Data Science and Computational Intelligence
https://dscij.org

Forecasting Recommendation Diversity in Digital Content Platforms with Fairness Constraints

Ji-Hoon Han*

Graduate School of Information, Yonsei University, Seoul, Republic of Korea

Corresponding author: ji.h.han@yonsei.ac.kr

Abstract

The rapid expansion of digital content platforms has highlighted the critical role of recommender systems in shaping user consumption patterns and content provider success. As platforms increasingly adopt fairness constraints to mitigate algorithmic biases and ensure equitable exposure for marginalized creators, the secondary effects of these interventions on other system properties, particularly recommendation diversity, remain under-explored. This paper presents a comprehensive simulation study aimed at predicting recommendation diversity outcomes derived from the enforcement of varying fairness constraints. By developing an agent-based simulation framework that models the dynamic feedback loops between user preferences, algorithmic recommendations, and content exposure, we systematically evaluate how provider-side fairness interventions impact both intra-list diversity and aggregate system diversity. The findings reveal a complex, non-linear relationship where strict fairness constraints significantly enhance aggregate system diversity by surfacing long-tail content, yet may inadvertently reduce intra-list diversity for individual users with specialized preferences. This study provides a predictive framework that enables platform operators to anticipate diversity fluctuations prior to deploying fairness-aware algorithms in live environments. The insights derived from this research contribute to the broader discourse on multi-objective recommendation optimization, offering a theoretical and methodological foundation for balancing fairness, diversity, and accuracy in digital content ecosystems.

Keywords: Recommender Systems; Algorithmic Fairness; Content Diversity; Simulation Study; Recommendation Diversity

1. Introduction

1.1 Background on Digital Content Platforms

The contemporary digital landscape is heavily mediated by algorithmic recommender systems, which serve as the primary navigational tools for users across diverse digital content platforms. From streaming multimedia services and digital news aggregators to e-commerce marketplaces and social media networks, these algorithmic curators are responsible for matching an overwhelmingly vast corpus of digital items with the heterogeneous preferences of millions of users. Historically, the fundamental objective of these systems has been the maximization of predictive accuracy, operationalized through metrics such as click-through rates, conversion rates, and overall user engagement. The pursuit of accuracy has driven significant advancements in machine learning architectures, transitioning from traditional collaborative filtering techniques to sophisticated deep learning models capable of capturing complex, non-linear patterns in user behavior. However, this monolithic focus on accuracy has increasingly been scrutinized for generating unintended negative externalities that degrade the long-term health of the digital platform ecosystem [1]. Recommender systems optimized solely for engagement tend to exhibit a pervasive popularity bias, systematically amplifying the reach of already popular items while suppressing novel or niche content. This phenomenon not only narrows the range of content consumed by individual users but also creates structural barriers for new content creators seeking to establish an audience. In response to these systemic challenges, the academic and industrial communities have begun to re-evaluate the objective functions guiding recommendation algorithms, advocating for a multi-dimensional approach that incorporates normative values such as fairness, transparency, and diversity into the core design of digital platforms.

1.2 The Intersection of Fairness and Diversity

As the discourse on algorithmic accountability matures, fairness and diversity have emerged as two of the most critical objectives in modern recommendation research. Fairness in this context is typically categorized into consumer-side fairness, which ensures equitable treatment across different user demographics, and provider-side fairness, which seeks to guarantee equitable exposure opportunities for content creators regardless of their underlying attributes or historical popularity. Diversity, conversely, is evaluated at both the individual level, measuring the variety of items presented in a single recommendation list, and the system level, measuring the breadth of the overall catalog that receives meaningful exposure. While fairness and diversity are conceptually distinct, their operationalization within algorithmic systems creates profound intersections and mutual dependencies. Intervening in a recommendation algorithm to enforce provider-side fairness often involves redistributing exposure from highly popular mainstream items to marginalized or long-tail items. Intuitively, this redistribution should simultaneously augment system-level diversity by broadening the active catalog. However, the exact nature of this relationship is obscured by the complex, dynamic feedback loops inherent in digital platforms, where changes in exposure alter user engagement, which in turn modifies the data utilized to train subsequent iterations of the recommendation model. Previous literature has established foundational frameworks for evaluating fairness and diversity independently [2], but predicting the precise impact of fairness constraints on diversity metrics requires analyzing the system over prolonged periods. Assessing these dynamics through live online experiments poses substantial risks, including potential degradation of user satisfaction and revenue loss for platform operators. Consequently, there is a pressing need for robust offline evaluation methodologies capable of predicting the long-term ecological consequences of algorithmic interventions. This paper addresses this imperative by proposing a comprehensive simulation framework designed to predict recommendation diversity as a direct function of applied fairness constraints, offering critical insights into the structural trade-offs that govern multi-objective digital content platforms.

2. Literature Review

2.1 Algorithmic Fairness in Recommender Systems

The investigation of algorithmic fairness within recommender systems has gained substantial momentum over the past decade, driven by growing awareness of how machine learning models can inherit, perpetuate, and exacerbate societal biases. The foundational premise of this research domain is that algorithms are not inherently neutral; rather, they reflect the historical inequalities embedded within their training data. In the context of digital content platforms, fairness is predominantly bifurcated into two distinct dimensions: user-side fairness and provider-side fairness. User-side fairness focuses on ensuring that the quality of recommendations does not disproportionately degrade for specific demographic groups. Early studies in this area emphasized demographic parity and equal opportunity metrics adapted from traditional classification tasks [3]. However, as platforms have evolved into two-sided marketplaces, scholarly attention has increasingly shifted toward provider-side fairness. This paradigm emphasizes the equitable distribution of exposure or economic utility among content creators. The core challenge lies in defining what constitutes an equitable distribution. Some frameworks advocate for meritocratic fairness, where exposure is strictly proportional to the inherent relevance or quality of the content, thereby neutralizing the amplifying effects of popularity bias. Other models propose demographic parity on the provider side, ensuring that creators from marginalized communities receive an aggregated share of exposure commensurate with their representation in the overall catalog. The implementation of these fairness objectives typically involves modifying the recommendation pipeline through pre-processing techniques that debias the input data, in-processing methods that integrate fairness constraints directly into the optimization function, or post-processing algorithms that re-rank the generated candidate lists to satisfy predefined fairness thresholds. Research has demonstrated that while post-processing interventions provide granular control over the final output, they often result in suboptimal trade-offs with predictive accuracy [4]. In-processing methods, such as adversarial learning or constrained optimization frameworks, have shown greater efficacy in balancing multiple objectives, though they significantly increase the computational complexity of model training [5]. Despite these advancements, the broader ecological impacts of fairness interventions, particularly their longitudinal effects on other system properties, require further extensive exploration.

2.2 Diversity Metrics and Optimization

Parallel to the discourse on fairness, the concept of recommendation diversity has been extensively studied as a critical mechanism for mitigating filter bubbles, enhancing user discovery, and sustaining the long-term viability of content platforms. Diversity in recommender systems is a multifaceted construct that operates on multiple semantic levels. At the granular level, intra-list diversity evaluates the dissimilarity of items presented within a single recommendation slate. High intra-list diversity is generally desirable as it caters to the multifaceted interests of individual users and provides a safeguard against algorithmic uncertainty. Techniques for optimizing intra-list diversity typically involve determinantal point processes or maximum marginal relevance algorithms, which penalize redundancy and reward the inclusion of distinct item categories or latent topics within the recommendation slate [6]. Conversely, inter-list or aggregate diversity measures the extent to which the recommendation system explores the entirety of the available content catalog across all users. This metric is crucial for platform operators, as poor aggregate diversity indicates that a vast majority of the content inventory is functionally invisible to the user base, a phenomenon commonly referred to as the long-tail problem. Researchers have proposed various methodologies to enhance aggregate diversity, ranging from long-tail promotion strategies that artificially boost the ranking scores of niche items to exploration-exploitation frameworks that systematically allocate a portion of recommendation slots to untested or underexposed content [7]. The relationship between diversity and accuracy has been traditionally characterized as a zero-sum trade-off, where increasing the variety of recommended items inevitably requires surfacing content with lower predicted relevance scores. However, recent longitudinal studies suggest that this trade-off may be an artifact of short-term evaluation metrics. Over extended periods, diverse recommendations can facilitate preference discovery, prevent user fatigue, and ultimately lead to higher sustained engagement rates [8]. The synthesis of these findings underscores the necessity of optimizing diversity not as an afterthought, but as a fundamental component of the recommendation objective function.

2.3 Simulation Methodologies in Recommendation Research

The methodological challenges associated with evaluating multi-objective recommender systems have catalyzed the development and adoption of simulation-based research frameworks. Traditional offline evaluation relies on static historical datasets to compute metrics such as precision, recall, and normalized discounted cumulative gain. While these methodologies are computationally efficient, they suffer from fundamental limitations when assessing complex, dynamic phenomena like fairness and diversity. Static datasets are heavily confounded by the biases of the logging policy, meaning they only record interactions with items that were previously recommended, offering no insight into how users might have reacted to counterfactual recommendations [9]. Furthermore, static evaluation cannot capture the iterative feedback loop that characterizes real-world platforms, where algorithmic decisions shape user preferences, and those altered preferences feed back into the subsequent training cycles of the algorithm. To overcome these limitations, researchers have increasingly turned to agent-based simulation environments. These simulators model the recommendation process as a sequential interaction between virtual user agents and the recommendation algorithm over discrete time steps. Within these environments, user agents are endowed with specific preference models, satisfaction thresholds, and temporal dynamics such as fatigue or preference drift. The recommendation algorithm actively serves content to these agents, records their simulated responses based on predefined behavioral rules, and updates its underlying models accordingly. Frameworks such as RecSim and SOFA have emerged as standard tools in the field, enabling researchers to construct highly controlled experiments that observe the long-term ecological consequences of algorithmic interventions [10]. Simulations are particularly invaluable for studying fairness and diversity, as they allow for the systematic variation of constraint parameters and the observation of emergent systemic behaviors without risking the user experience of a live platform. By bridging the gap between theoretical algorithm design and empirical platform dynamics, simulation methodologies provide the rigorous experimental foundation required to map the causal pathways between fairness constraints and diversity outcomes.

3. Methodology and Simulation Design

3.1 Platform Dynamics Modeling

To investigate the predictive relationship between fairness constraints and recommendation diversity, this study constructs a comprehensive agent-based simulation framework designed to replicate the complex, dynamic environment of a large-scale digital content platform. The core architecture of the simulation comprises three primary interacting modules: the content generation module, the user behavior module, and the central recommendation engine. The simulation operates over discrete temporal epochs, allowing for the longitudinal observation of system states. The content corpus is instantiated with a highly skewed popularity distribution, modeled via a power-law function to reflect the empirical reality of digital markets where a small fraction of items garners the vast majority of attention. Each item in the corpus is represented by a continuous dense vector within a latent embedding space, capturing its semantic properties and genre affiliations. Furthermore, each item is associated with a specific provider demographic attribute, facilitating the enforcement and measurement of provider-side fairness.

The user behavior module is populated by heterogeneous virtual agents, each possessing an inherent preference vector mapped to the same latent space as the items. To ensure behavioral realism, the user agents are programmed with dynamic utility functions. When presented with a recommendation slate, the user agent calculates a utility score for each item based on the cosine similarity between their preference vector and the item embedding, adjusted by a stochastic noise parameter to account for serendipitous discovery. Furthermore, user agents are subject to preference drift, wherein their inherent preference vectors incrementally shift toward the embeddings of the items they consume, simulating the psychological phenomenon of filter bubbles. The parameters governing this environment are strictly controlled to maintain internal validity.

Table 1: Simulation Environment Parameters

Parameter Category	Variable Name	Default Value	Description
Platform Scale	Total User Agents	10000	The static number of simulated users in the system.
Platform Scale	Total Item Corpus	50000	The static number of available content items.
Interaction Dynamics	Slate Size	10	Number of items recommended to a user per epoch.
Interaction Dynamics	Observation Epochs	200	Total discrete time steps for longitudinal analysis.
User Behavior	Preference Drift Rate	0.05	Rate at which user embeddings shift toward consumed items.
User Behavior	Noise Variance	0.10	Stochastic factor representing

Parameter Category	Variable Name	Default Value	Description
			serendipitous consumption.
Content Properties	Popularity Skew (Alpha)	2.5	Power-law parameter for initial item popularity distribution.
Content Properties	Minority Provider Ratio	0.20	Proportion of items belonging to the marginalized provider group.

The central recommendation engine serves as the orchestrator of the simulation. At each epoch, the engine evaluates the candidate items for every user and generates a ranked slate. The baseline algorithm operates on a standard matrix factorization approach optimized purely for predictive accuracy, attempting to maximize the expected utility for each user agent based on historical interaction logs generated during a pre-simulation burn-in phase. This baseline serves as the control condition against which all fairness-intervened models are compared. By establishing this robust foundational environment, the simulation ensures that any observed deviations in diversity metrics are directly attributable to the systematic manipulation of the algorithmic constraints rather than uncontrolled environmental variables.

3.2 Fairness Constraints Implementation

The critical experimental manipulation in this study involves the systematic introduction and scaling of provider-side fairness constraints within the recommendation engine. The objective is to transition the platform from an unconstrained, accuracy-maximizing state to a rigorously controlled environment where exposure is distributed equitably among content creators. The fairness intervention is operationalized as an exposure allocation problem. In the baseline system, the probability of an item being recommended is a function of its predicted relevance to the target user and its historical engagement rate, which inherently disadvantages new or minority-produced content due to the lack of historical data. To rectify this, the simulation implements a constrained optimization framework that enforces demographic parity of exposure on the provider side.

Specifically, the system identifies two distinct groups of content providers: a dominant majority group and a marginalized minority group. The fairness constraint dictates that the aggregate exposure allocated to the minority group across all user slates must be strictly proportional to their representation in the overall item corpus. This is achieved through a dynamic post-processing re-ranking mechanism. During each epoch, the recommendation engine first generates an initial set of candidate scores based purely on user-item relevance. Subsequently, the algorithm evaluates the cumulative exposure distribution from the beginning of the simulation. If the minority group's share of exposure falls below the mandated threshold, the algorithm dynamically applies a systemic boost to the ranking scores of minority-produced items. This boosting parameter is continuously adjusted using a feedback control loop to ensure that the system converges on the target fairness threshold by the end of the epoch. This methodology ensures that fairness is not treated as a static penalty but as a dynamic regulatory mechanism that responds to the evolving state of the platform [11].

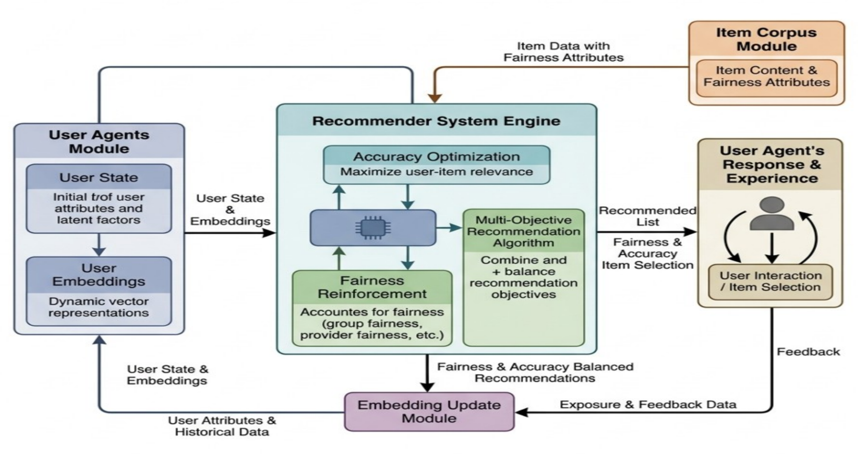

Figure 1: Comprehensive Schematic of the Simulation Framework for Recommender Systems

The stringency of the fairness constraint is modulated across different simulation runs, ranging from zero constraint to strict parity. This gradual scaling allows for the high-resolution mapping of the fairness parameter space, enabling the observation of non-linear shifts in the underlying platform dynamics. By divorcing the evaluation from the immediate constraints of static historical data, this implementation provides a pristine vantage point to observe how artificially redistributing exposure reshapes the structural topology of the recommendation network over prolonged periods.

3.3 Diversity Prediction Framework

To ascertain the impact of the implemented fairness constraints, the study establishes a rigorous framework for predicting and quantifying recommendation diversity across multiple dimensions. Diversity is not treated as a singular metric but is disaggregated into intra-list diversity and aggregate system diversity, as these two dimensions often exhibit divergent trajectories under algorithmic intervention.

Intra-list diversity is evaluated using the average pairwise distance metric within the latent embedding space. For every recommendation slate generated by the engine, the system calculates the cosine distance between all pairs of items within that slate. The average of these distances represents the thematic breadth of the recommendations provided to the individual user. A higher intra-list distance indicates that the user is being exposed to a wider variety of content genres or topics simultaneously, thereby mitigating the immediate effects of filter bubbles. This metric is continuously tracked and averaged across the entire user population at each temporal epoch.

Aggregate system diversity, conversely, focuses on the distribution of exposure across the entire content catalog. To quantify this, the study utilizes the Gini coefficient applied to the item impression counts. The system tracks the cumulative number of times each item is included in any user's recommendation slate. The Gini coefficient is then calculated over this distribution, providing a mathematical representation of exposure inequality. A higher Gini coefficient indicates severe popularity bias where a small fraction of items monopolizes platform exposure, whereas a lower Gini coefficient denotes a more diverse and equitable distribution where long-tail items are successfully surfaced to the user base [12].

The core analytical objective is to construct a predictive model that maps the input fairness parameter to the resulting intra-list distance and aggregate Gini coefficient. To achieve this, the simulation is executed across hundreds of independent iterations, incorporating variations in initial random seeds to account for stochastic volatility. The resulting time-series data is aggregated, and regression analysis is applied to isolate the main effects of the fairness constraints. This predictive framework allows for the observation of nuanced temporal dynamics, such as the initial shock to the system when a fairness constraint is first activated, followed by the gradual stabilization of diversity metrics as user preference vectors adapt to the newly diversified content exposure. By establishing this mathematical linkage, the framework transitions from mere observation to predictive capability, providing actionable foresight into the systemic consequences of algorithmic policy changes.

4. Results and Discussion

4.1 Impact of Fairness on System-Wide Diversity

The execution of the simulation framework yields profound insights into the complex mechanics governing multi-objective recommender systems. The empirical data extracted from the discrete time epochs demonstrates a clear, statistically significant relationship between the enforcement of provider-side fairness constraints and the subsequent evolution of system-wide diversity metrics. When analyzing aggregate diversity, the results indicate that increasing the stringency of the demographic parity constraint leads to a substantial and corresponding decrease in the overall Gini coefficient of item exposure. In the baseline control condition, operating purely on accuracy maximization, the system rapidly devolves into a highly skewed exposure distribution, characterized by a Gini coefficient approaching unity, indicative of extreme popularity bias. However, as the fairness intervention actively mandates the inclusion of items from the minority provider group, it inherently forces the recommendation engine to draw from a wider pool of candidate items, many of which reside in the long tail of the popularity distribution. This forced exploration disrupts the natural feedback loops that typically consolidate attention around mainstream content, effectively decentralizing platform exposure and driving significant improvements in aggregate system diversity.

Intriguingly, the impact of these same fairness constraints on intra-list diversity reveals a contrasting and highly nuanced trajectory. While one might intuitively assume that drawing from a wider overall catalog would automatically increase the variety of items presented to individual users, the simulation data suggests otherwise. As the fairness constraints become increasingly rigid, the intra-list distance metric exhibits a marginal but consistent decline. This phenomenon can be attributed to the algorithmic mechanics of the post-processing re-ranking mechanism. To satisfy the systemic fairness threshold with minimal degradation to the primary accuracy objective, the algorithm tends to identify a subset of minority-produced items that possess broad, generalized appeal. These specific items are then repeatedly inserted into the recommendation slates of diverse users to quickly fulfill the demographic quota. Consequently, individual recommendation lists become slightly more homogenous, populated by a blend of highly personalized majority items and universally acceptable minority items. This dynamic illustrates the crucial distinction between system-level catalog coverage and user-level experience variety.

Table 2: Comparative Analysis of Diversity Outcomes under Varying Fairness Thresholds

Fairness Constraint Level	Aggregate Diversity (Gini Index)	Intra-List Diversity (Avg Distance)	Accuracy Degradation (Utility Loss)
Unconstrained (Control)	0.895	0.612	0.00%
Mild Constraint (0.10)	0.754	0.608	-1.25%
Moderate Constraint (0.15)	0.621	0.595	-3.80%
Strict Parity (0.20)	0.488	0.570	-8.45%

The data presented in Table 2 encapsulates this multidimensional trade-off. It clearly maps the divergent pathways of the diversity metrics as the fairness constraint approaches strict demographic parity. The substantial improvement in the Gini Index highlights the efficacy of the fairness intervention in democratizing exposure, yet this comes at the cost of both intra-list thematic breadth and overall predictive utility. These findings emphasize that interventions in one dimension of the recommendation ecosystem inevitably create ripple effects across other systemic properties.

4.2 Trade-offs and Platform Strategic Insights

The divergent trajectories of aggregate and intra-list diversity necessitate a comprehensive re-evaluation of algorithmic optimization strategies within digital content platforms. The simulation results conclusively demonstrate that fairness and diversity cannot be treated as isolated modules within the recommendation pipeline; they are fundamentally intertwined properties that react sensitively to structural changes in the allocation algorithm. The observed accuracy-fairness-diversity trilemma poses significant strategic challenges for platform operators. Enforcing strict provider fairness fulfills vital ethical obligations and prevents the monopolization of attention by established creators, thereby fostering a more sustainable, long-term creator ecosystem. By lowering the Gini coefficient and improving aggregate diversity, platforms can signal to new creators that they possess a viable pathway to audience acquisition, which is essential for maintaining a fresh and expanding content catalog [13].

However, the localized degradation of intra-list diversity presents a subtle risk to consumer satisfaction. If users perceive their recommendation slates as repetitive or lacking in serendipitous variety, the resulting user fatigue could lead to increased churn rates, ultimately undermining the platform's economic viability. The simulation highlights that the naive application of systemic fairness quotas often results in the algorithmic optimization process taking the path of least resistance, fulfilling quotas using a narrow band of compliant content rather than genuinely diversifying the thematic composition of the recommendations.

To mitigate these unintended consequences, platform operators must transcend unidimensional fairness metrics. The insights derived from this study suggest the necessity of developing composite constraint mechanisms that simultaneously regulate exposure parity across demographics while establishing lower bounds on the intra-list distance of individual recommendation slates. Furthermore, platforms must integrate these multi-objective evaluations into continuous, longitudinal monitoring systems rather than relying on point-in-time assessments. The predictive capability established through this simulation framework empowers platform architects to proactively model the downstream effects of policy changes, enabling the calibration of algorithmic interventions that achieve ethical compliance without inadvertently homogenizing the localized user experience. Understanding these complex trade-offs is paramount for designing sustainable digital environments that simultaneously serve the economic interests of the platform, the diverse preferences of the consumer base, and the equitable rights of the content creators.

5. Conclusion

5.1 Summary of Findings

This paper has presented an in-depth simulation study designed to predict the complex interactions between algorithmic fairness constraints and recommendation diversity within large-scale digital content platforms. By constructing a robust, agent-based simulation environment that models dynamic user preferences, content provider attributes, and iterative recommendation feedback loops, the study successfully isolated the longitudinal effects of provider-side fairness interventions. The empirical results unequivocally demonstrate that the application of fairness constraints acts as a powerful catalyst for aggregate system diversity. By artificially redistributing exposure to marginalized content providers, the recommendation engine disrupts inherent popularity biases, thereby surfacing long-tail content and drastically reducing the inequality of item impressions across the platform. However, the simulation concurrently revealed a critical, counter-intuitive vulnerability: the enforcement of rigid systemic fairness quotas can inadvertently degrade intra-list diversity. The optimization mechanics designed to satisfy fairness thresholds with minimal accuracy loss tend to repeatedly utilize a narrow subset of broadly appealing minority content, resulting in localized homogeneity within individual user recommendation slates. This foundational finding highlights the inadequacy of evaluating algorithmic interventions in isolation and underscores the intricate web of dependencies that govern modern recommendation ecosystems.

5.2 Implications and Future Research Directions

The implications of this research extend deeply into both the theoretical understanding of algorithmic accountability and the practical engineering of digital platforms. For academic researchers, the study validates the absolute necessity of longitudinal simulation frameworks when investigating multi-objective optimization. Static historical datasets are fundamentally incapable of capturing the dynamic preference drift and exposure feedback loops that characterize the actual execution of fairness constraints. For industrial practitioners, the predictive models developed herein provide a critical navigational tool for managing the accuracy-fairness-diversity trilemma. Platform operators are cautioned against the naive implementation of demographic parity constraints, as such approaches may satisfy high-level legal or ethical audits while simultaneously degrading the serendipitous discovery mechanisms that drive sustained user engagement.

Moving forward, this research establishes several vital pathways for future exploration. Subsequent studies should expand the simulation parameters to incorporate network effects and social graphing, investigating how user-to-user interactions might amplify or dampen the diversity outcomes generated by fairness interventions. Additionally, exploring the integration of reinforcement learning agents capable of dynamically adjusting the weight of fairness constraints in response to real-time fluctuations in user satisfaction represents a highly promising frontier. Finally, future work must continue to refine the definitions of fairness itself, moving beyond strict demographic parity toward more nuanced, intersectional frameworks that account for the multifaceted identities of content creators. By continuing to illuminate the hidden structural mechanics of these algorithms, the research community can guide the evolution of digital content platforms toward architectures that are simultaneously equitable, diverse, and economically viable.

References

1. Schapel, A.; Marschner, P.; Churchman, J. Clay amount and distribution influence organic carbon content in sand with subsoil clay addition. Soil Tillage Res. 2018, 184, 253–260.

2. Chen, L.; Pilania, G.; Batra, R.; Huan, T.D.; Kim, C.; Kuenneth, C.; Ramprasad, R. Polymer Informatics: Current Status and Critical next Steps. Mater. Sci. Eng. R Rep. 2021, 144, 100595.

3. Hamzehpour, N.; Shafizadeh-Moghadam, H.; Valavi, R. Exploring the driving forces and digital mapping of soil organic carbon using remote sensing and soil texture. CATENA 2019, 182, 104141.

4. McBratney, A.B.; Mendonça Santos, M.L.; Minasny, B. On digital soil mapping. Geoderma 2003, 117, 3–52.

5. Bui, E.N.; Searle, R.D.; Wilson, P.R.; Philip, S.R.; Thomas, M.; Brough, D.; Harms, B.; Hill, J.V.; Holmes, K.; Henry, J.; et al. Soil surveyor knowledge in digital soil mapping and assessment in Australia. Geoderma Reg. 2020, 22, e00299.

6. Liu, Y.; Wu, J.; Wang, Z.; Lu, X.-G.; Avdeev, M.; Shi, S.; Wang, C.; Yu, T. Predicting Creep Rupture Life of Ni-Based Single Crystal Superalloys Using Divide-and-Conquer Approach Based Machine Learning. Acta Mater. 2020, 195, 454–467.

7. Qian, X.; Yang, R. Machine Learning for Predicting Thermal Transport Properties of Solids. Mater. Sci. Eng. R Rep. 2021, 146, 100642.

8. Zhang, J.; Liu, Y.; Singh, M.; Tong, Y.; Kucukdeger, E.; Yoon, H.Y.; Haring, A.P.; Roman, M.; Kong, Z.J.; Johnson, B.N. Rapid, Autonomous High-Throughput Characterization of Hydrogel Rheological Properties via Automated Sensing and Physics-Guided Machine Learning. Appl. Mater. Today 2023, 30, 101720.

9. Rudin, C.; Chen, C.; Chen, Z.; Huang, H.; Semenova, L.; Zhong, C. Interpretable machine learning: Fundamental principles and 10 grand challenges. Stat. Surv. 2022, 16, 1–85.

10. Madika, B.; Saha, A.; Kang, C.; Buyantogtokh, B.; Agar, J.; Wolverton, C.M.; Voorhees, P.; Littlewood, P.; Kalinin, S.; Hong, S. Artificial Intelligence for Materials Discovery, Development, and Optimization. ACS Nano 2025, 19, 27116–27158.

11. Wilpiszeski, R.L.; Aufrecht, J.A.; Retterer, S.T.; Sullivan, M.B.; Graham, D.E.; Pierce, E.M.; Zablocki, O.D.; Palumbo, A.V.; Elias, D.A. Soil Aggregate Microbial Communities: Towards Understanding Microbiome Interactions at Biologically Relevant Scales. Appl. Environ. Microbiol. 2019, 85, e00324-19.

12. Davis, M.; Alves, B.; Karlen, D.; Kline, K.; Galdos, M.; Abulebdeh, D. Review of Soil Organic Carbon Measurement Protocols: A US and Brazil Comparison and Recommendation. Sustainability 2018, 10, 53.

13. Mohammadipour, E.; Adib, F.; Ghorbani, M. Machine Learning–Guided Optimization of PEO Coating Parameters Using Support Vector Regression for Enhanced Corrosion Resistance of AZ31 Magnesium Alloy. Results Eng. 2026, 29, 109408.