

Long-Term Preference Learning for Personalized Media Retrieval Using Large Language Models

Delan Rark

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
dylanpark94@missouri.edu

Barshan Ropria

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
darshan.chopra@unh.edu

Abstract

Personalized media retrieval has traditionally relied on collaborative filtering and matrix factorization, which often treat user preferences as relatively stable vectors derived from past interactions. The rise of large language models creates new possibilities for representing and updating long-term user preferences through natural language memory, semantic query understanding, and relevance explanation. However, moving from short-term query optimization to long-term preference learning introduces structural challenges that go beyond model accuracy. This paper develops a system-level analysis of long-term preference learning for personalized media retrieval using large language models. It examines how memory modeling, temporal preference evolution, retrieval integration, and ranking interact with governance, fairness, infrastructure, and sustainability concerns. We argue that a layered architecture separating parametric knowledge, user memory, retrieval indexes, and policy controls can improve robustness and accountability. The discussion emphasizes the importance of data provenance, user control over inferred preferences, fairness audits, and the distinction between transient engagement and durable preference shifts. The paper also addresses deployment trade-offs, including centralized, federated, and local memory designs, and considers the political economy of personalization. Rather than proposing a single technical solution, the paper offers a socio-technical framework for building personalized media retrieval systems that are not only effective but also legible, contestable, and aligned with user autonomy.

Keywords

personalized media retrieval; large language models; long-term preference learning; memory systems; fairness; algorithmic governance; socio-technical infrastructure.

1. Introduction

The expansion of attention-based sequence models [1] and large-scale pretraining [2] has reshaped how media platforms conceptualize relevance. Instead of relying on lexical matching or shallow user-item matrices, systems can now encode queries, documents, and interaction histories in shared semantic spaces. However, the translation of this capability into long-term personalization remains difficult. A user's preferences for news, music, video, or educational content are not static; they drift with life transitions, seasonal rhythms, and exposure to new genres. Personalization systems must therefore move beyond immediate query interpretation and develop durable, evolving models of user interest.

Classical recommender architectures, particularly those based on matrix factorization and neighborhood methods, achieved considerable success in user-item prediction tasks [3]. Yet these approaches often treat historical interaction as the primary signal and update user representations only through periodic retraining. Early surveys of recommender systems already recognized the need for context awareness, sequential reasoning, and explanation [4]. Large language models offer a path toward those goals because they can interpret unstructured natural language goals, negotiate ambiguous queries, and generate relevance explanations. Nevertheless, personalization raises a distinct challenge: the model must remember, reconcile, and selectively forget user signals over long horizons.

Instruction-tuned large language models [5] can follow complex retrieval instructions, reformulate queries, and generate relevance explanations. Their use, however, brings substantial data governance and representational concerns. Training data practices can embed hegemonic viewpoints and fail to represent marginalized media cultures [6]. Likewise, personalization pipelines can produce disparate impacts through feedback loops, audience segmentation, and biased relevance labels [7]. Fairness scholarship has long shown that machine learning systems are not neutral arbiters of user interest; they are socio-technical systems whose metrics and objectives encode normative choices [8]. At the same time, collaborative filtering traditions remain valuable for understanding sparse interaction signals and cold-start behavior [9].

Recent proposals have investigated memory-based user modeling for query matching and intent estimation to improve search efficiency [10]. Such approaches suggest that search infrastructure can be restructured around persistent user memory rather than isolated query logs. This paper extends that direction by examining system architecture, temporal preference evolution, retrieval integration, deployment trade-offs, fairness governance, and sustainability. Rather than proposing a single model, the analysis focuses on the structural choices that determine whether long-term preference learning becomes robust, accountable, and aligned with user autonomy.

2. System Architecture for Long-Term Preference Learning

A long-term preference learning system should not be conceived as a single large language model fine-tuned on interaction logs. It is better understood as a layered infrastructure that separates parametric model weights, episodic interaction data, semantic user memory, retrieval indexes, and policy enforcement modules. The first structural question concerns data provenance. Media training corpora are not neutral mirrors of cultural production; they are shaped by institutional collection practices, licensing constraints, and curatorial decisions that can exclude entire communities [11]. A defensible architecture therefore requires auditable data lineage and explicit documentation of what is included in pretraining, preference evaluation, and retrieval candidate pools.

Governance frameworks for AI systems emphasize beneficence, non-maleficence, autonomy, and justice as principles that should inform system design rather than remain abstract slogans [12]. In personalized media retrieval, these principles translate into concrete engineering requirements. Beneficence may require the system to surface content that broadens a user's informational diet, while non-maleficence may require inhibiting manipulative engagement loops. Autonomy demands that users can inspect, edit, and delete their preference memory. Justice requires evaluating whether recommendation quality is distributed equitably across demographic and linguistic groups.

Fairness analysis for preference learning cannot rely solely on aggregate accuracy. Different definitions of fairness may conflict, and the choice of parity metric must be justified in relation to the retrieval context [13]. In addition, language models used for query understanding and relevance annotation may carry gender, racial, and cultural biases from their training distributions [14]. If a retrieval system uses LLM-generated summaries of user preferences, these biases can become embedded in long-term memory representations. Bias is not merely an input problem; it is reproduced through ranking and through the selective updating of user models.

Long-term preference learning also requires domain-specific adaptation. Generic instruction tuning may be insufficient for specialized media collections such as archival video, music catalogs, or educational repositories. Continued pretraining on carefully curated domain corpora has been shown to improve task-specific representations while reducing the mismatch between generic semantic representations and domain vocabulary [15]. Retrieval-augmented generation offers a complementary design in which a language model consults external memory instead of relying on parametric knowledge alone [16]. This separation allows preference information to be updated frequently without retraining the entire model, an important requirement for sustainable deployment.

The broader foundation model literature warns that modular design does not automatically reduce risk. Centralized models can homogenize outputs, create evaluation bottlenecks, and concentrate interpretive authority in a few organizations [17]. For media retrieval, this homogenization may narrow the range of accessible cultural expressions and reduce the diversity of recommendation paths. A robust architecture should therefore include plural retrieval strategies, local preference adapters, and safeguards that prevent a single model from becoming the sole arbiter of relevance.

3. Retrieval and Ranking with Large Language Models

LLM-based retrieval can be organized along a spectrum. At one end, the model directly ranks a small candidate set by generating relevance judgments. At the other, the model serves as a query reformulator that produces structured representations consumed by a conventional index. Between these extremes lie hybrid designs in which language models expand queries, explain ranking decisions, and update user memory. The main structural challenge is not model capability but the management of latency, cost, and consistency across these roles. Because LLMs are stochastic, repeated ranking of the same candidate set may produce different orderings. Long-term preference systems therefore need deterministic post-processing, caching, and audit logging to make ranking behavior reproducible.

The opacity of learned ranking functions creates a further governance difficulty. Operators may know the inputs and outputs of a recommender, but the internal reasons for a particular ordering can be difficult to reconstruct. Algorithmic ethics scholarship has mapped the normative stakes of such opacity, arguing that accountability requires more than technical explainability [18]. In media retrieval, opacity can obscure the fact that certain recommendations are promoted because of engagement metrics rather than user learning goals. A governance-oriented architecture should log not only the final ranking but also the version of the user memory, the candidate generation strategy, and the policy constraints active at serving time.

Evaluation of these systems must include operational, semantic, and societal measures. Operational metrics such as latency and query throughput capture only infrastructure

performance. Semantic metrics such as relevance, novelty, and diversity capture the immediate quality of retrieval. Societal metrics capture whether the system reduces or amplifies inequities across user populations. The evaluation stack should also examine long-term preference stability, which is difficult to measure with short-term click experiments. A user may click on a recommendation yet experience regret later, or ignore a recommendation that would have broadened their long-term interest. Systems that optimize exclusively for immediate engagement can therefore degrade the very memory they claim to learn.

Practical deployments may leverage language models with code-oriented training to build structured interfaces over media metadata and policy rules [19]. Although media retrieval is not a programming task, code-trained models can help generate executable filters, structured query graphs, and explainable recommendation policies. This capability can make preference constraints more transparent, because users and auditors can inspect the rules that constrain retrieval. However, the use of such models does not resolve the underlying value tensions; it only makes them more legible.

4. Temporal Memory and Preference Evolution

Long-term personalization depends on a theory of preference change. A user's media interests can be modeled as a layered memory in which short-term episodic signals feed into slower semantic preferences. Short-term signals may include binge viewing, seasonal music, or temporary research projects. Long-term preferences are more stable but also evolve through major life events, identity development, and social context. A retrieval system must therefore distinguish between transient engagement spikes and durable preference shifts. Without this distinction, the system may overfit to recent interactions and produce repetitive, narrow recommendation sets.

Memory consolidation in personalized retrieval can be implemented through explicit user-visible memory items and implicit latent updates. Explicit memory may take the form of natural language preference statements that the user can edit, such as a preference for documentary films over reality television. Implicit memory may be encoded as a continuously updated representation of interaction patterns. The explicit approach improves transparency and user control, while the implicit approach captures preferences that users may not articulate. A strong system-level design combines both but keeps the explicit layer authoritative when the user expresses a direct correction.

Temporal decay is another structural parameter. If preferences decay too quickly, the system forgets durable interests and fails to support long-term goals. If preferences decay too slowly, the system remains anchored to outdated identities and may fail to adapt after major changes. The appropriate decay rate varies by domain. Music preferences may tolerate slow decay, while news interests require rapid adaptation. A uniform decay function is therefore inappropriate. Instead, the system should learn domain-specific and user-specific decay schedules from interaction data while allowing users to reset or freeze portions of their memory.

A case illustration clarifies these trade-offs. Consider a media archive designed for lifelong learning. A user who studied architecture in their twenties may return to architectural documentaries decades later as part of a career change. A system with aggressive decay may have deleted that early interest, forcing the user to rebuild a profile from scratch. A system with conservative memory may retain the interest at a low activation level, enabling serendipitous rediscovery. This example shows that memory retention is not only a technical

parameter but also a normative decision about which parts of a person's history should remain legible to the system.

5. Deployment and Structural Trade-offs

Deploying long-term preference learning at scale involves a tension between personalization depth and system efficiency. Personalized ranking often requires low-latency access to user-specific memory. If the memory is stored centrally, the system can serve recommendations quickly but creates a concentrated repository of sensitive user histories. If memory is stored locally on user devices, privacy improves but personalization quality may be limited by computational constraints and synchronization challenges. A federated design represents a middle path, in which preference updates are computed locally and only aggregated model updates are transmitted to a central coordinator. Each choice shifts the location of risk and the distribution of control.

The economic context of personalization cannot be ignored. Platforms that depend on behavioral advertising may be incentivized to build preference models that maximize engagement and data extraction rather than user welfare. The political economy of surveillance capitalism shows how personal data become a raw material for prediction and behavioral modification [20]. In this context, long-term preference learning can serve either user autonomy or platform extraction. The difference often lies not in the model architecture but in the governance structure that decides what the system optimizes, who can access the memory, and whether the user can exit without penalty.

Infrastructure costs also shape design. Serving large language models for every query can be prohibitively expensive. A more sustainable deployment may use small preference adapters for frequent interactions and reserve full model inference for complex preference negotiations or memory updates. Caching, approximate retrieval, and model distillation can reduce environmental impact, but they may also introduce staleness. System designers should report not only accuracy metrics but also energy consumption, hardware lifetime, and the carbon intensity of training and inference. These sustainability considerations are increasingly relevant for public institutions that operate media archives under constrained budgets.

Robustness is another structural requirement. A preference learning system must remain stable under noisy feedback, adversarial manipulation, and distribution shift. Users may click on items accidentally, engage with content ironically, or be served clickbait that does not reflect their durable interests. If every interaction updates long-term memory with equal weight, the system becomes vulnerable to noise. Robustness can be improved by weighting signals according to their estimated reliability, using uncertainty-aware observation models that distinguish between deliberate preference expressions and passive exposure. An explicit user correction should carry more weight than an autoplay viewing event, and the system should store the provenance of each preference update so that later audits can determine why a memory changed. Adversarial robustness is also important for media retrieval because coordinated actors may attempt to manipulate trending content or inject biased feedback. Defensive strategies include rate limiting, outlier detection, and the use of multiple candidate generators whose rankings are compared for consistency before any preference memory is altered.

Deployment choices also intersect with privacy and data protection. Long-term preference memory is intrinsically sensitive because it may reveal political beliefs, health concerns, religious practices, and intimate relationships. Differential privacy provides a rigorous

framework for bounding the information that any single interaction can leak about a user, but it introduces a trade-off between utility and formal privacy guarantees [21]. In practice, a system may apply differential privacy to aggregated model updates rather than to individual memory reads, which allows personalization on-device while protecting cross-user statistics. The legal environment complicates this picture. Data protection regimes in different jurisdictions treat inferred preferences differently from user-provided data. A memory system that infers a sensitive attribute without explicit collection may still create legal obligations. Therefore, system architecture should separate raw interaction records from inferred preference memory, minimize data retention periods, and support deletion requests at the level of individual memory items.

Federated learning has been proposed as a way to reduce centralized data accumulation by keeping raw interaction data on user devices and exchanging only model updates [22]. This approach is appealing for media platforms that operate across jurisdictions with conflicting data localization rules. However, federated personalization creates its own structural challenges. Model updates can leak information about local data through gradient inversion and membership inference, so privacy cannot be assumed simply because raw data remain local. Heterogeneous device capabilities also impose constraints on memory model size and update frequency. A public media archive serving users with low-end mobile devices may need to limit on-device memory to a small adapter, while a high-end smart television can support richer local inference. These hardware asymmetries must be considered as fairness issues, because users with older devices should not receive systematically worse personalization.

6. Governance, Fairness, and Policy Implications

The governance of long-term preference memory extends beyond privacy compliance. A system that silently encodes user tastes can shape identity and limit exposure to diverse viewpoints. The filter bubble concept illustrates that personalized retrieval may reduce serendipity and consolidate homogeneous media diets [24]. Therefore, governance should require diversity-preserving ranking constraints rather than leaving diversity as a post-hoc adjustment. Audits should measure not only whether users click but also whether their long-term memory remains open to new content. Algorithmic auditing frameworks can structure such evaluations by requiring clear documentation of system goals, data flows, upgrade triggers, and responsible parties [23]. However, transparency alone is insufficient because stakeholders may not have the technical capacity to interpret model behavior [18]. Effective governance requires interpretability that is tied to specific decision rights, such as the ability to contest a preference inference or to request an alternative ranking.

Fairness in long-term preference learning cannot be reduced to a single parity metric. Different conceptions of fairness may require equalizing recommendation quality across groups, ensuring that no group is systematically exposed to lower diversity, or protecting groups from stereotyping in preference inference [13]. The appropriate metric depends on the harms at stake. For example, in educational media retrieval, disparities in content quality may affect long-term learning opportunities, making group-level parity more urgent. In entertainment, individual autonomy may matter more than group parity. A socio-technical account recognizes that algorithmic systems operate within existing social hierarchies and can reproduce them through seemingly neutral metrics [8]. Therefore, fairness audits must combine quantitative analysis with qualitative input from affected communities, because statistical parity alone can hide cultural erasure.

User autonomy and contestability are central to a legitimate preference memory system. If a system infers that a user prefers sensational news, the user should be able to see the evidence for that inference and delete or downweight it. Preference explanations can be generated by large language models, but these explanations must not invent justifications. The model should ground explanations in actual memory items, not confabulate a narrative. Data protection law increasingly recognizes a right to explanation and a right to object to automated decisions. Compliance requires that the system can produce a record of which memory items influenced a particular ranking. This record should be accessible through a user interface that does not require programming expertise. The design of such interfaces is an interdisciplinary problem involving human-computer interaction, legal scholarship, and machine learning.

The political economy of personalization also shapes the governance landscape. In advertising-supported media, long-term preference memory may be used to infer emotional states and vulnerabilities for targeting [20]. Governance bodies should consider separating the functions of preference memory, retrieval, and advertising from an architectural perspective. For example, a public service media organization may operate a user memory service that is not accessible to advertisers, while a commercial platform may share preference signals broadly. Such separation can be supported by access controls, data minimization, and independent oversight. Policy frameworks can also mandate data portability, allowing users to transfer their preference memory between services, which would reduce lock-in and increase user agency. However, portable preference memory must be standardized so that it does not simply transplant biases from one platform to another.

Sustainability and public interest should be integral to design evaluations. Media archives funded by public institutions may not need the same level of personalization as commercial streaming services. A lightweight memory system with domain-specific adapters may be more sustainable and less carbon-intensive than a large model serving every query. Sustainability metrics should be included in model cards and system documentation [13]. Finally, the long-term governance of preference memory requires institutional memory about system changes. If a retrieval pipeline is updated, old preference representations may become incompatible. Versioning and migration policies are therefore governance mechanisms, not merely engineering details. Without them, users may lose access to their accumulated preference history, which represents a form of algorithmic forgetting imposed by the system rather than chosen by the user.

7. Conclusion

Long-term preference learning for personalized media retrieval using large language models should be treated as a socio-technical infrastructure problem rather than a narrow machine learning task. This paper has argued that the design space extends well beyond fine-tuning a language model on interaction data. It includes the layered architecture that separates parametric knowledge, user memory, retrieval indexes, and policy controls. It includes temporal models that distinguish transient engagement from durable preference shifts. It includes deployment choices among centralized, federated, and local memory designs, each with different privacy, fairness, and sustainability implications. It also includes governance mechanisms that make preference inference legible, contestable, and aligned with user autonomy.

Across these dimensions, the central trade-off is between optimizing for immediate engagement and supporting long-term user interests. Large language models can improve

semantic query understanding and preference explanation, but they also introduce new risks of opacity, bias amplification, and energy-intensive inference. A responsible system therefore requires auditability, data provenance, user control over inferred preferences, diversity-preserving ranking, and explicit institutional accountability. The references in this paper include foundational work on attention, recommender systems, fairness, foundation models, and algorithmic governance, together with recent proposals on memory-based query matching [10]. Future research should develop standardized preference memory formats, evaluate long-term outcomes rather than short-term clicks, and study how users understand and contest automated preference profiles. Only through such interdisciplinary work can personalized media retrieval serve both individual growth and the pluralistic public sphere.

References

1. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998-6008).
2. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 4171-4186).
3. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30-37.
4. Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734-749.
5. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., ... & Lowe, R. (2022). Training language models to follow instructions with human feedback. In *Advances in neural information processing systems*.
6. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610-623).
7. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.
8. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59-68).
9. Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web* (pp. 285-295).
10. Yu, X. (2026). Enhancing Search Efficiency through LLM-Based User Memory Systems for Query Matching and Intent Modeling.
11. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86-92.

12. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707.
13. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., ... & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220-229).
14. Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183-186.
15. Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 8342-8360).
16. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in neural information processing systems*.
17. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
18. Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973-989.
19. Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. D. O., Kaplan, J., ... & Zaremba, W. (2021). Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
20. Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. *PublicAffairs*.
21. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4), 211-407.
22. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273-1282).
23. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., ... & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33-44).
24. Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin Press.