Why Do Recommenders Recommend? Three Waves of Research Perspectives on Recommender Systems

Annelien Smets[1,*], Christine Bauer[2]

[1]imec-SMIT, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium
[2]Department of Artificial Intelligence and Human Interfaces, University of Salzburg, Jakob-Haringer-Str 1, 5020 Salzburg, Austria

Abstract

Research on Recommender Systems (RSs) has evolved across disciplines, from a focus on technical optimization to a broader interest in these systems' societal role and impact. In this paper, we distinguish between two broad research orientations—technical and social—and identify three research waves that reflect shifting assumptions and research questions of both orientations. We discuss how assumptions within each research orientation have influenced the other over time, shaped by RSs becoming more deeply embedded in real-world, multi-sided applications. We argue that this evolving interplay has led to growing convergence around the central question: why do recommenders recommend? Addressing this question requires perspectives that span both technical and social domains, underscoring the importance of interdisciplinary collaboration. By charting this research evolution, this paper aims to support interdisciplinary exchange and collaboration in the field. It encourages researchers to explicitly acknowledge and revisit the assumptions that drive their research, and identifies which research questions arise more naturally and which may require additional effort.

Keywords

social orientation, technical orientation, social sciences, computer sciences, research field evolution

1. Introduction

The ever-increasing popularity of Recommender Systems (RSs) is commonly attributed to two primary value propositions. First, RSs assist end-users in managing information overload. Second, platform providers can leverage these systems to achieve business objectives such as increased sales, heightened engagement, and enhanced user retention [1, 2]. Although these value propositions are sometimes perceived as "diametrically opposed" [1], they share a commonality: the intended outcomes of RSs are framed in terms of their social implications, while achieving these goals often involves technical aspects, such as optimizing recommender algorithms. Over time, the relative emphasis on technical versus social aspects in RSs research has shifted. While much of the field has traditionally focused on technical developments, growing attention to social dimensions reflects RSs' increasing integration into everyday life. This shift has spurred research questioning the role and impact of RSs in domains such as online news [3, 4], creative industries [5, 6], tourism [7, 8], recruitment [9, 10], e-commerce [11, 12], health [13, 14], and many more. This growing societal relevance is also reflected in recent policy efforts, such as the EU's Digital Services Act, which explicitly addresses systemic risks linked to RSs [15].

These developments underscore the need to understand RSs not only as technical artifacts, but as systems with significant social consequences. As such, understanding RSs' real-world impact calls for approaches that move beyond disciplinary boundaries, combining technical and social research orientations [16, 17]. However, establishing such an interdisciplinary research agenda is challenging, as it requires a shared understanding of the diverse research traditions and assumptions.

In this paper, we distinguish between two broad orientations in RSs research: a *technical* and *social* orientation, encompassing disciplines such as computer science, psychology, media studies, economics,

Beyond Algorithms: Reclaiming the Interdisciplinary Roots of Recommender Systems Workshop (BEYOND 2025), September 26th, 2025, co-located with the 19th ACM Recommender Systems Conference, Prague, Czech Republic.
*Corresponding author.
annelien.smets@vub.be (A. Smets); christine.bauer@plus.ac.at (C. Bauer)
https://www.anneliensmets.be (A. Smets); https://christinebauer.eu (C. Bauer)
0000-0003-4771-7159 (A. Smets); 0000-0001-5724-1137 (C. Bauer)

and political science[1]. This distinction provides a useful lens to examine how different disciplinary assumptions have informed the questions, methods, and interpretations that shape RSs research. Rather than offering an exhaustive review, the paper identifies recurring patterns of how different orientations have approached RSs over time, structured along three broad research waves. The discussion sheds light on how these research perspectives have influenced one another, often implicitly, and how they have co-evolved with how RSs are deployed in real-world contexts. In doing so, the paper provides a basis for a deliberate reflection on how interdisciplinary research in this field can be framed and advanced.

2. Research Perspectives on Recommender Systems

Through an interpretative reading of how RSs research has evolved across technical and social orientations, we observe three broad waves of research. We use the term *waves* [2] deliberately: rather than clearly demarcated phases, they constitute overlapping currents of thought that continue to evolve and interact. One strand of research, characterized by a strong **technical orientation**, focuses primarily on optimizing RSs to narrowly defined performance metrics. Central to this line of work is the question: *what is a good recommendation*, a question that has led to diverse answers across different research perspectives. In contrast, research with a **social orientation** emphasizes the need to understand the role and impact of RSs. This line of work raises questions such as *what do* recommenders recommend, *what should* recommenders recommend, and *why do* recommenders recommend.

Table 1 presents an overview of the three research waves, indicating the interplay between RSs in real-world applications and the emerging research perspectives.

Table 1
Overview of three broad waves in RSs research, each associated with real-world application contexts and underlying research perspectives. The technical orientation outlines the evolving assumptions about what constitutes a good recommendation. The social orientation presents key questions that guide research regarding RSs' role and impact.

Research Wave	Applications	Technical Orientation *What is a good recommendation?*	Social Orientation *Understanding the role and impact of RSs*
Accuracy	Small pilots and first commercial applications	An accurate recommendation	What do recommenders recommend?
Beyond-Accuracy	Large-scale commercial applications	Depends on the user's context	What should recommenders recommend?
Multistakeholder	Embedded in multi-sided markets	Depends on the stakeholder	Why do recommenders recommend?

2.1. Early Days

In 1979, computer scientist Elaine Rich described the first RS in her doctoral research [18]. This system, known as 'Grundy the computer librarian', grouped users according to stereotypes and used this information to recommend relevant books. By the 1990s, RSs had emerged as a distinct research field, making it relatively new compared to the study of other information systems [19, 20].

During this first wave, taking a **technical orientation**, the primary focus was set on developing algorithmic techniques to improve recommendation accuracy; more specifically, addressing the question "how closely the recommender's predicted ratings are to the users' true ratings" [21]. The guiding

[1]While this distinction often aligns with disciplinary boundaries (i.e., technical research typically emerging from computer science and social research from the social sciences), we acknowledge that the boundary is not clear-cut. We aim to capture broad patterns rather than impose rigid categories.
[2]In this work, *waves* refers to epistemological shifts in research, not to commercial adoption cycles as in Martin et al. [1].

assumption was that ***the best recommendation is the most accurate one***. This view aligned well with 'Grundy the computer librarian' and the first applications developed in the 1990s, such as GroupLens [22], Tapestry [23], or Siteseer [24]. As a result, this research was characterized by a dominant single-user-focused paradigm, where RSs aimed to optimize outcomes for individual end-users—often based on assumed preferences rather than involving end-users directly in the research process. In this paradigm, systems attempt to predict which items end-users will like most and rank them accordingly, thereby assisting end-users in finding what they are looking for. It has been argued that this accuracy-oriented approach has led to a "hunt for algorithmic improvements" [25] with researchers abstracting away from domain-specific features and adopting a narrow focus on quantitative measures—particularly accuracy—without considering the broader impact of RSs on consumers, businesses, and society [26].

In **socially-oriented research**, the latter was the primary focus of the initial research strand on RSs that questioned ***what recommenders recommend***. For example, in economics and e-commerce, scholars studied the impact of algorithmic recommendations on sales diversity and concentration [27]. Hypothesizing that RSs reduce consumer search costs, empirical results suggest that online channels exhibit a significantly less concentrated sales distribution—implying that RSs may recommend items from the long tail of the popularity curve. Another extensively studied domain is online media, where researchers have investigated how algorithmic recommendations influence user exposure to content. Interest in this topic surged following concerns regarding the filter bubble problem, popularized by Pariser [28], assuming that RSs would mainly expose users to content resembling their prior consumption. Despite sustained criticism and conflicting empirical findings, this hypothesis remains central to many socially-oriented research inquiries to this day [4].

A recurring pattern in these lines of study is that socially-oriented research is often grounded in the assumption that RSs optimize for accuracy, and hence, are likely to recommend items the user has interacted with in the past, or steer users toward items other users have interacted with (e.g., popular items). While it is crucial to acknowledge the discrepancy between "general tendencies of algorithm families" and "what recommenders recommend" in practice [29], it is not surprising that most socially-oriented research has operated on assumptions about these 'general tendencies'. Social scientists have frequently criticized the "black-box nature" [30] of these algorithmic systems and therefore (had to) base their assumptions on the available literature at that time, which was highly dominated by the accuracy paradigm in research with a technical orientation. For example, academic literature has documented that content-based RSs are inclined to over-specialize and recommend items similar to those users have already interacted with [31] or collaborative filtering approaches to over-promote popular items [32].

This first wave, thus, reveals a notable tension when these two research orientations are considered together. Technically-oriented research was frequently seeking domain-agnostic approaches, representing an "academic abstraction" that neither sufficiently accounts for real-world implications nor accurately reflects real-world RS implementations [32, 26, 29]. Moreover, although it adopts a single-user-focused perspective, it typically relies on a narrow conception of the end-user: as a predictable individual whose preferences can be inferred from past behavior, without involving the user directly [33, 34]. However, it is precisely this paradigm that has informed many foundational hypotheses and assumptions in this first wave of socially-oriented research. By implicitly adopting technical framings of what RSs do and who the end-user is, this line of work risks building its critique on the same narrow premises it seeks to challenge. While a detailed discussion goes beyond the scope of this paper, it could be argued that this may (in part) account for the lack of empirical evidence supporting numerous socially-oriented research hypotheses about what recommenders recommend (e.g., the filter bubble hypothesis).

2.2. Large-Scale Commercial Applications

In the late 1990s to early 2000s, RSs experienced growing adoption in real-world (commercial) applications [35, 1]. In 1998, Amazon.com was the first to implement item-to-item collaborative filtering at a scale of millions of customers and catalog items [36], establishing itself as a pioneer in recognizing the commercial value of RSs. In 2006, the one-million-dollar Netflix Prize marked another significant milestone, cementing the role of RSs for commercial success.

In that period, a second research wave with a **technical orientation** began to move beyond the dominant accuracy paradigm. In 2006, McNee et al. [37] famously stated that "being accurate is not enough". They argued that the most accurate recommendations are not necessarily the most useful to end-users, and other quality indicators should also be considered. For example, recommending movies someone has already watched may be technically accurate, but it fails to consider that people generally do not wish to watch the same movie repeatedly. Consequently, these recommendations hold minimal value for the end-user, potentially leading to decreased user satisfaction with the service over time.

This has spurred a whole strand of research on "beyond-accuracy" objectives such as diversity, coverage, novelty, and serendipity [21]. These objectives better reflect the quality of the recommendations for the end-user. Thus, this research strand remains a highly single-user-focused paradigm. The main difference from the accuracy-focused paradigm (first wave) is that the user's context is now considered more relevant. For instance, while recommending a movie the user has already watched may be an accurate preference prediction, it is unhelpful [38, 26]; in other contexts, such as recurring online purchases, in contrast, recommending previously purchased items might align perfectly with the user's expectations [26]. This shift toward a more nuanced view reflects a growing recognition that ***the best recommendation depends on the user's context***. As a result, a greater focus on user-centered evaluation methods has emerged, involving users more directly in the research process and drawing partly on social science approaches to better understand user needs and behavior [39].

In **socially-oriented research**, shortly after the first surge in research assessing what recommenders recommend, a strong call echoed the importance of these beyond-accuracy objectives. For example, in the (news) media domain, a line of work emerged proposing "alternative recommender strategies" that reflect established media and democratic theories with a strong emphasis on beyond-accuracy metrics, particularly diversity [40, 41, 42]. By doing so, this line of work closely engages with policy and regulatory concerns, for example, by addressing ongoing discussions on network regulation and cultural policy for audiovisual content [43]. In other words, next to the descriptive line of research examining *what* recommenders recommend, this emerging line of academic work takes a more normative, often ethically grounded angle, discussing ***what recommenders should recommend***.

In this second wave, we observe both technical and socially-oriented research adopting a broader perspective on the recommendation problem, realizing that it is not just about predicting what items users will like the most. However, both fields continue to rely on a highly single-user-focused perspective, which neglects an important aspect: most RSs operate within multi-sided markets.

2.3. Embedding in Multi-Sided Markets

The peak of the platform economy (mid-2010s) marks a new paradigm shift, emphasizing that RSs operate in multi-sided markets; thus, more stakeholders are involved than solely the end-user [2, 44]. This period has been characterized by companies such as Booking.com, Facebook (now Meta), and Spotify, acquiring large market shares. These companies have in common that they act as an intermediary (or 'platform') between two or more stakeholder groups, facilitating interactions, transactions, or exchanges between them [45]. These groups may include stakeholders like platform owners, content providers, advertisers, and end-users. This configuration introduces new complexities in the design and evaluation of RSs. For example, the optimal outcome for the end-user may be at odds with the interests of other stakeholders [46], whose influence may be more decisive in the final design decisions.

As previously stated, the prevailing single-user-focused perspective in research with a **technical orientation** has faced criticism by scholars arguing that it represents an "academic abstraction" that fails to capture the complexity of real-world RSs in multi-sided markets [32, 47, 26]. In line with this critique, RSs have been redefined as a "multistakeholder environment" that moves beyond the dominant single-user-focused approach [44]. This paradigm considers the utility of all stakeholders involved in the recommendation process, defined as "any group or individual that can affect, or is affected by, the delivery of recommendations to users" [32]. This shift was largely inspired by social sciences, where economists have long discussed this shift from one-sided to multi-sided markets [48, 49]. This multistakeholder paradigm explicitly accounts for the various 'purposes' of a RS [2], such as showing relevant items to end-users while simultaneously increasing sales or time spent on a platform, which is of interest to platform owners. Consequently, in this paradigm, what constitutes ***the best recommendation depends on which stakeholder you ask***.

In real-world scenarios, balancing these varied interests poses computational challenges that researchers actively address through ongoing research efforts. For instance, researchers explore the use of contextual bandits in a multi-objective framework to drive recommendations in multistakeholder platforms [50, 51]. Moreover, scholars increasingly call for more extensive evaluation in RSs research, both in terms of evaluation methods (e.g., multi-method evaluations) and more informative metrics that capture the utility for diverse stakeholders and account for the different purposes of RSs [26, 52, 46]. This also includes considering RSs' long-term impacts and indirect effects, such as bias or fairness [9, 25].

In **socially-oriented research**, this platform pivot has led scholars to turn to the political economy of RSs to ask questions about which factors influence the development, operation, and impact of these systems in (online) markets [53]. This includes analyzing the power dynamics between the different RS stakeholders, as well as the implications for competition, market structure, and societal outcomes. Acknowledging the political economy of RSs increasingly informs the research hypotheses about the 'logics' of the RSs themselves, previously perceived as 'black boxes' [54]. This is evidenced by the fact that, in contrast to previous paradigms, socially-oriented research appears to be moving away from relying on these 'general tendencies' of algorithms, such as over-specialization. Instead, they question whether platforms' business models might stimulate certain tendencies. For example, a body of research examines how platforms engage in self-preferencing behavior [55]. For instance, empirical evidence suggests that products by Amazon's own brand receive significantly more 'frequently-bought-together' recommendations than third-party products on Amazon.com [56]. This type of self-preferencing behavior is considered an unfair practice under Europe's Digital Markets Act (DMA), which targets "gatekeeper platforms" with significant market power [57]. Although the DMA does not explicitly address RSs, its provisions on fairness, transparency, and competition could influence their deployment.

For socially-oriented research, this implies that the question is no longer just about understanding *what* recommenders recommend or what they *should* recommend, but also ***why do recommenders recommend***. This central question aligns closely with inquiries into the 'purpose' of RSs in technically-oriented research [2]. It inherently involves ethical considerations (e.g., fairness) and examines *for whom* recommendations are made (i.e., which stakeholder groups benefit or do not benefit) and with what intended purpose. While technically-oriented researchers focus on developing *metrics* and *techniques* to better capture, balance and evaluate these purposes, social scientists are primarily concerned with *understanding* these purposes and their relationship to platform ecosystems.

In this third wave, given the complexity of the systems studied and the markets they are embedded in, it is not surprising that we observe a greater convergence of these research orientations. However, this also means that researchers must be even more explicit about their assumptions, as these may cause a ripple effect across disciplines, shaping research questions, methods, and interpretations beyond their original context. Simultaneously, it offers excellent opportunities for a shared research agenda.

3. Discussion and Conclusions

In this paper, we explored how different research perspectives have shaped RSs research across technical and social orientations, aiming to foster reflection on interdisciplinary work in this field. In both orientations, we observe a shift from a single-user-focused perspective toward a more comprehensive view acknowledging RSs' embedding in multi-sided markets. Our discussion also highlights interplay between both orientations, showing how underlying assumptions have influenced one another over time. Initially, the accuracy-based, single-user-focused perspective rooted in technically-oriented research shaped assumptions and questions in socially-oriented work. More recently, however, the social orientation's view of RSs in multi-sided markets has begun to influence technically-oriented research. These transitions were partly driven by real-world applications, which for both research orientations sharpened the need to understand the complexities of contemporary RSs.

We observe that this leads to a growing convergence of both research orientations, with an increasing focus on the question of **why do recommenders recommend**. Embracing a technical orientation, this question reflects the purpose of RSs and informs the design and evaluation of these systems, while for socially-oriented research, it helps unpack RSs' role in mediating user interactions and their real-world impact. This convergence underscores the value of interdisciplinary engagement, as addressing this central question would benefit from insights drawn from both technical and social perspectives. By charting this evolution, this paper aims to contribute to such an interdisciplinary research agenda in at least two ways. First, it aims to demonstrate the importance of being explicit about the assumptions underpinning research paradigms and encourage researchers to revisit these assumptions, even across research orientations and disciplines. Second, it seeks to identify these research orientations' 'natural tendencies' and indicate which questions may need additional effort.

First, highlighting how assumptions influence research paradigms in various disciplines underscores the need to acknowledge these assumptions explicitly. In that regard, it is important to note that each wave has a continued relevance, regardless of the emergence of a new one. Taking a technical orientation, the academic hunt for algorithmic improvements on accuracy metrics remains a relevant topic for scientific progress in this field. However, one should realize that in this first wave, there is a widening gap between theory and contemporary practice, where state-of-the-art algorithms may underperform in real-world settings. Similarly, in socially-oriented research, understanding the political economy of RSs does not make concerns about what recommenders recommend obsolete. However, approaching this from a 'first wave understanding of RSs' may result in significant abstraction from reality, impacting research hypotheses and interpretations of findings. Consequently, researchers should transparently outline the assumptions that guide their research design and hypotheses to invite critical evaluation and adaptability to evolving contexts. Moreover, it presents the opportunity to revisit certain assumptions in light of more recent findings or real-world applications.

Although making underlying assumptions explicit may seem straightforward, some of the most fundamental premises are frequently taken for granted. This is often because they are deeply ingrained within the dominant research paradigm of a particular discipline. There are various ways to promote transparency around these assumptions, such as pre-registering experiments [58]. Pre-registration involves publicly sharing the research plan, including hypotheses, methods, and analysis, before conducting the study. Journals such as *ACM Transactions on Recommender Systems (TORS)* support this through their registered reports format, where study protocols are peer-reviewed in principle before data collection. By encouraging pre-registration and similar practices, researchers can emphasize the importance of communicating about the fundamental premises that underpin their work.

Second, by delineating these research paradigms, this discussion may help the community identify which research questions emerge less naturally and thus need additional effort. Although growing, many research questions in this 'third wave' are pursued by only a few researchers. One reason could be the need for methodological innovation, which may face reviewer resistance and lead to difficulties publishing findings [26]. Consequently, conferences and journals should continue efforts recruiting diverse reviewers and actively promote research exploring innovative methodological approaches and perspectives. Another challenge is that the nature of these questions may require access to real-world data and/or online evaluations. Given their increasing impact on markets and societies, governance bodies have taken steps to facilitate research by allowing "vetted researchers" to request data from very large online platforms [15, 59]. However, it remains unclear to what extent this will be sufficient to address the critical inquiries that researchers seek to explore. To this end, researchers should familiarize themselves with the procedures and engage in open dialogue to share challenges and best practices.

The evolution sketched in this paper illustrates the potential for RSs research to move toward a more comprehensive understanding of RSs and their real-world impact. While the field has grown substantially and demonstrated its versatility through contributions from diverse disciplines and application domains, true integration requires deliberate engagement with both technical and social orientations. Each brings distinct strengths—whether in system development and evaluation or in examining social context, purpose, and impact. Moving forward, sustained collaboration is essential to bridge assumptions, methods, and priorities to foster shared ground for interdisciplinary research.

Acknowledgments

This publication was supported by the Research Foundation Flanders (FWO) under grant number S006323N, and by the Excellence in Digital Sciences and Interdisciplinary Technologies (EXDIGIT) project, funded by Land Salzburg under grant number 20204-WISS/263/6-6022.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] F. J. Martin, J. Donaldson, A. Ashenfelter, M. Torrens, R. Hangartner, The big promise of recommender systems, AI Magazine 32 (2011) 19–27. doi:10.1609/aimag.v32i3.2360.

[2] D. Jannach, G. Adomavicius, Recommendations with a purpose, in: Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16, Association for Computing Machinery, New York, NY, USA, 2016, pp. 7–10. doi:10.1145/2959100.2959186.

[3] J. Möller, D. Trilling, N. Helberger, B. van Es, Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity, Information, Communication & Society 21 (2018) 959–977. doi:10.1080/1369118X.2018.1444076.

[4] L. Michiels, J. Leysen, A. Smets, B. Goethals, What are filter bubbles really? a review of the conceptual and empirical work, in: Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '22 Adjunct, Association for Computing Machinery, New York, NY, USA, 2022, pp. 274–279. doi:10.1145/3511047.3538028.

[5] D. Hesmondhalgh, R. Campos Valverde, D. B. V. Kaye, Z. Li, The Impact of Algorithmically Driven Recommendation Systems on Music Consumption and Production: A Literature Review, Report, Centre for Data Ethics and Innovation; Department for Digital, Culture, Media & Sport, UK, 2023. URL: https://www.gov.uk/government/publications/research-into-the-impact-of-streaming-services-algorithms-on-music-consumption/the-impact-of-algorithmically-driven-recommendation-systems-on-music-consumption-and-production-a-literature-review.

[6] T. Bonini, A. Gandini, "First week is editorial, second week is algorithmic"': Platform gatekeepers and the platformization of music curation, Social Media + Society 5 (2019). doi:10.1177/2056305119880006.

[7] D. Massimo, F. Ricci, Building effective recommender systems for tourists, AI Magazine 43 (2022) 209–224. doi:10.1002/aaai.12057.

[8] A. Banerjee, P. Banik, W. Wörndl, A review on individual and multistakeholder fairness in tourism recommender systems, Frontiers in Big Data 6 (2023). doi:10.3389/fdata.2023.1168692.

[9] Y. Wang, W. Ma, M. Zhang, Y. Liu, S. Ma, A survey on the fairness of recommender systems, ACM Transactions on Information Systems 41 (2023). doi:10.1145/3547333.

[10] Y. Mashayekhi, N. Li, B. Kang, J. Lijffijt, T. De Bie, A challenge-based survey of e-recruitment recommendation systems, ACM Computing Surveys 56 (2024). doi:10.1145/3659942.

[11] K. Hosanagar, D. Fleder, D. Lee, A. Buja, Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation, Management Science 60 (2014) 805–823. doi:10.1287/mnsc.2013.1808.

[12] G. Adomavicius, J. C. Bockstedt, S. P. Curley, J. Zhang, Effects of online recommendations on consumers' willingness to pay, Information Systems Research 29 (2018) 84–102. doi:10.1287/isre.2017.0703.

[13] R. De Croon, L. Van Houdt, N. N. Htun, G. Štiglic, V. Vanden Abeele, K. Verbert, Health recommender systems: Systematic review, Journal of Medical Internet Research 23 (2021). doi:10.2196/18035.

[14] C. A. Figueroa, H. Torkamaan, A. Bhattacharjee, H. Hauptmann, K. W. Guan, G. Sedrakyan, Designing health recommender systems to promote health equity: A socioecological perspective, Journal of Medical Internet Research 27 (2025). doi:10.2196/60138.

[15] European Union, 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), 2022. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32022R2065.

[16] D. Jannach, P. Pu, F. Ricci, M. Zanker, Recommender systems: Trends and frontiers, AI Magazine 43 (2022) 145–150. doi:10.1002/aaai.12050.

[17] J. Stray, A. Halevy, P. Assar, D. Hadfield-Menell, C. Boutilier, A. Ashar, C. Bakalar, L. Beattie, M. Ekstrand, C. Leibowicz, C. Moon Sehat, S. Johansen, L. Kerlin, D. Vickrey, S. Singh, S. Vrijenhoek, A. Zhang, M. Andrus, N. Helberger, P. Proutskova, T. Mitra, N. Vasan, Building human values into recommender systems: An interdisciplinary synthesis, ACM Transactions on Recommender Systems 2 (2024). doi:10.1145/3632297.

[18] E. Rich, User modeling via stereotypes, Cognitive Science 3 (1979) 329–354. doi:10.1016/S0364-0213(79)80012-9.

[19] F. Ricci, L. Rokach, B. Shapira, Introduction to Recommender Systems Handbook, in: F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), Recommender Systems Handbook, Springer US, Boston, MA, USA, 2011, pp. 1–35. doi:10.1007/978-0-387-85820-3_1.

[20] B. Smyth, People who liked this also liked... A publication analysis of three decades of recommender systems research, ACM Transactions on Recommender Systems (2025). doi:10.1145/3742442.

[21] M. Kaminskas, D. Bridge, Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems, ACM Transactions on Interactive Intelligent Systems 7 (2016) 1–42. doi:10.1145/2926720.

[22] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, J. Riedl, GroupLens: an open architecture for collaborative filtering of netnews, in: Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work, CSCW '94, Association for Computing Machinery, New York, NY, USA, 1994, pp. 175–186. doi:10.1145/192844.192905.

[23] D. Goldberg, D. Nichols, B. M. Oki, D. Terry, Using collaborative filtering to weave an information tapestry, Communications of the ACM 35 (1992) 61–70. doi:10.1145/138859.138867.

[24] J. Rucker, M. J. Polanco, Siteseer: personalized navigation for the Web, Communications of the ACM 40 (1997) 73–76. doi:10.1145/245108.245125.

[25] D. Jannach, O. S. Shalom, J. A. Konstan, Towards more impactful recommender systems research, in: Proceedings of the 1st Workshop on the Impact of Recommender Systems, co-located with 13th ACM Conference on Recommender Systems (ACM RecSys 2019), volume 2462 of *CEUR Workshop Proceedings*, Aachen, Germany, 2019. URL: https://ceur-ws.org/Vol-2462/short6.pdf.

[26] D. Jannach, C. Bauer, Escaping the McNamara Fallacy: toward more impactful recommender systems research, AI Magazine 41 (2020) 79–95. doi:10.1609/aimag.v41i4.5312.

[27] E. Brynjolfsson, Y. Hu, D. Simester, Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales, Management Science 57 (2011) 1373–1386. doi:10.1287/mnsc.1110.1371.

[28] E. Pariser, The filter bubble: what the Internet is hiding from you, Penguin Press, New York, NY, USA, 2011. doi:10.5555/2029079.

[29] D. Jannach, M. Jugovac, Measuring the business value of recommender systems, ACM Transactions on Management Information Systems 10 (2019). doi:10.1145/3370082.

[30] F. Pasquale, The black box society: The secret algorithms that control money and information, Harvard University Press, Cambridge, MA, USA, 2015. doi:10.5555/2717112.

[31] Z. Abbassi, S. Amer-Yahia, L. V. Lakshmanan, S. Vassilvitskii, C. Yu, Getting recommender systems to think outside the box, in: Proceedings of the Third ACM Conference on Recommender Systems, RecSys '09, Association for Computing Machinery, New York, NY, USA, 2009, pp. 285–288. doi:10.1145/1639714.1639769.

[32] H. Abdollahpouri, Popularity bias in recommendation: a multi-stakeholder perspective, Phd thesis,

University of Colorado, 2020. URL: http://arxiv.org/abs/2008.08551.
[33] S. S. Anand, B. Mobasher, Contextual recommendation, in: B. Berendt, A. Hotho, D. Mladenic, G. Semeraro (Eds.), From Web to Social Web: Discovering and Deploying User and Content Profiles, Springer, Berlin, Heidelberg, 2007, pp. 142–160. doi:10.1007/978-3-540-74951-6_8.
[34] A. Sun, Take a fresh look at recommender systems from an evaluation standpoint, in: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 2629–2638. doi:10.1145/3539618.3591931.
[35] J. B. Schafer, J. Konstan, J. Riedl, Recommender systems in e-commerce, in: Proceedings of the 1st ACM Conference on Electronic Commerce, EC '99, Association for Computing Machinery, New York, NY, USA, 1999, pp. 158–166. doi:10.1145/336992.337035.
[36] B. Smith, G. Linden, Two decades of recommender systems at Amazon.com, IEEE Internet Computing 21 (2017) 12–18. doi:10.1109/MIC.2017.72.
[37] S. M. McNee, J. Riedl, J. A. Konstan, Being accurate is not enough: how accuracy metrics have hurt recommender systems, in: CHI '06 Extended Abstracts on Human Factors in Computing Systems, CHI EA '06, Association for Computing Machinery, New York, NY, USA, 2006, pp. 1097–1101. doi:10.1145/1125451.1125659.
[38] A. V. Bodapati, Recommendation systems with purchase data, Journal of Marketing Research 45 (2008) 77–93. doi:10.1509/jmkr.45.1.077.
[39] B. P. Knijnenburg, M. C. Willemsen, Evaluating recommender systems with user experiments, in: F. Ricci, L. Rokach, B. Shapira (Eds.), Recommender Systems Handbook, Springer US, Boston, MA, USA, 2015, pp. 309–352. doi:10.1007/978-1-4899-7637-6_9.
[40] N. Helberger, K. Karppinen, L. D'Acunto, Exposure diversity as a design principle for recommender systems, Information, Communication & Society 21 (2018) 191–207. doi:10.1080/1369118X.2016.1271900.
[41] S. Vrijenhoek, M. Kaya, N. Metoui, J. Möller, D. Odijk, N. Helberger, Recommenders with a mission: Assessing diversity in news recommendations, in: Proceedings of the 2021 Conference on Human Information Interaction and Retrieval, CHIIR '21, Association for Computing Machinery, New York, NY, USA, 2021, pp. 173–183. doi:10.1145/3406522.3446019.
[42] J. K. Sørensen, J.-H. Schmidt, An algorithmic diversity diet?: Questioning assumptions behind a diversity recommendation system for PSM, in: RIPE@2016: Public Service Media in a Networked Society, 2016.
[43] K. Irion, P. Valcke, Cultural diversity in the digital age: EU competences, policies and regulations for diverse audiovisual and online content, in: E. Psychogiopoulou (Ed.), Cultural Governance and the European Union: Protecting and Promoting Cultural Diversity in Europe, Palgrave Macmillan UK, London, UK, 2015, pp. 75–90. doi:10.1057/9781137453754_7.
[44] R. D. Burke, H. Abdollahpouri, B. Mobasher, T. Gupta, Towards multi-stakeholder utility evaluation of recommender systems, UMAP (Extended Proceedings) 1618 (2016). URL: https://ceur-ws.org/Vol-1618/SOAP_paper2.pdf.
[45] D. S. Evans, R. Schmalensee, Matchmakers: The new economics of multisided platforms, Harvard Business Review Press, Boston, MA, USA, 2016.
[46] E. Zangerle, C. Bauer, Evaluating recommender systems: survey and framework, ACM Computing Surveys 55 (2022). doi:10.1145/3556536.
[47] H. Abdollahpouri, G. Adomavicius, R. Burke, I. Guy, D. Jannach, T. Kamishima, J. Krasnodebski, L. Pizzato, Multistakeholder recommendation: Survey and research directions, User Modeling and User-Adapted Interaction (2020) 127–158. doi:10.1007/s11257-019-09256-1.
[48] J.-C. Rochet, J. Tirole, Platform competition in two-sided markets, Journal of the European Economic Association 1 (2003) 990–1029. doi:10.1162/154247603322493212.
[49] T. Eisenmann, G. Parker, M. W. Van Alstyne, Strategies for two-sided markets, Harvard Business Review 84 (2006) 92. URL: https://hbr.org/2006/10/strategies-for-two-sided-markets.
[50] R. Mehrotra, N. Xue, M. Lalmas, 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, 2020, pp. 3224–3233. doi:10.1145/3394486.3403374.

[51] O. Jeunen, B. Goethals, Top-k contextual bandits with equity of exposure, in: Proceedings of the 15th ACM Conference on Recommender Systems, RecSys ’21, Association for Computing Machinery, New York, NY, USA, 2021, pp. 310–320. doi:10.1145/3460231.3474248.

[52] R. Burke, G. Adomavicius, T. Bogers, T. Di Noia, D. Kowald, J. Neidhardt, O. Özgöbek, M. S. Pera, N. Tintarev, J. Ziegler, De-centering the (traditional) user: Multistakeholder evaluation of recommender systems, International Journal of Human-Computer Studies 203 (2025). doi:10.1016/j.ijhcs.2025.103560.

[53] H. Ranaivoson, A. Smets, P. Ballon, Challenges and opportunities for recommender systems in media markets, in: U. Rohn, M. B. Rimscha, T. Raats (Eds.), De Gruyter Handbook of Media Economics, De Gruyter, Berlin, Boston, 2024, pp. 215–228. doi:10.1515/9783110793444-015.

[54] A. Smets, J. Hendrickx, P. Ballon, We’re in this together: a multi-stakeholder approach for news recommenders, Digital Journalism 10 (2022) 1813–1831. doi:10.1080/21670811.2021.2024079.

[55] Y. Kittaka, S. Sato, Y. Zennyo, Self-preferencing by platforms: A literature review, Japan and the World Economy 66 (2023). doi:10.1016/j.japwor.2023.101191.

[56] N. Chen, H.-T. Tsai, Steering via algorithmic recommendations, RAND Journal of Economics 55 (2024) 501–518. doi:10.1111/1756-2171.12481.

[57] European Union, Regulation (EU) 2022/1925 of the European Parliament and of the Council of 16 November 2022 on Contestable and Fair Markets in the Digital Sector (Digital Markets Act), 2022. URL: https://eur-lex.europa.eu/eli/reg/2022/1925/oj.

[58] J. P. Simmons, L. D. Nelson, U. Simonsohn, Pre-registration: Why and how, Journal of Consumer Psychology 31 (2021) 151–162. doi:10.1002/jcpy.1208.

[59] M. Vermeulen, Researcher access to platform data: European developments, Journal of Online Trust and Safety 1 (2022). doi:10.54501/jots.v1i4.84.