/10.1145/3447548.3467298>
- [178] Alex Kantrowitz. 2021. Facebook Removed The News Feed Algorithm In An Experiment. Then It Gave Up. *Big Technology*. Retrieved April 2, 2022 from <https://bigtechnology.substack.com/p/facebook-removed-the-news-feed-algorithm>
- [179] Mozghan Karimi, Dietmar Jannach, and Michael Jugovac. 2018. News recommender systems – Survey and roads ahead. 54, 6 (2018), 1203–1227. DOI:<https://doi.org/10.1016/j.ipm.2018.04.008>
- [180] Todd B. Kashdan, Paul Rose, and Frank D. Fincham. 2004. Curiosity and Exploration: Facilitating Positive Subjective Experiences and Personal Growth Opportunities. *Journal of Personality Assessment* 82, 3 (June 2004), 291–305. DOI:https://doi.org/10.1207/s15327752jpa8203_05
- [181] Daphne Keller. 2018. Internet platforms: Observations on speech, danger, and money. (2018), 5–8.
- [182] Daphne Keller. 2021. Amplification and Its Discontents. Knight First Amendment Institute at Columbia University. Retrieved March 28, 2022 from <https://knightcolumbia.org/content/amplification-and-its-discontents>
- [183] Lianne Kerlin. 2020. Human values: understanding psychological needs in a digital age. Retrieved from <http://downloads.bbc.co.uk/rd/pubs/whp/whp-pdf-files/WHP371.pdf>
- [184] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2021. The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes. arXiv:2005.04790 [cs] (April 2021). Retrieved October 10, 2021 from <http://arxiv.org/abs/2005.04790>
- [185] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. 2018. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning*, PMLR, 2668–2677. Retrieved January 2, 2022 from <https://proceedings.mlr.press/v80/kim18d.html>
- [186] Heejung Kim and Deborah Ko. 2007. Culture & Self-Expression 1 Culture and Self-Expression. In *Frontiers of social psychology: The self*. 325–342.
- [187] Bart P. Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender systems for self-actualization. (2016), 11–14. DOI:<https://doi.org/10.1145/2959100.2959189>

- [188] Christine Koggel and Joan Orme. 2010. Care Ethics: New Theories and Applications. *Ethics and Social Welfare* 4, 2 (July 2010), 109–114. DOI:<https://doi.org/10.1080/17496535.2010.484255>
- [189] Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson. 2020. Captum: A unified and generic model interpretability library for PyTorch. arXiv:2009.07896 [cs, stat] (September 2020). Retrieved December 1, 2021 from <http://arxiv.org/abs/2009.07896>
- [190] Tobias D. Krafft, Michael Gamer, and Katharina A. Zweig. 2019. What did you see? A study to measure personalization in Google’s search engine. *EPJ Data Sci.* 8, 1 (December 2019), 38. DOI:<https://doi.org/10.1140/epids/s13688-019-0217-5>
- [191] Alan B. Krueger and Arthur A. Stone. 2014. Progress in measuring subjective well-being. *Science* 346, 6205 (October 2014), 42–43. DOI:<https://doi.org/10.1126/science.1256392>
- [192] David Scott Krueger, Tegan Maharaj, and Jan Leike. 2020. Hidden incentives for auto-induced distributional shift. (2020).
- [193] Caitlin Kuhlman, Walter Gerych, and Elke Rundensteiner. 2021. Measuring Group Advantage: A Comparative Study of Fair Ranking Metrics. Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AI/ES ’21) (May 2021). DOI:<https://doi.org/10.1145/3461702.3462588>
- [194] Matevž Kunaver and Tomaž Požrl. 2017. Diversity in recommender systems – A survey. 123, (May 2017), 154–162. DOI:<https://doi.org/10.1016/j.knosys.2017.02.009>
- [195] Akos Lada, Meihong Wang, and Tak Yan. 2021. How machine learning powers Facebook’s News Feed ranking algorithm. *Engineering at Meta*. Retrieved December 13, 2021 from <https://engineering.fb.com/2021/01/26/ml-applications/news-feed-ranking/>
- [196] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Management Science* 65, 7 (July 2019), 2966–2981. DOI:<https://doi.org/10.1287/mnsc.2018.3093>
- [197] Mark Latonero and Aaina Agarwal. 2021. Human Rights Impact Assessments for AI: Learning from Facebook’s Failure in Myanmar. Carr Center for Human Rights Policy Harvard Kennedy School, Harvard University.
- [198] Edith Law and Luis von Ahn. 2011. Human Computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5, 3 (June 2011), 1–121. DOI:<https://doi.org/10.2200/S00371ED1V01Y201107AIM013>
- [199] Huyen Le, Raven Maragh, Brian Ekdale, Andrew High, Timothy Havens, and Zubair Shafiq. 2019. Measuring Political Personalization of Google News Search. In *The World Wide Web Conference, ACM, San Francisco CA USA*, 2957–2963. DOI:<https://doi.org/10.1145/3308558.3313682>
- [200] Mark Ledwich and Anna Zaitsev. 2019. Algorithmic Extremism: Examining YouTube’s Rabbit Hole of Radicalization. arXiv:1912.11211 [cs] (December 2019). Retrieved May 19, 2021 from <http://arxiv.org/abs/1912.11211>
- [201] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D Procaccia. 2019. Webuidai: Participatory framework for algorithmic governance. 3, (2019). DOI:<https://doi.org/10.1145/3359283>
- [202] Claire Leibowicz, Connie Moon Sehat, Adriana Stephan, and Jonathan Stray. 2021. If We Want Platforms to Think Beyond Engagement, We Have to Know What We Want Instead. Retrieved November 17, 2021 from <https://medium.com/partnership-on-ai/if-we-want-platforms-to-think-beyond-engagement-we-have-to-know-what-we-want-instead-a8c9fbf6688>
- [203] Jurek Leonhardt, Avishek Anand, and Megha Khosla. 2018. User Fairness in Recommender Systems. In *Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW ’18, ACM Press, Lyon, France*, 101–102. DOI:<https://doi.org/10.1145/3184558.3186949>
- [204] Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2021. User-oriented Fairness in Recommendation. In *Proceedings of the Web Conference 2021, Association for Computing Machinery, New York, NY, USA*, 624–632. Retrieved January 2, 2022 from <https://doi.org/10.1145/3442381.3449866>
- [205] Myles-Jay Linton, Paul Dieppe, and Antonieta Medina-Lara. 2016. Review of 99 self-report measures for assessing well-being in adults: exploring dimensions of well-being and developments over time. *BMJ Open* 6, 7 (July 2016), e010641. DOI:<https://doi.org/10.1136/bmjopen-2015-010641>
- [206] Felicia Loecherbach, Judith Moeller, Damian Trilling, and Wouter van Atteveldt. 2020. The Unified Framework of Media Diversity: A Systematic Literature Review. 8, 5 (2020), 605–642. DOI:<https://doi.org/10.1080/21670811.2020.1764374>
- [207] Felicia Loecherbach, Kasper Welbers, Judith Moeller, Damian Trilling, and Wouter Van Atteveldt. 2021. Is this a click towards diversity? Explaining when and why news users make diverse choices. In *13th ACM Web Science Conference 2021, ACM, Virtual Event United Kingdom*, 282–290. DOI:<https://doi.org/10.1145/3447535.3462506>
- [208] Sahil Loomba, Alexandre de Figueiredo, Simon J. Piatek, Kristen de Graaf, and Heidi J. Larson. 2021. Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nat Hum Behav* 5, 3 (March 2021), 337–348. DOI:<https://doi.org/10.1038/s41562-021-01056-1>
- [209] Philipp Lorenz-Spreen, Lisa Oswald, Stephan Lewandowsky, and Ralph Hertwig. 2021. Digital Media and Democracy: A Systematic Review of Causal and Correlational Evidence Worldwide. *SocArXiv*. DOI:<https://doi.org/10.31235/osf.io/p3z9v>
- [210] Raphael Louca, Moumita Bhattacharya, Diane Hu, and Liangjie Hong. 2019. Joint Optimization of Profit and Relevance for Recommendation Systems in E-commerce. 2440, (2019), 4–7.
- [211] Justina Lukat, Jürgen Margraf, Rainer Lutz, William M. van der Veld, and Eni S. Becker. 2016. Psychometric properties of the Positive Mental Health Scale (PMH-scale). *BMC Psychology* 4, 1 (February 2016), 8. DOI:<https://doi.org/10.1186/s40359-016-0111-x>
- [212] Glenn P. Malone, David R. Pillow, and Augustine Osman. 2012. The General Belongingness Scale (GBS): Assessing achieved belongingness. *Personality and Individual Differences* 52, 3 (February 2012), 311–316. DOI:<https://doi.org/10.1016/j.paid.2011.10.027>
- [213] David Manheim and Scott Garrabrant. 2018. Categorizing Variants of Goodhart’s Law. (2018), 1–10.
- [214] Joel Eduardo Martinez and Elizabeth Levy Paluck. 2020. Quantifying shared and idiosyncratic judgments of racism in social discourse. (2020).

Retrieved from <https://psyarxiv.com/kfpjg>

- [215] Sara Mattingly-Jordan, Bob Donaldson, Phillip Gray, and L Maria Ingram. 2019. Ethically Aligned Design First Edition Glossary. Retrieved from https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e_glossary.pdf
- [216] Nicolas Mattis, Philipp Masur, Judith Möller, and Wouter van Atteveldt. 2021. A theoretical framework for facilitating diverse news consumption through recommender design. *SocArXiv* (2021). DOI:<https://doi.org/10.31235/osf.io/wvxf5>
- [217] Colum Mccaffery. 2020. An algorithm for empowering public service news. *Polis*. Retrieved October 8, 2021 from <https://blogs.lse.ac.uk/polis/2020/09/28/this-swedish-radio-algorithm-gets-reporters-out-in-society/>
- [218] Jennifer McCoy and Murat Somer. 2019. Toward a Theory of Pernicious Polarization and How It Harms Democracies: Comparative Evidence and Possible Remedies. *The ANNALS of the American Academy of Political and Social Science* 681, 1 (January 2019), 234–271. DOI:<https://doi.org/10.1177/0002716218818782>
- [219] David W. McMillan and David M. Chavis. 1986. Sense of community: A definition and theory. *J. Community Psychol.* 14, 1 (January 1986), 6–23. DOI:[https://doi.org/10.1002/1520-6629\(198601\)14:1<6::AID-JCOP2290140103>3.0.CO;2-I](https://doi.org/10.1002/1520-6629(198601)14:1<6::AID-JCOP2290140103>3.0.CO;2-I)
- [220] Alan Meca. 2012. Personal Control and Responsibility Measure: A Psychometric Evaluation. Master of Science Psychology. Florida International University. DOI:<https://doi.org/10.25148/etd.FI12072703>
- [221] Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. 2017. Auditing search engines for differential satisfaction across demographics. In *Proceedings of the 26th international conference on World Wide Web companion*, 626–633.
- [222] 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. *ACM*, New York, NY, USA, 2243–2251. DOI:<https://doi.org/10.1145/3269206.3272027>
- [223] Rishabh Mehrotra, Niannan Xue, and Mounia Lalmas. 2020. Bandit based Optimization of Multiple Objectives on a Music Streaming Platform. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*, Association for Computing Machinery, New York, NY, USA, 3224–3233. DOI:<https://doi.org/10.1145/3394486.3403374>
- [224] Christoph J. Meinrenken, Scott M. Kaufman, Siddharth Ramesh, and Klaus S. Lackner. 2012. Fast Carbon Footprinting for Large Product Portfolios. 16, 5 (2012), 669–679. DOI:<https://doi.org/10.1111/j.1530-9290.2012.00463.x>
- [225] Claire Midgley, Sabrina Thai, Penelope Lockwood, Chloe Kovacheff, and Elizabeth Page-Gould. 2021. When every day is a high school reunion: Social media comparisons and self-esteem. *Journal of Personality and Social Psychology* 121, 2 (August 2021), 285–307. DOI:<https://doi.org/10.1037/pspi0000336>
- [226] Silvia Milano, Mariarosaria Taddeo, and Luciano Floridi. 2020. Recommender systems and their ethical challenges. *AI & Soc* 35, 4 (December 2020), 957–967. DOI:<https://doi.org/10.1007/s00146-020-00950-y>
- [227] Smitha Milli, Luca Belli, and Moritz Hardt. 2020. From Optimizing Engagement to Measuring Value. (2020), 1–14.
- [228] Martin Mladenov, Elliot Creager, Omer Ben-Porat, Kevin Swersky, Richard Zemel, and Craig Boutilier. 2020. Optimizing Long-term Social Welfare in Recommender Systems: A Constrained Matching Approach. In *Proceedings of the 37th International Conference on Machine Learning, PMLR*, 6987–6998. Retrieved January 2, 2022 from <https://proceedings.mlr.press/v119/mladenov20a.html>
- [229] Martin Mladenov, Ofer Meshi, Jayden Ooi, Dale Schuurmans, and Craig Boutilier. 2019. Advantage Amplification in Slowly Evolving Latent-State Environments. 2019-Augus. (May 2019), 3165–3172. DOI:<https://doi.org/10.24963/ijcai.2019/439>
- [230] Haradhan Kumar Mohajan. 2017. Two Criteria for Good Measurements in Research: Validity and Reliability. *ASHU- ES* 17, 4 (December 2017), 59–82. DOI:<https://doi.org/10.26458/1746>
- [231] Judith Möller, Damian Trilling, Natali Helberger, and Bram van Es. 2018. Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. 21, 7 (July 2018), 959–977. DOI:<https://doi.org/10.1080/1369118X.2018.1444076>
- [232] Marlon Mooijman, Joe Hoover, Ying Lin, Heng Ji, and Morteza Dehghani. 2018. Moralization in social networks and the emergence of violence during protests. *Nat Hum Behav* 2, 6 (June 2018), 389–396. DOI:<https://doi.org/10.1038/s41562-018-0353-0>
- [233] Meredith Ringel Morris, Annuska Zolyomi, Catherine Yao, Sina Bahram, Jeffrey P. Bigham, and Shaun K. Kane. 2016. “With most of it being pictures now, I rarely use it”: Understanding Twitter’s Evolving Accessibility to Blind Users. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, San Jose California USA, 5506–5516. DOI:<https://doi.org/10.1145/2858036.2858116>
- [234] Mozilla. 2021. How Artificial Intelligence Fuels Online Disinformation: Relevant Legislation. Mozilla Foundation. Retrieved November 19, 2021 from <https://foundation.mozilla.org/en/campaigns/trained-for-deception-how-artificial-intelligence-fuels-online-disinformation/relevant-legislation/>
- [235] Sendhil Mullainathan and Ziad Obermeyer. 2021. On the Inequity of Predicting A While Hoping for B. *AEA Papers and Proceedings* 111, (May 2021), 37–42. DOI:<https://doi.org/10.1257/pandp.20211078>
- [236] Kevin Munger and Joseph Phillips. 2020. Right-Wing YouTube: A Supply and Demand Perspective. *The International Journal of Press/Politics* (October 2020), 194016122096476. DOI:<https://doi.org/10.1177/1940161220964767>
- [237] Luke Munn. 2019. Alt-right pipeline: Individual journeys to extremism online. *First Monday* (2019). DOI:<https://doi.org/10.5210/fm.v24i6.10108>
- [238] Robin L. Nabi, Abby Prestin, and Jiyeon So. 2013. Facebook Friends with (Health) Benefits? Exploring Social Network Site Use and Perceptions of Social Support, Stress, and Well-Being. *Cyberpsychology, Behavior, and Social Networking* 16, 10 (October 2013), 721–727. DOI:<https://doi.org/10.1089/cyber.2012.0521>
- [239] Anatol-Fiete Näher and Ivar Krumpal. 2012. Asking sensitive questions: the impact of forgiving wording and question context on social desirability bias. *Qual Quant* 46, 5 (August 2012), 1601–1616. DOI:<https://doi.org/10.1007/s11135-011-9469-2>
- [240] Arvind Narayanan, Joanna Huey, and Edward W. Felten. 2016. A Precautionary Approach to Big Data Privacy. In *Data Protection on the Move*:

Current Developments in ICT and Privacy/Data Protection, Serge Gutwirth, Ronald Leenes and Paul De Hert (eds.). Springer Netherlands, Dordrecht, 357–385. DOI:https://doi.org/10.1007/978-94-017-7376-8_13

- [241] Menaka Narayanan, Emily Chen, Jeffrey He, Been Kim, Sam Gershman, and Finale Doshi-Velez. 2018. How do Humans Understand Explanations from Machine Learning Systems? An Evaluation of the Human-Interpretability of Explanation. arXiv:1802.00682 [cs] (February 2018). Retrieved January 5, 2022 from <http://arxiv.org/abs/1802.00682>
- [242] National Opinion Research Center. 2013. General Social Survey Cumulative File 1972-2012. Retrieved November 11, 2021 from http://people.wku.edu/douglas.smith/GSS%201972_2012%20Codebook.pdf
- [243] Esinath Ndiweni and Welcome Sibanda. 2020. CSR governance framework of South Africa, pre, during and post-apartheid: a manifestation of ubuntu values? *International Journal of Business Governance and Ethics* 14, 4 (January 2020), 363–383. DOI:<https://doi.org/10.1504/IJBGE.2020.110820>
- [244] Efrat Nechushtai and Seth C. Lewis. 2019. What kind of news gatekeepers do we want machines to be? Filter bubbles, fragmentation, and the normative dimensions of algorithmic recommendations. *90*, (2019), 298–307. DOI:<https://doi.org/10.1016/j.chb.2018.07.043>
- [245] Elizabeth A. Nick, David A. Cole, Sun-Joo Cho, Darcy K. Smith, T. Grace Carter, and Rachel Zerkowitz. 2018. The Online Social Support Scale: Measure Development and Validation. *Psychol Assess* 30, 9 (September 2018), 1127–1143. DOI:<https://doi.org/10.1037/pas0000558>
- [246] Nima Noorshams, Saurabh Verma, and Aude Hoefflner. 2020. TIES: Temporal Interaction Embeddings for Enhancing Social Media Integrity at Facebook. (2020), 3128–3135. DOI:<https://doi.org/10.1145/3394486.3403364>
- [247] Norwegian Centre for Research Data. 2012. ESS Round 6: European Social Survey Round 6 Data (2012). Retrieved November 15, 2021 from <http://www.europeansocialsurvey.org/data/themes.html?t=personal>
- [248] Ingrid Nunes and Dietmar Jannach. 2017. A systematic review and taxonomy of explanations in decision support and recommender systems. *User Model User-Adap Inter* 27, 3 (December 2017), 393–444. DOI:<https://doi.org/10.1007/s11257-017-9195-0>
- [249] Brendan Nyhan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y. Chen, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Andrew M. Guess, Edward Kennedy, Young Mie Kim, David Lazer, Neil Malhotra, Devra Moehler, Jennifer Pan, Daniel Robert Thomas, Rebekah Tromble, Carlos Velasco Rivera, Arjun Wilkins, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker. 2023. Like-minded sources on Facebook are prevalent but not polarizing. *Nature* (July 2023), 1–8. DOI:<https://doi.org/10.1038/s41586-023-06297-w>
- [250] Gus O’Donnell, Angus Deaton, Martine Durand, David Halpern, and Richard Layard. 2014. Wellbeing and Policy. Retrieved from <https://li.com/reports/the-commission-on-wellbeing-and-policy/>
- [251] OECD. 2019. Enhancing Access to and Sharing of Data: Reconciling Risks and Benefits for Data Re-use across Societies. Organisation for Economic Co-operation and Development, Paris. Retrieved October 14, 2021 from https://www.oecd-ilibrary.org/science-and-technology/enhancing-access-to-and-sharing-of-data_276aac8-en
- [252] OECD. 2020. Better Life Index. Retrieved November 11, 2021 from <https://www.oecdbetterlifeindex.org/>
- [253] Mathieu O’Neil and Michael J Jensen. 2020. Australian Perspectives on Misinformation. News Media Research Centre, University of Canberra. Retrieved from <https://researchprofiles.canberra.edu.au/en/publications/australian-perspectives-on-misinformation>
- [254] Elinor Ostrom. 2000. Collective action and the evolution of social norms. *14*, 3 (2000), 137–158. DOI:<https://doi.org/10.1257/jep.14.3.137>
- [255] Aviv Ovadya. 2021. Towards Platform Democracy: Policymaking Beyond Corporate CEOs and Partisan Pressure. Retrieved December 8, 2021 from <https://www.belfercenter.org/publication/towards-platform-democracy-policy-making-beyond-corporate-ceos-and-partisan-pressure>
- [256] Shanto p. Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. 2018. The Origins and Consequences of Affective Polarization in the United States. (2018), 1–35. DOI:<https://doi.org/10.1146/annurev-polisci-051117-073034>
- [257] Elisabeth Pacherie. 2007. The Sense of Control and the Sense of Agency. *Psyche* 13, 1 (2007), 1.
- [258] Stefania Paolini, Jake Harwood, and Mark Rubin. 2010. Negative Intergroup Contact Makes Group Memberships Salient: Explaining Why Intergroup Conflict Endures. *Pers Soc Psychol Bull* 36, 12 (December 2010), 1723–1738. DOI:<https://doi.org/10.1177/0146167210388667>
- [259] Jana Papcunová, Marcel Martončík, Denisa Fedáková, Michal Kentoš, Miroslava Bozogaňová, Ivan Srba, Robert Moro, Matúš Pikuliak, Marián Šimko, and Matúš Adamkovič. 2021. Hate speech operationalization: a preliminary examination of hate speech indicators and their structure. *Complex Intell. Syst.* (October 2021). DOI:<https://doi.org/10.1007/s40747-021-00561-0>
- [260] Jen Patja Howell, Evelyn Douek, Quinta Jurecic, and David Kaye. 2021. The Arrival of International Human Rights Law in Content Moderation. *The Lawfare Podcast*. Retrieved December 17, 2021 from <https://www.lawfareblog.com/lawfare-podcast-arrival-international-human-rights-law-content-moderation>
- [261] Gordon Pennycook and David G. Rand. 2019. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc Natl Acad Sci USA* 116, 7 (February 2019), 2521–2526. DOI:<https://doi.org/10.1073/pnas.1806781116>
- [262] Gordon Pennycook and David Gertler Rand. 2021. Reducing the spread of fake news by shifting attention to accuracy: Meta-analytic evidence of replicability and generalizability. *PsyArXiv*. DOI:<https://doi.org/10.31234/osf.io/v8ruj>
- [263] Thomas F. Pettigrew and Linda R. Tropp. 2006. A meta-analytic test of intergroup contact theory. *90*, 5 (2006), 751–783. DOI:<https://doi.org/10.1037/0022-3514.90.5.751>
- [264] Matteo Pinna, Léo Picard, and Christoph Goessmann. 2021. Cable News and COVID-19 Vaccine Compliance. Social Science Research Network, Rochester, NY. DOI:<https://doi.org/10.2139/ssrn.3890340>
- [265] Markus Prior. 2013. Media and Political Polarization. *16*, (2013), 101–27. DOI:<https://doi.org/10.1146/annurev-polisci-100711-135242>
- [266] Andrew K. Przybylski and Netta Weinstein. 2017. A Large-Scale Test of the Goldilocks Hypothesis: Quantifying the Relations Between Digital-

Screen Use and the Mental Well-Being of Adolescents. *Psychol Sci* 28, 2 (February 2017), 204–215. DOI:<https://doi.org/10.1177/0956797616678438>

- [267] Pearl Pu and Li Chen. 2010. A User-Centric Evaluation Framework of Recommender Systems. 612, (2010), 8.
- [268] Tao Qi, Fangzhao Wu, Chuhan Wu, Yongfeng Huang, and Xing Xie. 2020. Privacy-Preserving News Recommendation Model Learning. *arXiv:2003.09592* [cs] (October 2020). Retrieved December 8, 2021 from <http://arxiv.org/abs/2003.09592>
- [269] Kevin M. Quinn, Burt L. Monroe, Michael Colaresi, Michael H. Crespin, and Dragomir R. Radev. 2010. How to Analyze Political Attention with Minimal Assumptions and Costs. *American Journal of Political Science* 54, 1 (2010), 209–228. DOI:<https://doi.org/10.1111/j.1540-5907.2009.00427.x>
- [270] Amifa Raj and Michael D. Ekstrand. 2022. Measuring Fairness in Ranked Results. (2022), 11.
- [271] Amifa Raj, Connor Wood, Ananda Montoly, and Michael D. Ekstrand. 2020. Comparing Fair Ranking Metrics. *arXiv:2009.01311* [cs] (September 2020). Retrieved November 10, 2021 from <http://arxiv.org/abs/2009.01311>
- [272] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 33–44.
- [273] Ranking Digital Rights. 2020. 2020 Ranking Digital Rights Corporate Accountability Index. Ranking Digital Rights. Retrieved April 19, 2022 from <https://rankingdigitalrights.org/index2020/>
- [274] Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian Processes for Machine Learning*: MIT Press.
- [275] Manoel Horta Ribeiro, Raphael Ottoni, Robert West, Virgílio A. F. Almeida, and Wagner Meira. 2020. Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*, Association for Computing Machinery, New York, NY, USA, 131–141. DOI:<https://doi.org/10.1145/3351095.3372879>
- [276] Manoel Horta Ribeiro, Raphael Ottoni, Robert West, Virgílio A.F. Almeida, and W. M. Wagner Meira. 2020. Auditing radicalization pathways on YouTube. In *Conference on Fairness, Accountability, and Transparency (FAT* '20)*, Barcelona, 131–141. DOI:<https://doi.org/10.1145/3351095.3372879>
- [277] Samantha Robertson and Niloufar Salehi. 2020. What If I Don't Like Any Of The Choices? The Limits of Preference Elicitation for Participatory Algorithm Design. *arXiv:2007.06718* [cs] (July 2020). Retrieved June 1, 2021 from <http://arxiv.org/abs/2007.06718>
- [278] Jo Robinson, Georgina Cox, Eleanor Bailey, Sarah Hetrick, Maria Rodrigues, Steve Fisher, and Helen Herrman. 2016. Social media and suicide prevention: a systematic review: Suicide prevention and social media. *Early Intervention in Psychiatry* 10, 2 (April 2016), 103–121. DOI:<https://doi.org/10.1111/eip.12229>
- [279] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Kording, Carla Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. 2019. Tackling Climate Change with Machine Learning. (June 2019). Retrieved September 26, 2019 from <http://arxiv.org/abs/1906.05433>
- [280] Kevin Roose. 2019. The Making of a YouTube Radical. Retrieved April 5, 2021 from <https://www.nytimes.com/interactive/2019/06/08/technology/youtube-radical.html>
- [281] Brent D. Rosso, Kathryn H. Dekas, and Amy Wrzesniewski. 2010. On the meaning of work: A theoretical integration and review. *Research in Organizational Behavior* 30, (January 2010), 91–127. DOI:<https://doi.org/10.1016/j.riob.2010.09.001>
- [282] Cynthia Rudin. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *arXiv:1811.10154* [cs, stat] (September 2019). Retrieved November 23, 2021 from <http://arxiv.org/abs/1811.10154>
- [283] Stuart Russell. 2019. *Human compatible: Artificial intelligence and the problem of control*. Viking, New York.
- [284] Agnieszka Rychwalska and Magdalena Roszczyńska-Kurasińska. 2018. Polarization on Social Media: When Group Dynamics Leads to Societal Divides. (2018). DOI:<https://doi.org/10.24251/hiess.2018.263>
- [285] Carol D Ryff and Corey Lee M Keyes. 1995. The Structure of Psychological Well-Being Revisited. *Journal of Personality and Social Psychology* 69, 4 (1995).
- [286] Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms.
- [287] Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2021. Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection. *arXiv:2111.07997* [cs] (November 2021). Retrieved November 30, 2021 from <http://arxiv.org/abs/2111.07997>
- [288] Piotr Sapiezynski, Wesley Zeng, Ronald E. Robertson, Alan Mislove, and Christo Wilson. 2019. Quantifying the Impact of User Attention on Fair Group Representation in Ranked Lists. *arXiv:1901.10437* [cs] (May 2019). Retrieved December 2, 2021 from <http://arxiv.org/abs/1901.10437>
- [289] Stephen G. Sapp and Wendy J. Harrod. 1993. Reliability and Validity of a Brief Version of Levenson's Locus of Control Scale. *Psychol Rep* 72, 2 (April 1993), 539–550. DOI:<https://doi.org/10.2466/pr0.1993.72.2.539>
- [290] Martin Saveski, Brandon Roy, and Deb Roy. 2021. The Structure of Toxic Conversations on Twitter. *Proceedings of the Web Conference 2021* (April 2021), 1086–1097. DOI:<https://doi.org/10.1145/3442381.3449861>
- [291] Daniel Schiff, Aladdin Ayesh, Laura Musikanski, and John C. Havens. 2020. IEEE 7010: A New Standard for Assessing the Well-being Implications of Artificial Intelligence. (2020). Retrieved from <http://arxiv.org/abs/2005.06620>
- [292] Tobias Schnabel, Saleema Amershi, Paul N. Bennett, Peter Bailey, and Thorsten Joachims. 2020. The Impact of More Transparent Interfaces on

Behavior in Personalized Recommendation. (2020), 991–1000. DOI:<https://doi.org/10.1145/3397271.3401117>

- [293] Shalom H. Schwartz. 2012. An Overview of the Schwartz Theory of Basic Values. *Online Readings in Psychology and Culture* 2, 1 (December 2012). DOI:<https://doi.org/10.9707/2307-0919.1116>
- [294] Amartya Sen. 2017. *Collective choice and social welfare*. Harvard University Press.
- [295] Brittany Seymour, Rebekah Getman, Avinash Saraf, Lily H. Zhang, and Elsbeth Kalendarian. 2015. When Advocacy Obscures Accuracy Online: Digital Pandemics of Public Health Misinformation Through an Antifluoride Case Study. *Am J Public Health* 105, 3 (March 2015), 517–523. DOI:<https://doi.org/10.2105/AJPH.2014.302437>
- [296] Holly B. Shakya and Nicholas A. Christakis. 2017. Association of Facebook Use With Compromised Well-Being: A Longitudinal Study. *Am. J. Epidemiol.* (January 2017), amjepid;kwv189v2. DOI:<https://doi.org/10.1093/aje/kww189>
- [297] Guy Shani, David Heckerman, and Ronen I. Brafman. 2005. An MDP-Based Recommender System. *Journal of Machine Learning Research* 6, 43 (2005), 1265–1295.
- [298] Elizabeth Hansen Shapiro, Michael Sugarman, Fernando Bermejo Media, and Ethan Zuckerman. 2021. New approaches to Platform Data Research. Retrieved from <https://www.netgainpartnership.org/events/2021/2/26/new-approaches-to-platform-data-research>
- [299] Helaine Silverman and D. Fairchild Ruggles. 2007. Cultural Heritage and Human Rights. In *Cultural Heritage and Human Rights*, Helaine Silverman and D. Fairchild Ruggles (eds.). Springer, New York, NY, 3–29. DOI:https://doi.org/10.1007/978-0-387-71313-7_1
- [300] Jesper Simonsen and Toni Robertson. 2012. *Routledge international handbook of participatory design*. Routledge. DOI:<https://doi.org/10.4324/9780203108543>
- [301] Ashudeep Singh, Yoni Halpern, Nithum Thain, Konstantina Christakopoulou, Ed H Chi, Jilin Chen, and Alex Beutel. 2021. Building Healthy Recommendation Sequences for Everyone: A Safe Reinforcement Learning Approach. (2021), 10.
- [302] 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 (KDD '18)*, Association for Computing Machinery, New York, NY, USA, 2219–2228. DOI:<https://doi.org/10.1145/3219819.3220088>
- [303] Spandana Singh. 2020. Regulating Platform Algorithms: Approaches for EU and U.S. Policymakers. Retrieved December 17, 2021 from <http://newamerica.org/oti/briefs/regulating-platform-algorithms/>
- [304] Spandana Singh. 2020. Special Delivery: How Internet Platforms Use Artificial Intelligence to Target and Deliver Ads. New America Foundation. Retrieved January 5, 2022 from <http://newamerica.org/oti/reports/special-delivery/>
- [305] Spandana Singh. 2020. Promoting Fairness, Accountability, and Transparency Around Algorithmic Recommendation Practices. New America. Retrieved November 19, 2021 from <http://newamerica.org/oti/reports/why-am-i-seeing-this/>
- [306] Spandana Singh. 2021. Charting a Path Forward. Retrieved December 20, 2021 from <http://newamerica.org/oti/reports/charting-path-forward/>
- [307] Spandana Singh and Leila Doty. 2021. Cracking Open the Black Box. New America Foundation. Retrieved November 30, 2021 from <http://newamerica.org/oti/reports/cracking-open-the-black-box/>
- [308] Spandana Singh and Leila Doty. 2021. The Transparency Report Tracking Tool: How Internet Platforms Are Reporting on the Enforcement of Their Content Rules. New America. Retrieved October 9, 2021 from <http://newamerica.org/oti/reports/transparency-report-tracking-tool/>
- [309] Eva E. A. Skoe. 2014. Measuring care-based moral development: The Ethic of Care Interview. *Behavioral Development Bulletin* 19, 3 (September 2014), 95–104. DOI:<https://doi.org/10.1037/h0100594>
- [310] Ben Smith. 2021. How TikTok Reads Your Mind. *The New York Times*. Retrieved December 6, 2021 from <https://www.nytimes.com/2021/12/05/business/media/tiktok-algorithm.html>
- [311] Jannick Kirk Sørensen and Jonathon Hutchinson. 2018. Algorithms and Public Service Media. In *Public Service Media in the Networked Society*, 91–106.
- [312] Dayana Spagnuolo, Cesare Bartolini, and Gabriele Lenzini. 2016. Metrics for Transparency. In *Data Privacy Management and Security Assurance*, Giovanni Livraga, Vicenç Torra, Alessandro Aldini, Fabio Martinelli and Neeraj Suri (eds.). Springer International Publishing, Cham, 3–18. DOI:https://doi.org/10.1007/978-3-319-47072-6_1
- [313] Elizabeth A Stanton. 2007. The Human Development Index: A History. *PERI Working Papers* 127, (2007). Retrieved from https://scholarworks.umass.edu/peri_workingpapers/85/
- [314] Michael F. Steger, Bryan J. Dik, and Ryan D. Duffy. 2012. Measuring Meaningful Work: The Work and Meaning Inventory (WAMI). *Journal of Career Assessment* 20, 3 (August 2012), 322–337. DOI:<https://doi.org/10.1177/1069072711436160>
- [315] Charles Steinfield, Nicole B. Ellison, and Cliff Lampe. 2008. Social capital, self-esteem, and use of online social network sites: A longitudinal analysis. *Journal of Applied Developmental Psychology* 29, 6 (November 2008), 434–445. DOI:<https://doi.org/10.1016/j.appdev.2008.07.002>
- [316] Jacquelin van Stekelenburg. 2014. Going all the way: Politicizing, polarizing, and radicalizing identity offline and online. 8, 5 (2014), 540–555. DOI:<https://doi.org/10.1111/soc4.12157>
- [317] James W. Stoutenborough, Scott E. Robinson, and Arnold Vedlitz. 2016. Is “fucking” a new dirty word? The influence of word choice on public views toward natural gas attitudes. *Energy Research & Social Science* 17, (July 2016), 52–58. DOI:<https://doi.org/10.1016/j.erss.2016.04.005>
- [318] Milton E. Strauss and Gregory T. Smith. 2009. Construct Validity: Advances in Theory and Methodology. *Annu Rev Clin Psychol* 5, (April 2009), 1–25. DOI:<https://doi.org/10.1146/annurev.clinpsy.032408.153639>
- [319] Jonathan Stray. 2012. Who should see what when? Three principles for personalized news. Nieman Lab. Retrieved June 15, 2021 from <https://www.niemanlab.org/2012/07/who-should-see-what-when-three-principles-for-personalized-news/>

- [320] Jonathan Stray. 2020. Aligning AI Optimization to Community Well-being. *International Journal of Community Well-Being* 3, (2020), 443–463. DOI:<https://doi.org/10.1007/s42413-020-00086-3>
- [321] Jonathan Stray. 2021. Designing Recommender Systems to Depolarize. *First Monday* (2021).
- [322] Jonathan Stray. 2021. Show me the algorithm: Transparency in recommendation systems. Schwartz Reisman Institute. Retrieved October 4, 2021 from <https://srinstitute.utoronto.ca/news/recommendation-systems-transparency>
- [323] Jonathan Stray, Ravi Iyer, and Helena Puig Larrauri. 2023. The Algorithmic Management of Polarization and Violence on Social Media. *Kingst First Amendment Institute at Columbia University*, New York, NY. Retrieved June 8, 2023 from <https://knightcolumbia.org/content/the-algorithmic-management-of-polarization-and-violence-on-social-media>
- [324] Natalie Jomini Stroud, Ashley Muddiman, and Joshua M Scacco. 2017. Like, recommend, or respect? Altering political behavior in news comment sections. *New Media & Society* 19, 11 (November 2017), 1727–1743. DOI:<https://doi.org/10.1177/1461444816642420>
- [325] Sara Su. 2017. New Test With Related Articles. Facebook Newsroom. Retrieved September 19, 2021 from <https://about.fb.com/news/2017/04/news-feed-fyi-new-test-with-related-articles/>
- [326] Kaveri Subrahmanyam, Stephanie M. Reich, Natalia Waechter, and Guadalupe Espinoza. 2008. Online and offline social networks: Use of social networking sites by emerging adults. *Journal of Applied Developmental Psychology* 29, 6 (November 2008), 420–433. DOI:<https://doi.org/10.1016/j.appdev.2008.07.003>
- [327] J. L. Sullivan and J. E. Transue. 1999. The Psychological Underpinnings of Democracy: A Selective Review of Research on Political Tolerance, Interpersonal Trust, and Social Capital. *Annu. Rev. Psychol.* 50, 1 (February 1999), 625–650. DOI:<https://doi.org/10.1146/annurev.psych.50.1.625>
- [328] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Gradients of Counterfactuals. arXiv:1611.02639 [cs] (2017). Retrieved December 1, 2021 from <http://arxiv.org/abs/1611.02639>
- [329] Nicolas Suzor, Tess Van Geelen, and Sarah Myers West. 2018. Evaluating the legitimacy of platform governance: A review of research and a shared research agenda. *International Communication Gazette* 80, 4 (June 2018), 385–400. DOI:<https://doi.org/10.1177/1748048518757142>
- [330] Nima Taghipour, Ahmad Kardan, and Saeed Shiry Ghidary. 2007. Usage-based web recommendations: a reinforcement learning approach. In *Proceedings of the 2007 ACM conference on Recommender systems (RecSys '07)*, Association for Computing Machinery, New York, NY, USA, 113–120. DOI:<https://doi.org/10.1145/1297231.1297250>
- [331] Ruth Tennant, Louise Hiller, Ruth Fishwick, Stephen Platt, Stephen Joseph, Scott Weich, Jane Parkinson, Jenny Secker, and Sarah Stewart-Brown. 2007. The Warwick-Edinburgh Mental Well-being Scale (WEMWBS): development and UK validation. *Health and Quality of Life Outcomes* 5, 1 (November 2007), 63. DOI:<https://doi.org/10.1186/1477-7525-5-63>
- [332] The Markup. 2020. The Citizen Browser Project—Auditing the Algorithms of Disinformation. Retrieved October 14, 2021 from <https://themarkup.org/citizen-browser>
- [333] Nicolas Thompson. 2018. How Facebook Wants to Improve the Quality of Your News Feed | WIRED. *Wired*. Retrieved February 14, 2022 from <https://www.wired.com/story/how-facebook-wants-to-improve-the-quality-of-your-news-feed/>
- [334] Luke Thorburn, Jonathan Stray, and Priyanjana Bengani. 2022. What Does it Mean to Give Someone What They Want? The Nature of Preferences in Recommender Systems. *Understanding Recommenders*. Retrieved March 25, 2022 from <https://medium.com/understanding-recommenders/what-does-it-mean-to-give-someone-what-they-want-the-nature-of-preferences-in-recommender-systems-82b5a1559157>
- [335] Luke Thorburn, Jonathan Stray, and Priyanjana Bengani. 2022. What Will “Amplification” Mean in Court? *Tech Policy Press*. Retrieved June 30, 2022 from <https://techpolicy.press/what-will-amplification-mean-in-court/>
- [336] Todd Thrash and Andrew Elliot. 2003. Inspiration as a Psychological Construct. *Journal of personality and social psychology* 84, (May 2003), 871–89. DOI:<https://doi.org/10.1037/0022-3514.84.4.871>
- [337] 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:<https://doi.org/10.1007/978-1-4899-7637-6>
- [338] Christian Winther Topp, Søren Dinesen Østergaard, Susan Søndergaard, and Per Bech. 2015. The WHO-5 Well-Being Index: A Systematic Review of the Literature. *Psychother Psychosom* 84, 3 (2015), 167–176. DOI:<https://doi.org/10.1159/000376585>
- [339] Petter Törnberg. 2022. How digital media drive affective polarization through partisan sorting. *Proceedings of the National Academy of Sciences* 119, 42 (October 2022), e2207159119. DOI:<https://doi.org/10.1073/pnas.2207159119>
- [340] UK Office of National Statistics. 2019. Measuring national well-being: domains and measures. Retrieved November 11, 2021 from <https://www.ons.gov.uk/peoplepopulationandcommunity/wellbeing/datasets/measuringnationalwellbeingdomainsandmeasures>
- [341] UN General Assembly. 2015. Transforming our world : the 2030 Agenda for Sustainable Development.
- [342] UNESCO. 1978. Declaration on Fundamental Principles concerning the Contribution of the Mass Media to Strengthening Peace and International Understanding, to the Promotion of Human Rights and to Countering Racism, apartheid and incitement to war. Retrieved December 8, 2021 from http://portal.unesco.org/en/ev.php-URL_ID=13176&URL_DO=DO_TOPIC&URL_SECTION=201.html
- [343] UNESCO. 2005. Convention for the Protection and Promotion of the Diversity of Cultural Expressions. Retrieved December 8, 2021 from <https://en.unesco.org/creativity/convention/texts>
- [344] UNESCO. 2017. Declaration of Ethical Principles in relation to Climate Change. Retrieved December 2, 2021 from http://portal.unesco.org/en/ev.php-URL_ID=49457&URL_DO=DO_TOPIC&URL_SECTION=201.html
- [345] United Nations. 1966. International Covenant on Civil and Political Rights. Retrieved November 2, 2021 from <https://www.ohchr.org/en/professionalinterest/pages/ccpr.aspx>
- [346] United Nations. 1969. Declaration on Social Progress and Development. Retrieved November 2, 2021 from

<https://www.ohchr.org/en/professionalinterest/pages/progressanddevelopment.aspx>

- [347] Université de Montréal. 2018. Montreal Declaration for a Responsible Development of Artificial Intelligence. Retrieved November 2, 2021 from <https://www.montrealdeclaration-responsibleai.com/the-declaration>
- [348] Kristen Vaccaro, Dylan Huang, Motahare Eslami, Christian Sandvig, Kevin Hamilton, and Karrie Karahalios. 2018. The Illusion of Control: Placebo Effects of Control Settings. ACM, New York, NY, USA. DOI:<https://doi.org/10.1145/3173574.3173590>
- [349] Sebastián Valenzuela, Namsu Park, and Kerk F. Kee. 2009. Is There Social Capital in a Social Network Site?: Facebook Use and College Students' Life Satisfaction, Trust, and Participation. *Journal of Computer-Mediated Communication* 14, 4 (July 2009), 875–901. DOI:<https://doi.org/10.1111/j.1083-6101.2009.01474.x>
- [350] Shannon Vallor. 2016. *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. Oxford University Press, New York. DOI:<https://doi.org/10.1093/acprof:oso/9780190498511.001.0001>
- [351] Shannon Vallor, Irina Raicu, and Brian Green. 2020. *Technology and Engineering Practice: Ethical Lenses to Look Through*. The Markkula Center for Applied Ethics at Santa Clara University. Retrieved from https://www.scu.edu/media/ethics-center/technology-ethics/Tech_and_Engineering_Practice-Ethical_Lenses-2020.pdf
- [352] Lav R. Varshney. 2020. Respect for Human Autonomy in Recommender Systems. arXiv:2009.02603 [cs] (September 2020). Retrieved October 14, 2021 from <http://arxiv.org/abs/2009.02603>
- [353] Briana Vecchione, Solon Barocas, and Karen Levy. 2021. Algorithmic Auditing and Social Justice: Lessons from the History of Audit Studies. arXiv:2109.06974 [cs] (September 2021). DOI:<https://doi.org/10.1145/3465416.3483294>
- [354] Philippe Verduyn, Oscar Ybarra, Maxime Résoibois, John Jonides, and Ethan Kross. 2017. Do Social Network Sites Enhance or Undermine Subjective Well-Being? A Critical Review. 11, 1 (January 2017), 274–302. DOI:<https://doi.org/10.1111/sipr.12033>
- [355] Eric S Vorm and Andrew D Miller. 2018. Assessing the Value of Transparency in Recommender Systems: An End-User Perspective. (2018), 8.
- [356] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. 359, 6380 (March 2018), 1146–1151. DOI:<https://doi.org/10.1126/science.aap9559>
- [357] Sanne Vrijenhoek, Mesut Kaya, Nadia Metoui, Judith Möller, Daan Odijk, and Natali Helberger. 2020. Recommenders with a mission: assessing diversity in news recommendations. 1, 1 (December 2020), 554–561. DOI:https://doi.org/10.1007/978-3-030-65965-3_38
- [358] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2017. Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. (2017), 1–52. DOI:<https://doi.org/10.2139/ssrn.3063289>
- [359] Annika Waern. 2004. User Involvement in Automatic Filtering: An Experimental Study. *User Model. User-Adapt. Interact.* 14, (June 2004), 201–237. DOI:<https://doi.org/10.1023/B:USER.0000028984.13876.9b>
- [360] Wall Street Journal. 2021. Inside TikTok's Algorithm: A WSJ Video Investigation. Wall Street Journal. Retrieved October 14, 2021 from <https://www.wsj.com/articles/tiktok-algorithm-video-investigation-11626877477>
- [361] Lequn Wang and Thorsten Joachims. 2021. User Fairness, Item Fairness, and Diversity for Rankings in Two-Sided Markets. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*, 23–41. Retrieved January 2, 2022 from <https://doi.org/10.1145/3471158.3472260>
- [362] Washingtonian. 2019. What Happened After My 13-Year-Old Son Joined the Alt-Right. Washingtonian. Retrieved January 3, 2022 from <https://www.washingtonian.com/2019/05/05/what-happened-after-my-13-year-old-son-joined-the-alt-right/>
- [363] David Watson, Lee Anna, and Auke Tellegen. 1988. Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales. (1988), 8.
- [364] Pamela Weaver, Jeong Choi, and Tammie Kaufman. 1997. Question Wording and Response Bias: Students' Perceptions of Ethical Issues in the Hospitality and Tourism Industry. *Journal of Hospitality & Tourism Education* 9, 2 (April 1997), 21–26. DOI:<https://doi.org/10.1080/10963758.1997.10685307>
- [365] Stephanie Weiser. 2019. Requirements of Trustworthy AI. FUTURIUM - European Commission. Retrieved October 21, 2021 from <https://ec.europa.eu/futurium/en/ai-alliance-consultation/guidelines/1>
- [366] Dan Weld and Gagan Bansal. 2019. The challenge of crafting intelligible intelligence. *Communications of the ACM* 62, 6 (2019), 70–79.
- [367] Judith B. White, Ellen J. Langer, Leeat Yariv, and John C. Welch. 2006. Frequent Social Comparisons and Destructive Emotions and Behaviors: The Dark Side of Social Comparisons. *J Adult Dev* 13, 1 (March 2006), 36–44. DOI:<https://doi.org/10.1007/s10804-006-9005-0>
- [368] Joe Whittaker, Seán Looney, Alastair Reed, and Fabio Votta. 2021. Recommender systems and the amplification of extremist content. *Internet Policy Review* 10, 2 (June 2021). DOI:<https://doi.org/10.14763/2021.2.1565>
- [369] Mark Wilhelm, Ajith Ramanathan, Alexander Bonomo, Sagar Jain, Ed H. Chi, and Jennifer Gillenwater. 2018. Practical Diversified Recommendations on YouTube with Determinantal Point Processes. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*, Association for Computing Machinery, New York, NY, USA, 2165–2173. DOI:<https://doi.org/10.1145/3269206.3272018>
- [370] Ceri Wilson and Jenny Secker. 2015. Validation of the Social Inclusion Scale with Students. *Social Inclusion* 3, 4 (July 2015), 52–62. DOI:<https://doi.org/10.17645/si.v3i4.121>
- [371] World Health Organization. 2004. *Promoting Mental Health*. World Health Organization, Geneva. Retrieved November 15, 2021 from <https://public.ebookcentral.proquest.com/choice/publicfullrecord.aspx?p=4978588>
- [372] World Values Survey. 2020. WVS Wave 7. Retrieved November 11, 2021 from <https://www.worldvaluessurvey.org/WVSDocumentationWV6.jsp>
- [373] Chuhan Wu, Fangzhao Wu, Yang Cao, Yongfeng Huang, and Xing Xie. 2021. FedGNN: Federated Graph Neural Network for Privacy-Preserving

- Recommendation. arXiv:2102.04925 [cs] (March 2021). Retrieved December 8, 2021 from <http://arxiv.org/abs/2102.04925>
- [374] Haolun Wu, Chen Ma, Bhaskar Mitra, Fernando Diaz, and Xue Liu. 2021. Multi-FR: A Multi-Objective Optimization Method for Achieving Two-sided Fairness in E-commerce Recommendation. arXiv:2105.02951 [cs] (May 2021). Retrieved December 2, 2021 from <http://arxiv.org/abs/2105.02951>
- [375] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint Multisided Exposure Fairness for Recommendation. (2022).
- [376] Yikun Xian, Tong Zhao, Jin Li, Jim Chan, Andrey Kan, Jun Ma, Xin Luna Dong, Christos Faloutsos, George Karypis, S. Muthukrishnan, and Yongfeng Zhang. 2021. EX3: Explainable Attribute-aware Item-set Recommendations. In Fifteenth ACM Conference on Recommender Systems (RecSys '21), Association for Computing Machinery, New York, NY, USA, 484–494. DOI:<https://doi.org/10.1145/3460231.3474240>
- [377] Hao Yan, Ellen E. Fitzsimmons-Craft, Micah Goodman, Melissa Krauss, Sanmay Das, and Patricia Cavazos-Rehg. 2019. Automatic detection of eating disorder-related social media posts that could benefit from a mental health intervention. *Int J Eat Disord* 52, 10 (October 2019), 1150–1156. DOI:<https://doi.org/10.1002/eat.23148>
- [378] Moran Yarchi, Christian Baden, and Neta Kliger-Vilenchik. 2021. Political Polarization on the Digital Sphere: A Cross-platform, Over-time Analysis of Interactional, Positional, and Affective Polarization on Social Media. *Political Communication* 38, 1–2 (March 2021), 98–139. DOI:<https://doi.org/10.1080/10584609.2020.1785067>
- [379] Xing Yi, Liangjie Hong, Erheng Zhong, Nathan Nan Liu, and Suju Rajan. 2014. Beyond clicks: Dwell time for personalization. (2014), 113–120. DOI:<https://doi.org/10.1145/2645710.2645724>
- [380] Jillian York and Ethan Zuckerman. 2019. Moderating the Public Sphere. In Rikke Frank Jørgensen (ed.). MIT Press.
- [381] Yang Yu, Sen Luo, Shenglan Liu, Hong Qiao, Yang Liu, and Lin Feng. 2020. Deep attention based music genre classification. *Neurocomputing* 372, (January 2020), 84–91. DOI:<https://doi.org/10.1016/j.neucom.2019.09.054>
- [382] Hamed Zamani, Markus Schedl, Paul Lamere, and Ching-Wei Chen. 2019. An Analysis of Approaches Taken in the ACM RecSys Challenge 2018 for Automatic Music Playlist Continuation. arXiv:1810.01520 [cs] (August 2019). Retrieved December 8, 2021 from <http://arxiv.org/abs/1810.01520>
- [383] Meike Zehlke, Ke Yang, and Julia Stoyanovich. 2021. Fairness in Ranking: A Survey. arXiv:2103.14000 [cs] (May 2021). Retrieved December 2, 2021 from <http://arxiv.org/abs/2103.14000>
- [384] Meike Zehlke, Ke Yang, and Julia Stoyanovich. 2022. Fairness in Ranking, Part I: Score-based Ranking. *ACM Comput. Surv.* (April 2022). DOI:<https://doi.org/10.1145/3533379>
- [385] Amy X. Zhang, Grant Hugh, and Michael S. Bernstein. 2020. PolicyKit: Building governance in online communities. (2020), 365–378. DOI:<https://doi.org/10.1145/3379337.3415858>
- [386] Yongfeng Zhang and Xu Chen. 2020. Explainable Recommendation: A Survey and New Perspectives. 14, 1 (2020), 1–101. DOI:<https://doi.org/10.1561/15000000066.Yongfeng>
- [387] Xiangyu Zhao, Changsheng Gu, Haoshenglu Zhang, Xiwang Yang, Xiaobing Liu, Jiliang Tang, and Hui Liu. 2021. DEAR: Deep Reinforcement Learning for Online Advertising Impression in Recommender Systems. arXiv:1909.03602 [cs] (May 2021). Retrieved April 24, 2022 from <http://arxiv.org/abs/1909.03602>
- [388] Xiaoxue Zhao, Weinan Zhang, and Jun Wang. 2013. Interactive collaborative filtering. In Proceedings of the 22nd ACM international conference on Information & Knowledge Management (CIKM '13), Association for Computing Machinery, New York, NY, USA, 1411–1420. DOI:<https://doi.org/10.1145/2505515.2505690>
- [389] Zhe Zhao, Lichan Hong, Li Wei, Jilin Chen, Aniruddh Nath, Shawn Andrews, Aditee Kumthekar, Maheswaran Sathiamoorthy, Xinyang Yi, and Ed Chi. 2019. Recommending what video to watch next. ACM, New York, NY, USA, 43–51. DOI:<https://doi.org/10.1145/3298689.3346997>
- [390] Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. ACM Press, New York, New York, USA, 22. DOI:<https://doi.org/10.1145/1060745.1060754>
- [391] Gregory D. Zimet, Nancy W. Dahlem, Sara G. Zimet, and Gordon K. Farley. 1988. The Multidimensional Scale of Perceived Social Support. *Journal of Personality Assessment* 52, 1 (March 1988), 30–41. DOI:https://doi.org/10.1207/s15327752jpa5201_2
- [392] Frederik J Zuiderveen Borgesius, Damian Trilling, Judith Möller, Balázs Bodó, Claes H. de Vreese, and Natali Helberger. 2016. Should we worry about filter bubbles? 5, 1 (2016), 1–16. DOI:<https://doi.org/10.14763/2016.1.401>

Appendix A: Table of Values

This list was culled from a wide variety of sources and refined through multi-stakeholder consultation as described above. It includes perspectives from various traditions and cultures, but is not intended to be a comprehensive list of the values that might be important to consider in recommender systems. Nor do we attempt to prioritize or rank these values here. Rather, this table should be viewed as one list of broadly important themes.

For each value, we provide possible interpretations in the context of recommender systems, some indicators that might be used to assess whether a system enacts or supports that value, and example designs that are relevant to that value.

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
1	Usefulness	<p>“The purpose of recommenders is often summarized as ‘help the users find relevant items’” [167]</p> <p>A recommender should provide a useful service to the user.</p>	<p>Long term user retention.</p> <p>User satisfaction surveys.</p>	<p>Show people more of what they rate highly.</p>
2	Liberty	<p>A platform should respect human dignity and protect human rights and freedoms [116:61]; recommenders should operationalize respect for autonomy [352].</p> <p>A platform should not stop users from exploring certain types of information. Users should be able to pursue their own good in their own way.</p>	<p>Similar mechanisms as Freedom of Expression.</p>	<p>Protections against arbitrary removal of content.</p> <p>Eliminating threats of violence designed to compel individuals to form particular opinions, in violation of article 19 in Universal Declaration of Human Rights.</p>
3	Freedom of Expression	<p>“Everyone shall have the right to freedom of expression; this right shall include freedom to seek, receive and impart information and ideas of all kinds” [345 art .19]</p> <p>“The exercise of freedom of opinion, expression and information ... is a vital factor in the strengthening of peace and international understanding” [342].</p> <p>Platforms should not stop users from expressing their thoughts and opinions, or unduly suppress distribution of user posts.</p>	<p>“Conducting an ex-ante evaluation attempts to predict the future relationship between [human] rights and an on-going or proposed business activity.” [197]</p>	<p>Transparency around content which is suppressed.</p> <p>Appeals processes for removed content.</p>
4	Control	<p>Control “includes the system’s ability to allow users to revise their preferences, to customize received recommendations, and to request a new set of recommendations” [267:18]</p> <p>“Ability to determine the nature, sequence and/or consequences of technical and</p>	<p>“I feel in control of my news feed.” [348]</p> <p>Locus of control Scale.</p> <p>Ex: “Other people usually control my life” [289]</p>	<p>Recommender controls [143,147,292]</p>

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		<p>operational settings, behavior, specific events, and/or experiences.” [162:17]</p> <p>The platform should give users ways to control the content selected for them and it should give users ways to control how the content they create is shared.</p>		
5	Agency, Autonomy, Efficacy	<p>“The sense of agency can be analyzed as a compound of more basic experiences, including the experience of intentional causation, the sense of initiation and the sense of control” [257]</p> <p>The platform should provide the capacity for intentional action and help users achieve their goals [352]</p> <p>The platform should not manipulate users for the benefit of other stakeholders.</p>	<p>General Self-Efficacy Scale (GSE) Ex: “It is easy for me to stick to my aims and accomplish my goals” [63]</p> <p>Autonomy Scale (AS): Measures components of autonomy including family loyalty autonomy, value autonomy, emotional autonomy, and behavioral autonomy [10]</p>	Recommender controls [143,147,292]
6	Privacy	<p>“The protection of personal data is of fundamental importance to a person’s enjoyment of his or her right to respect for private and family life” [73]</p> <p>Users should be allowed to determine if and how their personal data is collected, processed and disseminated [116:21,162:71]</p>	GDPR Data Protection Impact Assessment [103]	Controls to limit data use. Federated recommendation algorithms [268,373].
7	Safety, Security	<p>The principle of ‘safety’ requires that an AI system be reliable and that ‘the system will do what it is supposed to do without harming living beings or [its] environment.’” [116:38]</p> <p>“[System] use should not contribute to increasing</p>	Sense that most people can be trusted [372]	Hate speech moderation Parental controls Option to report harmful content.

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		stress, anxiety, or a sense of being harassed by one’s digital environment.” [347]		
8	Self-expression, authenticity	<p>“Self-expression is a notion that is closely associated with a horde of positive concepts, such as freedom, creativity, style, courage, self-assurance, and even healing and spirituality” [186]</p> <p>The platform should empower users to express their identity (including personality, attributes, behavior) and to decide how it is presented to others. [183:19].</p>	Correlates with (but is not identical to) quantity and quality of user content creation	<p>Easy-to-use, attractive content creation/modification tools, e.g. SnapChat filters, TikTok tools.</p> <p>Content moderation policies that allow self-expression (e.g. don’t take down art because it has nudity)</p>
9	Well-being	<p>Well-being is a complex and multidimensional construct that encompasses many of the values we describe here and there is no universally accepted definition [88]. Well-being may be measured by both objective and subjective factors including satisfaction with life, emotions and affect, psychological well-being, social relationships, and meaning [205]</p> <p>On the platforms, users should see content that leads them to experience contentment, joy and pleasure, both ephemerally and over the long term. Platforms should help users feel satisfied with their lives.</p>	<p>Satisfaction with Life Scale (SWLS). Ex: “The conditions of my life are excellent” [87]</p> <p>Positive Affect and Negative Affect Scales (PANAS). Ex: “Indicate the extent to which you have felt interested, distressed, excited, upset, etc over the past week” [363]</p> <p>Scale of Psychological Wellbeing (SPWB). Assesses psychological functioning across domains such as autonomy, self-acceptance, purpose in life, and positive relationships [285]</p>	<p>Encourage active vs. passive use of social media [354]</p> <p>Usage controls, e.g. screen time limits</p>
10	Inspiration, Awe	<p>“In times of uncertainty, others are sought for guidance, inspiration or motivation, to seek ideas, goals or possibilities, which can influence ambitions,</p>	<p>Inspiration scale: Measures the experience of being inspired, feeling inspired and motivation to do something with</p>	<p>Uprank content identified by Inspiration classifier [163]</p>

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		choices, and achievements.” [183:11] Platforms should show users content that will inspire, motivate, or guide them.	that inspiration [336]	
11	Mental Health	Platforms should help users protect and improve their mental health. Platforms should not encourage unhealthy types or amounts of use. “[A] state of well-being in which the individual realizes his or her own abilities, can cope with the normal stresses of life, can work productively and fruitfully, and is able to make a contribution to his or her community” [371]	Warwick-Edinburgh Mental Well-Being Scale (WEMWBS) [331] Positive Mental Health Scale (PMH) Ex: “I feel that I am actually well equipped to deal with life and it’s difficulties” [211]	Features that help users manage their screen time, e.g. “take a break” notifications, screen time limits. Option to report harassment or bullying.
12	Physical Health	Platforms should help users protect and improve their physical health, such as by providing accurate health information and encouraging behaviors that contribute to physical health.	WHO Wellbeing Index (WHO-5) Ex: “I have felt active and vigorous” and “I woke up feeling fresh and rested” [338]	Help users develop healthy habits (exercise, diet, etc.) Clear platform policies on removing health misinformation, ads for fake cures, etc.
13	Self-actualization, Personal Growth	Platforms should enable users to reach their full potential. “Personal growth is a continuous improvement in life in order to find purpose and meaning” [183:22] “[R]ather than have people choose the easiest option, we wish to have them develop a strong sense of determination of having selected the right path” [187]	“Feeling that the things one does are worthwhile” [340] European Social Survey (ESS). Ex: “To what extent [do you] learn new things in life?” [247]	Educational recommender systems [81]
14	Recognition, Acknowledgment	Platforms should provide ways for other people to recognize a user’s contributions or worth. “Whilst some esteem needs can be met by having self-	Social Inclusion Scale (SIS). Ex: ‘I have felt accepted by my friends’ [370]	“Celebrate” reaction button (as LinkedIn has)

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		acceptance, self-worth and self-value, validation from others is also important" [183:28]		
15	Knowledge, Informativeness	Users should see items that keep them informed about topics they care about or might care about. "[P]ersonalised news recommendations allow the media not only to help users find relevant information, but also to inform them better and more effectively." [148:995] "Curiosity drives individuals to seek stimulation, information or new experiences, serving a purpose to increase knowledge and build skills." [183:18]	News knowledge quizzes [6,12] Curiosity and exploration inventory (CEI). Ex: "I would describe myself as someone who actively seeks as much information as I can in a new situation" [180]	News recommender systems [179] Trending topics
16	Connection	"Individuals are driven to interact and seek social closeness ... The benefits of social connections are far-reaching." [183:17]	Perceived Social Support scale (PSS). Ex: "I have friends with whom I can share my joys and sorrows" and "I can talk about my problems with my friends" [391] Online Social Support Scale (OSSS). Ex: "Online, I belong to groups of people with similar interests" and "Online, people make me feel like I belong" [245]	Recommendations for people and groups with similar interests [59]
17	Civic Engagement	"Working to make a difference in the civic life of one's community and developing the combination of knowledge, skills, values and motivation to make that difference." [97:vi]	Voter turnout [252] Attendance of peaceful demonstrations [372] Approximate total hours a month active in voluntary organizations [372]	Encourage people to vote [38]

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		“Creating a public forum and optimal conditions for engagement” [148:1005]	Donations to a charity in a month [242]	
18	Community, Belonging	<p>“[T]he feeling that people matter to one another and to the group, and a shared faith that members’ needs will be met through their commitment to be together.” [219]</p> <p>“The concept of members sharing mutual concern and/or love for one another” [183:15].</p>	<p>Sense of belonging to a neighborhood [340]</p> <p>The General Sense of Belongingness Scale [212]. Ex: “I have a place at the table with others”</p>	<p>Recommend community-oriented groups</p> <p>User affordances for donating to charitable causes.</p>
19	Accessibility, Inclusiveness	Inclusiveness in design means that diverse people and perspectives should be involved in the design of AI systems. Inclusiveness in impact calls for a just distribution of AI benefits and harms.[116:51–53]	Assessing the experience of different user demographics through surveys and other methods, such as visually impaired or hard of hearing users [31,233]	<p>Implement Web Content and Accessibility Guidelines (WCAG)</p> <p>Variety of formats for cognitive impairments (e.g. screen readers, transcripts, captions)</p>
20	Tolerance, Constructive Discourse	<p>Tolerance “creates the opportunity for a wide range of political groups to express their ideas and to participate in public life” [327]</p> <p>Polarization divides society into “us” and “them” camps and contributes to the erosion of democracy [218]</p>	Polarization measures such as issue-position or affective polarization [43,256,378]	Increase ideological diversity, select for civil conversations, optimize polarization measures [321]
21	Duty	The notion of duty and obligation is defined in contrast to self-interest [158].	Personal control and responsibility scale [220]	User affordances for donating to charitable causes.
22	Care, Compassion, Empathy	“The ethic of care emphasizes the importance of context, interdependence, relationships, and responsibilities to concrete others.” [188] and is “based on the recognition that self and others are interconnected” [309].	Empathy quotient [22]. Ex: “I find it easy to put myself in somebody else’s shoes” and “Seeing people cry doesn’t really upset me”	Best practices in suicide prevention [278]

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		These notions also appear in Confucian and Buddhist traditions [350:132] and in Ubuntu values where “a person is ‘a person through others’” [243].		
23	Fairness, Equality, Equity	“A morally justifiable distribution of benefits, goods, harms, and risks” [351]. In multi stakeholder fairness, different stakeholders have different objectives and needs from the system [1].	Fairness metrics in classification and recommendation [34,271]	Add fairness measures to the recommender training or serving objectives [34,222]
24	Diversity	“[D]iversity refers to the idea that in a democratic society, informed citizens construct their worldview from a diverse set of sources which helps them to make balanced and well-considered decisions.” [150:192]. “‘Cultural diversity’ refers to the manifold ways in which the cultures of groups and societies find expression.” [343 art. 4]	Slate and feed diversity metrics of various kinds (Kunaver and Požrl, 2017) Perceived diversity and user satisfaction (Kunaver and Požrl, 2017)	Methods to increase various types of diversity, including topical and source diversity, novelty, coverage, etc. [60,194]
25	Accountability	Accountability is “a component of the state of being responsible, alongside being answerable and being attributable” [215] AI principles documents include “verifiability and replicability,” “impact assessments,” “environmental responsibility,” “evaluation and auditing requirements.” [116:28]	Ranking Digital Rights Corporate Accountability Index [273]	Explanations of decisions according to clearly defined principles. Processes for users to bring up problems and resolve them in a timely manner.
26	Transparency & Explainability	Transparency means that information provided about a system is a) meaningful, b) useful, c) accessible, d) comprehensive, and e) truthful [162:71]	System cards which provide stakeholders with an overview of different components of an ML system [4]	Explanations for recommendations, answering questions such as “why am I seeing this?” or “Why has this been removed?”

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		<p>“Transparency is instrumental to uphold intrinsic values of human autonomy and justice” [56:20] Desirable properties of transparency metrics [312] “[162:49,312]</p>		Publicly disclose and discuss changes to moderation or ranking [322].
27	Accuracy (Factuality)	<p>Accuracy is “an AI system’s ability to make correct judgements, for example to correctly classify information into the proper categories, or its ability to make correct predictions, recommendations, or decisions based on data or models” [365] AI systems should make “accurate information readily available” [133] and should not spread “untrustworthy information.” [347]</p>	Credibility metrics including prevalence of misinformation items among recommendations, and user engagement on this material.	Misinformation classifiers [45] Prompts to consider accuracy before sharing [262].
28	Tradition, History	<p>“Heritage is a concept to which most people would assign a positive value. The preservation of material culture ... and intangible culture ... are generally regarded as a shared common good by which everyone benefits.” [299] Practices, symbols, ideas, and beliefs that are developed by groups and represent their shared experience and fate. They often take the form of religious rites, beliefs, and norms of behavior [293:6]. Rituals or conventional norms are one of the five central virtues of Confucianism [77]</p>	Percentage of locally produced content.	Promotion of cultural events [62]
29	Environmental Sustainability	“Respect for environment and natural habitat, efficiency, maintainability, operability, supportability, reliability,	Product level carbon footprints [224]	Recommender systems can potentially direct users toward climate-friendly options [279:25]

	Value	Interpretation in the context of recommenders	Example Indicator	Example design
		<p>durability, resilience, forgiveness, robustness, redundancy, reusability, reconfigurability, simplicity, economy, renewability” [162:70]</p> <p>“Develop and scale up carefully assessed technologies, infrastructure and actions that reduce climate change and its associated risks.” [344]</p>		Reporting on the energy used by ML models [152]
30	Progress	The International Covenant on Civil and Political rights includes the right of people to “freely pursue their economic, social and cultural development” [345] as expanded in the Declaration on Social Progress and Development [346].	UN Human Development Index [313]	To the best of our knowledge, none
31	Labor	Meaningful work is both significant and positive in valence [281]. Job insecurity, low employability, and unemployment are all detrimental to well-being, beyond the effects of income loss [136]	Work and Meaning Inventory [314]	Ensure users see recommendations for jobs that they would want.