Beyond Personalization: Recommender Systems and Digital Humanism

Digital Humanism #4

Julia Neidhardt

julia.neidhardt@tuwien.ac.at

July 9th, 2026

Before You Opened Social Media Today,...

- **Recommender Systems (RSs)** had already decided what would were likely to see:
  - The first videos,
  - The first post,
  - The first songs,
  - The first products.
- You didn't choose from everything on the platform, you chose from what was shown to you.
- RSs are today among the most widely deployed forms of AI and influence billions of decisions every day

Recommendations reportedly drive **over 70% of YouTube watch time** and generate more viewership than subscriptions or search
(Mohan, cited in Quartz, 2018; YouTube)

What This Session is For

- This session is not a technical tutorial on recommender algorithms.
- Instead, we discuss RSs **from a Digital Humanism perspective**:
  - What kinds of choices do these systems make visible?
  - Which values are built into recommendation and ranking?
  - Who benefits, who is disadvantaged, and who gets to decide?
  - How can RSs better serve human autonomy, fairness, and democratic society?

Part 1: Recommender Systems Basics

How platforms decide what you see

Something to Keep in Mind

- Before we discuss algorithms, **think about this**:
  - What was the first recommendation you encountered today?
  - Why do you think it appeared?
  - What information might the platform have used?
  - How much of that decision belonged to you?
  - Have you ever been shown something that felt awkwardly specific, almost as if the platform knew too much?

What are Recommender Systems?

- **RSs are software systems** that
  - estimate what users are likely to value,
  - rank available options,
  - decide which items become visible.
- They learn from clicks, purchases, ratings, watch time, many other behavioral signals.
- Commercial systems often optimise user engagement but this is a design choice rather than a technical necessity.
- RSs decide which options become visible.

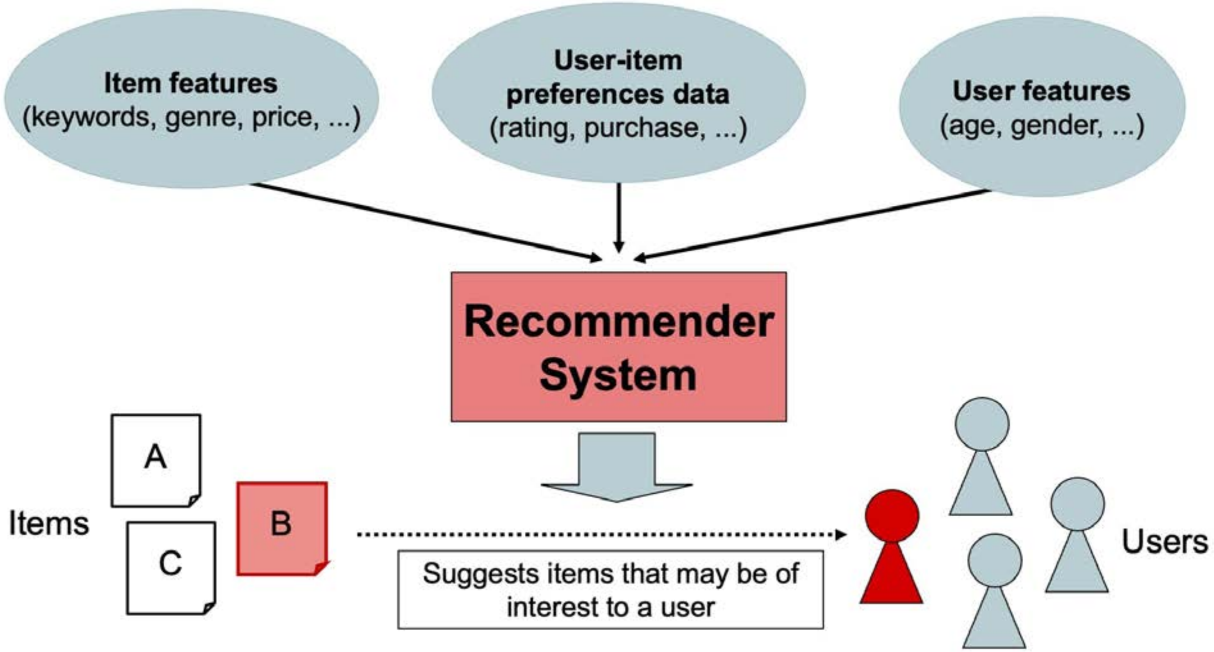

[Ricci et al., 2011]

Where Do We Encounter RSs?

- Recommender systems power many of the digital services we use every day:
  - social media platforms,
  - video and music streaming,
  - online shopping,
  - news services,
  - entertainment platforms.

- RSs are no longer only tools for finding relevant content. They increasingly determine which information, products and opportunities receive attention.

[Ricci et al., 2011]

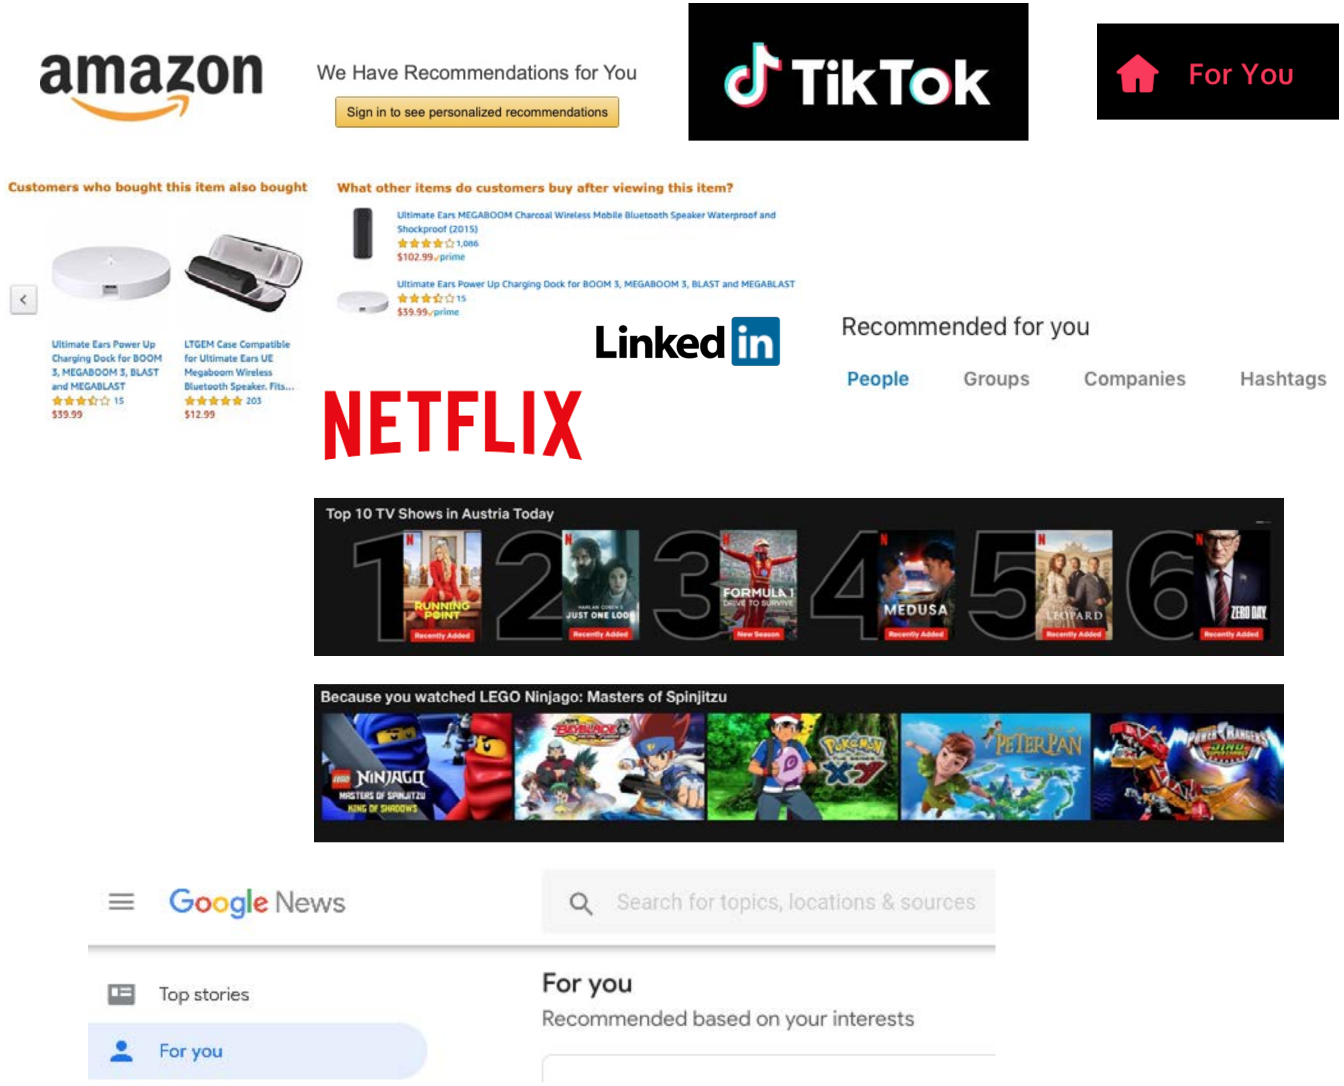

How Recommender Systems Work

- Most RSs combine several approaches:
  - **Collaborative filtering:** recommend items liked by similar users.
  - **Content-based recommendation:** recommend items similar to those a user has previously enjoyed.
  - **Hybrid methods:** combine multiple approaches to improve performance.
- Different methods but with the same goal: Estimate which items are most relevant for each user.
- Every recommendation reflects assumptions about what is worth showing.

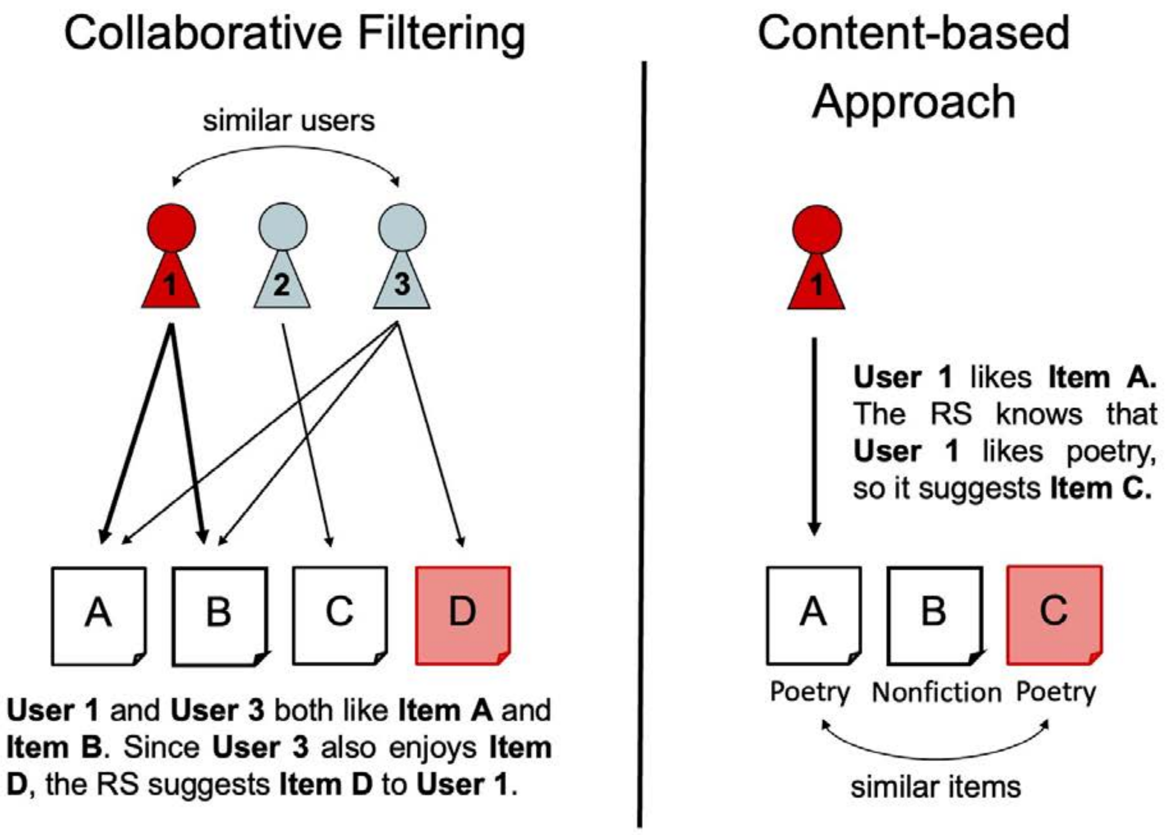

[Ricci et al., 2011]

Recommendation as Prediction

- The recommendation task can be viewed as a prediction problem.
- For each user and candidate item, the system estimates how likely a desired outcome is for example:
  - clicking, watching, purchasing, learning, long-term satisfaction.
- Items are assigned prediction scores and ranked accordingly.
- The central challenge is that users interact with only a tiny fraction of all available items, so the system must infer preferences from sparse data.

user × item matrix

	Dune	Arrival	Blade Runner	Ex Machina
You	★	?	★	?
Anna	★		★	★
Ben	★	★		★
Clara		★		

? = predict the blanks from the filled cells

Explicit Signals	Implicit Signals
★★★★★ rating	Watched to the end
Like	Paused
Save	Re-watched
Follow	skipped immediately

[Ricci et al., 2011; Aggarwal, 2016]

Learning Hidden Preferences

- Rather than memorizing every interaction, RSs learn compact representations of users and items.
- These representations (called **latent representations**) capture underlying preferences that are not directly observable.
- Users with similar behavior are represented close together in this learned space, making it possible to recommend nearby items.
- Modern RSs learn increasingly rich representations using large-scale machine learning models.

[Koren, Bell & Volinsky, 2009]

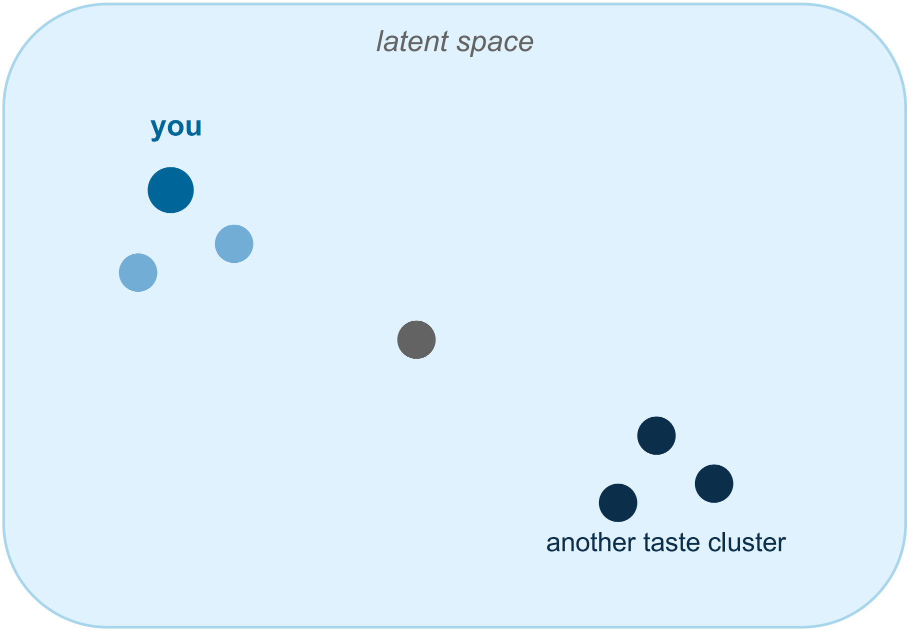

near = similar taste

The Cold-Start Problem

- Recommendations become easier once a system has observed user behavior.
- The challenge is **starting from little or no information**:
  - new users have no interaction history,
  - new items have not yet been evaluated.
- To address this, platforms combine signals such as:
  - onboarding preferences,
  - item descriptions,
  - popularity,
  - exploration of new content.
- The system **gradually replaces general assumptions with observations** about an individual user.

Exploit: show what already works	***Familiar ⟷ Novel***	**Explore: try something new**

[Schein et al., 2002]

From Millions of Items to Your Feed

- Recommendation is a **multi-stage filtering process**.
  1. retrieve a small set of promising candidates,
  2. rank these candidates using more sophisticated models,
  3. re-rank the results to account for objectives such as diversity, safety, freshness, or business rules,
  4. present the final selection to the user.

- Only a **tiny fraction of all available content** appears in your feed. Most available information never reaches the user!

[Covington et al., 2016]

What Does a Platform Optimize For?

- Although RSs differ in their implementation, they all do the following:

Observe clicks, views, dwell → **Model** infer preferences/make predictions → **Rank** score candidates → **Show** deliver the feed

- The critical question is therefore not only how systems recommend, but what they are designed to optimize.

- Different optimization objectives produce different recommendation outcomes.

[Stray et al., 2023]

Part 2: Inside the Feed

How platforms decide what you see and what they optimize.

Every Recommendation Follows the Same Process

- Every time you open an app, thousands of candidate items compete for your attention.
- For every candidate item the system estimates the probability of different user actions (click, watch to the end, like, comment, ignore).
- These predictions are combined into a ranking score.
- The highest-ranked items become your feed.

Gather candidates → **Predict** P(like), P(share)… → **Score** weighted sum → **Sort & filter** rank + demotions

[Covington et al., 2016]

Not All Signals Are Equally Informative

- Different interactions tell a recommender system different things about a user's interests.

User interaction	What it may indicate
View	Exposure only
Like	Positive reaction
Watch to the end	Sustained interest
Save	Future intention
Share	High perceived value
Rewatch	Strong interest

- Recommendation systems learn not only from what we do, but also from which actions we take.

[Hu, Koren & Volinsky, 2008]

TikTok: Learning Your Interests

- TikTok relies **primarily on behavior**, not on who you follow.

- Particularly informative signals include:
  - watch time,
  - video completion,
  - rewatches,
  - quick swipes.

- Because every interaction provides immediate feedback, the system can adapt surprisingly quickly.

- The underlying objective is not simply clicks, it is long-term retention.

[NYT, 2021 (Algo 101); WSJ, 2021; TikTok Newsroom, 2020]

YouTube: Optimizing Watch Time

- YouTube predicts which videos you are most likely to watch for longer.
- Since 2012, the main optimization objective has shifted from clicks to **expected watch time.**
- Autoplay further extends viewing sessions by continuously suggesting the next video.
- Recommendations reportedly generate more viewership than subscriptions or search

[Covington et al., 2016; Mohan, 2018]

Instagram: Different Feeds, Different Objectives

- Like other platforms, Instagram does not use one recommendation algorithm.

- **Different ranking models** power
  - Feed
  - Stories
  - Explore
  - Reels
  - Search
- because each **supports different user goals**.

[Mosseri, 2023]

Different Platforms, Different Signals

- Although the signals differ, every platform predicts future behavior and ranks accordingly.

Platform	Primary Signals
TikTok	watch time, completion, retention
YouTube	watch time
Instagram	watch time, likes, shares
Facebook	weighted engagement

[Facebook Papers, 2021; X open-sourced algorithm, 2023]

What Research Shows

- Studies consistently **report two observations**, i.e., RSs can:
  - learn user preferences very quickly
  - continuously adapt as user behavior changes
- **Examples:**
  - TikTok: Feeds can become highly specialised within a short time.
  - YouTube: A substantial share of videos users later regretted watching had been recommended by the platform.

[WSJ, 2021; Mozilla RegretsReporter, 2021]

Feedback Loops

- Recommendations are not static.
- Each interaction becomes new data, influencing the next recommendation.
- Over time, this feedback loop can reinforce even small preferences.

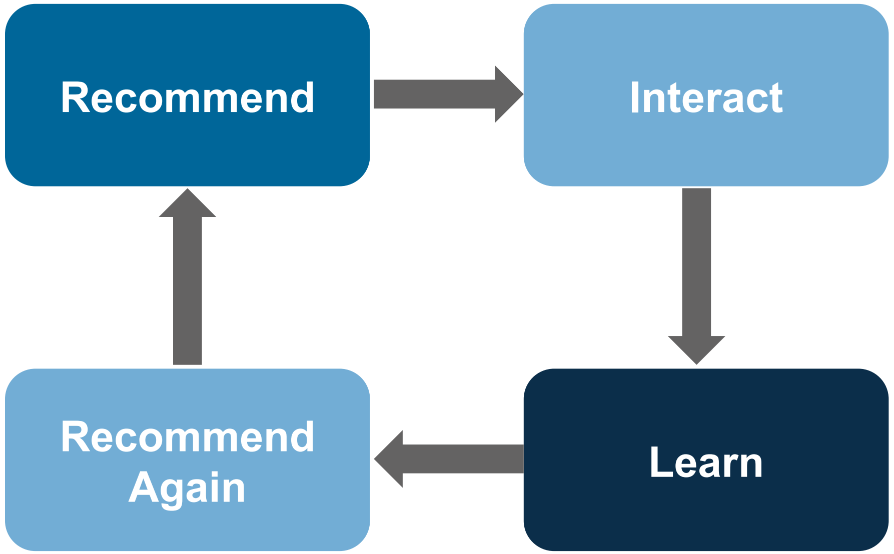

[Mansoury et al., 2020]

Are We Trapped in Filter Bubbles?

- Early concerns suggested that personalization isolates people from opposing viewpoints.
- Current research paints a more nuanced picture.
- Both algorithms and individual choices influence what we ultimately see.
- Algorithms influence exposure, but they interact with people's own choices and broader social dynamics.

[Bakshy et al., 2015; Guess et al., 2023; Haidt & Bail, 2023]

When Engagement Amplifies Content

- RSs do typically not optimize for outrage or extremism. They optimize for engagement.

- **As a result**:
  - emotionally charged content often receives more attention,
  - existing interests can become reinforced,
  - users may be exposed to increasingly similar or more extreme content.

- Current evidence suggests that recommender systems are more likely to amplify existing preferences than to create new ones.

[Brady et al., 2017; Ribeiro et al., 2020, Haidt & Bail, 2023]

User behavior
↓
Prediction
↓
Recommendation
↓
More interaction

Why Engagement Becomes the Default

- Platforms optimize engagement because it is:
  - easy to measure,
  - available immediately,
  - closely linked to advertising revenue,
  - useful for improving recommendation models.

- Long-term outcomes, such as learning, well-being, or informed decision-making, are much harder to observe and optimize.

[Milli et al., 2023; Stray et al., 2023]

Interface Design

- Interface design also influences user behavior.

- **Examples include:**
  - infinite scroll,
  - autoplay,
  - notifications.

- These features encourage more interactions, providing additional data for recommendation algorithms.

[Mathur et al., 2019]

Example: Spotify

- Spotify combines several recommendation techniques to help users discover music.

- At the same time,
  - popular artists receive disproportionate attention,
  - emerging artists are harder to discover.

- **Recommendation therefore involves balancing**
  - relevance,
  - discovery,
  - diversity.

[Mehrotra et al., 2018]

News Feeds

- RSs generally rank content according to predicted engagement rather than its truthfulness.

- **As a result**,
  - emotionally charged information often spreads faster,
  - misinformation and clickbait may receive greater visibility.

- Recommendation systems do not evaluate the societal value of information, they optimize measurable user behavior.

[Vosoughi, Roy & Aral, 2018]

Mental Health and Vulnerability

- The same optimization that maximizes engagement **may also intensify risks for vulnerable users**.

- Research has linked RSs to concerns such as
  - compulsive use,
  - sleep disruption,
  - social comparison,
  - exposure to harmful content.

- Current evidence suggests caution rather than simple conclusions.

[cf. Haidt & Bail, 2023]

Discussion: Designing the Ranking System

- Imagine you are designing a RSs.

- Assume your goal is to maximize engagement.

- Which user action should receive the greatest weight?
  - Like
  - Share
  - Comment
  - Watch time
  - Angry reaction

- What kind of content would your ranking system reward?

Part 3: Consequences of Recommendation

When recommendation affects people, markets, and society

Beyond Accuracy: What Should We Optimize?

- Accurate recommendations are not automatically good recommendations.
  - They may be repetitive.
  - They may reduce discovery.
  - They may create unfair exposure.
  - They may ignore long-term human outcomes.

- The central question is: Good for whom?

Accuracy	**Novelty**	**Fairness**
Diversity	**Serendipity**	**Long-term value**

[Kaminskas & Bridge, 2017]

Where Does Bias Come From?

- **Bias** refers to systematic patterns in data, models, or human behavior that produce outcomes that are consistently skewed in a particular direction.

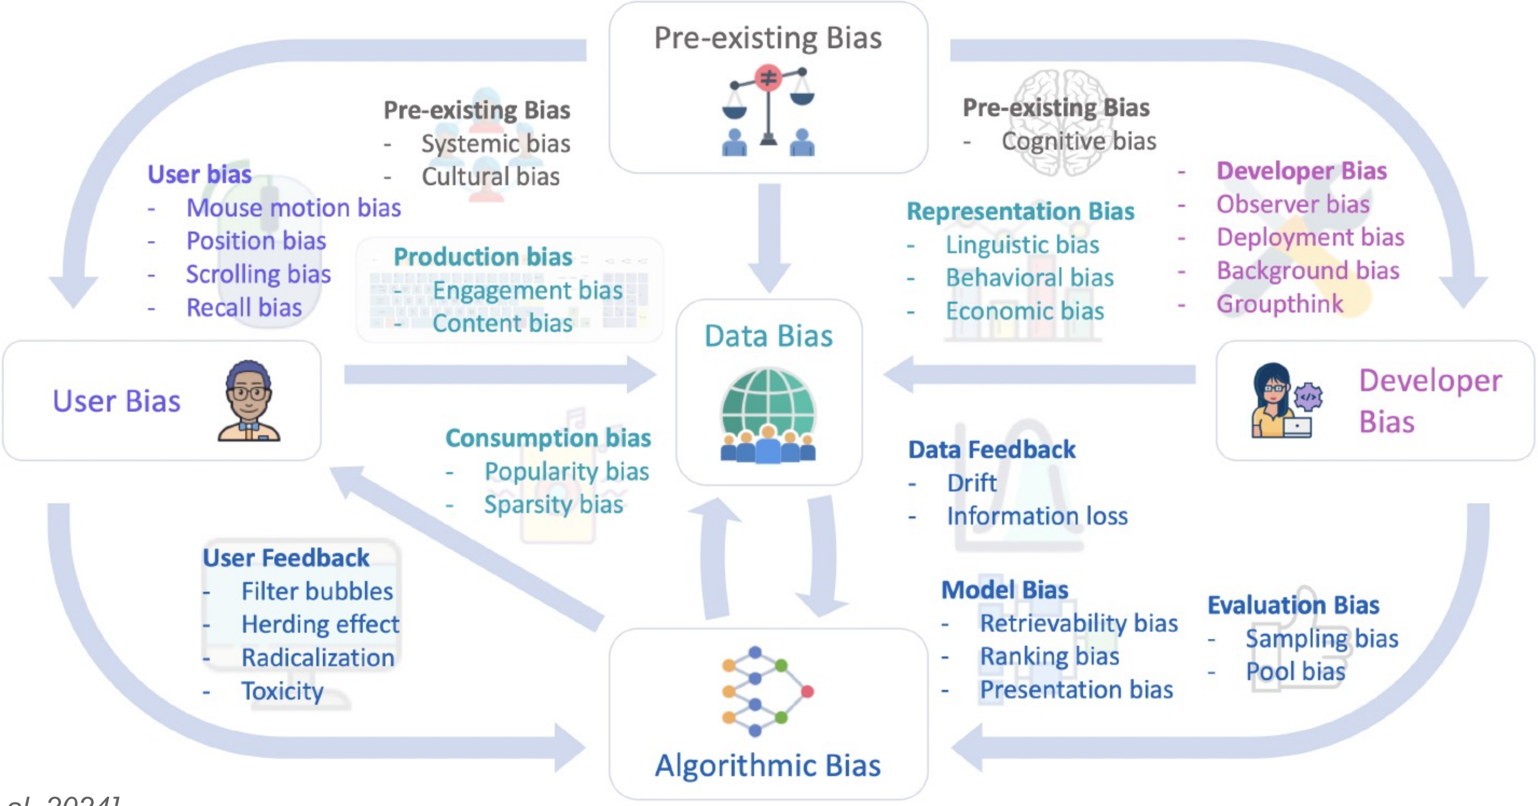

[Baeza-Yates et al. 2024]

Popularity

- Popular items receive **more visibility**. More visibility creates **more interactions**. More interactions create even **more visibility**.
- New creators, niche interests, and minority voices may remain underexposed.
- Popularity can become self-reinforcing.

[Abdollahpouri et al., 2019]

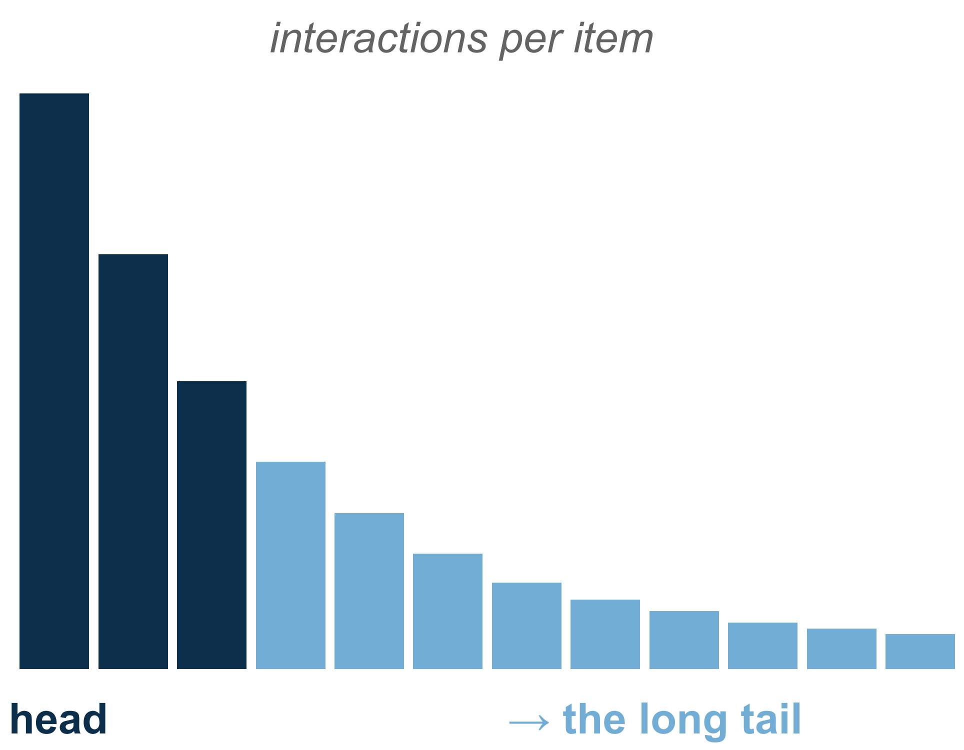

Fairness for Whom?

- Recommendation is multi-sided:

Consumers Useful recommendation	**Providers** Fair visibility
Communities Balanced representation	**Society** Public-interest outcomes,

- Who gets left out? RSs can systematically disadvantage some groups.
  - Fairness depends on whose perspective we consider

[Burke et al., 2024]

Design Choices

Design choice	Observed effect
Emotion-first ranking	More polarisation and outrage
Repeated recommendation	Echo-chamber effects, less viewpoint diversity
Popularity-based ranking	Boosts misinformation, clickbait, low quality
High speed and virality	Undermines deliberation, rewards reaction

[Haidt & Bail, 2023]

Hidden Human Labor

- RSs feel automated, but they still **depend on human work**:
  - labeling data,
  - moderating harmful content,
  - evaluating outputs,
  - handling edge cases.

- Automation often reorganizes human labour rather than eliminating it.

[Gray & Suri, 2019; Fox et al., 2023]

Data Colonialism

- A few companies control
  - data,
  - platforms,
  - infrastructure,
  - access to users.

- Value is created globally, but power and profits remain concentrated.

- Who benefits from personalization?

[Couldry & Mejias, 2024]

Consequential Recommendation

- Recommendation increasingly affects access to opportunities, not only attention.
  - hiring
  - healthcare
  - education
  - finance
  - justice

- Personalization can become opaque gatekeeping.

Persuasion, Manipulation and Autonomy

- Influence exists on a spectrum:

Persuasion	**Nudging**	**Exploitation**	**Manipulation**
open appeal to reasons	shaping defaults	leveraging weakness	bypassing agency

acceptable ⟶ *unacceptable?*

- At what point does influence become manipulation?

[Susser, Roessler & Nissenbaum, 2019]

Measuring Diversity, Novelty and Fairness

- Accuracy tells us whether recommendations are relevant.
- But we also care about:
  - **Diversity**: Are recommendations varied?
  - **Novelty**: Are users discovering new content?
  - **Fairness**: Who benefits and who is left out?
- What we measure determines what we optimize.

[Castells et al., 2015; Kaminskas & Bridge, 2017]

Transparency, Explainability and Contestability

- Human-centered recommender systems should allow people to:
  - understand how recommendations are produced,
  - see why something was recommended,
  - challenge or change the system’s behavior.

Principle	Question
Transparency	How does the system work?
Explainability	Why was this recommended?
Contestability	What can I do if I disagree?

[Tintarev & Masthoff, 2015]

Discussion: Who Should Decide?

- Who should decide what recommender systems optimize?
  - platform designers?
  - users?
  - regulators?
  - civil society?
  - affected communities?

- Should this decision be technical, democratic, or both?

Discussion: What Would You Build?

- Imagine you are designing a recommender system for a public library.

- Should it optimize for
  - the books people already like?
  - books that broaden horizons?
  - local authors?
  - underrepresented voices?
  - learning?
  - something else?

Part 4: The Generative Turn

LLMs, conversational recommendation, and agentic systems

From Ranking to Interaction

Traditional recommender systems

User
↓
Retrieve candidates
↓
Rank
↓
Top-10 list

Generative recommender systems

User request
↓
Retrieve relevant items
↓
Generate response
↓
Refine through dialogue

- Ranking does not disappear. It becomes one component of a conversational recommendation process.

[Jannach et al., 2021]

What Large Language Models Add

- Traditional recommenders learn mainly from interactions. Large language models additionally **allow users to express preferences in natural language**.

- Examples:
  - *Something similar to Arrival but less depressing.*
  - *A vegetarian restaurant within walking distance.*
  - *A beginner-friendly book about climate change.*

- LLMs allow users to express preferences more naturally.

[Wagne et al., 2026]

Example

- User Request:
  - *I'm visiting Vienna with two children.*
  - *It might rain.*
  - *We have three hours and a budget of €50.*

- Traditional RS
  - ranked attractions

- Generative RS
  - suitable attractions
  - indoor alternative
  - suggested route
  - explanation of trade-offs

- The recommendation becomes contextual rather than simply personalized.

Classical vs. Generative Recommendation

Classical	Generative
ranks items	generates responses around items
predicts interactions	processes natural-language requests
limited explanations	generated justifications and comparisons
fixed interface	iterative refinement

- The core pipeline remains, but recommendation becomes language-mediated and iterative.

[Wagne et al., 2026]

New Capabilities, New Challenges

- Generative RSs can:
  - generate persuasive explanations for recommendations,
  - produce personalised natural-language responses based,
  - use stored interaction history to maintain conversational context, and
  - perform user-authorised tasks or actions through integrated tools or services.
- New Challenges
  - reduced transparency,
  - greater persuasive power,
  - loss of user agency,
  - privacy concerns through richer conversations.

[Wagne et al., 2026]

- User:
  - *Find a quiet café near the museum that is open after 6 pm.*
- Possible failure
  - recommends a café that closed permanently
  - invents opening hours
  - cites reviews that do not exist
- The response appears plausible, but may not be correct.

The Next Step: Agentic Recommendation

- Agentic recommender systems can connect recommendations to follow-up actions.
- They may support sequences such as:

Search
↓
Compare
↓
Reserve
↓
Update calendar

- Recommendation moves from presenting options to supporting multi-step tasks.

[Maragheh & Deldjoo, 2025]

AI Slop and Synthetic Content

- Generative AI makes synthetic content cheap and abundant.
- Recommenders then rank and may learn from this machine-made content.
- Over time, this can reduce quality, blur provenance, and weaken trust in human-created work.

[Shumailov et al., 2024]

model quality over generations

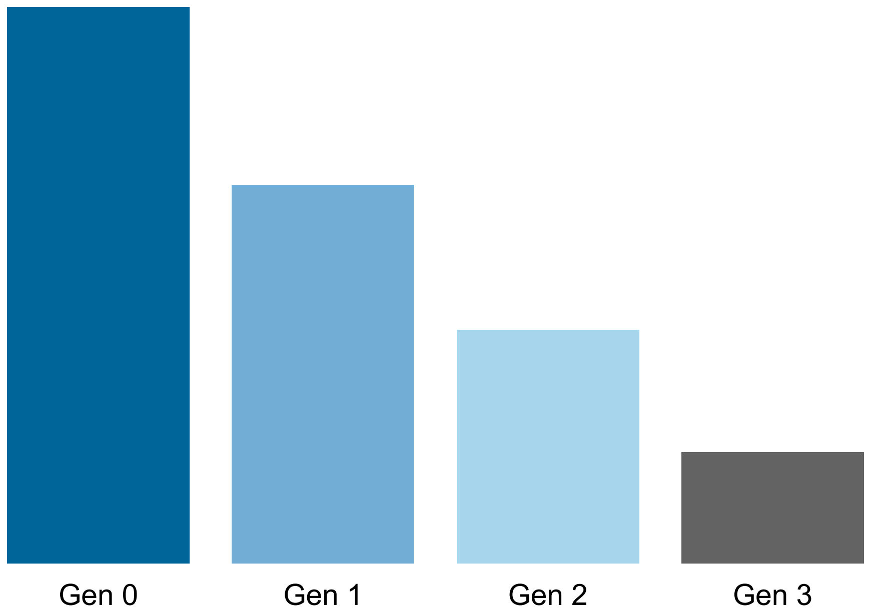

diversity & quality decay each generation

Part 5: Designing for Human Values

Human-centered design, Digital Humanism, and responsible RSs

From Optimization to Human Values

- **Human-centered RSs** should support people, not only optimize engagement.

- This means designing for:
  - informed decision-making,
  - user autonomy,
  - meaningful transparency,
  - user control,
  - diverse stakeholders,
  - long-term outcomes.

- Human-centered design asks what people need from the system not only what the system can predict.

[Stray et al, 2023, Nalis & Neidhardt, 2023; Knees, Neidhardt & Nalis, 2024]

Digital Humanism

- **Digital Humanism** asks how digital technologies can serve human and democratic values.

- For recommender systems, this means asking:
  - Who decides what is optimized?
  - Who benefits from personalization?
  - Who is disadvantaged?
  - How much control should users have?
  - Which values should not be reduced to metrics?

- Digital Humanism turns recommendation from a technical problem into a societal design question.

[Vienna Manifesto, 2019; Prem et al., 2023]

From Personalization to Empowerment

- Future RSs should give users more control over personalization.

- This includes:
  - transparent objectives,
  - editable preference models,
  - user-controlled recommendation strategies,
  - meaningful explanations,
  - portability of preference profiles.

- The goal is not to optimize users' behavior, but to support their agency.

[Kazienko & Cambria, 2024; Prem et al., 2023]

Responsible Recommender Systems

- Responsibility is more than fairness and bias.

- It concerns the whole system lifecycle:
  - objectives,
  - data,
  - models,
  - deployment,
  - monitoring,
  - governance,
  - societal impacts.

- **Responsibility is a socio-technical property**, not only a metric.

[Kazienko & Cambria, 2024]

Governance

- Governance becomes part of system design.

- **Key questions**:
  - Who defines optimization objectives?
  - Who is accountable for harms?
  - Which decisions should remain under human control?
  - How can transparency and oversight be ensured?

- Design choices and governance structures must be considered together.

[Kazienko & Cambria, 2024]

Looking ahead

- **Future RSs should**:
  - support human agency,
  - strengthen democratic discourse,
  - reduce information asymmetries,
  - enable participatory personalization,
  - evaluate long-term societal outcomes.

- This requires collaboration across computer science, social sciences, law, and the humanities.

[Kazienko & Cambria, 2024; Prem et al., 2023]

Recommender Systems Beyond Personalization

- RSs are not only personalization tools. They are digital infrastructures that affect individuals, markets, and society.
- Digital Humanism asks us to build systems that support:
  - human agency,
  - democratic values,
  - fairness,
  - societal well-being.
- The challenge is not only to build more intelligent systems, but more responsible ones.

https://dl.acm.org/doi/10.1145/3828551

Conversational Recommender Systems Using Generative Models (Gen-CRS): A Literature Review

AHMADOU WAGNE, TU Wien, Vienna, Austria
THOMAS KOLB, TU Wien, Vienna, Austria
ASHMI BANERJEE, Technical University of Munich, Munich, Germany
FATEMEH NAZARY, Polytechnic University of Bari, Bari, Italy
JULIA NEIDHARDT, TU Wien, Vienna, Austria
YASHAR DELDJOO, Department of Electrical Engineering and Information Technology, Polytechnic University of Bari, Bari, Italy

Generative models are profoundly impacting how conversational recommender systems (CRS) are conceptualized, designed, and deployed in practice, enabling mixed-initiative, context-aware, and tool-augmented recommendation workflows that go beyond rigid traditional pipelines. This survey provides a comprehensive review of Generative Conversational Recommender Systems (Gen-CRS) from 2018–2025. We organize the field into three analytical layers: (i) architectural and methodological foundations, covering unified, modular, and agentic system designs, strategies for adapting foundation models, and methods for recommendation and response generation, (ii) knowledge and data foundations, detailing how item-level information (structured catalogs, knowledge graphs, unstructured and multi-modal content), user-level signals (short-term preferences, long-term profiles, and persona representations), and dialogue corpora and logs are integrated into generative workflows, and (iii) evaluation methodologies, structured around what is evaluated (output types), which quality dimensions are measured (e.g., task effectiveness, conversational quality, efficiency, user trust, and ethical concerns), how evaluation is conducted (offline metrics, simulation, user studies, online testing), who evaluates (humans, automated metrics, LLM-based judges), and which stakeholders are considered (consumers, item providers, platforms and society). Our analysis highlights both the capabilities and the risks introduced by generative components, including challenges in grounding, catalog fidelity, controllability, safety, and bias. We argue for responsible and transparent system design, emphasizing validated knowledge integration and evaluation protocols that reflect the full complexity of conversational interactions and system behaviors, and we outline open research challenges that must be addressed to develop reliable and trustworthy Gen-CRS. This survey is also informed by, and partly the result of, tutorial feedback we collected at ACM RecSys 2025 in Prague [42].[1]

CCS Concepts: • **Information systems** → **Recommender systems**; • **Human-centered computing** → *Interactive systems and tools*; • **Computing methodologies** → Artificial Intelligence.

Additional Key Words and Phrases: Generative AI, LLM, Conversational Recommender System, Dialogue System

[1] https://www.youtube.com/watch?v=oWLr23-wRRE

Authors' Contact Information: Ahmadou Wagne, TU Wien, Vienna, Austria; e-mail: ahmadou.wagne@tuwien.ac.at; Thomas Kolb, TU Wien, Vienna, Austria; e-mail: thomas.kolb@tuwien.ac.at; Ashmi Banerjee, Technical University of Munich, Munich, Germany; e-mail: ashmi.banerjee@tum.de; Fatemeh Nazary, Polytechnic University of Bari, Bari, BA, Italy; e-mail: fatemeh.nazary@poliba.it; Julia Neidhardt, TU Wien, Vienna, Vienna, Austria; e-mail: julia.neidhardt@tuwien.ac.at; Yashar Deldjoo, Department of Electrical Engineering and Information Technology, Polytechnic University of Bari, Bari, Italy; e-mail: deldjooy@acm.org.

Tutorial RecSys 2025 and IJCAI 2026 (forthcoming):
https://recsys-lab.at/gen-conv-recsys-tutorial-ijcai-ecai-2026/