

A Fairness-Aware Multi-Objective Recommendation Framework for Mitigating Creator Popularity Bias and Balancing User Utility and Engagement

Mr. A. Srinivasan,

Research Scholar & Assistant Professor, School of Intelligent Systems and Data Science,
KPR College of Arts Science and Research, Coimbatore,
srinivasarar91@gmail.com

Dr. G. Satyavathy,

Professor & Head, Department of Computer Science with Data Analytics,
KPR College of Arts Science and Research, Coimbatore
satyavathy.g@kprcas.ac.in

Abstract

Recommendation systems have become a central mechanism for organizing digital content, products, creators, and services. Although conventional recommender models are commonly optimized for accuracy and short-term engagement, their ranking behavior can systematically favor already popular creators and reduce the visibility of less-established creators. This review examines the emerging research problem of creator popularity bias from a multi-objective perspective, where user utility, engagement, creator fairness, exposure diversity, and platform sustainability must be considered simultaneously. The paper synthesizes major recommendation paradigms, fairness definitions, popularity-bias mechanisms, exposure-based evaluation, multi-objective optimization, re-ranking, learning-to-rank, and reinforcement-learning approaches. Based on the reviewed literature, a conceptual framework named Fairness-Aware Multi-Objective Recommendation (FAMOR) is proposed. The framework combines personalized relevance with creator exposure constraints and engagement-aware utility optimization. It uses creator-level fairness indicators, popularity-adjusted exposure, diversity measures, and user satisfaction as complementary objectives rather than treating fairness as a post-processing add-on. The review identifies important challenges involving fairness definitions, cold-start creators, feedback loops, dynamic preferences, creator heterogeneity, computational cost, and the risk of reducing perceived recommendation quality. Future research should investigate context-sensitive fairness, causal exposure modeling, long-term creator welfare, explainability, and online adaptive optimization. The proposed framework provides a structured foundation for designing recommender systems that preserve useful personalization while promoting more equitable opportunities for creators.

Keywords: recommender systems; creator fairness; popularity bias; multi-objective optimization; user utility; engagement; exposure fairness; learning-to-rank; recommendation diversity; responsible AI

1. Introduction

Recommender systems are widely used to reduce information overload by selecting and ranking items that are likely to be useful or interesting to users. In social media, creator platforms, streaming services, marketplaces, and user-generated-content ecosystems, recommendation algorithms also determine which creators receive attention. This makes recommendation a distribution mechanism: the algorithm does not merely predict preferences; it influences the allocation of visibility, traffic, and economic opportunities.

A major concern is popularity bias. Popular items and creators tend to receive more interactions, stronger historical signals, and greater algorithmic visibility. These advantages can reinforce one another through feedback loops. A creator who is initially visible can obtain more impressions, clicks, likes, and followers, which then provide stronger signals for subsequent ranking. Less-popular creators may therefore receive fewer opportunities to generate the very interactions required to improve their future ranking.

Traditional evaluation has frequently emphasized accuracy-oriented metrics such as precision, recall, normalized discounted cumulative gain, or click-through rate. These metrics remain important because recommendations that users do not find relevant have limited value. However, optimizing only for immediate engagement can overlook distributional effects across creators. A ranking that maximizes clicks may repeatedly show content from a narrow group of highly popular creators even when comparable content from emerging creators could provide similar utility. Fairness in recommendation is consequently a multi-sided problem. Users seek relevant, diverse, and trustworthy content. Creators seek reasonable and non-discriminatory opportunities for exposure. Platforms seek engagement, retention, quality, and sustainable ecosystems. These objectives can conflict. Increasing exposure for underrepresented creators may reduce a short-term engagement metric in some contexts, while aggressive fairness constraints may introduce irrelevant content. The research challenge is therefore not simply to maximize fairness, but to identify a principled balance among competing objectives.

This review focuses specifically on the intersection of creator fairness, popularity bias, user utility, and engagement. It synthesizes concepts and methods from recommender systems, fair ranking, multi-objective optimization, and responsible machine learning. It also develops a conceptual framework that can be used as a basis for empirical research.

2. Review Objectives and Methodological Scope

This article is a structured narrative review rather than a systematic meta-analysis. Its purpose is to organize the conceptual and technical landscape around creator popularity bias and fairness-aware recommendation. The review addresses five questions:

- How does popularity bias emerge and persist in modern recommendation pipelines?
- What forms of fairness are relevant when recommendation affects creator exposure?
- How can user utility and engagement be balanced against creator-level fairness objectives?
- Which algorithmic approaches are most suitable for fairness-aware multi-objective recommendation?
- What research gaps remain for practical deployment and rigorous evaluation?

The discussion covers collaborative filtering, content-based recommendation, hybrid models, learning-to-rank, re-ranking, fairness-aware recommendation, multi-objective optimization, and sequential or reinforcement-learning approaches. The review treats engagement as an outcome that may be valuable but potentially myopic when used as the sole optimization target.

3. Recommendation Systems and the Emergence of Popularity Bias

3.1 Conventional recommendation pipeline

A modern recommendation pipeline typically contains candidate generation, scoring, ranking, and post-processing stages. Candidate generation identifies a manageable set of potentially relevant items. A scoring model then estimates relevance or expected interaction probability. A ranking component orders candidates, and post-processing may impose business, safety, diversity, or fairness constraints. Popularity can enter at every stage through historical interaction counts, popularity-derived features, retrieval priors, or implicit feedback.

Collaborative filtering is particularly sensitive to exposure patterns because interactions are not observed uniformly. If users mostly encounter popular creators, the system collects more interaction data for those creators. Matrix factorization, neural collaborative filtering, and related methods can consequently learn strong representations for popular entities while producing weaker estimates for creators with sparse data.

3.2 The popularity feedback loop

Popularity bias can be represented as a reinforcing cycle: exposure generates interactions; interactions become training signals; stronger signals increase predicted relevance; predicted relevance increases future exposure. This creates a distinction between observed preference and algorithmically mediated preference. A high interaction rate does not necessarily mean that users would prefer an item if it had equal exposure to alternatives.

The problem is particularly significant for creator platforms because exposure is often concentrated. A small group of creators can accumulate a large fraction of impressions, while a long tail receives limited visibility. Long-tail content may be high quality but difficult to discover because quality and popularity are not equivalent.

Stage	Mechanism	Potential effect
Initial exposure	Ranking favors established creators	Unequal visibility
Interaction collection	Popular creators receive more clicks/views	Denser historical data
Model training	Interaction volume becomes a relevance signal	Popularity correlated with predicted relevance
Ranking	High-scoring creators appear repeatedly	Exposure concentration
Feedback	More exposure generates more interactions	Popularity reinforcement

4. Creator Fairness in Recommendation

Fairness in recommendation can be defined from multiple perspectives. User-side fairness concerns whether users receive equitable quality of service. Item-side fairness concerns whether items are treated fairly. Creator-side fairness focuses on the distribution of exposure or opportunity across content producers. Group fairness may compare predefined creator groups, while individual fairness may seek similar treatment for creators with similar relevant characteristics.

4.1 Exposure fairness

Exposure is especially useful for creator platforms because ranking position has unequal visibility. An item at the first position is generally more likely to be seen than one near the bottom. A fairness objective can therefore measure expected exposure rather than simply counting recommendations. Position-based exposure models assign higher attention weights to higher-ranked positions, allowing the system to estimate how much visibility each creator receives.

4.2 Opportunity versus outcome

A key distinction is between equality of exposure and equality of outcomes. Creators may receive similar exposure but obtain different engagement because of content quality, audience fit, or timing. Conversely, equalizing engagement may be unrealistic because creators have different audience sizes. A practical fairness objective may therefore focus on opportunity: creators should have a reasonable chance to reach users who are likely to value their content.

4.3 Fairness constraints

Fairness can be incorporated through minimum exposure constraints, exposure regularization, popularity-adjusted scores, group-based quotas, or constrained ranking. However, rigid constraints can be harmful when applied without considering relevance. A fair recommender should avoid substituting arbitrary equality for useful personalization. The goal is better described as constrained optimization under relevance and quality requirements.

5. User Utility and Engagement

User utility refers to the value a recommendation provides to the user. It can include relevance, satisfaction, novelty, diversity, discovery, trust, and task completion. Engagement, by contrast, often measures observable actions such as clicks, watch time, likes, comments, shares, saves, or session duration. Engagement is measurable and operationally important, but it is not always identical to user benefit.

Short-term engagement may reward sensational, repetitive, or highly familiar content. A user can click an item without finding it valuable, while a useful recommendation may generate little immediate interaction. This motivates the inclusion of multiple utility dimensions. A recommendation framework should distinguish immediate behavioral response from longer-term satisfaction where possible.

Objective	Example metric	Strength	Risk if optimized alone
Relevance	NDCG / Recall	Captures ranking quality	Can reinforce popularity
Engagement	CTR / watch time	Easy to measure	May encourage myopic behavior
User satisfaction	Rating / survey / return behavior	Closer to perceived value	Costly or sparse
Diversity	Intra-list diversity	Reduces repetition	May lower relevance

Creator fairness	Exposure parity / Gini	Captures distribution	May ignore quality
Novelty	Novelty score	Supports discovery	Can reduce familiarity

6. Multi-Objective Recommendation

The central premise of multi-objective recommendation is that no single metric adequately represents system quality. A generic optimization problem can be written as:

$$\text{Maximize } F = \alpha R + \beta U + \gamma E + \delta D + \epsilon F_c - \lambda C$$

where R represents relevance, U user utility, E engagement, D diversity or novelty, F_c creator-fairness performance, C represents undesirable costs or constraint violations, and the coefficients control the trade-offs. In practice, these objectives can be normalized and dynamically weighted.

A weighted-sum approach is simple, but it assumes that objectives can be combined meaningfully on a common scale. Alternative strategies include Pareto optimization, constrained optimization, lexicographic ranking, scalarization with adaptive weights, and multi-objective reinforcement learning. The choice depends on whether fairness is treated as a hard requirement, a soft objective, or a context-dependent preference.

6.1 Pareto perspective

A Pareto-efficient recommendation policy is one for which improving one objective would require sacrificing another. This perspective is useful because there may not be a single universally optimal balance between engagement and creator fairness. Instead, the platform can select a policy from a Pareto frontier according to governance requirements, product context, and acceptable trade-offs.


6.2 Dynamic weighting

Fixed weights can become inappropriate as user interests, creator populations, or platform objectives change. Dynamic weighting can respond to context, such as new-user onboarding, creator cold-start conditions, content scarcity, or periods of high popularity concentration. However, dynamic policies require careful monitoring to prevent unstable ranking behavior.

7. Proposed Fairness-Aware Multi-Objective Recommendation Framework (FAMOR)

Based on the reviewed concepts, this paper proposes a conceptual Fairness-Aware Multi-Objective Recommendation (FAMOR) framework. The framework is designed for platforms in which creators compete for user attention and where popularity bias can emerge from feedback-driven ranking.

Figure 1. Fairness-Aware Multi-Objective Recommendation (FAMOR) Framework



7.1 Framework architecture

The framework contains six logical layers: (1) data and interaction modeling, (2) candidate generation, (3) popularity-aware relevance scoring, (4) fairness-aware multi-objective ranking, (5) exposure monitoring, and (6) feedback and adaptive calibration.

- Data layer: user preferences, item characteristics, creator identifiers, historical interactions, temporal context, and exposure logs.
- Candidate layer: retrieve relevant candidates using collaborative, content-based, or hybrid retrieval.
- Scoring layer: estimate personalized relevance while explicitly modeling popularity and uncertainty.
- Fair ranking layer: optimize relevance, utility, engagement, diversity, and creator exposure jointly.
- Monitoring layer: calculate exposure concentration, creator coverage, fairness gaps, and user-quality metrics.
- Adaptive layer: update weights or constraints using offline evaluation and controlled online experimentation.

A simplified conceptual flow is: User Context → Candidate Generation → Relevance & Popularity Modeling → Multi-Objective Fair Ranking → Exposure Allocation → User/Creator Feedback → Model Calibration. The key design principle is that fairness is incorporated before the final ranking is served, while monitoring remains continuous after deployment.

7.2 Popularity-adjusted relevance

Popularity should not simply be removed from the model because it can contain legitimate information about quality or broad demand. Instead, the framework treats popularity as one feature among several and estimates whether high predicted engagement is partly explained by unequal prior exposure. A popularity-adjusted score can reduce excessive dependence on raw interaction volume while retaining useful information.

7.3 Creator exposure regularization

Let e_c denote expected exposure for creator c and q_c denote a target exposure level based on eligible audience, content quality, or historical opportunity. A regularization term can penalize large deviations:

$$L_{\text{fair}} = \sum_c w_c |e_c - q_c|$$

The target q_c should not necessarily be equal for every creator. It may be conditioned on relevant audience size, content quality, category, creator eligibility, or uncertainty. This avoids simplistic equal-exposure rules that could overwhelm relevance.

7.4 Engagement-aware utility

The framework separates immediate engagement from broader utility. For example, expected click probability may be combined with estimated satisfaction, diversity, novelty, and repeat-use indicators. The resulting objective encourages recommendations that are both attractive and useful rather than simply attention-maximizing.

7.5 Cold-start creator protection

New creators have limited historical interaction data. The framework can allocate controlled exploration opportunities to reduce uncertainty. Exploration should be quality-aware and targeted to users with a plausible interest in the creator's content. This is preferable to indiscriminate exposure because it protects both user experience and creator opportunity.

8. Algorithmic Approaches for Fair Recommendation

8.1 Fair re-ranking

Re-ranking is often the most practical intervention because it can operate after an existing recommendation model generates candidates. A re-ranker can trade a small amount of predicted engagement for improved creator coverage or exposure diversity. The approach is modular and can be evaluated without replacing the entire recommendation stack.

8.2 Fair learning-to-rank

Fairness can also be integrated into learning-to-rank objectives. Pairwise or listwise losses can be augmented with exposure penalties, group fairness constraints, or regularization. This allows the model to learn trade-offs directly, although it increases complexity and requires careful training data construction.

8.3 Constrained optimization

Constrained optimization treats selected fairness conditions as requirements. For example, the ranking process may require minimum creator coverage or limit concentration within a recommendation slate. This approach is appropriate when governance policies establish non-negotiable fairness boundaries.

8.4 Reinforcement learning

Sequential recommendation creates an opportunity to optimize long-term outcomes. A reinforcement-learning agent can consider how current exposure decisions affect future user satisfaction, creator participation, and content supply. Multi-objective reinforcement learning can represent multiple reward components, but the method requires robust reward design and safeguards against unintended optimization.

8.5 Explainable fairness

Explainability is important because creators and platform managers may need to understand why exposure changes. A useful explanation could distinguish relevance, user preference, popularity adjustment, and fairness constraints. Such explanations can support auditing and make algorithmic trade-offs more transparent.

9. Evaluation Framework

A credible evaluation should report both recommendation quality and distributional outcomes. Evaluating only engagement can conceal popularity concentration. Conversely, reporting only fairness can conceal deterioration in user experience.

Dimension	Representative measures	Purpose
Accuracy	Precision@K, Recall@K, NDCG@K	Measure relevance
Engagement	CTR, watch time, saves	Measure behavioral response
Utility	Satisfaction, retention, task success	Assess broader user value
Creator exposure	Exposure Gini, exposure variance	Measure concentration
Fairness	Group exposure gap, parity measures	Compare creator groups
Coverage	Creator/item coverage@K	Measure breadth
Diversity	Intra-list diversity	Measure recommendation variety
Novelty	Novelty@K	Measure discovery
Long-term health	Creator retention, catalog diversity	Assess ecosystem sustainability

Offline evaluation should use temporal splits when possible to reduce leakage and simulate future recommendation. Popular and long-tail creators should be analyzed separately. Results should also be segmented by user activity level, content category, creator age, and platform context. Online evaluation should use controlled experiments with guardrail metrics, because fairness improvements may interact with engagement and satisfaction in nonlinear ways.

10. Research Gaps and Challenges

10.1 Defining the correct fairness target

There is no universal creator-fairness definition. Equal exposure, proportional exposure, opportunity-based exposure, and quality-adjusted exposure can lead to different outcomes. Future research should justify fairness targets using platform-specific evidence rather than selecting a metric solely for convenience.

10.2 Causality and exposure

Observed interactions are confounded by exposure. A creator may appear unpopular because users did not see the content, rather than because users rejected it. Causal and counterfactual approaches are needed to estimate what engagement might have been under alternative exposure policies.

10.3 Fairness versus quality

A major practical question is how much short-term utility a platform should trade for improved creator opportunity. The answer may vary by context. A mature framework should identify acceptable trade-off regions instead of assuming that any fairness gain is beneficial.

10.4 Dynamic ecosystems

Creator populations and user preferences evolve. A fairness-aware system must therefore be adaptive. Static fairness constraints can become stale and may unintentionally disadvantage creators in rapidly changing categories.

10.5 Gaming and strategic behavior

Creators may adapt to algorithmic incentives. If the system rewards certain exposure-related behaviors, creators may optimize for those signals. Future research should evaluate robustness against manipulation and consider incentive-compatible fairness mechanisms.

10.6 Privacy and governance

Creator-level auditing can require sensitive platform data. Fairness research therefore needs privacy-preserving measurement, access controls, transparent governance, and careful treatment of protected or sensitive attributes where applicable.

11. Discussion

The literature suggests that popularity bias is not merely a model defect; it is an ecosystem-level phenomenon produced by the interaction of ranking, exposure, behavior, and feedback. This observation has an important implication: fairness interventions should be evaluated at the level of the recommendation process rather than only at the level of model predictions.

The proposed FAMOR framework emphasizes this systems perspective. Its principal contribution is conceptual integration. Instead of treating creator fairness as an isolated constraint, the framework places it alongside user utility and engagement within a common optimization architecture. Popularity is modeled explicitly so that the system can distinguish demand signals from exposure-induced signals. Exposure monitoring then provides feedback for calibration.

The most appropriate implementation will vary across domains. A music platform may prioritize listening satisfaction and catalog diversity, while a short-video platform may emphasize watch time and creator opportunity. An e-commerce marketplace may require seller fairness, product quality, and transaction outcomes. Consequently, the framework should be instantiated with domain-specific objectives and governance rules.

Importantly, fairness should not mean indiscriminate promotion of low-quality content. The desired principle is fair opportunity conditional on relevance and eligibility. A creator should receive exposure where there is evidence that their content can provide value to a suitable audience. This approach aligns creator welfare with user-centered recommendation rather than positioning them as opposing goals.

12. Future Research Directions

- Develop causal estimators that separate intrinsic creator quality from exposure-generated popularity.
- Study Pareto-frontier methods for selecting explicit trade-offs between engagement and creator fairness.
- Develop user-specific fairness policies that adapt to interests rather than applying uniform exposure rules.
- Investigate multi-objective reinforcement learning for long-term user satisfaction and creator ecosystem health.
- Create standardized benchmarks containing exposure logs, creator metadata, and long-tail distributions.
- Explore fairness-aware foundation and generative recommendation models with transparent exposure controls.
- Design explainable dashboards for creators and platform governance teams.
- Study incentive-compatible algorithms that remain robust when creators adapt strategically.
- Evaluate fairness interventions across different content domains and cultural contexts.
- Combine creator fairness with privacy, safety, diversity, and sustainability objectives in a unified governance framework.

13. Conclusion

Recommendation systems increasingly function as attention-allocation mechanisms, making creator exposure an important algorithmic outcome. Popularity bias can emerge through feedback loops in which historical visibility produces interactions, interactions influence ranking, and ranking produces additional visibility. Optimizing only for immediate engagement can therefore amplify concentration and reduce opportunities for emerging or long-tail creators.

This review argues for a multi-objective approach in which relevance and engagement remain essential but are balanced with user utility, diversity, novelty, and creator fairness. The proposed Fairness-Aware Multi-Objective Recommendation (FAMOR) framework integrates popularity-aware scoring, exposure-aware ranking, cold-start exploration, multi-objective optimization, and continuous monitoring. The framework does not assume that all creators should receive identical exposure; instead, it emphasizes fair opportunity subject to relevance, quality, and user context.

Future empirical studies should test the framework using realistic exposure logs and online experiments, with particular attention to causal effects and long-term ecosystem outcomes. Such research can help move recommender

systems from purely engagement-centered optimization toward responsible personalization that benefits users, creators, and platforms simultaneously.

References

- Abdollahpouri, H., Burke, R., & Mobasher, B. (2017). Controlling popularity bias in learning-to-rank recommendation. *Proceedings of the 11th ACM Conference on Recommender Systems*.
- Abdollahpouri, H., Burke, R., & Mobasher, B. (2019). Managing popularity bias in recommender systems with personalized re-ranking. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Burke, R. (2017). Multisided fairness for recommendation. *FATREC Workshop on Fairness, Accountability, and Transparency in Recommender Systems*.
- Ekstrand, M. D., Burke, R., & Diaz, F. (2018). Fairness and discrimination in recommendation and retrieval. *Proceedings of the 11th ACM International Conference on Web Search and Data Mining*.
- Fabbri, F., Bonchi, F., Boratto, L., & Castillo, C. (2021). Exposure inequality in recommender systems: A review and research agenda. *ACM Computing Surveys*.
- Ge, Y., Liu, S., Fu, Z., Tan, J., Li, Z., & others. (2021). Towards long-term fairness in recommendation. *Proceedings of the ACM International Conference on Information and Knowledge Management*.
- Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*.
- Joachims, T., Swaminathan, A., & Schnabel, T. (2017). Unbiased learning-to-rank with biased feedback. *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*.
- Jannach, D., Zanker, M., Felfernig, A., & Friedrich, G. (2010). *Recommender Systems: An Introduction*. Cambridge University Press.
- Kunaver, M., & Požrl, T. (2017). Diversity in recommender systems: A survey. *Knowledge-Based Systems*, 123, 154–162.
- Liu, Y., Sun, W., & others. (2023). Fairness-aware recommendation: A survey of recent advances and open challenges. *ACM Transactions on Recommender Systems*.
- Patro, G. K., Biswas, A., Ganguly, N., & others. (2020). FairRec: Two-sided fairness for personalized recommendations in two-sided platforms. *Proceedings of The Web Conference*.
- Rastegarpanah, B., Rudra, K., & others. (2019). Fighting fire with fire: Mitigating algorithmic bias with a bias-aware recommender system. *Proceedings of the 13th ACM Conference on Recommender Systems*.
- Ricci, F., Rokach, L., & Shapira, B. (Eds.). (2022). *Recommender Systems Handbook* (3rd ed.). Springer.
- Singh, A., & Joachims, T. (2018). Fairness of exposure in rankings. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Yao, S., & Huang, B. (2017). Beyond parity: Fairness objectives for collaborative filtering. *Advances in Neural Information Processing Systems*.
- Zhu, Z., & Caverlee, J. (2022). Fairness-aware recommender systems: A survey and future directions. *ACM Transactions on Information Systems*.

Figure 1. Conceptual FAMOR Framework

User Context	Creator / Content Data	Exposure History	Popularity Signals	Platform Constraints
Candidate Generation	Hybrid Retrieval	Relevance Estimation	Popularity Adjustment	Quality Filtering
		FAMOR MULTI-OBJECTIVE RANKING		
User Utility	Engagement	Creator Exposure Fairness	Diversity / Novelty	Long-Term Ecosystem
Exposure Monitoring → Feedback → Calibration → Re-ranking → Continuous Evaluation				

Figure 1 summarizes the proposed architecture. The system jointly considers user context, creator/content characteristics, historical exposure, popularity signals, and platform constraints. The central ranking layer optimizes multiple objectives, while exposure monitoring provides feedback for subsequent calibration.