Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

Content Exploration in Recommendation Systems: Balancing Discovery and Efficiency

Abhishek Kumar
Meta, USA

ARTICLE INFO

Received: 05 Nov 2025

Revised: 10 Dec 2025

Accepted: 19 Dec 2025

ABSTRACT

The content exploration problem is one of the most unresolved issues in recommendation systems. At the time the new items, creators, or types of content are being introduced into a platform, the system does not have enough data on interaction to approximate the actual utility or relevance of the item or content. The classic collaborative filtering and engagement alone models are biased towards popular items and do not allow new content to be exposed, as well as the overall ecosystem is less diverse. This article examines algorithmic and architectural approaches to content exploration—including uncertainty-aware value estimation, contextual bandits, reinforcement learning, and hybrid representation models—that enable efficient discovery while maintaining recommendation quality. The article addresses trade-offs between exploration cost, user satisfaction, and creator fairness, proposing strategies to align exploration with long-term platform and user objectives. Representation learning leveraging multimodal content features enables intelligent inference of item relevance despite sparse interaction data. Hybrid architectures integrating exploration throughout candidate generation and ranking pipelines enable sophisticated coordination of multiple modeling components. Multi-objective optimization formally tackles conflicting objectives in the area of engagement, diversity, fairness, and long-term value. Evaluation systems should be able to represent the opinions of different stakeholders via metrics related to user satisfaction, the diversity of content, fairness by creators, and the health of the platform.

Keywords: Content Exploration, Cold Start Problem, Recommendation Systems, Multi-Objective Optimization, Creator Fairness

1. Introduction

Recommender systems are digital ecosystem discovery engines, presenting items in huge and ever-growing catalogs. The exponential growth of user-generated content across platforms such as YouTube, TikTok, Spotify, and various e-commerce marketplaces has intensified the challenge of effective content discovery. YouTube's recommendation system, which generates over 200 million hours of watch time daily and drives approximately 70% of total viewing time on the platform, exemplifies the scale and impact of modern recommendation architectures [1]. The platform processes recommendations for over one billion users, with the candidate generation network reducing millions of videos to hundreds of candidates and the ranking network providing precise predictions for the final dozen recommendations presented to users. The cold start problem, however, that is, the case where new material has little to no historical information about its interaction, is a significant obstacle to good personalization. Such discovery engines that encourage invisible information are thus critical to sustaining diversity, equity, and novelty in suggestions.

Their inherent difficulty is the need to find an equilibrium between exploration and exploitation: adding new items without compromising the user experience and general metrics of engagement.

Most large-scale recommenders focus on short-term interaction metrics optimized under the basis on logged user interactions, resulting in popularity bias and a lack of representation of emerging content. This systematic bias results in a vicious circle, in which popular items will receive more impressions, whereas new material has a difficult time gaining a presence, no matter how good or relevant it may be. The deep neural network architecture employed by platforms like YouTube uses implicit feedback mechanisms, where watch time serves as the primary training signal, inherently favoring content with established engagement histories over novel items [1]. Research on conversational recommendation systems reveals that language models exhibit systematic biases in recommendation generation, with popularity bias being the most prevalent issue affecting content distribution fairness [2]. These biases manifest through disproportionate representation of popular items in generated recommendations, creating barriers for emerging content regardless of quality or user relevance.

In large-scale platforms, over 30–50% of new daily impressions are exposed to recently uploaded or low-interaction content. Even a modest improvement in exploration efficiency of 1–2% can lead to measurable increases in content diversity and creator retention. Effective cold start handling thus plays a critical role not only in user experience but also in sustaining healthy ecosystem growth. The economic implications extend beyond user satisfaction to encompass creator retention, platform differentiation, and long-term marketplace sustainability. Analysis of recommendation system biases demonstrates that unintended disparities in content exposure significantly impact both user satisfaction and creator engagement, with mitigation strategies requiring careful calibration to balance multiple stakeholder objectives [2].

Despite advances in contextual bandits, Bayesian uncertainty estimation, and embedding-based similarity models, these approaches often underperform when data is sparse or feedback is delayed. Moreover, exploration policies are rarely aligned with platform-level goals such as creator growth or user retention. This article addresses these limitations by exploring algorithmic and modeling strategies that integrate exploration strategies with the core ranking and value estimation pipelines of recommendation systems.

Component	Function	Scale Characteristics	Bias Manifestation
Candidate Generation Network	Reduces the catalog to the candidate set	Processes millions of videos for hundreds of candidates	Favors established content with interaction history
Ranking Network	Provides precise predictions	Final dozen recommendations	Implicit feedback mechanisms privilege watch time
Language Model Systems	Conversational recommendation	Multiple model architectures	Popularity bias in top quintile items
Mitigation Strategies	Balance stakeholder objectives	Platform-wide calibration	Requires careful multi-objective tuning

Table 1: Recommendation System Architecture and Bias Characteristics [1, 2]

2. Representation Learning for Cold Start Content

The cold start problem fundamentally stems from the absence of behavioral signals that traditional collaborative filtering relies upon. Representation learning offers a pathway to circumvent this limitation by leveraging content features—textual descriptions, visual attributes, metadata, and

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

contextual information—to infer latent item relevance before substantial user interaction data accumulates. Neural collaborative filtering approaches demonstrate substantial improvements over traditional matrix factorization methods, with multi-layer perceptron architectures achieving Hit Ratio improvements of 4.5% and Normalized Discounted Cumulative Gain gains of 4.8% on benchmark datasets [3]. These neural architectures learn nonlinear user-item interaction functions through deep neural networks, replacing the inner product operation in matrix factorization with learnable interaction layers that capture complex patterns in user preferences and item characteristics.

Contemporary methods use multimodal embedding algorithms that combine the information of various feature spaces. Transformer-based text encoders like BERT or its variants derive the semantic content of titles, descriptions, and tags, including the thematic content and linguistic cues. Convolutional neural network or vision transformer-based visual embeddings are computed on the thumbnail image or video frame, or a product photograph, and encode aesthetic and compositional information. The generalized matrix factorization component within neural collaborative filtering frameworks provides a linear pathway that preserves collaborative filtering capabilities while neural layers capture nonlinear relationships, achieving optimal performance through ensemble combinations that weight these components appropriately [3]. The experimental findings on MovieLens and Pinterest datasets indicate that this hybrid architecture outperforms pure deep learning methods by 2-3 percent without making a significant sacrifice in computational efficiency, which can be used in large-scale deployment.

The communication of these dissimilar modalities has its opportunities and challenges. The first methods of fusion combine feature vectors in batch before processing, allowing cross-modal interactions at the expense of signal strength. Late fusion methods process each modality independently before combining predictions, preserving modality-specific patterns but missing interaction effects. Attention-based fusion mechanisms offer a middle ground, dynamically weighting modality contributions based on learned relevance patterns. Research on language model-driven conversational recommendation reveals systematic biases in how different modalities influence recommendation generation, with popularity bias emerging as the most prevalent issue across recommendation tasks [4]. Analysis of large language models, including ChatGPT, GPT-3.5, and specialized recommendation models, demonstrates that these systems exhibit significant popularity bias, with items in the top popularity quintile receiving recommendation frequencies 2.5 to 3.5 times higher than their actual popularity distribution would suggest, even when prompted to emphasize diversity.

Transfer learning strategies further enhance cold start performance by leveraging knowledge from related domains or platforms. Pre-trained models capture generalizable patterns about content quality, topic relevance, and aesthetic preferences that transfer across contexts. The neural collaborative filtering framework enables effective transfer learning through pre-training embedding layers on large datasets before fine-tuning on target domains, reducing data requirements by approximately 40% while maintaining recommendation accuracy [3]. Fine-tuning these models on platform-specific data adapts representations to local user preferences while retaining broader semantic understanding.

Representation-based methods necessarily rely on the quality of features and their correspondence to relevance. High-dimensional embeddings represent subtle similarities but have a chance of overfitting to superficial characteristics as opposed to actual relevance. Empirical analysis demonstrates that embedding dimensionality significantly impacts model performance, with optimal dimensions ranging from 64 to 256 depending on dataset characteristics and sparsity levels [4]. Furthermore, temporal dynamics require continuous model updates as content trends, creator styles, and user preferences evolve.

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

Approach	Architecture Type	Key Capability	Performance Characteristic
Neural Collaborative Filtering	Hybrid generalized matrix factorization and multi-layer perceptron	Captures linear and nonlinear interactions	Outperforms pure deep learning by modest margins
Multimodal Embedding	Text encoders, vision transformers, metadata fusion	Synthesizes diverse feature spaces	Optimal dimensions vary by dataset sparsity
Transfer Learning	Pre-trained embeddings with fine-tuning	Domain knowledge transfer	Reduces data requirements substantially
Language Model Systems	Large-scale transformer architectures	Natural language understanding	Exhibits systematic popularity bias

Table 2: Neural Collaborative Filtering and Representation Learning Approaches [3, 4]

3. Exploration Algorithms and Policy Design

Exploration algorithms provide the decision-theoretic framework for balancing the introduction of novel content against the exploitation of known preferences. Classical multi-armed bandit formulations model content selection as choosing among alternatives with uncertain rewards, where exploration reduces uncertainty and exploitation maximizes immediate expected value. Contextual bandit algorithms have demonstrated substantial improvements in recommendation system performance, with LinUCB achieving click-through rate improvements of 12.5% over context-free methods and consistent performance gains across different pool sizes ranging from 10 to 250 articles in news recommendation scenarios [5]. These algorithms solve the exploration-exploitation dilemma by keeping confidence boundaries on the values of items, and the way to choose items that trade off expected reward with the reduction of uncertainty, producing real-time systems with millions of user interactions.

The ε-greedy is the simplest exploration policy: a random selection with probability ε; otherwise, an item with the highest estimated value. Despite its simplicity, ε-greedy exploration suffers from uniform randomization that wastes opportunities on clearly irrelevant items. Adaptive variants that decay ε over time or condition it on uncertainty estimates improve efficiency but retain fundamental limitations in targeting exploration effort. Experimental evaluations on Yahoo! Front Page Today Module demonstrate that ε-greedy approaches underperform contextual methods by approximately 7-10%, particularly when dealing with high-dimensional feature spaces where contextual information enables more efficient exploration through similarity-based generalization [5].

Thompson sampling offers a more principled Bayesian approach by sampling item values from posterior distributions reflecting current uncertainty. Items with high mean estimates or high variance receive proportionally more exploration, naturally balancing confidence and potential. This probabilistic framework elegantly addresses the exploration-exploitation trade-off without requiring explicit hyperparameter tuning for exploration intensity. Upper Confidence Bound algorithms, including LinUCB for contextual settings, explicitly quantify uncertainty through confidence intervals and select items maximizing the sum of estimated value plus an exploration bonus. The LinUCB algorithm constructs a confidence ellipsoid around parameter estimates using ridge regression, with the exploration bonus term scaled by a parameter α that controls exploration aggressiveness [5]. The UCB framework provides theoretical optimality guarantees with regret bounds growing

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/ Research Article

logarithmically with time, and naturally reduces exploration as confidence increases through accumulated observations.
Reinforcement learning approaches treat recommendation as a sequential decision problem where current actions influence future states and rewards. Policy gradient methods, actor-critic architectures, and deep Q-learning enable optimization of long-term value rather than immediate engagement. The Deep Reinforcement Learning framework for recommendations with negative feedback demonstrates that incorporating explicit negative signals improves recommendation quality by 3-5% compared to systems using only positive feedback [6]. This temporal perspective naturally aligns with platform objectives around user retention, lifetime value, and ecosystem health. The actor-critic architecture combines a deep Q-network for value estimation with a policy network for action selection, achieving convergence after approximately 50,000 training iterations on e-commerce datasets containing millions of user-item interactions. However, the complexity of state spaces, delayed rewards, and off-policy learning from logged data present substantial implementation challenges, with training requiring 8-12 hours on GPU clusters for production-scale systems.
The choice of exploration algorithm involves trade-offs between computational efficiency, sample complexity, and adaptability to non-stationary environments. Contextual bandits like LinUCB achieve near-optimal performance while maintaining computational complexity linear in the number of features, making them suitable for real-time serving with latency requirements below 50 milliseconds [5].

Algorithm	Exploration Mechanism	Computational Profile	Optimization Target
ε-greedy	Uniform random selection	Low complexity, easily scalable	Immediate reward maximization
LinUCB	Confidence bounds with contextual features	Linear feature complexity	Balances exploration and exploitation
Thompson Sampling	Bayesian posterior sampling	Maintains probability distributions	Naturally balances uncertainty and reward
Deep Reinforcement Learning	Actor-critic architecture	High complexity, extensive training	Long-term sequential value

Table 3: Exploration Algorithm Characteristics and Performance [5, 6]

4. Hybrid Architectures and Multi-Objective Optimization

The process of effective content exploration cannot be done as a separate module but must be incorporated into the larger recommended architecture. The hybrid frameworks are modeling paradigms that are merged to take advantage of their strengths and counter the weaknesses of each other. Neural collaborative filtering evidences the usefulness of hybrid buildings by joining the components of generalized matrix factorization and multi-layer perceptron to obtain Hit Ratio enhancements of 4.5% and 4.8% of Pinterest and MovieLens datasets, as compared to traditional matrix factorization [7]. These architectures are based on both linear and nonlinear modeling features, and the generalized matrix factorization feature of these architectures captures collaborative filtering signals, also referred to as inner products, and neural layers are acquired by training on deep learning of the complex user-item interaction functions.

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

Two-stage architectures distinguish between candidate generation and ranking steps, so that each activity can be handled specially. The candidate generation stage uses effective retrieval schemes such as collaborative filtering, content-based matching, and exploration-specific sampling to produce a manageable set of potentially relevant items out of the entire catalog. The ranking stage applies sophisticated models incorporating user context, item features, and cross-feature interactions to produce final orderings. Neural collaborative filtering architectures typically employ embedding layers with dimensions ranging from 8 to 64, with experimental results showing that 32-dimensional embeddings achieve optimal performance on datasets containing 100,000 to 1 million user-item interactions [7]. Exploration can be introduced at either or both stages: broadening candidate diversity during retrieval or adjusting scores during ranking to favor novel items. The framework uses embedding-dimension-linear computational complexity to predict recommendations, and with a zero-milliseconds-per-user-item-pair serving time, can effectively serve users in real time.

Multi-objective optimization explicitly balances conflicting objectives such as engagement, diversity, novelty, and fairness. Scalarization methods are those methods that combine objectives by means of weighted sums, allowing easy optimization, but necessitating precise tuning of trade-off parameters. Multi-objective Pareto-efficient approaches demonstrate significant advantages over single-objective methods, with the Linear Scalarization for Multiobjective Ranking (LSMOR) algorithm achieving diversity improvements of 15-20% while maintaining recommendation accuracy within 3-5% of single-objective baselines [8]. Pareto optimization identifies non-dominated solutions representing different compromise points, allowing dynamic selection based on context or policy. The Pareto-Efficient Ranking (PER) method generates approximately 8-12 distinct Pareto-optimal configurations on MovieLens datasets, each representing different trade-offs between accuracy metrics like Normalized Discounted Cumulative Gain and diversity metrics such as intra-list distance and catalog coverage. Constrained optimization formulations maximize primary objectives subject to minimum requirements on secondary goals, such as maintaining diversity above a threshold while maximizing engagement.

Value decomposition techniques separate immediate engagement signals from longer-term impact indicators. Models predicting both immediate clicks and downstream retention or satisfaction enable principled trade-offs between short-term metrics and platform health. The neural collaborative filtering framework supports multi-task learning through shared embedding layers with task-specific prediction heads, reducing training time by 30-40% compared to separate models while improving generalization performance [7]. Exploration bonuses can be calibrated against these multi-horizon value estimates to ensure novel content receives appropriate exposure without excessively sacrificing current user experience.

Position bias correction represents another critical consideration in exploration strategies. Items presented in prominent positions receive disproportionate attention regardless of intrinsic quality. Multi-objective approaches inherently address position bias by optimizing for both relevance and diversity simultaneously, with experimental results showing a 25-30% reduction in popularity bias compared to single-objective accuracy-focused methods [8].

Component	**Design Pattern**	**Integration Strategy**	**Optimization Approach**
Two-Stage Pipeline	Candidate generation and ranking separation	Specialized handling per stage	Linear complexity in embedding dimensions
Neural Collaborative Filtering	Generalized matrix factorization with neural layers	Ensemble combinations	Multi-task learning with shared embeddings

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

Pareto-Efficient Ranking	Non-dominated solution identification	Dynamic policy selection	Multiple trade-off configurations
Linear Scalarization	Weighted objective combination	Straightforward gradient optimization	Careful trade-off parameter tuning

Table 4: Hybrid Architecture Components and Multi-Objective Methods [7, 8]

5. Evaluation Metrics and Trade-off Analysis

Assessing exploration strategies requires metrics capturing diverse stakeholder interests beyond traditional accuracy measures. User satisfaction manifests through engagement metrics, including click-through rates, watch time, and session length, but also through retention, return visits, and survey-based satisfaction scores. Short-term engagement may decrease during aggressive exploration, while long-term satisfaction improves through increased diversity and discovery of niche interests. Research on novelty and diversity metrics demonstrates that recommendation quality depends critically on balancing accuracy with diversity, with frameworks incorporating both relevance ranking and novelty discovery achieving superior user satisfaction across multiple evaluation dimensions [9]. These frameworks reveal that traditional accuracy metrics like precision and recall provide incomplete assessments of recommendation quality, as they fail to capture user preferences for discovering unexpected yet relevant items.

Content diversity metrics quantify the breadth of items shown and consumed. Coverage measures the fraction of catalog items receiving impressions, while entropy quantifies the distribution of attention across items. Gini coefficients capture inequality in exposure or engagement, with lower values indicating a more equitable distribution. Experimental analyses on MovieLens datasets of 100,000 ratings show that diversity-conscious ranking algorithms gain new scores of 25-40% with the same relevance at the same accuracy-optimal baselines [9]. The Intent-Aware Diversity metric, which takes into account the similarity distances between items as well as their relevance probabilities, is more finely-tuned than the simple dissimilarity metrics because it takes into account the fact that dissimilar but unhelpful recommendations do not offer much user value. Temporal diversity is the measure of the speed at which recommendations get updated, between familiarity and novelty.

Creator fairness metrics address whether new or marginalized creators receive equitable opportunities for exposure. Exposure rates conditioned on content quality, demographic attributes, or platform tenure reveal systemic biases. Equal opportunity metrics assess whether items of comparable quality receive comparable exposure regardless of creator characteristics. Research on multi-sided fairness in recommendation systems demonstrates that neighborhood-based collaborative filtering algorithms can be enhanced with fairness constraints that balance provider-side equity with consumer-side utility [10]. The balanced neighborhoods approach achieves C-fairness values above 0.85 on a 0-1 scale while maintaining recommendation accuracy within 3-5% of unconstrained baselines, where C-fairness measures the calibration between provider representation in the catalog and their representation in recommendations. Experiments on LastFM music recommendation datasets show that fairness-enhanced algorithms reduce the standard deviation of item exposure across provider groups by 35-45%, substantially improving equity without severely compromising personalization quality.

The cold start penalty quantifies how quickly new items reach performance parity with established content. Metrics tracking engagement rates, conversion rates, or quality scores as functions of item age or impression count characterize learning curves. Good exploration plans help in making sure that convergence is made faster so that new content does not have to spend longer durations at a competitive disadvantage. Trade-off analysis tells about conflicts between conflicting goals. More aggressive exploration increases diversity and improves cold start performance, but may reduce immediate engagement. Research demonstrates that diversity-optimized recommendations can decrease immediate click-through rates by 5-8% while improving session-level engagement by 10-15%

and long-term retention by 12-18% [9]. Personalizing exploration intensity to user preferences represents one strategy for managing this heterogeneity.

Platform-level metrics, including creator retention, content upload rates, and marketplace liquidity capture ecosystem health effects. Multi-sided fairness approaches demonstrate that provider-side equity improvements of 30-40% can be achieved with consumer-side utility losses limited to 2-4%, indicating that fairness and personalization objectives are not strictly adversarial [10].

Conclusion

There are the cold start problem and the popularity bias of engagement-based optimization, which pose challenges to diversity and equity in the ecosystem. To overcome these, it is necessary to go beyond siloed exploration heuristics and integrate the ideas of representation learning, principled exploration algorithms, and multi-objective optimization. Representation learning that uses the multimodal content features allows learning item relevance with sparse interaction data, and transformer text encoders, vision models, and metadata embeddings have rich semantic representations that can scale across data. Pretrained models and transfer learning help to accelerate cold start behaviour by transferring knowledge in related domains, but the quality and relevance of learned representations are critically dependent on feature engineering, model architecture, and ongoing adaptation to changing content and user preferences. Decision-theoretic approaches to exploration based on both bandit theory and reinforcement learning give decision-theoretic approaches to balancing discovery and exploitation, and Thompson sampling, contextual bandits, and policy learning techniques all offer different trade-offs between computational efficiency, sample complexity, and optimality. Hybrid architectures that combine exploration during candidate generation and ranking pipelines allow a greater level of coordination of various modeling components, whereas multi-objective optimization explicitly deals with conflicting objectives around engagement, diversity, fairness, and long-term value. The value decomposition and position bias correction techniques enhance the quality of counterfactual reasoning and forecasts of long-term impacts. Assessment systems should reflect the views of various stakeholders by using measures that measure user satisfaction, content diversity, fairness to creators, and ecosystem health, and trade-off analysis shows a conflict between immediate interaction and the sustainable nature of the ecosystem in the long term. Future research directions include exploration strategies that can accommodate personal user preferences towards novelty, better off-policy evaluation techniques of counterfactual exploration policies, and cross-domain and cross-platform cold start mechanisms. The unification of causal inference methods can yield more plausible responsibility of the engagement patterns on each quality of content over presentation impacts, whereas exploration strategies that openly maximize creator variety and equity in the marketplace can serve as a viable approach to creating healthier online ecosystems. The ongoing evolution of content platforms toward increasingly personalized and diverse catalogs ensures that exploration will remain a central challenge in recommendation systems, requiring synthesis of insights from machine learning, decision theory, and platform economics, coupled with careful attention to the multi-sided nature of recommendation marketplaces.

References

[1] Paul Covington, et al., "Deep Neural Networks for YouTube Recommendations," ACM Digital Library, 2016. Available: https://dl.acm.org/doi/10.1145/2959100.2959190

[2] Tianshu Shen, et al., "Towards Understanding and Mitigating Unintended Biases in Language Model-driven Conversational Recommendation," ScienceDirect, 2023. Available: https://www.sciencedirect.com/science/article/abs/pii/S0306457322002400

[3] Xiangnan He et al., "Neural Collaborative Filtering," arxiv, 2017, Available: https://arxiv.org/abs/1708.05031

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

[4] Tianshu Shen, et al., "Towards Understanding and Mitigating Unintended Biases in Language Model-driven Conversational Recommendation," ScienceDirect, 2023. Available: https://www.sciencedirect.com/science/article/abs/pii/S0306457322002400
[5] Lihong Li, et al., "A Contextual-Bandit Approach to Personalized News Article Recommendation," 2012. Available: https://arxiv.org/abs/1003.0146
[6] Xiangyu Zhao, et al., "Recommendations with Negative Feedback via Pairwise Deep Reinforcement Learning," ACM Digital Library, 2018. Available: https://doi.org/10.1145/3219819.3219886
[7] Xiangnan He, et al., "Neural Collaborative Filtering," arxiv, 2017. Available: https://arxiv.org/abs/1708.05031
[8] Marco Tulio Ribeiro, et al., "Multiobjective Pareto-Efficient Approaches for Recommender Systems," ACM Digital Library, 2014. Available: https://doi.org/10.1145/2629350
[9] Saúl Vargas, Pablo Castells, "Rank and Relevance in Novelty and Diversity Metrics for Recommender Systems," ACM Digital Library, 2011. Available: https://dl.acm.org/doi/10.1145/2043932.2043955'
[10] Robin Burke, "Balanced Neighborhoods for Multi-sided Fairness in Recommendation," Proceedings of Machine Learning Research, 2018, pp. 202-214. Available: https://proceedings.mlr.press/v81/burke18a.html