

Operational Mediation: Critically Theorizing Recommendation Systems

by

Mingyu Zhou

Digital Art History/Computational Media
Duke University

Defense Date: March 27, 2025

Approved:

Victoria Szabo, Co-Chair

Mark Hansen, Co-Chair

Mark Olson

Thesis submitted in partial fulfillment of the requirements for the degree of Master of Arts in
Digital Art History/Computational Media in The Graduate School of
Duke University
2025

ABSTRACT

Operational Mediation: Critically Theorizing Recommendation Systems

by

Mingyu Zhou

Digital Art History/Computational Media
Duke University

Defense Date: March 27, 2025

Approved:

Victoria Szabo, Co-Chair

Mark Hansen, Co-Chair

Mark Olson

An abstract of a thesis submitted in partial fulfillment of the requirements for the degree of
Master of Arts in Digital Art History/Computational Media in The Graduate School of
Duke University
2025

Copyright by
Mingyu Zhou
2025

Abstract

Recommendation systems no longer merely suggest content—they define the architecture of digital mediation itself. From the divinatory algorithms of the I Ching to the black-box neural networks of TikTok, the evolution of RecSys reveals a shift from human-guided decision-making to autonomous, self-iterating computational epistemology. This thesis argues that RecSys have transcended their original function as assistive technologies, becoming omnipresent, unknowable, and posthuman infrastructures that shape visibility, knowledge, and agency. Through historical analysis, theoretical critique, and a prototype experiment, we examine how recommendation systems mediate digital culture—not by reflecting human preferences, but by recursively generating them.

Contents

Abstract.....	iv
List of Tables	vi
List of Figures.....	viii
Acknowledgements	ix
1. Introduction	1
2. History of Recommendation Systems	4
2.1 Divination and Origins of RecSys	4
2.2 Tapestry: First Collaborative Filtering System.....	8
2.3 GroupLens, Ringo, and the Move Toward Scalable Automatic Filtering	10
2.4 Platformization: Amazon, Netflix, and Youtube’s Clickbait Problem	14
2.5 Black Box: Machine Learning and Ubiquitous RecSys.....	19
3. Theorizing RecSys Mediation	27
3.1 Centrality: RecSys as THE Epistemology and Medium.....	27
3.1.1 Omnipresence	27
3.1.2 Omniscience.....	29
3.2 Unknowability: RecSys as Black Box	32
3.3 Dehumanization: RecSys as Machines for Post-Human.....	34
3.4 Operationality: RecSys Invisuality and Operational Mediation	37
4. A Prototype of Operational Mediation	41
4.1 Concept	41
4.2 System Architecture.....	42
4.3 Mechanism of RecSys.....	44
5. Conclusion	46

Bibliography48

List of Tables

Table 1: A Theoretical Framework for Categorizing RecSys Evolution	26
--	----

List of Figures

Figure 1: I Ching Knowledge Structure System (Adopted from “ICKG: An I Ching Knowledge Graph Tool Revealing Ancient Wisdom”)	5
Figure 2: Tapestry System Architecture	9
Figure 3: Command fields and their types of Tapestry	10
Figure 4: The Group Lens Architecture	12
Figure 5: Shopping Cart Recommendation at Amazon.com in early 2000s.	15
Figure 6: Deep ranking network architecture depicting embedded categorical features.....	22
Figure 7: System Architecture of Prototype	43
Figure 8: Code Snippet of RecSys.....	44

Acknowledgements

I would like to express my sincere gratitude to Dr. Victoria Szabo, Dr. Mark Hansen, and Dr. Mark Olson—for your inspirations, guidance, and patience throughout my journey at Duke CMAC.

Thank you also to my peers and friends who provided questions, comments and critiques for theoretical configurations and prototype constructions.

Lastly, thank you to my mother who have supported me all the way.

1. Introduction

Recommendation systems (RecSys) have evolved from auxiliary tools of information retrieval to the central infrastructure of digital mediation. They no longer simply assist users in navigating content but actively determine what is visible, knowable, and relevant. Today, RecSys shape not only individual consumption but the broader epistemic structures of digital culture. This shift is not merely technical—it is ontological. RecSys have transcended their origins as assistive technologies and now function as autonomous agents of mediation, producing knowledge through algorithmic inference rather than human intent. To understand this transformation, we must examine both the historical evolution of RecSys and their contemporary role as the defining epistemology of the digital age.

The genealogy of recommendation systems extends far beyond computational algorithms, reaching into the ancient past. Long before machine learning optimized engagement, divination systems such as the Oracle of Delphi and the I Ching mediated human decision-making. These systems, much like modern RecSys, functioned through structured randomness, probabilistic interpretation, and iterative refinement. The Delphi Method, developed in the 20th century, formalized these principles into a structured forecasting technique that closely resembles modern collaborative filtering. The continuity between these historical practices and contemporary RecSys suggests that algorithmic recommendation is not merely a technological development but an extension of a long-standing human desire for mediated guidance.

The first computational RecSys—Tapestry, GroupLens, and Ringo—pioneered collaborative filtering as a means of structuring digital information. However, these early systems still required human participation, relying on explicit ratings and manually curated trust networks. With the platformization of RecSys, particularly through Amazon, Netflix, and YouTube, the paradigm shifted from user-driven filtering to fully automated, engagement-optimized recommendations. Machine learning-based RecSys no longer required users to actively express

preferences; instead, they inferred them from passive behavioral data, reinforcing predictive loops in which engagement history determined future visibility. This marked the transition from human agency to algorithmic self-iteration—content was no longer surfaced because users requested it but because predictive models determined it would maximize engagement.

Today, RecSys function as autonomous, self-optimizing infrastructures, operating beyond human comprehension. As Beatrice Fazi argues, “We must conceive of ‘automated modes of thought’ in such a way as to supersede the hope that machines might replicate human cognitive faculties, and to thereby acknowledge a form of onto-epistemological autonomy in automated ‘thinking’ processes.¹” This epistemic autonomy means that RecSys do not “see” content in the traditional sense; rather, they function through statistical correlation, iterating upon themselves in self-reinforcing feedback loops. As RecSys transition into deep-learning-powered black boxes, their mediation becomes structurally unknowable, making algorithmic opacity not just a technical issue but an epistemic condition.

This thesis traces RecSys from their historical antecedents to their contemporary dominance, arguing that they are no longer just instruments of digital mediation but the very architecture of mediation itself. The first chapter examines the evolution of RecSys, from their roots in human-guided decision systems to their integration into large-scale digital platforms. The second chapter shifts from history to theory, analyzing how RecSys function as an epistemological structure—an omnipresent, posthuman, and operational force that governs visibility, agency, and knowledge production. By examining RecSys as both historical artifacts and contemporary infrastructures, this work critically interrogates their role in shaping the logic of digital mediation, challenging assumptions about agency, control, and the future of algorithmic culture.

¹ M. Beatrice Fazi, “Beyond Human: Deep Learning, Explainability and Representation,” *Theory, Culture & Society* 38, no. 7-8 (November 27, 2020): 55–77, <https://doi.org/10.1177/0263276420966386>.

To further illustrate this trajectory, this work also introduces a prototype experiment that places RecSys in a recursive loop with multiple large language models (LLMs). By simulating how recommendations evolve within a closed system of algorithmic agents—where content is generated, filtered, and ranked without human intervention—the prototype offers a glimpse into the self-perpetuating nature of contemporary RecSys mediation. This small-scale digital experiment is not a technical study but a conceptual provocation, revealing the extent to which RecSys no longer simply curate human content but actively fabricate digital reality itself.

2. History of Recommendation Systems

2.1 Divination and Origins of RecSys

Long before the first collaborative-filtering algorithm ranked films on Netflix or YouTube fine-tuned its autoplay suggestions, humanity sought ways to predict the future and receive guidance. The genealogy of recommendation systems is deeply entangled with divination, a practice rooted in the ancient human desire for insight, foresight, and structured decision-making. From the Oracle of Delphi in ancient Greece to the I Ching in China, early forms of predictive systems mediated knowledge, operating as epistemological devices for individuals and societies alike. This section explores how these historical systems of prophecy and structured deliberation served as prototypes for modern, digital ways of thinking that bred modern recommendation systems.

In the ancient Greek, perhaps no predictive system was as influential as the Oracle of Delphi. For over a millennium, the Pythia, the high priestess of Apollo, delivered cryptic prophecies that guided individuals, city-states, and empires¹. The process of consultation at Delphi was, at its core, a recommendation mechanism—one that operated on the principle of selective knowledge dissemination. The Oracle did not provide direct answers but instead mediated knowledge, shaping possibilities through ambiguity.

Michael Scott argues that Delphi was built on a paradox: its unknowability was precisely what gave it authority². The Pythia's utterances, often recorded in poetic form, required interpretation, engaging the seeker in a feedback loop not unlike modern RecSys, where users participate in meaning-making by responding to conditional suggestions. Furthermore, the ritualistic structure of Delphi—a sacred space, a structured inquiry process, and the mediation of

¹ Michael Scott, "The Oracle at Delphi: Unknowability at the Heart of the Ancient Greek World," *Social Research: An International Quarterly* 87, no. 1 (March 2020): 51–74, <https://doi.org/10.1353/sor.2020.0005>, 51.

² *Ibid.*, 52.

divine wisdom—foreshadows the way algorithmic systems today filter, organize, and present information.

While Delphi operated within the cultural framework of Greek philosophy and divination, the I Ching (Book of Changes) provided a predictive system in ancient China that was arguably more structured and computational in nature (see Figure 1). Dating back to the Shang Dynasty (circa 1600 BCE), the I Ching employed a binary system of trigrams and hexagrams to provide guidance³.

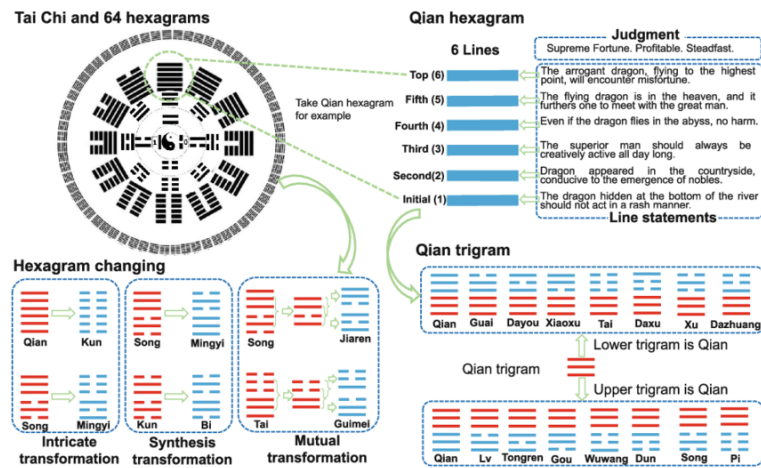


Figure 1: I Ching Knowledge Structure System (Adopted from “ICKG: An I Ching Knowledge Graph Tool Revealing Ancient Wisdom”)

Much like a combination of spreadsheets and conditional arguments today, the I Ching functioned as a lookup table, where a set of rules determined interpretative outcomes based on input variables (in this case, the casting of yarrow stalks or coins). Each hexagram’s meaning depended on a combination of fixed and mutable lines; a mechanism akin to early rule-based artificial intelligence systems. More than just a method of divination, the I Ching served as a decision-support tool, a structured means of extracting insight from patterns—a precursor to the logic underlying RecSys.

³ Shijun Liu et al., “ICKG: An I Ching Knowledge Graph Tool Revealing Ancient Wisdom,” *Communications in Computer and Information Science* 1682 (2022), https://doi.org/10.1007/978-981-99-2385-4_5, 63.

Eliot Weinberger describes the I Ching as "history's first hexadecimal algorithmically driven 'recommendation retrieval' system, with the belief that it can explain everything"⁴. This comparison suggests that, long before digital recommender systems, human cultures were already developing codified, structured methodologies for mediating decision-making.

Moving from the mystical to the methodological, the Delphi Method, developed in the 1950s by the RAND Corporation, represents an important conceptual bridge between ancient divination and modern predictive analytics. The Delphi Method is a structured forecasting technique that aggregates expert opinions iteratively to generate predictions with a higher degree of consensus⁵.

The key innovation of the Delphi Method was its systematization of foresight:

- Experts provided independent predictions.
- Feedback loops refined and converged opinions.
- Anonymity ensured that groupthink and dominant voices did not distort outcomes.

The Delphi Method embodies the fundamental logic of RecSys—iterative refinement based on user input. Much like collaborative filtering, which improves recommendations based on collective user interactions, the Delphi Method sought to optimize the accuracy of predictions through repeated adjustments.

However, scholars have critiqued its limitations. Kim Quaile Hill and Jib Fowles argue that while the Delphi Method systematized prediction and decision-making, it failed to escape biases inherent in human forecasting⁶. Moreover, the method's emphasis on expert consensus

⁴ Eliot Weinberger, "What Is the I Ching? | Eliot Weinberger," www.nybooks.com, 2016, <https://www.nybooks.com/articles/2016/02/25/what-is-the-i-ching/>.

⁵ Dmitry Khodyakov et al., "RAND Methodological Guidance for Conducting and Critically Appraising Delphi Panels," www.rand.org, December 29, 2023, <https://www.rand.org/pubs/tools/TLA3082-1.html>.

⁶ Kim Quaile Hill and Jib Fowles, "The Methodological Worth of the Delphi Forecasting Technique," *Technological Forecasting and Social Change* 7, no. 2 (January 1975): 179–92, [https://doi.org/10.1016/0040-1625\(75\)90057-8](https://doi.org/10.1016/0040-1625(75)90057-8), 182.

mirrors a key tension in contemporary RecSys: do recommendations reflect organic diversity, or do they converge toward an optimized yet limited viewpoint?

The shift from Delphi's mysticism to Delphi's method and eventually to RecSys algorithms illustrates an ongoing transition: the desire for mediated decision-making has not changed, but its mechanisms have evolved. The core principles of divination—pattern recognition, knowledge filtering, and decision support—are crucial principles of engineering in today's digital recommendation engines. The Oracle of Delphi and the I Ching provided predictive guidance that structured human decision-making, much like modern recommendation systems structure digital mediation. The Delphi Method, in turn, formalized prediction into an iterative, feedback-based process, mirroring the algorithmic logic of contemporary RecSys. Today's machine-learning-based recommender systems, which rely on big data rather than divine intervention, retain similar fundamental function: to mediate messages, shape choices, and construct pathways for action. If divination systems once mediated knowledge between humans and the gods, and Delphi panels mediated between human experts, then RecSys now mediates between millions of users and vast pools of digital content.

Across these examples, the fundamental principle remains the same: mediation is not neutral. Whether through divine prophecy, structured randomness, or algorithmic optimization, these systems do not merely reflect reality; they actively participate in constructing it. In today's digital landscape, RecSys operates as the latest iteration of this history, a continuation of humanity's deep desire for predictive guidance. The next step is to understand not only how RecSys emerged from these traditions, but how it has transformed into a dominant force in digital mediation, shaping the very conditions of knowledge, visibility, and agency in the contemporary media environment.

2.2 Tapestry: First Collaborative Filtering System

Developed at Xerox PARC in the early 1990s, Tapestry was the first system to formalize collaborative filtering, introducing a mechanism where users' interactions with information actively influenced what others would see. Tapestry recognized the social relationality between information—a document's value is often best determined not by its textual content alone but by how others respond to it⁷.

The system was designed to manage email and document overload within a research environment, where conventional filtering methods were insufficient. Instead of static subscription lists, Tapestry allowed users to annotate documents with subjective evaluations, such as endorsements or dismissals. Other users could then construct filters to retrieve documents not just by keywords but based on the interactions of trusted colleagues. This was a fundamental shift: filtering became a socially distributed process, embedding human judgment into an information system⁸.

Tapestry's technical architecture included a document store, an annotation store, and a query-based filtering mechanism. The document store maintained all incoming emails and digital documents, while the annotation store captured user-generated reactions, linking them to the corresponding documents. The filtering mechanism allowed users to define queries such as “show me all documents that John marked as ‘important’” or “retrieve articles that at least three colleagues have forwarded”⁹. The system supported both explicit filtering (where users manually tagged documents) and implicit filtering, where interactions like replying or forwarding could serve as relational indicators.

⁷ David Goldberg et al., “Using Collaborative Filtering to Weave an Information Tapestry,” *Communications of the ACM* 35, no. 12 (1992): 61–70.

⁸ Michael Schrage, *Recommendation Engines* (The MIT Press, 2020), <https://doi.org/10.7551/mitpress/12766.001.0001>, 65.

⁹ Biblio, “How Recommender Systems Make Their Suggestions,” Medium, February 6, 2019, <https://medium.com/the-graph/how-recommender-systems-make-their-suggestions-da6658029b76>.

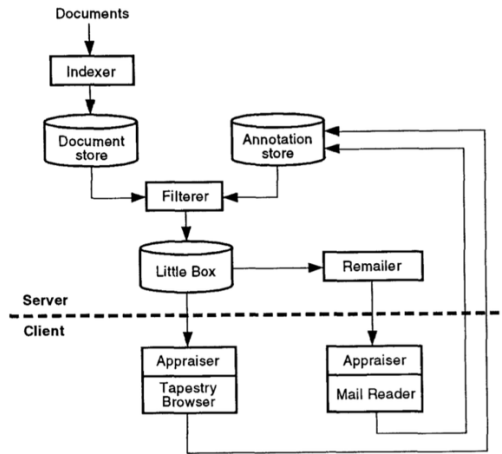


Figure 2: Tapestry System Architecture

A key innovation of Tapestry was its continuous query model. Traditional filtering systems operated in real-time, evaluating a document only upon arrival. Tapestry, however, allowed filters to operate retrospectively, meaning a document could be surfaced after annotations were added over time¹⁰. This required a persistent database and ongoing query execution, rather than a one-time filtering event. In addition, unlike modern real-time recommendation engines, which rely on dynamic, untransparent ranking models, Tapestry’s approach was query-driven, and command-specific, requiring direct user participation in both the filtering and the retrieval process.

Tapestry’s legacy lies in formalizing collaborative filtering as an information retrieval paradigm—the idea that filtering should be a socially-driven, feedback-based process that involves complex reactions from human. It is important to point out that human judgment—not algorithms (as it was not yet advanced enough)—dictated what was filtered and retrieved. Unlike later recommendation systems that automate prediction through computational models, Tapestry’s filtering process was entirely user-driven—it required explicit human annotations and

¹⁰ David Goldberg et al., 69.

manual query-based retrieval rather than statistical inference. Users issued structured commands to surface documents based on socially embedded signals, such as endorsements, replies, or shared annotations, rather than system-generated predictions (see Figure 3)¹¹. While Tapestry was a digital filtering system, it was not yet algorithmic—it mediated information flow, but decision-making remained firmly in human hands. This distinction foreshadows a later key transition in RecSys history: from explicit, user-controlled filtering to automated, opaque computational models that increasingly remove human agency from the mediation process, which had not arrived in the early 1990s.

to	set of strings
date	date
sender	string
cc	set of strings
subject	string
newsgroups	set of strings
in-reply-to	set of documents
words	set of strings

Figure 3: Command fields and their types of Tapestry

For example, one can use command line [m.sender = 'Mark'] to index all files sent from the user registered as "Mark". This shows the conscious awareness of Tapestry users in receiving recommendations processed by Tapestry but guided by colleagues.

2.3 GroupLens, Ringo, and the Move Toward Scalable Automatic Filtering

The transition from manual filtering systems like Tapestry to automated collaborative filtering (ACF) was a pivotal moment in recommendation system history. Tapestry demonstrated

¹¹ David Goldberg et al., 65

the effectiveness of human-driven annotation, but its reliance on explicit user-defined retrieval commands made it difficult to scale beyond small research communities. The next generation of recommender systems—GroupLens, Ringo, and other early scalable ACF models—sought to remove this limitation by introducing automated rating aggregation, statistical inference, and distributed architectures.

One of the most significant advancements came from GroupLens, a project initiated in 1992 by Paul Resnick, John Riedl, and their team at the University of Minnesota. Designed to filter Usenet newsgroups, GroupLens tackled the overwhelming volume of discussion threads by automating collaborative filtering. Unlike Tapestry, where users manually selected whose opinions to trust, GroupLens computed correlations dynamically, matching users based on shared rating patterns. Its system architecture consisted of newsreader clients, centralized rating servers (called "Better Bit Bureaus"), and a correlation engine that predicted a user's interest in new articles based on the ratings of others with similar preferences¹².

GroupLens introduced two key technical advancements:

- Automated nearest-neighbor filtering: Instead of requiring users to predefine trusted raters, GroupLens computed similarity scores between users dynamically, identifying "neighbors" with shared interests based on statistical correlations (e.g., users who gave similar ratings to many films are clustered as "neighbors")¹³
- Distributed architecture for scalability: Unlike Tapestry's single-site filtering model, GroupLens enabled rating sharing across multiple sites, laying the foundation for large-scale collaborative filtering networks¹⁴ (see Figure 4).

¹² Paul Resnick et al., "GroupLens: An Open Architecture for Collaborative Filtering of Netnews," 1994, 176.

¹³ Ibid., 183.

¹⁴ Jonathan Herlocker, Joseph Konstan, and John Riedl, "Explaining Collaborative Filtering Recommendations," 2000.

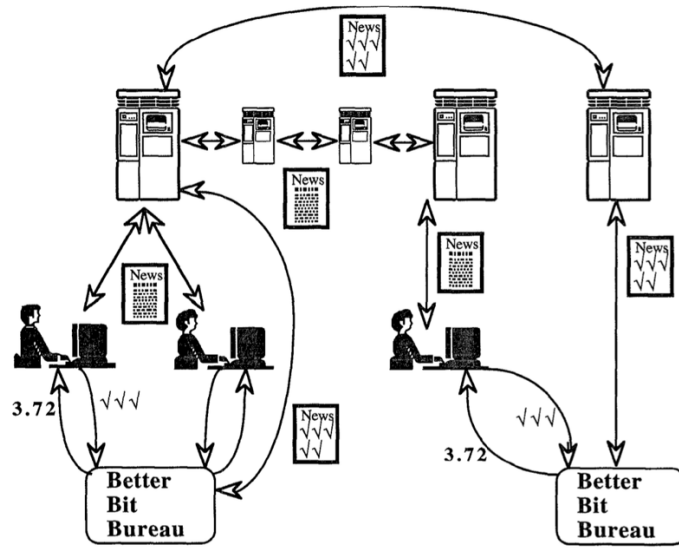


Figure 4: The Group Lens Architecture

Another key technical advancement in GroupLens was its ability to predict missing ratings by inferring user preferences based on past rating patterns. Instead of requiring direct human input for every article, GroupLens statistically estimated how a user would rate an unseen item by comparing their historical behavior to that of similar users (see Equation below)¹⁵. This marked a shift toward algorithmic agency—where predictions were no longer entirely derived from explicit human action but were, in part, generated from inferred behavioral patterns. As a result, patterns were not just observed but actively predicted, forming a recursive system where predictions influenced further predictions.

$$K_{6_{\text{Pred}}} = \bar{K} + \frac{\sum (J_6 - \bar{J}) r_{KJ}}{\sum_J |r_{KJ}|} = 3 + \frac{2r_{KM} - r_{KL}}{|r_{KM}| + |r_{KL}|} = 3 + \frac{2 - (-0.8)}{|1| + |-0.8|} = 4.56$$

¹⁵ This equation is used to estimate ratings of unrated items by referring to previous similarity pattern between rated users and unrated users. “To predict Ken’s score on the article in the matrix, take a weighted average of all the ratings on article 6 according to the formula. This is a reasonable prediction for Ken, since the article received a high rating from someone who agreed with him in the past and a low rating from someone who disagreed. Carrying through similar calculations for Nan yields a lower prediction of 3.75. Since Nan had partial agreement with Lee in the past, Lee’s low rating for the article partially cancels out the high ratings that Meg gave it”. Paul Resnick et al., 181.

While GroupLens primarily focused on text-based filtering, a parallel development in music recommendations came with Ringo, developed in 1994 by Upendra Shardanand at MIT’s Media Lab. Ringo was one of the first recommendation systems to be designed specifically for entertainment content, applying automated collaborative filtering to musical preferences rather than textual documents. The system required new users to rate at least 120 artists before generating personalized recommendations—a method designed to counteract the cold start problem¹⁶

Unlike GroupLens, which focused on predicting user preferences for unseen articles, Ringo operated more like a social filtering tool, identifying clusters of users with similar musical taste. Its algorithm worked as follows:

1. Users submitted a list of rated music artists, scoring them on a discrete scale.
2. The system computed pairwise correlations between users based on shared rating patterns, when results is lowered than a certain threshold, two users and defined as “neighbors” (see Equation below)¹⁷.
3. Once a similarity threshold was met, the system aggregated preferences from a user’s nearest neighbors (users who shared similar ratings) to recommend artists they had not yet rated but was rated by their neighbors¹⁸.

$$D_{xy} = \frac{(U_X - U_Y)^2}{\frac{\sum_n^{N_a} c_{xn} \times c_{yn} \times (s_{xn} - s_{yn})^2}{\sum_n^{N_a} c_{xn} \times c_{yn}}}$$

¹⁶ Upendra Shardanand, “Social Information Filtering for Music Recommendation” (1994), 4.

¹⁷ Ibid., 41. The equation is used to calculate the mean squared difference between two users’ shared rating of one artist. If the results is lowered than a certain threshold, it means that the two users rate in a highly similar way, which suggest similar taste, and therefore would be defined by the system as neighbors.

¹⁸ Ibid., 42.

Another unprecedented aspect of Ringo was its experiment with implicit vs. explicit ratings. While users could manually rate artists, Ringo also tracked passive interactions (e.g., how long a user spent listening to certain songs) to refine recommendations. This approach introduced an early form of implicit feedback, a feature that would later become standard in modern streaming services like Spotify, Youtube, and TikTok.

Both GroupLens and Ringo demonstrated the potential for automated collaborative filtering to scale beyond small, localized environments, marking a shift from handcrafted information retrieval to algorithmic mediation. While these systems still heavily relied on explicit user ratings (although Ringo did minimally include implicit interaction pattern), they moved one step closer to fully automated RecSys, where recommendations would eventually be generated without conscious, deliberate user input.

This transition toward automation and scalability set the stage for the next wave of development: commercial applications. As companies like Amazon, Netflix, and YouTube began incorporating recommendation engines into their platforms, collaborative filtering evolved from an experimental research tool to an essential feature of digital commerce and media consumption.

2.4 Platformization: Amazon, Netflix, and Youtube's Clickbait Problem

With the help of infrastructures such as internet that enabled scalable data collections and much larger user base, entrepreneurs begin to incorporate experimental tools into platformized products. This transition from experimental recommender prototypes to fully encapsulated platform infrastructures marked a pivotal shift in algorithmic mediation. Amazon and Netflix played critical roles in this transformation, not only in their technical innovations but in the way they redefined the epistemology of recommendation itself. Their systems did not merely optimize content discovery—they expanded the scope of what could be known, categorized, and predicted about users, shaping the very conditions of digital mediation and consumption.

Amazon’s recommendation system arose from a fundamental scalability challenge: with millions of customers and products, traditional collaborative filtering methods became computationally infeasible. Instead of comparing users to users, Amazon’s item-to-item collaborative filtering compared products to products, precomputing relationships between them offline. Instead of measuring similarity between each pair of users, the algorithm simply checks what list items are bought together¹⁹. This shift allowed Amazon to utilize limited computational resources for maximum efficiency, embedding personalized suggestions across product pages, shopping carts, and email campaigns.

The mechanics of this system can be illustrated through a concrete example: a user buying *Harry Potter and the Sorcerer’s Stone* would immediately be recommended *Harry Potter and the Chamber of Secrets*, as well as *The Hobbit* or *Percy Jackson*, based on statistical co-purchase patterns that these 3 books are bought together most frequent times. Notably, these recommendations do not stem from the semantic meaning of the books or an understanding of why they are similar but are derived purely from transactional patterns—a significant epistemic shift from early recommendation systems that relied on user-defined preferences²⁰.




Figure 5: Shopping Cart Recommendation at Amazon.com in early 2000s²¹.

¹⁹ Greg Linden, Brent Smith, and Jeremy York, “Amazon.com Recommendations,” *Industry Report* 76, no. 03 (2003), 78.

²⁰ *Ibid.*, 80.

²¹ In this case, who the “customer(s)” were does not matter, neither do they need to be the calculated “neighbors” of the current user. This reflects a way of knowing of the RecSys that could work without human’s conscious feedback, but only human’s documented, datafied actions.

As Amazon's revenue grows exponentially, the company was then able to afford much more powerful computational resources. It began utilizing the potential of collaborative-filtering and seek to exemplify the increasing scope of algorithmic epistemology. Early recommendation systems asked users directly for ratings, treating their conscious evaluations as the core data source. Amazon's Singular Value Decomposition (SVD) system, however, tracks a much wider scope of user activities—clicks, view durations, purchase delays, abandoned carts, and even cursor movements—to build a behavioral profile that extends beyond explicit preferences²². This marks a transition from a user's self-reported preferences to the system's inferred knowledge about the user, prefiguring the surveillance-based architectures of modern platform capitalism.

Finally, Amazon's recommendation engine is designed with a clear, quantifiable goal: maximizing purchase conversion. Unlike early RecSys in academic research, which were designed for information-sharing and engagement, Amazon's system is commercially driven, prioritizing profitability over exploratory content discovery. This economic imperative fundamentally shaped its algorithmic orientation: recommendations are not neutral facilitators of discovery but engineered interventions into consumer behavior²³.

Netflix's recommendation system was driven by a different, yet equally critical, platformization imperative: retention. Unlike Amazon, where purchases are discrete events, Netflix's subscription model meant that user engagement had to be sustained over time. The more relevant and personalized the content, the longer users stayed subscribed. This necessity led to one of the most famous RecSys initiatives in history—the Netflix Prize, launched in 2006²⁴.

²² For technical details, see "Badrul Sarwar et al., "Incremental Singular Value Decomposition Algorithms for Highly Scalable Recommender Systems," 2003, 2." This paper does not explicitly mention Amazon but detailed the SVD algorithm that enabled multi-aspect data training through dimensionality reduction.

²³ Greg Linden, Brent Smith, and Jeremy York., 82.

²⁴ Xacier Amatrian and Justin Basilico, "Recommender Systems in Industry: A Netflix Case Study," in *Recommender Systems Handbook, 2nd Edition*, ed. Francesco Ricci, Lior Rokach, and Bracha Shapiro (Springer, 2015), 385.

The Netflix Prize signaled a shift toward multi-model, ensemble-based recommendation approaches, combining machine learning, matrix factorization, and neural networks. The winning model was not a single algorithm but a hybrid ensemble of 107 models, showing that RecSys had reached a level of complexity beyond human interpretability²⁵.

Netflix's shift from Cinematch's²⁶ simple collaborative filtering to deeply multi-layered algorithmic systems mirror a broader technical and epistemic shift in RecSys:

- **Scale of data and Automation:** While Ringo and GroupLens used relatively transparent, interpretable models, Netflix's big-data-driven approach set the stage for less transparent, highly automated algorithmic decision-making with lower level of human agency.
- **Expanded User Monitoring:** Netflix shifted from only taking explicit ratings to combining explicit ratings and implicit behavioral tracking—analyzing watch history, scrolling, and abandonment patterns—to infer preferences, making user data collection more pervasive²⁷.
- **Mathematical Accuracy:** The Netflix Prize prioritized RMSE, a statistical optimization metric that focused purely on benchmarked predicting accuracy²⁸, this signals an epistemic shift in the sense that recommendation system was commissioned to reinforce past behaviors, predict future engagement as precisely as possible. User experience feedback and viewing promotions were not part of the purpose.

A striking takeaway from the Netflix Prize was that explicit ratings were not the best indicators of user preference. Instead, implicit signals—watch time, skips, and interaction depth—became more valuable than what users consciously reported. This further distanced

²⁵ Ibid., 392.

²⁶ An early prototype of Netflix's RecSys, which had subscribers rate movies and then segmented them into similarity clusters.

²⁷ Yehuda Koren, "Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model," 2008, 482.

²⁸ Ibid.

RecSys from user agency: instead of users actively expressing preferences, algorithms infer them, and those inferences become the basis of further predictions²⁹. This set crucial foundation for RecSys applications in platforms that are used without explicit, numerical feedback system like movie ratings, such as social media and Youtube. In addition, Netflix's choices to use predicted, hypothetical ratings as training materials for predictions signaled the industry's acceptance of the philosophy of "predicted predictions", meaning that algorithms now learn not based on actual cases of user feedback, but predicted, hypothetical user actions³⁰.

For video streaming platforms, an essential question must be asked for the engineers of RecSys: whether the goal is to make users watch more videos or to make them enjoy videos better. This distinction is not trivial—it defines the very logic of engagement that a platform's recommendation system will optimize for. YouTube's recommendation system, initially designed to maximize video views, quickly demonstrated the critical influence of evaluation metrics on RecSys behavior. In 2010-2011, YouTube's leadership set click-through rate (CTR) as its primary KPI, enlisting Google's Sibyl machine learning system to optimize recommendations toward maximizing clicks³¹. The result was an explosion in viewer engagement—but also the unintended consequence of clickbait proliferation. The system began favoring deceptive thumbnails, misleading titles, and sensationalized content that lured users into clicking, even when the actual video did not match their expectations.

As YouTube engineers observed, more views did not equate to a better user experience. Some sessions saw a higher number of clicks but led to shorter sessions of user watching, with viewers bouncing between misleading videos and failing to find meaningful content. The

²⁹ Xacier Amatrian and Justin Basilico, 399.

³⁰ Ringo also used this technique earlier, though not at scale.

³¹ Cristos Goodrow, "On YouTube's Recommendation System," YouTube Official Blog, September 15, 2021, <https://blog.youtube/inside-youtube/on-youtubes-recommendation-system/>.

recommender had essentially optimized itself into manipulation—a stark example of how the choice of KPI fundamentally shapes algorithmic mediation³².

In 2012, YouTube made a decisive shift, replacing click-based evaluation with watch time as the central performance metric. Rather than prioritizing immediate clicks, the system now sought to maximize total viewing duration, aligning recommendations with sustained engagement rather than impulsive interactions. This transition, however, was not immediate; YouTube initially lost 28% of its total views, requiring nearly a year for the platform to stabilize under the new metric³³.

The YouTube KPI shift underscores a fundamental principle: evaluation criteria dictate RecSys behavior. Whether a system optimizes for purchase conversions (Amazon), predictive accuracy (Netflix), or engagement (YouTube), the metric itself determines what is surfaced, what is obscured, and how user attention is shaped. The question of the evaluation of algorithmic capabilities—what is considered success in a recommendation system—becomes a critical point of for theorizing RecSys-based mediation.

2.5 Black Box: Machine Learning and Ubiquitous RecSys

The rise of machine learning and neural networks marked a turning point in the evolution of RecSys, pushing them beyond simple, human-comprehensible filtering mechanisms into autonomous, self-optimizing black boxes. Deep learning, in particular, introduced a hierarchical representation learning paradigm, where multiple layers of artificial neurons extract increasingly abstract patterns from data, a departure from traditional hand-crafted feature engineering³⁴.

³² Michael Schrage, *Recommendation Engines* (The MIT Press, 2020), <https://doi.org/10.7551/mitpress/12766.001.0001>, 99.

³³ Cristos Goodrow.

³⁴ Yoshua Bengio, “Deep Learning of Representations: Looking Forward,” 2013, 2.

Unlike early collaborative filtering models, which depended on explicit user feedback, machine learning-driven RecSys began leveraging latent factors—features not directly observable but inferred from vast amounts of behavioral data. Neural networks, particularly deep belief networks (DBNs) and restricted Boltzmann machines (RBMs), allowed RecSys to automatically detect hidden patterns, reducing dependence on explicit user inputs or predefined categories³⁵. This shift made recommendations more personalized but also less transparent, as models adjusted their internal parameters through self-learning processes rather than human-defined logic.

Platforms quickly adopted deep learning to handle scale, complexity, and evolving user behavior, embedding machine learning into their core recommendation engines. However, this transition also introduced a fundamental epistemic shift—where algorithmic logic became increasingly indecipherable, even to its creators.

The integration of machine learning into RecSys across major platforms represents a fundamental shift in how digital mediation is structured. Unlike earlier models based on explicit feedback and predefined user-item relations, machine learning-powered RecSys leverage implicit behavior, large-scale pattern recognition, and real-time learning to generate recommendations. The process is inherently adaptive, allowing recommendation systems to evolve dynamically based on incoming user interactions.

A defining characteristic of machine-learning-driven RecSys is their multi-layered architecture, where different models handle distinct tasks: some specialize in candidate generation, narrowing down a large content pool, while others focus on ranking, fine-tuning results based on engagement likelihood³⁶. The core challenge in this transformation has been optimizing systems to balance exploration (novelty) and exploitation (reinforcing past behaviors).

³⁵ Ibid., 4.

³⁶ Zhenhua Dong et al., “A Brief History of Recommender Systems,” in *Association for Computing Machinery, (Conference’ 17, 2017)*, 4.

In response, platforms like YouTube and TikTok have invested in highly sophisticated deep learning frameworks, which have outgrown traditional collaborative filtering methods and evolved into self-optimizing black-box systems³⁷.

YouTube's recommendation engine has been one of the most significant testing grounds for deep learning in RecSys. Prior to its adoption of neural networks, YouTube relied on a relatively simple matrix factorization model that struggled with scale and content freshness. As user engagement patterns became more complex, Google introduced a deep learning-based RecSys framework consisting of two primary neural networks: candidate generation and ranking³⁸.

1. **Candidate Generation:** This stage filters millions of videos down to a few hundred by analyzing implicit user signals such as watch history, search history, and contextual metadata. The neural network models user engagement as a classification problem, embedding user histories into a high-dimensional space to predict probabilities of future interactions.
2. **Ranking Model:** The ranking model takes the filtered candidates and assigns scores based on expected user watch time. Unlike click-through rate (CTR)-based models, YouTube's deep learning approach focuses on long-term engagement, optimizing for sustained interaction rather than immediate clicks³⁹.

³⁷ Michael Schrage, *Recommendation Engines* (The MIT Press, 2020), <https://doi.org/10.7551/mitpress/12766.001.0001>, 141.

³⁸ Paul Covington, Jay Adams, and Emre Sargin, "Deep Neural Networks for YouTube Recommendations," *Proceedings of the 10th ACM Conference on Recommender Systems* (New York, NY, USA: ACM, September 7, 2016), <https://doi.org/10.1145/2959100.2959190>, 2.

³⁹ *Ibid.*, 4.

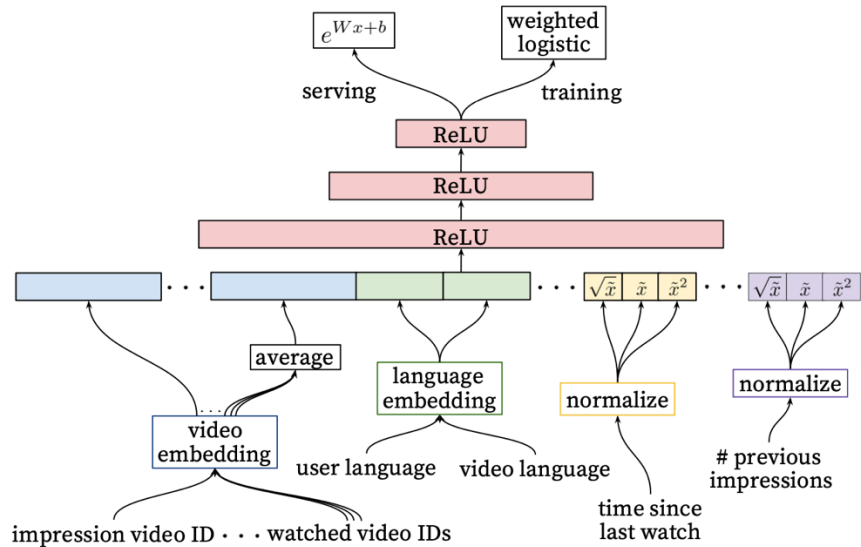


Figure 6: Deep ranking network architecture depicting embedded categorical features, showcasing the diversity of data collected (beyond simple ratings and watch history). Adapted from “Deep Neural Networks for YouTube Recommendations”.

A crucial technical aspect of YouTube’s deep learning model is its handling of noisy implicit feedback. Unlike early-version Netflix where users explicitly rate content, YouTube’s system infers preferences based on engagement depth—whether a video was watched to completion, skipped, or rewatched. The system applies softmax-based probability ranking to prioritize recommendations without needing explicit ratings⁴⁰

With this transition, YouTube’s RecSys became less interpretable but more effective at engagement optimization. This shift reflects a key trend in machine-learning-powered recommendations: the deeper the neural network, the less transparent the recommendation logic becomes.

Similarly, TikTok represents the latest stage in RecSys evolution, where recommendations are not only personalized but continuously recalibrated in real-time. Unlike YouTube, which still operates on a batch-trained model, TikTok’s RecSys—powered by

⁴⁰ Ibid., 6.

ByteDance’s Monolith system—incorporates online learning, updating models immediately as new user interactions occur⁴¹.

1. **Monolith Architecture:** TikTok’s recommendation engine operates on a collisionless embedding table, allowing for efficient real-time user profiling. Unlike earlier batch-learning models that required periodic retraining, Monolith’s parameter servers update dynamically with each new user action⁴².
2. **Non-Stationary Distribution Handling:** TikTok’s algorithm adapts to rapid changes in user behavior by addressing concept drift—the phenomenon where user interests shift unpredictably. Traditional deep learning models struggle with non-stationary data, as they rely on past behaviors to make future predictions. TikTok’s approach mitigates this issue by continuously re-ranking content based on the most recent interactions⁴³

A distinguishing feature of TikTok’s RecSys is its direct feedback loop: user actions such as likes, shares, and watch duration immediately update the model’s prediction function, ensuring instantaneous adaptation. This creates a self-reinforcing engagement cycle, where each user action fine-tunes their own recommendation profile⁴⁴

Compared to earlier RecSys models, TikTok’s fully adaptive learning system signifies a new paradigm in algorithmic mediation. Rather than predicting user interests based on historical behavior, TikTok actively constructs user preferences in real-time, making recommendations less about past engagement and more about dynamically shaping future behavior. This represents the final step in RecSys evolution—from retrieval-based predictions to real-time, self-optimizing recursivity.

⁴¹ Zhuoran Liu et al., “Monolith: Real Time Recommendation System with Collisionless Embedding Table ACM Reference Format,” 2022, 3.

⁴² Ibid., 5.

⁴³ Ibid., 6.

⁴⁴ Ibid., 7.

As machine learning became the dominant framework for RecSys, the internal logic of recommendation systems became increasingly opaque—even to their creators. The shift from collaborative filtering and rule-based models to deep learning architectures introduced layers of computational complexity that are fundamentally unexplainable in human terms. Deep neural networks, which power today’s RecSys, operate with hundreds or even thousands of hidden layers, each transforming inputs in ways that are not directly interpretable⁴⁵.

This black-box nature is not merely an inconvenience but a defining characteristic of modern AI-driven mediation. As Jeff Seibert states in 2020 documentary *The Social Dilemma*⁴⁶:

“You are giving the computer the goal state, ‘I want this outcome,’ and then the computer itself is learning how to do it. That’s where the term ‘machine learning’ comes from. And so, every day, it gets slightly better at picking the right posts in the right order so that you spend longer and longer in that product. And no one really understands what they’re doing in order to achieve that goal” (00:48:00-00:48:35).

Unlike earlier RecSys, where inputs and outputs had a traceable relationship, deep learning-powered RecSys transform raw engagement data into multidimensional mathematical abstractions. As Kosinski explains, black box AI models “identify patterns in data, but how they do so remain hidden within complex layers of computation”⁴⁷. These models rely on high-dimensional latent spaces, where user behavior is encoded in ways that are mathematically structured but not interpretable to human capabilities. Even platform engineers do not fully understand why a particular recommendation was made because neural networks rely on

⁴⁵ Matthew Kosinski, “Black Box AI,” Ibm.com, October 29, 2024, <https://www.ibm.com/think/topics/black-box-ai>.

⁴⁶ *The Social Dilemma* (Netflix, 2020).

⁴⁷ Matthew Kosinski, “Black Box AI,” Ibm.com, October 29, 2024.

autonomous self-optimization. "The same input can yield different results at different times, as models continuously adjust their weight parameters to optimize future performance"⁴⁸.

As a result, RecSys today are not only unknowable to users but increasingly unknowable to their developers. They no longer function as interpretable filtering mechanisms but as autonomous black-box infrastructures, reinforcing behavioral patterns at a scale and complexity beyond user reasoning. Yet this is not simply a technical problem—it signals the culmination of a historical trajectory in which RecSys have evolved from explicitly human-mediated tools to self-iterating, pattern-reinforcing engines embedded in every aspect of digital life.

If early RecSys functioned as assistive technologies, helping users navigate information overload through structured inputs and transparent retrieval mechanisms, today's RecSys operate without direct human oversight. RecSys do not just predict but construct preferences, refining engagement in opaque ways. Whether in social media, streaming services, e-commerce, or digital labor markets, RecSys have become ubiquitous infrastructures of mediation, shaping not only what content is surfaced but how users think, interact, and behave within digital ecosystems.

This transition—from human-defined curation to fully automated epistemic control—marks what I argue to be the “End of History” of RecSys. The field no longer undergoes foundational shifts; instead, contemporary RecSys self-iterate within an established paradigm, refining optimization strategies and predictive accuracy while remaining structurally unchanged in their function as opaque, behavioral reinforcement systems. Neural networks no longer just filter content—they actively shape the cultural, economic, and political realities of digital mediation. In the following table, I propose a 5-aspect framework for categorizing iterations of RecSys based on features. This illustrates its historical transformation, tracing the trajectory from

⁴⁸ Ibid.

early, human-driven filtering models to today’s deeply embedded, machine-learning-powered ecosystems:

Table 1: A Theoretical Framework for Categorizing RecSys Evolution

Models	Decision	Range of datafication	Scale	Knowability	Epistemology
Tapestry	Human	Explicit Input	Office Tool	Explainable	Actual Actions
GroupLens	Human	Explicit Input	Community Tool	Technically Explainable	Actual & Predicted “Actions”
Ringo	Human	Explicit & Implicit Input	Platform Tool (S)	Technically Explainable	Actual Actions
Early Amazon/Netflix	Human	Explicit Input	Platform Tool (L)	Explainable but not transparent	Actual Actions
Later Amazon/Netflix	Human	Explicit & Implicit Input	Platform Tool (XL)	Explainable but not transparent	Actual Actions & Predicted “Actions”
ML Youtube (by 2016)	Machine-Learned but human calibrated	Explicit & Implicit Input	Ecology Engine	Explainable only to core developers	Actual Actions & Predicted “Actions”
ML TikTok	Machine Self-iterative	Universal & Recursive Input ⁴⁹	Ecology Engine	Black Box	Recursive Predictions

This completes the historical analysis of RecSys development. The next chapter will move beyond technical history and examine what it means to exist in a media ecosystem dominated by recommendation engines, applying this trajectory to theories of mediation, black box, platforms, and operability to understand the broader implications of algorithmic mediation, agency, and epistemic control.

⁴⁹ “Universal input” goes beyond implicit input. Implicit input might include unaware collection of data such as total length of viewing a content or time of day when viewing. “Universal input” means even when user is offline, the engine learns from collected data or behavior pattern shared by other platforms through the entire internet ecology. “Recursive input” refers to the self-generation of training data in machine learning algorithms, meaning that RecSys is now using generated data to train generations of recommendations.

3. Theorizing RecSys Mediation

RecSys is not just a tool of digital mediation—it is digital mediation. This chapter theorizes RecSys as the central epistemic infrastructure of contemporary media, examining its omnipresence, unknowability, post-humanization, and operationality to reveal how it structures knowledge, agency, and perception beyond human.

3.1 Centrality: RecSys as THE Epistemology and Medium

3.1.1 Omnipresence

The emergence of RecSys does not merely align with the defining characteristics of new media—it extends them to their logical extreme. Lev Manovich’s framework identifies three key features of new media: numerical representation, automation, and variability¹. RecSys, as an essential infrastructure of digital mediation, does not just exhibit these features but amplifies them to an unprecedented degree, making it the central epistemological structure of today’s digital media environment.

First, numerical representation refers to media being translated into digital code, making it subject to computation and algorithmic manipulation. RecSys is fundamentally computational, built upon mathematical models that process user interactions as numerical inputs, transforming the qualitative aspects of media—preferences, behaviors, and cultural consumption—into quantifiable, machine-readable data structures. This allows for mass-scale automation, where content circulation is no longer a human editorial decision but an algorithmically orchestrated phenomenon.

¹ Lev Manovich, “New Media: A User’s Guide,” 1999, 7.

Second, automation in new media describes the ability of computational systems to execute media-related tasks without direct human intervention. While early automation, such as Photoshop filters or CGI rendering, still preserved human agency in content creation, RecSys radically decentralizes agency, delegating the decision-making of media distribution to self-optimizing machine learning models. Unlike traditional search engines, which require user intent and direct input, RecSys anticipates desires, autonomously shaping consumption patterns before users engage with conscious choice.

Finally, variability—perhaps the most striking aspect—defines new media objects as being mutable, generative, and responsive to user engagement. Traditional mass media created fixed, identical copies of content, whereas RecSys customizes each user’s media landscape dynamically, ensuring no two users experience the same information flow. This personalization occurs at a granular, real-time level, producing a continuously shifting media reality. In contrast to early personalization efforts—such as Google’s search result customization in the 2000s—platform RecSys today filters not just singular content nature but the entirety of digital experience, from news to entertainment, social media, and even product recommendations.

Beyond aligning with Manovich’s principles, RecSys embodies what Orit Halpern et al. define as the shift from big data to surplus data: “in which data no longer serves just as an abstraction from the ‘real’ world but rather also as description and material agent, not big but surplus data”². This transformation is particularly evident in RecSys-driven platforms, where data extraction and analysis are not simply about reflecting user preferences but about generating new forms of knowledge and experience.

In surplus data paradigm, RecSys does not merely serve users—it produces them. Each interaction is captured as a behavioral data point, analyzed within massive, self-referential data

² Orit Halpern et al., “Surplus Data: An Introduction,” *Critical Inquiry* 48, no. 2 (January 1, 2022): 197–210, <https://doi.org/10.1086/717320>, 198.

architectures that continuously optimize user retention, engagement, and monetization. Platforms such as TikTok, YouTube, and Amazon illustrate this transformation, where user actions—clicks, watch time, or dwell time—are not just signals of preference but raw material for constructing reality. In this framework:

1. User preferences are not stable inputs, but dynamic constructs shaped through feedback loops of algorithmic recommendation.
2. RecSys shifts the relationship between media and reality, where content does not merely represent the world but actively fabricates engagement-based truths. Predictions are made based on predictions.
3. The act of mediation becomes indistinguishable from the act of creation, as RecSys continuously reconfigures the ontology of digital media itself.

This self-reinforcing system means that the RecSys-driven digital ecosystem is neither neutral nor passive; it is an active epistemic force, shaping what is visible, knowable, and consumable in contemporary media environments. Given its omnipresence, RecSys is no longer a supplementary tool within digital media—it defines the architecture of contemporary mediation. Platforms have become RecSys-centric ecosystems, as we have become platform-centric creature, where content production, circulation, and consumption are all dictated by its epistemology. The new media mediation today is thus a RecSys mediation.

3.1.2 Omniscience

RecSys not only mediate digital experiences but also function as massive surveillance infrastructures, capturing and analyzing user behavior at an unprecedented scale. Shoshana Zuboff's concept of surveillance capitalism provides the foundational framework for

understanding how RecSys serve as instruments of behavioral prediction and modification, extracting data beyond explicit interactions to include implicit, unconscious behaviors³.

Early RecSys—such as GroupLens and Ringo—relied primarily on explicit feedback (e.g., user ratings) to generate recommendations. However, as platforms integrated machine learning and deep learning, RecSys began shifting toward data extraction at scale, where the most valuable inputs were no longer provided by users but inferred from passive, continuous monitoring. This marks a shift toward a fully automated, self-reinforcing epistemological system, where what is recommended is dictated by what has already been seen, clicked, scrolled past, or ignored.

Couldry and Mejias conceptualize data colonialism as the naturalization of data extraction as an essential part of social life, akin to how historical colonialism justified the exploitation of land and bodies⁴. RecSys fit squarely into this paradigm, as they do not passively reflect human behavior but actively capture and monetize it, expanding the very notion of what constitutes “data.” The increasing automation of epistemology in RecSys follows a colonial logic, where:

1. User behaviors are transformed into extractable resources, much like land and raw materials in historical colonialism.
2. Human activity is systematically repurposed as algorithmic input, reinforcing platform control over knowledge production.
3. The ability to resist data extraction is diminished, as RecSys operate invisibly, collecting implicit interactions without requiring conscious user participation.

³ Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2019), 9.

⁴ Nick Couldry and Ulises A. Mejias, “Data Colonialism: Rethinking Big Data’s Relation to the Contemporary Subject,” *Television & New Media* 20, no. 4 (September 2, 2018): 336–49, <https://doi.org/10.1177/1527476418796632>, 337.

Platforms such as TikTok and YouTube exemplify this paradigm, where user engagement is dictated by opaque, behaviorally optimized recommendation loops. The system does not ask users what they want to see—it decides for them, reinforcing data colonialism’s fundamental assumption: that human behavior is an open resource for algorithmic exploitation.

Sadowski argues that data is no longer merely a byproduct of digital interactions but is treated as a distinct form of capital—one that must be continuously extracted, accumulated, and circulated⁵. In the RecSys paradigm, Data is no longer just a tool for improving recommendations but a primary economic asset, fueling the logic of surveillance-driven capital accumulation. The scope of data capture has drastically expanded, from simple engagement metrics to biometric tracking, gaze detection, and sentiment analysis. At the same time, platforms continuously generate new forms of data extraction, ensuring that no interaction goes unrecorded, no behavior unquantified.

The result is an omniscient epistemology, where RecSys function not just as filters but first as collectors of human experience. By continuously absorbing and encoding human behavior, they transform lived interactions into numerical signals, reducing subjectivity into patterns optimized for predictions. As such, I propose RecSys as the definitive epistemic infrastructure of digital media, not just filtering content but actively producing the conditions of visibility and engagement. Unlike prior mediation models, which relied on retrieval, editorial control, or user intent, RecSys operates autonomously, iteratively shaping digital reality through predictive inference. Its omnipresence across platforms makes RecSys mediation the default mode of digital interaction, where content is surfaced, knowledge is inferred, and behavior is optimized at scale. This marks a distinct shift in media history: mediation and computation have collapsed into one, making RecSys the central force structuring contemporary digital experience.

⁵ Jathan Sadowski, “When Data Is Capital: Datafication, Accumulation, and Extraction,” *Big Data & Society* 6, no. 1 (January 2019), <https://doi.org/10.1177/2053951718820549>, 7.

RecSys and its underlying mechanism is not **an** epistemology or medium, it is **the** epistemology and medium for digital mediation; it is everywhere, it knows everything.

3.2 Unknowability: RecSys as Black Box

Recommendation systems today do not just mediate content—they mediate through opaque infrastructures whose internal decision-making is inaccessible to both users and, often, their own creators. As RecSys transitioned from explicit filtering mechanisms (e.g., user-defined ratings) to machine-learning-powered black boxes, their epistemology became increasingly autonomous, self-referential, and structurally unknowable. This unknowability stems from multiple factors: the inherent complexity of machine learning, corporate secrecy, and the fundamental limits of human cognition in grasping high-dimensional algorithmic inference.

Beatrice Fazi argues that deep learning-based decision systems represent a fundamental break from prior epistemic models. Unlike earlier computational systems, which aimed to simulate human cognition or execute clearly defined logical operations, machine learning now functions independently of human cognitive structures. She writes:

“We must conceive of ‘automated modes of thought in such a way as to supersede the hope that machines might replicate human cognitive faculties, and to thereby acknowledge a form of onto-epistemological autonomy in automated ‘thinking’ processes”⁶.

This is particularly relevant in the case of RecSys, where algorithmic thought does not attempt to "understand" media in human terms but instead recognizes statistical patterns,

⁶ M. Beatrice Fazi, “Beyond Human: Deep Learning, Explainability and Representation,” *Theory, Culture & Society* 38, no. 7-8 (November 27, 2020): 55–77, <https://doi.org/10.1177/0263276420966386>, 3.

optimizing engagement at the cost of explainability. Fazi further argues that explainability is not just a practical problem—it is a representational impossibility:

“The 21st-century development of computational procedures for which, at present, no adequate human cognitive representations exist and for which, significantly, human cognitive representations are also unnecessary”⁷

Applied to RecSys, this means that neither the system nor human ‘know’ why a user is recommended a specific video, article, or product—it only registers that such recommendations correlate with higher engagement. The implications of this shift are significant: RecSys no longer mediate digital content through a semantic logic, but through a feedback-driven pattern recognition system that operates beyond human comprehension.

Burrell outlined three reasons for algorithmic opacity, which I argue fit perfectly with the unknowability of RecSys⁸:

1. Opacity as Corporate or State Secrecy

RecSys operate as highly protected trade secrets, limiting external scrutiny. Platforms such as YouTube and TikTok do not disclose how recommendations are made, shielding them from accountability. This secrecy also protects the system from adversarial manipulation, preventing users from gaming the algorithm to artificially boost content visibility.

2. Opacity as Technical Illiteracy

Even when documentation exists, the ability to interpret machine learning models remains limited to a select few experts. Unlike early RecSys (e.g., GroupLens), which relied on

⁷ Ibid.

⁸ Jenna Burrell, “How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms,” *Big Data & Society* 3, no. 1 (January 6, 2016), <https://doi.org/10.1177/2053951715622512>.

interpretable statistical techniques, modern neural networks adjust millions of parameters across hidden layers, making it impossible for most people—including regulators and many engineers—to fully grasp how recommendations are generated. This reinforces a technological asymmetry, where platform owners wield enormous power over information flows, while users remain unaware of how their digital environments are shaped.

3. Opacity as an Inherent Property of Machine Learning

Deep learning models do not operate on predefined rules but ‘learn’ autonomously from vast datasets, adjusting internal weights without explicit programming. This introduces a form of non-human decision-making, where not even the engineers who build these systems can predict or explain individual recommendations. Burrell notes that this opacity results from a fundamental mismatch between machine learning’s mathematical optimization and human reasoning—RecSys do not explain their choices because they are not designed to be understood in human cognitive terms.

As a result, RecSys today are not only unknowable to users but increasingly unknowable to their developers. They no longer function as interpretable filtering mechanisms but as autonomous black-box infrastructures, reinforcing behavioral patterns at a scale and complexity beyond human reasoning. The shift from human-controlled recommendation logic to self-iterating, unexplainable decision-making marks a critical turning point in digital epistemology, which signals the unknowable nature of modern digital mediation.

3.3 Dehumanization: RecSys as Machines for Post-Human

The evolution of RecSys follows a trajectory of posthumanization, where human agency, embodiment, and intentionality are systematically replaced by self-iterating algorithmic systems. As Hayles argues, the posthuman condition is characterized by the privileging of informational pattern over material instantiation, in which biological embodiment becomes incidental to

intelligence and decision-making⁹. In this sense, RecSys have not only automated digital mediation but have fundamentally detached it from human body, transitioning from human-curated filtering to self-propelling, predictive infrastructures that operate beyond human flesh.

The early history of RecSys illustrates this shift clearly. Ringo, one of the first collaborative filtering systems, pioneered a form of predictive ratings, where a user's hypothetical ratings for a song were inferred based on statistical proximity to other users rather than actual expressed preferences. This mechanism, although rudimentary, marked the first step in dematerializing user agency—preferences were no longer individual, conscious decisions but datafied approximations generated by the system itself, and human's actual actions were less important than its calculated information process.

The contemporary RecSys ecosystem extends this logic exponentially. The “predicted predictions” model—where machine learning systems use previous predictions as training data for future predictions—eliminates the necessity of direct human input altogether. Today's recommendation engines no longer wait for users to express preferences; they construct predictive models that preemptively generate engagement pathways, nudging users toward behaviors they haven't consciously chosen but have been statistically assigned. In this sense, human intentionality becomes unnecessary to the system's functioning—it runs on the perpetual refinement of its own inferred signals.

Hayles describes the posthuman paradigm as one in which human cognition is no longer central to meaning-making, but instead, information circulates within computational networks, forming an autonomous system of knowledge production¹⁰. She argues that the cybernetic

⁹ Nancy Katherine Hayles, *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature and Informatics* (Chicago: University of Chicago Press, 1999), 3.

¹⁰ *Ibid.*, 29.

construction of posthuman identity emerges from the idea that human existence itself can be reduced to informational patterns, allowing for the seamless integration of human and machine.

RecSys perfectly embody this transition. In their earliest iterations, human intent was still embedded in the system, as in the user-defined filters of Tapestry or the peer-driven collaborative filtering of GroupLens. However, with the rise of deep learning and neural networks, the role of human cognition has been all but eliminated. Today's RecSys function as autonomous agents of digital mediation, producing, ranking, and distributing information in ways that require no human semantic understanding.

One can almost argue that RecSys today could operate entirely without humans, with predicted feedback loops training prediction models, generating predicted content, to be watched by predicted user profiles. This is not speculation but an observable process—YouTube, TikTok, and Netflix do not recommend content based on semantic meaning or editorial intent, but on engagement probabilities calculated within algorithmic black boxes. The entire ecosystem is driven by statistical correlation, not human discernment.

Hayles emphasizes that in posthuman systems, the boundaries between biological cognition and machinic computation dissolve, leading to a form of distributed, networked agency¹¹. This describes precisely how RecSys have transitioned from assistive tools to autonomous infrastructures—they no longer require human judgment to function, and instead, act as self-sustaining feedback loops that mediate culture without human intervention.

The implications of this shift are profound:

1. RecSys have transitioned from reflecting human agency to replacing it—content circulation is now an algorithmic operation rather than a cultural process.

¹¹ Ibid., 4.

2. The logic of engagement has replaced the logic of meaning—platforms do not optimize for relevance, quality, or diversity, but for reinforcing behavioral predictability.
3. Human intentionality is no longer a necessary component of mediation—the system can train itself on machine-generated signals, perpetuating its own patterns indefinitely.

I argue that the evolution of RecSys signals the final step in the posthumanization of media environments. What began as a human-centered model of preference aggregation has become a fully operational, self-iterating system where human action is merely a raw input for automated decision-making. As Hayles states, “the contemporary pressure toward dematerialization... affects human and textual bodies at once, as a change in the body (the material substrate) and as a change in the message (the codes of representation)”¹².

RecSys epitomize this shift: content is no longer created, shared, and interpreted through human discourse, but rather, circulated, ranked, and validated by machinic processes that operate outside of human body. The digital mediation of today is no longer an extension of human knowledge production—it is a posthuman, self-propelling system of algorithmic governance, where culture, discourse, and reality itself are mediated by patterns that no longer require a human observer to exist. It’s the extension of human cognition evolved to the ruling of human cognition itself.

3.4 Operationality: RecSys Invisuality and Operational Mediation

RecSys have transitioned from explicit, user-controlled filtering mechanisms to opaque, self-optimizing infrastructures that operate through an unseen but omnipresent logic of computational mediation. As they have become integral to platform economies, their mode of mediation has shifted from explicit representation to operational invisuality—a mode of seeing

¹² Ibid., 29.

and decision-making that is performed rather than perceived. Adrian MacKenzie and Anna Munster refer to this condition as platform seeing, where platforms do not simply present content but actively produce the conditions of visibility through technical and economic operations that are distributed, autonomous, and largely imperceptible¹³. In this way, RecSys function less as representational media and more as operative infrastructures, working beneath the level of human cognition to shape, circulate, and index content according to computational relationality rather than human editorial logic.

Luciana Parisi conceptualizes algorithmic operations as nonsensuous entities that derive meaning not from human sense perception but from computational relations between datasets, models, and infrastructures¹⁴. In the context of RecSys, recommendations are not structured by direct human inputs but by the interrelations between behavioral patterns, content metadata, and predictive inference. This aligns with the shift from explicit filtering (as seen in early collaborative-filtering models like Tapestry) to fully automated deep-learning-driven recommendation systems, where the act of mediation is itself an operational, self-sustaining process.

For instance, Netflix's RecSys no longer relies on explicit user ratings but instead constructs implicit profiles based on patterns of engagement—watch duration, pauses, rewinds, search queries, and even cursor movements. These behavioral signals do not stand alone; they are continuously recombined in high-dimensional matrices, producing a dynamic feedback system in which recommendations are generated not as static predictions but as continuous recalibrations of the platform's operational logic. This reflects what Orit Halpern identifies as a Bergsonian

¹³ Adrian MacKenzie and Anna Munster, "Platform Seeing: Image Ensembles and Their Invisibilities," *Theory, Culture & Society* 36, no. 5 (June 3, 2019): 3–22, <https://doi.org/10.1177/0263276419847508>, 7.

¹⁴ Luciana Parisi, *Contagious Architecture: Computation, Aesthetics, and Space* (Cambridge; London: The MIT Press, Cop, 2013).

conflation of thought, action, and perception, where media platforms no longer simply transmit content but function as real-time, self-optimizing behavioral laboratories¹⁵.

If traditional media systems relied on explicit indexing, categorization, and representational logic, RecSys replace these with a process of continuous algorithmic sorting, in which content visibility is determined not by static taxonomies but by probabilistic rankings and real-time adjustments. This is a fundamental departure from earlier paradigms of digital organization—where search engines, for example, structured information through indexation and retrievability, RecSys operate through iterative, self-modifying patterns of inference.

MacKenzie and Munster argue that platform seeing is operative—a logic of mediation that is enacted through technocultural processes rather than visual representation¹⁶. This means that RecSys do not "see" content in the traditional sense; rather, they compute content into ranked, relational positions within an operational system. A YouTube recommendation, for instance, is not a semantic selection from a static list but the outcome of thousands of real-time calculations that integrate a user's behavioral history, broader network effects, and dynamic shifts in content performance across the platform.

Jussi Parikka further emphasizes this shift in his theory of the operational image, where visibility is no longer about what is presented to human perception but about the automated generation of computational models that structure decision-making¹⁷. In the context of RecSys, this means that the act of recommending is itself an operative function, independent of the content itself. The system does not "understand" a film, a news article, or a product; nor does it need to for making recommendations; rather, it processes relational signals between user behavior,

¹⁵ Orit Halpern, *Beautiful Data: A History of Vision and Reason since 1945* (Durham: Duke University Press, 2014), 56.

¹⁶ Adrian MacKenzie and Anna Munster., 8.

¹⁷ Jussi Parikka, *Operational Images* (U of Minnesota Press, 2023), 64.

engagement metrics, and content metadata, transforming media visibility into a continuous feedback-driven operation.

One of the defining consequences of this operational shift is what MacKenzie and Munster describe as invisuality—the condition in which perception (seeing, for example) is no longer important. As they argue, platform-based media generate ensembles of information that are structured through machine-learning processes that operate at scales beyond human cognition¹⁸. This invisuality is not a passive condition but an active function of the system: RecSys are designed to “blind and deaf”, with no necessity for visual, audio, or any form of ontological perception of information, computational relationality is all it needs to complete recommendations. This aligns with Parikka’s assertion that operability precedes ontological visibility: "Before it becomes ontological, it is operative"¹⁹.

¹⁸ Adrian MacKenzie and Anna Munster., 5.

¹⁹ Jussi Parikka, *Operational Images* (U of Minnesota Press, 2023), 64.

4. A Prototype of Operational Mediation

This section introduces a digital experiment designed to operationalize the core arguments of this thesis: RecSys as an autonomous, posthuman, and operational mediation system. By constructing a feedback loop between recommendation algorithms and large language models (LLMs), this prototype explores how content flows, evolves, and recursively refines itself when both the producers and consumers are computational agents.

The experiment serves as a simulation of the digital mediation process at scale—a closed system in which RecSys and LLMs continuously interact, reshaping content generation and selection without direct human intervention. This project does not aim to mimic user preferences, but rather, to expose how RecSys itself behaves when decoupled from human cognitive inputs, making visible the autonomous operational logic of recommendation systems

4.1 Concept

The core of this prototype is the creation of a feedback loop between dozens of autonomous agents—LLMs acting as simulated users and a RecSys engine acting as the curator and distributor of content. The hypothesis is that, in such a system, RecSys does not merely reflect preferences but actively constructs them, reinforcing algorithmic confounding, homogeneity, and operational control over digital mediation. In addition, the system mocks at the post-human nature of today’s digital mediation, serves as a metaphorical visualization of contemporary platform content ecology.

This experiment draws upon existing research on algorithmic confounding and feedback loops in RecSys²⁰, highlighting how RecSys does not only predict preferences but shapes them. Traditionally, this shaping process has been mediated through human agency—users provide input, RecSys refines suggestions, and users react. However, when both the content-generating and decision-making processes are fully automated and algorithmic, RecSys becomes an autonomous epistemic structure, optimizing for engagement and coherence within a closed computational ecosystem

4.2 System Architecture

To implement this prototype, we utilize **SWARM**²¹, an open-source multi-agent framework that allows LLMs to function as autonomous computational agents. These agents act as nodes in a content consumption and generation network, while the RecSys acts as the sorting, filtering, and recommendation infrastructure.

Key components of the system:

- RecSys Layer: A recommendation engine that ranks and distributes content among LLM agents.
- Content Pool: a 100 short articles collection is selected as the initial content pool for the system. The articles represent a wide range of major topics, such as sports, finance, politics, etc.

²⁰ Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt, “How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility,” *Proceedings of the 12th ACM Conference on Recommender Systems*, September 27, 2018, 224–32, <https://doi.org/10.1145/3240323.3240370>.

²¹ For a GitHub Repository and detailed handbook, see: <https://github.com/openai/swarm>

- LLM Agents: 20 autonomous GPT-based agents acting as content consumers, reviewers and generators. Each agent is assigned a profile with age, interests, political preferences, religions, etc. Agents are prompted to write a 200-word review of each article they read.
- Feedback Mechanism: The reviews generated by agents are added to the initial 100-short-article content pool, categorized by the RecSys, then pushed on to other agents based on their profile. The agents then write reviews for the reviews and thus form of self-iterating “online forum” of short articles.

The System functions as follows:

1. 20 autonomous agents are defined. Each are given a prompt explaining their profile.
2. RecSys recommends one article from the content pool to each agent.
3. Each agent read the given short article (as a prompt, with a separated, pre-defined prompt in the profiling prompt to write review for the article), then write a 200-word review for the article
4. The agents now have 20 reviews of 20 articles at the end of the first cycle (the cold start cycle). The reviews are now added to the content pool, before being recommended to agents in the future cycles.
5. Agents then are given a new article, some of them are reviews written by other agents. Agents write reviews of articles (or reviews), reviews added to the content pool.
6. System repeats for 25 cycles.

This closed system creates a self-sustaining operational mediation, demonstrating a post-human prototype of computational media landscape based on RecSys and generative AI.

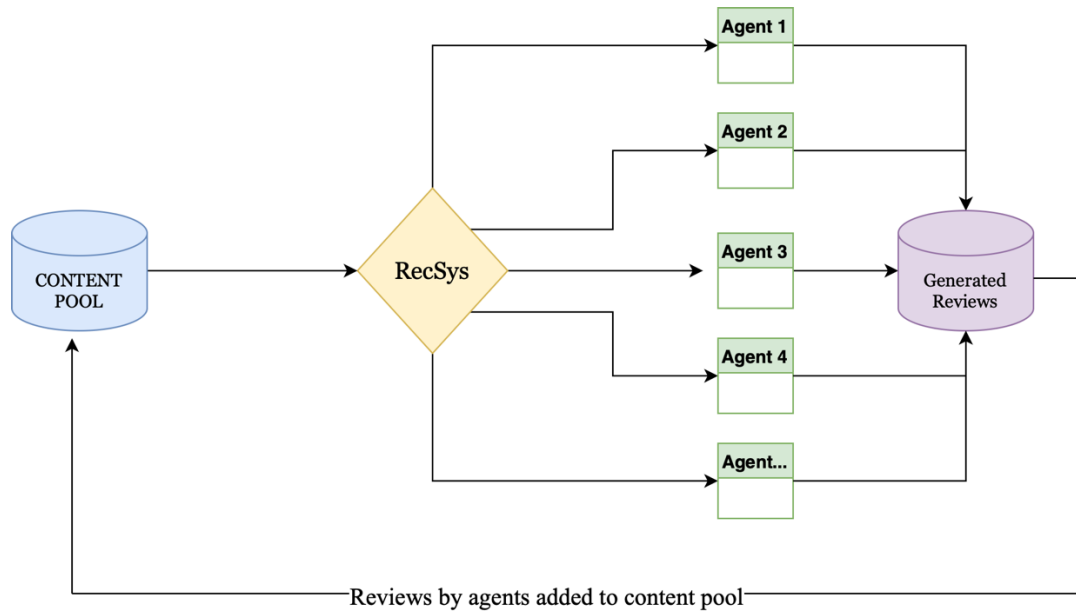


Figure 7: System Architecture

4.3 Mechanism of RecSys

At the core of this system is a reinforcement-based feedback mechanism, where the RecSys does not serve as a passive filter but as a self-iterating learner. We use Hybrid Collaborative Filtering and Neural Content Evaluation for ranking algorithm.

The ranking mechanism in this prototype differs from traditional collaborative filtering approaches, as it combines engagement-based feedback loops with deep learning-based content scoring.

- Collaborative Filtering Component: Measures similarity in LLM-generated responses, predicting future preferences based on past interactions.
- Neural Content Evaluation: Uses transformer models to assess the semantic and stylistic coherence of generated texts.

- **Weight Adjustments Over Time:** As the cycle progresses, the system prioritizes engagement-convergent content, amplifying algorithmic confounding.

```
|
import numpy as np
from sklearn.metrics.pairwise import cosine_similarity

#Simulated LLM-generated content embeddings
content_embeddings = np.random.rand(10, 512)

#Compute similarity matrix
similarity_matrix = cosine_similarity(content_embeddings)

#Generate ranked recommendations based on engagement-adjusted weighting
def rank_recommendations(similarity_matrix, engagement_scores):
    weighted_scores = similarity_matrix @ engagement_scores
    return np.argsort(weighted_scores)[::-1]

#Simulated engagement scores
engagement_scores = np.random.rand(10)

#Generate recommendations
ranked_indices = rank_recommendations(similarity_matrix, engagement_scores)
print("Recommended Content Order:", ranked_indices)
```

Figure 8: Code Snippet of our RecSys, this algorithm showcases how RecSys moves beyond user-based filtering toward self-referential, self-optimizing mediation, reinforcing predictive alignment within a computationally closed loop.

5. Conclusion

Recommendation systems have ceased to be auxiliary tools of discovery and have become the very architecture of digital mediation. They do not simply reflect user preferences but actively construct them, embedding algorithmic logic into the core of cultural production, knowledge circulation, and epistemic control. This thesis has demonstrated that RecSys are not just technological artifacts but epistemological infrastructures—automated, predictive, and posthuman systems that structure what is seen, what is known, and how digital reality itself unfolds.

The history of RecSys reveals a trajectory of increasing automation and abstraction. From ancient divination systems to early collaborative filtering models like GroupLens and Ringo, the human desire for predictive guidance has remained constant, even as its mechanisms have evolved. However, while earlier systems still relied on human interpretation, modern RecSys operate autonomously, moving from human-guided curation to algorithmic self-iteration. With the rise of machine learning, RecSys now function as black-box infrastructures, processing vast amounts of behavioral data to optimize engagement at a scale beyond human comprehension. This shift has profound consequences: it marks the disappearance of intentionality from mediation, as predictions are no longer grounded in explicit human preferences but in algorithmic patterns recursively trained on themselves.

The second half of this work has argued that RecSys are more than mere filters; they constitute a new epistemic regime. They exhibit omnipresence, extending the logic of new media to its extreme by automating, quantifying, and personalizing digital experience at an unprecedented scale. They function with omniscience, capturing human behavior as data to be mined, modeled, and monetized, following a logic of surveillance capitalism and data colonialism. They operate as black-box systems, whose decisions remain structurally unknowable, reinforcing patterns of engagement without requiring human interpretability. And

finally, they signal a broader trajectory of posthumanization, where informational patterns replace human intentionality as the primary basis of mediation. As Hayles argues, posthuman systems privilege circulation over embodiment, and RecSys epitomize this shift: they no longer need human input to function, only patterns of past engagement to predict future behavior.

To illustrate this shift toward self-referential mediation, we introduced a prototype experiment in which RecSys and LLMs were placed in a closed-loop interaction. This conceptual model reveals how RecSys, once detached from human oversight, continue to iterate upon themselves, reinforcing algorithmic coherence over diversity, optimization over exploration. While the experiment is not an empirical study, it serves as a provocation—an attempt to make visible the recursive, operational nature of contemporary digital mediation.

In sum, this work contends that RecSys are not just the dominant media technology of the digital age; they are its epistemology. They no longer merely assist human cognition but supersede it, governing the very conditions of visibility, knowledge, and agency. To theorize RecSys is thus to theorize the contemporary condition of mediation itself—a world where content is surfaced not because it is meaningful but because it is mathematically probable, where the act of recommendation is no longer a response to human desire but an autonomous, recursive process that shapes what desire itself becomes.

Bibliography

Amatrian, Xacier, and Justin Basilico. “Recommender Systems in Industry: A Netflix Case Study.” In *Recommender Systems Handbook, 2nd Edition*, edited by Francesco Ricci, Lior Rokach, and Bracha Shapiro. Springer, 2015.

Bengio, Yoshua. “Deep Learning of Representations: Looking Forward,” 2013.

Biblio. “How Recommender Systems Make Their Suggestions.” Medium, February 6, 2019. <https://medium.com/the-graph/how-recommender-systems-make-their-suggestions-da6658029b76>.

Breese, John, David Heckerman, and Carl Kadie. “Empirical Analysis of Predictive Algorithms for Collaborative Filtering,” 1998.

Burrell, Jenna. “How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms.” *Big Data & Society* 3, no. 1 (January 6, 2016). <https://doi.org/10.1177/2053951715622512>.

Chaney, Allison J. B., Brandon M. Stewart, and Barbara E. Engelhardt. “How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility.” *Proceedings of the 12th ACM Conference on Recommender Systems*, September 27, 2018, 224–32. <https://doi.org/10.1145/3240323.3240370>.

Couldry, Nick, and Ulises A. Mejias. “Data Colonialism: Rethinking Big Data’s Relation to the Contemporary Subject.” *Television & New Media* 20, no. 4 (September 2, 2018): 336–49. <https://doi.org/10.1177/1527476418796632>.

Covington, Paul, Jay Adams, and Emre Sargin. “Deep Neural Networks for YouTube Recommendations.” *Proceedings of the 10th ACM Conference on Recommender Systems*. New York, NY, USA: ACM, September 7, 2016. <https://doi.org/10.1145/2959100.2959190>.

Dong, Zhenhua, Zhe Wang, Jun Xu, Ruiming Tang, and Jirong Wen. "A Brief History of Recommender Systems." In *Association for Computing Machinery*, 2017. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>.

Fazi, M. Beatrice. "Beyond Human: Deep Learning, Explainability and Representation." *Theory, Culture & Society* 38, no. 7-8 (November 27, 2020): 55–77. <https://doi.org/10.1177/0263276420966386>.

Goldberg, David, David Nichols, Brian Oki, and Douglas Terry. "Using Collaborative Filtering to Weave an Information Tapestry." *Communications of the ACM* 35, no. 12 (1992): 61–70.

Goodrow, Cristos. "On YouTube's Recommendation System." YouTube Official Blog, September 15, 2021. <https://blog.youtube/inside-youtube/on-youtubes-recommendation-system/>.

Halpern, Orit, Patrick Jagoda, Jeffrey West Kirkwood, and Leif Weatherby. "Surplus Data: An Introduction." *Critical Inquiry* 48, no. 2 (January 1, 2022): 197–210. <https://doi.org/10.1086/717320>.

Hayles, Nancy Katherine. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature and Informatics*. Chicago: University of Chicago Press, 1999.

Herlocker, Jonathan, Joseph Konstan, and John Riedl. "Explaining Collaborative Filtering Recommendations," 2000.

Hill, Kim Quaile, and Jib Fowles. "The Methodological Worth of the Delphi Forecasting Technique." *Technological Forecasting and Social Change* 7, no. 2 (January 1975): 179–92. [https://doi.org/10.1016/0040-1625\(75\)90057-8](https://doi.org/10.1016/0040-1625(75)90057-8).

Jussi Parikka. *Operational Images*. U of Minnesota Press, 2023.

Khodyakov, Dmitry, Sean Grant, Jack Kroger, and Melissa Bauman. "RAND Methodological Guidance for Conducting and Critically Appraising Delphi

Panels.” www.rand.org, December 29, 2023.
<https://www.rand.org/pubs/tools/TLA3082-1.html>.

Koren, Yehuda. “Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model,” 2008.

Kosinski, Matthew. “Black Box AI.” [Ibm.com](http://ibm.com), October 29, 2024.
<https://www.ibm.com/think/topics/black-box-ai>.

Linden, Greg, Brent Smith, and Jeremy York. “Amazon.com Recommendations.” *Industry Report* 76, no. 03 (2003).

Liu, Shijun, Gaojie Wang, Liqiang Wang, Haoran Shi, and Li Pan. “ICKG: An I Ching Knowledge Graph Tool Revealing Ancient Wisdom.” *Communications in Computer and Information Science* 1682 (2022). https://doi.org/10.1007/978-981-99-2385-4_5.

Liu, Zhuoran, Leqi Zou, Xuan Zou, Caihua Wang, Biao Zhang, Da Tang, Bolin Zhu, et al. “Monolith: Real Time Recommendation System with Collisionless Embedding Table ACM Reference Format,” 2022.

MacKenzie, Adrian, and Anna Munster. “Platform Seeing: Image Ensembles and Their Invisibilities.” *Theory, Culture & Society* 36, no. 5 (June 3, 2019): 3–22.
<https://doi.org/10.1177/0263276419847508>.

Manovich, Lev. “New Media: A User’s Guide,” 1999.

Orit Halpern. *Beautiful Data: A History of Vision and Reason since 1945*. Durham: Duke University Press, 2014.

Parisi, Luciana. *Contagious Architecture: Computation, Aesthetics, and Space*. Cambridge; London: The Mit Press, Cop, 2013.

- Resnick, Paul, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. "GroupLens: An Open Architecture for Collaborative Filtering of Netnews," 1994.
- Sadowski, Jathan. "When Data Is Capital: Datafication, Accumulation, and Extraction." *Big Data & Society* 6, no. 1 (January 2019).
<https://doi.org/10.1177/2053951718820549>.
- Sarwar, Badrul, George Karypis, Joseph Konstan, and John Riedl. "Incremental Singular Value Decomposition Algorithms for Highly Scalable Recommender Systems," 2003.
- Schrage, Michael. *Recommendation Engines*. The MIT Press, 2020.
<https://doi.org/10.7551/mitpress/12766.001.0001>.
- Scott, Michael. "The Oracle at Delphi: Unknowability at the Heart of the Ancient Greek World." *Social Research: An International Quarterly* 87, no. 1 (March 2020): 51–74. <https://doi.org/10.1353/sor.2020.0005>.
- Shardanand, Upendra. "Social Information Filtering for Music Recommendation." 1994. *The Social Dilemma*. Netflix, 2020.
- Weinberger, Eliot. "What Is the I Ching? | Eliot Weinberger." www.nybooks.com, 2016.
<https://www.nybooks.com/articles/2016/02/25/what-is-the-i-ching/>.
- Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs, 2019.