

Multisided Fairness in Recommender Systems

Reza Shafiloo, Maria Stratigi, Kostas Stefanidis

ADBIS 2026

September 2026, Orléans, France

Agenda


- Context and Foundations of Fairness in Rec Sys [20mins]
- Multisided Fairness in Rec Sys
 - Stakeholders [10mins]
 - Fairness Definitions [10mins]
 - Solutions [20mins]
- Evaluation [10mins]
- Multisided Fairness: Limited Item Availability [10mins]
- Open Challenges [10mins]

Algorithmic Fairness: Why?


We live in a world where decisions are assisted or taken by algorithmic systems driven by large amounts of data

From simple, or not that simple, **personal** decisions:


Where to eat?



Who to date?



What to read, watch, buy..?



What are the news?



*What job to take? What school to attend?
Who to follow? ...?*

Algorithmic Fairness: Why?

To more complex ones:

- Education, school admission
- Law enforcement, sentencing decisions
- Job recruitment
- Insurance, Credit
- Housing
- Pricing of goods and services
- ...

How much can/should we trust such systems?

And this concern has not been without reason:

A steady stream of empirical findings has shown that data-driven methods can unintentionally both encode **existing human biases** and **introduce new ones**.

A Snapshot of the Frontiers of Fairness in Machine Learning
By Alexandra Chouldechova, Aaron Roth
Communications of the ACM, May 2020, Vol. 63 No. 5, Pages 82-89

Fairness matters! Unfair systems can cause real-world harm

- Allocative harms: When decision-making systems in criminal justice, health care, etc. are discriminatory, they create allocative harms, which are caused when a system withholds certain groups an opportunity or a resource.
- Representational harms: When systems reinforce the subordination of some groups along the lines of identity, race, class, gender, etc., they create stereotype perpetuation and cultural denigration.

Examples: Jobs ads

Facebook Inc. disproportionately shows certain types of job ads to men and women. Job ads were more likely to present job ads to users if their gender identity reflected the concentration of that gender in a particular position or industry.

(study led by University of Southern California researchers)

Example:

Ads for delivery driver job listings that had similar qualification requirements but for different companies.

- The ads did not specify a specific demographic.
- One was an ad for Domino's pizza delivery drivers, the other for Instacart drivers.
- Instacart has more female drivers but Domino's has more male drivers.

Facebook targeted the Instacart delivery job to more women and the Domino's delivery job to more men.

Examples: Music Streaming Recommendations

Diversity problem:

E.g., Spotify's prime 10 most streamed artists of 2020, for instance, solely two are ladies

- When a recommendation service gives suggestions, it does so by leveraging the historical data.
- That creates a suggestions loop, **amplifying existing bias** and **reducing diversity**.


Examples: Facial Recognition Technology



Last year, he was accused of reaching into a vehicle, grabbing a cellphone from a man and damaging it.

Officials concluded Oliver **had been misidentified** as the perpetrator and dismissed the case.

Detroit Police used **facial recognition technology** in the investigation.



- Facial recognition systems have been used by police forces **for more than two decades**.
- While the technology works relatively well on white men, **the results are less accurate for other demographics** (Recent studies by M.I.T. and the National Institute of Standards and Technology)
- In part because of a **lack of diversity in the images** used to train the algorithm.

What is Algorithmic Fairness?

Fairness: lack of discrimination (bias)

Bias may come from:

- **The actual data**
 - If a survey contains biased questions.
 - If some specific population is misrepresented in the input data.
 - If the data itself is a product of a historical process that operated to the disadvantage of certain groups.
- **The algorithm**
 - Reflecting, for example, commercial or other preferences of its designers.
 - Data processing (**technical bias**).
 - Feedback loop (**amplified bias**).

What is Algorithmic Fairness?

Individual fairness: Similar people should be treated similarly.

Examples:

- Fairness through awareness
- Fairness through unawareness
- Counterfactual fairness

Many others:

E.g., based on diversity and coverage

Group fairness: All groups should be treated similarly.


I.e., similar outcomes across demographic groups (e.g., women vs. men)

Examples:

- Demographic (or statistical) parity
- Equal opportunity
- Equalized odds
- Conditional statistical parity
- Treatment equality

Impossibility of Fairness

Fairness involves mathematical and sociotechnical trade-offs



Mathematical impossibility

- (Some) fairness metrics are mutually incompatible and cannot be satisfied simultaneously
- Trade-offs: improving one may harm another
 - Choose which fairness definition to prioritize based on context and goals

Conceptual impossibility

- Formal fairness definitions require abstraction and simplification
- But fairness is socially situated, it depends on context, history, power, and values
 - No definition is universally correct

Achieving Fairness

Methods for achieving fairness in recommenders can be distinguished as:

Pre-processing: Transform the data so that any underlying bias is removed

In-processing: Modify existing or introduce new recommendations algorithms that produce fair outputs

Post-processing: Treat the recommendation algorithms as black boxes and modify their output to ensure fairness



Making the data *fairer* will result in a *less unfair* output.

Working with the algorithm offers greater control over fairness behavior.
Balance predictive and fairness performance.

Apply fairness adjustments on the output, with no changes to the data or algorithm.
Minimal interventions to preserve predictive performance.

Agenda

- Context and Foundations of Fairness in Rec Sys [20mins]
- Multisided Fairness in Rec Sys
 - Stakeholders [10mins]
 - Fairness Definitions [10mins]
 - Solutions [20mins]
- Evaluation [10mins]
- Multisided Fairness: Limited Item Availability [10mins]
- Open Challenges [10mins]

The Problem

Most recommender systems are designed to optimize the outcome for a consumer.

Examples:

- Amazon recommends products to consumers
- Netflix recommends movies to viewers
- Spotify recommend songs to listeners.




Traditional evaluation metrics also focus on consumer-centric metrics

- Accuracy
- Precision
- Recall
- NDCG
- User satisfaction

What happens to everyone else affected by the recommendation?

Recommendations Affect More Than Consumers

Recommendations create consequences for multiple stakeholders



A recommender is not simply an information retrieval task!

It is also an allocation mechanism among different stakeholders for:


- Visibility
- Exposure
- Economic opportunities

Example: Amazon recommends a laptop to a consumer

The recommendation affects:

- **Consumer:** Receives a relevant product that matches their needs
- **Provider (Seller):** Gains visibility, potential clicks and increased sales
- **Item (Laptop):** Receives recommendation exposure and a higher chance of being discovered and purchased
- **Platform (Amazon):** Generates revenue, improves user engagement and influences market dynamics

What is a Multisided Recommender System?



Why Fairness Becomes Necessary


Do all stakeholders benefit equally? **No!**

Example:


- Consumers: lower quality recommendations
- Providers: Low visibility
- Items: No exposure
- Platform: Reduced sustainability

Recommendation systems often create unequal outcomes.


Item Popularity Unfairness



Provider Unfairness




Stakeholders Conflicts



Fairness per Stakeholder

Fairness must be defined with respect to the outcomes that matter to each stakeholder.



Consumer Fairness

Consumer Fairness focuses on ensuring that recommendation quality is distributed fairly across users and user groups.

Why it matters

- Recommendation influences opportunities and choices
- Some users may systematically receive lower quality recommendations.
- Historical data may disadvantage certain groups

Consumer Fairness seeks to reduce disparities in recommendation outcomes

Individual Fairness

Core idea:

Similar users should receive similar recommendations

Problem

Two identical users may receive very different recommendation quality due to noise, bias, or data sparsity.

Example:

Two job seekers with similar qualifications should receive similar job recommendations

Group Fairness

Core idea:

Recommendations should be fair across groups

Groups can be defined by:

- Gender
- Age
- Religion
- Other protected attribute

Example

A movie recommendation should not offer lower-quality suggestions to older users compared to younger users.

Recommendation Quality Fairness

Core idea:

Consumers should receive comparable recommendation quality regardless of their characteristics and activity

Consumers should not be systematically disadvantaged because they :

- Are new users
- Have fewer interactions
- Have fewer preferences

Example

Active users receive much better quality recommendations than users that have no previous interactions

Provider Fairness

Providers depend on recommendations for visibility and economic opportunities

Why it matters

- Recommendations directly influence provider visibility and opportunities
- Visibility may be concentrated among fewer providers, limiting exposures for others
- Fair treatment encourages providers participation and market diversity

Recommendation influence who gets attention and opportunities

Examples:

- Sellers on Amazon
- Restaurants on food delivery platforms
- Companies posting jobs

Exposure Fairness

Core idea:

Providers should receive fair visibility in recommendation lists

Problem

Popular providers often dominate recommendations

Example

Provider A has 80% exposure

Provider B has 15%

Provider C has 5%

Fairness Goal

Exposure should not be concentrated on a few providers

Opportunity Fairness

Core idea:

Providers should have equal opportunity to reach consumers

Unlike exposure fairness, the focus is on matching opportunities

- Exposure fairness → How much visibility
- Opportunity fairness → Who gets access to relevant consumers

Example

A small seller should have the same opportunities to reach consumers as large retailers.

Item Fairness

Items compete for a limited number of recommendation positions. Some items may receive large exposure while others remain unseen.

Item Fairness aims to distribute recommendation opportunities equitably

Why it matters:

Recommendation exposure is valuable and limited resource.

Popular items often dominate recommendation lists, while niche items remain less visible.

Unequal exposure can reduce diversity.

Popularity Bias Mitigation

Core idea:

Popularity bias can create unequal recommendation opportunities across items

Popularity bias occurs when few popular items receive more recommendation than others. As they gain more interactions, they are more likely to be recommended again, creating a cycle of increased attention.

This may lead to:

- Less popular items receiving little to no visibility.
- Users being presented with a limited range of options.
- Exposure be concentrated to a small subset of items

Long-tail Promotion

Core idea:

Long-tail promotion increases opportunities for less visible items

Long-tail promotion aims to increase the visibility of less popular, niche or infrequent recommended items. Rather than focusing on the highly popular items, the recommender system actively supports the discovery of a wider range of content

Why it matters:

- Help users discover niche and diverse items
- Creates opportunities for under-exposed items

Coverage Fairness

Core idea:

Coverage fairness evaluates whether recommendation opportunities are distributed across the providers' catalog

A fair system should provide exposure to a range of items in a provider's catalog rather than repeatedly recommend the same items.

Why it matters:

- Encourages broader utilization of item catalog.
- Prevents recommendation opportunities from being concentrate.
- Increases the chance that users receive a wider range of relevant items.

Platform Fairness

The platform acts as an intermediary between consumers and providers and controls the recommendation process.

Platforms often pursue objectives such as revenue generation, user retention, marketplace growth, and long-term sustainability.

The platform's role is to balance the differing objectives of stakeholders while ensuring a healthy ecosystem

Why it matters:

- Recommendation decisions shape the entire marketplace
- Optimizing one stakeholder may penalize the others
- Platforms must balance relevance, fairness and business objectives
- Long-term ecosystem health depends on maintaining stakeholders participation

Marketplace Sustainability

Core idea:

Sustainable marketplaces require continued participation from a diverse set of stakeholders

Marketplace sustainability focuses on maintaining a healthy and viable ecosystem over time. A sustainable market should continue to attract consumers, providers and items while ensuring the recommendation opportunities remain meaningful for all participants.

Why it matters:

- Providers are more likely to participate when they receive opportunities.
- Consumers benefit from diverse marketplaces.
- Sustainable ecosystems lead to long-term platform growth.

Stakeholder Balance

Core idea:

Platforms should balance the objectives of multiple stakeholders rather than maximizing a single stakeholders outcome

Stakeholder Balance focuses on managing the trade-offs between competing stakeholders' objectives. Improving outcomes for one stakeholder may reduce the benefits of others. This makes the recommendation a multi-objective decision problem.

Why it matters:

- Recommendation systems operate in a multisided ecosystem
- Balanced outcome help maintain trust, participation and long-term sustainability

Agenda

- Context and Foundations of Fairness in Rec Sys [20mins]
- Multisided Fairness in Rec Sys
 - Stakeholders [10mins]
 - Fairness Definitions [10mins]
 - Solutions [20mins]
- Evaluation [10mins]
- Multisided Fairness: Limited Item Availability [10mins]
- Open Challenges [10mins]

FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms

FairRec is one of the first algorithms to explicitly optimize **fairness for both consumers and providers** in two-sided marketplaces.

Consumer Fairness: Envy-Free up to One Item: Consumers should not strongly prefer another user's recommendation list after removing one item.

Provider Fairness: Maximin Share Exposure: Every provider receive at least a minimum guaranteed level recommendation exposure whenever possible.

Methodology

The key idea is to transform recommendation into a fair allocation problem. Instead of simply ranking items, FairRec allocates recommendation opportunities while satisfying fairness constraints.

Instead of recommending the top-k items directly, FairRec allocates recommendation slots while ensuring providers receive a minimum level of exposure.

The recommendation process is formulated as a constrained allocation of indivisible goods, where:

- recommendation slots are the resources,
- consumers receive bundles of items,
- providers receive exposure through allocated items

FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms

FairRec is one of the first algorithms to explicitly optimize **fairness for both consumers and providers** in two-sided marketplaces.

Consumer
recommen

Provider F
level recom

Methodolo

The key id
items, Fair

Instead of
ensuring prov

Key Contributions

- One of the first algorithms to address **two-sided fairness** in personalized recommendation.
- Introduces formal fairness guarantees for **both consumers and providers**.
- Bridges recommender systems and **fair allocation theory**.
- Provides theoretical fairness guarantees while preserving recommendation quality.

ther user's

aranteed

ply ranking

s.

s while

The recommendation process is formulated as a constrained allocation of indivisible goods, where:

- recommendation slots are the resources,
- consumers receive bundles of items,
- providers receive exposure through allocated items

CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems

CPFair introduces a joint optimization framework that improves fairness for **both consumers and providers** while preserving recommendation quality.

Consumer Fairness: Reduce disparities in recommendation quality between consumer groups (e.g., active vs. inactive users).

Provider Fairness: Reduce disparities in recommendation exposure between provider groups (e.g., popular vs. long-tail providers).

Methodology

CPFair re-ranks the recommendation list by jointly optimizing relevance and fairness.

The framework computes fairness disparities from both perspectives:

- DCF (Deviation from Consumer Fairness): Measures differences in recommendation quality across consumer groups.
- DPF (Deviation from Producer Fairness): Measures differences in recommendation exposure across provider groups.

Instead of enforcing fairness as a hard constraint, CPFair formulates an optimization problem that jointly balances: recommendation accuracy, consumer fairness, provider fairness.

The objective is solved using a mixed-integer optimization formulation, followed by an efficient greedy algorithm to produce the final re-ranked list.

CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems

CPFair introduces a joint optimization framework that improves fairness for **both consumers and providers** while preserving recommendation accuracy.

Consumer
(users).

Provider
(providers).

Methodology

CPFair re-ranks

The framework

Key Contributions

- Jointly optimizes **consumer and provider fairness** within a single framework.
- Introduces fairness disparity measures (DCF and DPF) to quantify two-sided unfairness.
- Uses **optimization-based re-ranking**, making it independent of the underlying recommender model.
- Demonstrates that bilateral fairness can be improved without significantly reducing recommendation quality.

- DCF (Deviation from Consumer Fairness): Measures differences in recommendation exposure across consumer groups.
- DPF (Deviation from Producer Fairness): Measures differences in recommendation exposure across provider groups.

Instead of enforcing fairness as a hard constraint, CPFair formulates an optimization problem that jointly balances: recommendation accuracy, consumer fairness, provider fairness.

The objective is solved using a mixed-integer optimization formulation, followed by an efficient greedy algorithm to produce the final re-ranked list.

Promoting Two-sided Fairness with Adaptive Weights for Providers and Customers in Recommendation

Most two-sided fairness methods apply fairness **after recommendation generation** through re-ranking. However, this does not address biases learned during model training. This paper proposes an **adaptive weighting framework** that incorporates fairness directly into the learning process by dynamically adjusting the importance of different consumers and providers during training.

Consumer Fairness: Improve recommendation quality for disadvantaged customer groups. Reduce disparities in recommendation performance across consumers.

Provider Fairness: Improve exposure for under-served providers. Prevent recommendation opportunities from concentrating among dominant providers.

Methodology

Instead of treating every interaction equally during training, the model **assigns adaptive weights** to consumers and providers. Stakeholders experiencing poorer outcomes receive **larger training weights**, encouraging the model to pay more attention to them.

After each training iteration, the framework evaluates:

- recommendation quality across consumers
- exposure across providers

Rather than using fixed loss weights, the framework **updates the importance of each stakeholder dynamically**.

- Consumers receiving poorer recommendations receive larger training weights.
- Providers receiving less exposure also receive larger training weights.

The recommender is retrained using the updated weights. The model gradually learns representations that improve outcomes for disadvantaged stakeholders while preserving recommendation quality.

Key Contributions

- Introduces an **adaptive weighting strategy** for two-sided fairness.
- Simultaneously optimizes **consumer** and **provider** fairness during training.
- Eliminates the need for post-processing re-ranking.
- Automatically adjusts fairness priorities as training progresses.

Towards Platform Profit-aware Fairness in Personalized Recommendation

This work proposes a framework that jointly considers **consumer relevance, provider fairness, and platform profit**.

Consumer: Maintain personalized recommendation quality.

Provider: Improve fairness in provider exposure. Reduce visibility imbalance among providers.

Platform: Maximize platform profit while preserving fairness.

Methodology

The framework extends the recommendation objective by incorporating **platform profit** alongside relevance and fairness. Rather than optimizing only recommendation accuracy, the algorithm searches for recommendations that maximize a **combined objective**.

The framework evaluates how recommendation exposure is distributed across providers. Providers receiving insufficient exposure receive higher priority during optimization.

Each recommendation contributes differently to platform revenue (e.g., commissions, product prices, advertising value). The framework estimates the expected profit associated with recommending each item.

The final recommendation list is selected by jointly optimizing: recommendation relevance, provider fairness, platform profit.

Instead of maximizing one objective, the algorithm searches for a balanced solution among the three objectives.

Towards Platform Profit-aware Fairness in Personalized Recommendation

This work proposes a framework that jointly considers **consumer relevance, provider fairness, and platform profit**.

Consumer: Maintain personalized recommendation quality.

Provider:

Platform:

Method

The fra

Rather

objecti

The fra

exposu

Key Contributions

- One of the first recommender systems to explicitly include **platform profit** as an optimization objective.
- Introduces a **three-sided recommendation framework** considering consumers, providers, and the platform.
- Demonstrates that fairness and profitability can be optimized jointly rather than independently.
- Formulates recommendation as a **multi-objective optimization problem**.

fairness.

ize a **combined**

sufficient

Each recommendation contributes differently to platform revenue (e.g., commissions, product prices, advertising value). The framework estimates the expected profit associated with recommending each item.

The final recommendation list is selected by jointly optimizing: recommendation relevance, provider fairness, platform profit.

Instead of maximizing one objective, the algorithm searches for a balanced solution among the three objectives.

Regret-aware Re-ranking for Guaranteeing Two-sided Fairness and Accuracy in Recommender Systems

Most provider fairness methods allocate exposure based on the **current recommendation list**. However, providers that repeatedly receive insufficient exposure may accumulate **unfairness over time**.

Regret-Aware Re-ranking introduces the concept of **provider regret**, ensuring that providers who have historically received less exposure are prioritized in future recommendations.

Consumer Fairness: Preserve recommendation accuracy and user satisfaction.

Provider Fairness: Reduce cumulative provider regret. Ensure providers receive their deserved recommendation opportunities over time.

Methodology

Instead of considering only **current exposure**, the algorithm maintains a history of provider exposure. Providers with **higher accumulated regret** receive higher priority during re-ranking.

For each provider, the algorithm compares: 1. expected recommendation exposure 2. actual recommendation exposure

The difference represents the provider's **regret**. Higher regret indicates that the provider has consistently received fewer recommendation opportunities than expected.

The recommendation list is adjusted by jointly considering: recommendation relevance, current provider exposure, accumulated provider regret.

Providers with larger regret receive increased priority.

After recommendations are delivered, provider regret is updated and carried forward to future recommendation rounds.

Regret-aware Re-ranking for Guaranteeing Two-sided Fairness and Accuracy in Recommender Systems

Most provider fairness methods allocate exposure based on the **current recommendation list**. However, providers that repeatedly receive insufficient exposure may accumulate **unfairness over time**.

Regret-Aware Re-ranking introduces the concept of **provider regret**, ensuring that providers who have historically received less exposure are prioritized in future recommendations.

Consumer Fair

Provider Fair

Methodology

Instead of con
regret receive

For each prov
The difference
opportunities tr

Key Contributions

- Introduces **provider regret** as a new fairness objective.
- Considers **historical fairness**, rather than only current recommendation lists.
- Simultaneously optimizes consumer accuracy and provider fairness.
- Extends two-sided fairness to sequential recommendation scenarios.

The recommendation list is adjusted by jointly considering: recommendation relevance, current provider exposure, accumulated provider regret.

Providers with larger regret receive increased priority.

After recommendations are delivered, provider regret is updated and carried forward to future recommendation rounds.

OKRA: An Explainable, Heterogeneous, Multi-stakeholder Job Recommender System

OKRA proposes a **multi-stakeholder job recommender system** that combines personalized recommendations with explainability to support fair and trustworthy hiring decisions.

Candidate (Consumer) Fairness: Recommend jobs based on relevant qualifications and preferences. Improve equal access to employment opportunities.

Recruiter / Employer (Providers) Fairness: Recommend qualified candidates while avoiding systematic bias. Increase transparency in recommendation decisions.

Explainability: Provide understandable reasons for recommendations to support trust and accountability.

Methodology

Unlike previous methods that modify ranking or optimization objectives, OKRA models the recruitment ecosystem as a **heterogeneous graph**.

The graph connects different entity types (candidates, jobs, recruiters, skills, companies), allowing recommendations to be generated by exploiting multiple relationship types.

A graph-based recommendation model learns embeddings for each entity by exploiting the heterogeneous relationships within the recruitment ecosystem.

The learned representations are used to recommend:

- jobs to candidates,
- candidates to recruiters,
- while considering the needs of multiple stakeholders.

Rather than returning only recommendation scores, OKRA also provides explanations based on graph paths and entity relationships. Examples include: matching skills, previous experience, company requirements, shared qualifications.

These explanations help stakeholders understand why recommendations were made.

OKRA: An Explainable, Heterogeneous, Multi-stakeholder Job Recommender System

OKRA proposes a **multi-stakeholder job recommender system** that combines personalized recommendations with explainability to support fair and trustworthy hiring decisions.

Candidate (Consumer) Fairness: Recommend jobs based on relevant qualifications and preferences. Improve equal access to employment opportunities.

Recruiter
recommen
Explaina

Methodol

Unlike pre
graph.

The graph
exploiting

A graph-b
recruitment

The learned

- ## Key Contributions
- Introduces a **heterogeneous graph-based recommender** for recruitment.
 - Supports **multiple stakeholders** within a single recommendation framework.
 - Integrates **explainability** directly into the recommendation process.
 - Demonstrates how transparent recommendations can support fairer hiring decisions.

- jobs to candidates,
- candidates to recruiters,
- while considering the needs of multiple stakeholders.

Rather than returning only recommendation scores, OKRA also provides explanations based on graph paths and entity relationships. Examples include: matching skills, previous experience, company requirements, shared qualifications.

These explanations help stakeholders understand why recommendations were made.

ProFairRec: Provider Fairness-aware News Recommendation

News recommendation models are trained on historical click data, which is often biased toward **popular news providers**. As a result, recommendation models learn these biases and repeatedly favor dominant providers, reducing the visibility of smaller or minority news publishers. ProFairRec addresses this issue by learning provider-fair representations during model training.

Provider Fairness:

- Reduce provider bias inherited from historical click data.
- Ensure that recommendation opportunities are more evenly distributed across news providers.
- Improve exposure for **minority or less popular providers** while maintaining recommendation accuracy.

Methodology

The central idea is to **separate provider-specific information from recommendation-relevant information**.

1. For each news article, the model learns:

- a provider-fair representation from the article content,
- a provider-biased representation from the provider identity.

2. An adversarial network attempts to predict the news provider from the **provider-fair representation**. The recommendation model is trained to **prevent** this prediction. As a result, the provider-fair representation contains recommendation-relevant information while removing provider identity as much as possible.

3. The model encourages the provider-fair and provider-biased representations to be **orthogonal** (independent). This improves the separation between useful recommendation signals and provider-specific bias.

4. Only the **provider-fair representations** of news and users are used for matching and ranking candidate news.

ProFairRec: Provider Fairness-aware News Recommendation

News recommendation models are trained on historical click data, which is often biased toward **popular news providers**. As a result, recommendation models tend to recommend popular news providers, leading to the visibility of smaller, minority news publishers.

Provider Fairness

- Reduce provider bias
- Ensure fairness
- Improve recommendation quality

Methodology

The central idea is to create a **provider-fair representation** of news providers, which is orthogonal to the provider-biased representation.

1. For each news provider, we create two representations:

- a **provider-biased representation** (capturing popularity)
- a **provider-fair representation** (capturing recommendation-relevant information)

2. An adversarial model is trained to **prevent** this prediction. As a result, the provider-fair representation contains recommendation-relevant information while removing provider identity as much as possible.

3. The model encourages the provider-fair and provider-biased representations to be **orthogonal** (independent). This improves the separation between useful recommendation signals and provider-specific bias.

4. Only the **provider-fair representations** of news and users are used for matching and ranking candidate news.

Key Contributions

- First work to explicitly study **provider fairness in news recommendation**.
- Introduces a **representation learning framework** for provider fairness.
- Combines adversarial learning with orthogonal regularization to remove provider bias.
- Designed as a **model-agnostic framework**, allowing integration with many existing news recommendation models.

The Many Facets of Fairness in Recommender Systems

The Many Facets of Fairness in Recommender Systems

Motivation


Existing fairness-aware recommender systems typically optimize fairness for a **single stakeholder**, most commonly consumers or providers. However, recommender systems affect **multiple stakeholders simultaneously**, and improving fairness for one stakeholder may negatively impact others. This work proposes a **unified framework** for studying fairness across **consumers, providers, and items** within the same recommendation process.

Research Questions

1. Can multiple stakeholder fairness objectives be evaluated within a common framework?
2. How do fairness interventions for one stakeholder affect the others?
3. Which fairness trade-offs emerge when multiple stakeholders are considered simultaneously?

Key Idea

Fairness should be evaluated as a property of the entire recommendation ecosystem rather than from a single stakeholder's perspective.




Methodology

- 1 Input: Historical Interactions**
User-item interactions and contextual information.
- 2 Base Recommendation**
A baseline recommender generates candidate items for each user.
- 3 Fairness Evaluation Layer**
Evaluate the candidate list from three stakeholder perspectives.

Key Idea

Unlike previous methods, **fairness is evaluated jointly**, allowing the framework to observe interactions among stakeholder objectives.


- 4 Fairness-aware Re-ranking**
Re-rank candidates to balance fairness across consumers, providers, and items.




Findings and Insights

Observed Trade-offs Among Multiple Fairness Objectives

Legend: ↑ Improves ↓ Decreases — May have little / mixed effect



- ### Key Contributions
- First unified study of **consumer, provider, and item fairness**.
 - Demonstrates that fairness objectives interact rather than operate independently.
 - Provides a common evaluation framework for multisided fairness.
 - Highlights the importance of analyzing fairness trade-offs instead of optimizing a single stakeholder.



Key Insight

No single objective can be optimized without affecting others. A balance must be achieved based on the application's priorities.



Takeaway

Fairness in recommender systems is inherently multi-dimensional. Enhancing fairness for one stakeholder may lead to compromises for others. Careful design and tuning are essential to achieve an appropriate balance.

In Summary


Paper	Consumers	Providers	Items	Platform	Core Idea
FairRec	✓	✓	✗	✗	Fair allocation of recommendation opportunities
CPFair	✓	✓	✗	✗	Joint fairness-aware re-ranking
Adaptive Weights	✓	✓	✗	✗	Adaptive fairness-aware training
Profit-aware	✓	✓	✗	✓	Joint optimization including platform profit
Regret-aware	✓	✓	✗	✗	Compensate historically under-exposed providers
OKRA	✓	✓	✗ [*]	✗	Explainable heterogeneous graph recommendation
Many Facets	✓	✓	✓	✗	Unified consumer–provider–item fairness framework

* **OKRA** recommends jobs (items), but the paper's fairness objectives are centered on **candidates and recruiters/employers**, rather than explicitly defining or optimizing **item fairness**. Therefore, I would leave the **Items** column blank.

Agenda

- Context and Foundations of Fairness in Rec Sys [20mins]
- Multisided Fairness in Rec Sys
 - Stakeholders [10mins]
 - Fairness Definitions [10mins]
 - Solutions [20mins]
- Evaluation [10mins]
- Multisided Fairness: Limited Item Availability [10mins]
- Open Challenges [10mins]

Fairness Evaluation




Consumer-side Fairness Evaluation

Do all users receive comparable recommendation quality?


Typical evaluation metrics

- NDCG difference across user groups
- Performance gap or prediction error gap
- Deviation from consumer fairness
- Min-max ratio for individual-level quality
- Entropy of received limited items over time

Group Fairness Comparison



Individual Fairness Comparison



Provider-side Fairness Evaluation

Are exposure and visibility distributed fairly across providers?

Typical evaluation metrics

- Exposure disparity or exposure ratio
- Disparate visibility across provider groups
- Gini index, entropy, variance, or coefficient of variation
- Max-min fairness and maximin share
- Fraction of providers receiving enough exposure




Item-side Fairness Evaluation

Does provider fairness also reach the items inside each inventory?

Typical evaluation metrics

- Long-tail coverage and catalog coverage
- Novelty or average percentage of long-tail items
- Entropy of item exposure
- Deviation from item fairness
- Exposure comparison across item groups

Why provider fairness is not enough?




Item-side metrics reveal whether exposure is concentrated on a small subset of items.

Platform and system-level Evaluation

Can fairness improve without harming broader platform goals?

Typical evaluation metrics

- Revenue, profit, gains, price, spending, or pay-per-click
- Engagement and long-term participation
- Cumulative delivery distribution for scarce items
- Joint reporting of utility and fairness metrics
- Trade-off analysis across stakeholders



Challenge: Improving fairness for one stakeholder may reduce utility or fairness for another. Therefore, recent work reports several fairness and utility metrics together.

Agenda

- Context and Foundations of Fairness in Rec Sys [20mins]
- Multisided Fairness in Rec Sys
 - Stakeholders [10mins]
 - Fairness Definitions [10mins]
 - Solutions [20mins]
- Evaluation [10mins]
- Multisided Fairness: Limited Item Availability [10mins]
- Open Challenges [10mins]

Fairness Under Limited Item Availability


Many fairness-aware recommender systems assume that items are always available.

In many real-world settings, items are limited. In this case, recommendation becomes not only a **ranking problem**, but also an **allocation problem**.

Allocating a scarce item to one consumer may reduce the chance for others to receive it.

Therefore, fairness must consider:

- who receives the opportunity
- when they receive it
- how allocations affect both consumers and providers over time



SoCRATe: Compensating Users When Items Are Scarce

Main idea

- Limited items create competition and unequal recommendation quality.
- SoCRATe compensates users across multiple sessions for past recommendation loss.

How SoCRATe works

- Measures the gap between delivered and optimal recommendations.
- Tracks cumulative loss across sessions.
- Prioritizes users with higher past loss.
- Allocates scarce items using preference-driven or round-robin compensation.


Why it is important

- Prevents the same users from being repeatedly disadvantaged.
- Loss-based prioritization reduces disparity, especially with round-robin allocation.

Why Consumer History Matters?

To allocate scarce items fairly, consumers can be viewed in three broad groups:

- **Loyal users**
 - users with repeated positive interactions with a provider or category
 - they may expect some continuity or reward for their past engagement
- **Neutral users**
 - users with no strong historical preference or repeated interaction pattern
 - they should still receive fair access without being overshadowed by highly active users
- **Newcomers**
 - new or low-history users with limited interaction data
 - they are often disadvantaged in recommendation because the system knows less about them



Key fairness challenge:

- If the system favors only loyal users, **newcomers may be excluded**.
- If it ignores loyalty completely, it may **fail to reflect meaningful user-provider relationships**.
- A fair system should **balance past interactions with equal opportunity**.

EQUALI: A Framework for Fair Allocation of Limited Items

Main idea

- EQUALI is designed for situations where relevant items are scarce.
- It does not only ask: “Who is most likely to like this item?”
- It also asks: “Who has been underserved so far?”

How EQUALI works

- Uses the recommendation score from the base recommender system.
- Looks at past delivery history.
- Checks whether some groups received too many or too few limited items.
- Gives better chances to underserved groups.
- Updates the delivery history after each session.

Why it is important

- It moves fairness-aware recommendation closer to real-world scenarios.
- It recognizes that limited resources create competition.
- It shows that fairness should be studied as a dynamic allocation problem, not only as a static ranking problem.

Agenda

- Context and Foundations of Fairness in Rec Sys [20mins]
- Multisided Fairness in Rec Sys
 - Stakeholders [10mins]
 - Fairness Definitions [10mins]
 - Solutions [20mins]
- Evaluation [10mins]
- Multisided Fairness: Limited Item Availability [10mins]
- Open Challenges [10mins]

Open Challenges

Balancing Multiple Fairness Objectives

- Balance competing fairness goals across consumers, providers, items, platforms.
- Improving fairness for one group may reduce utility or exposure for another.
- Design algorithms that manage these trade-offs while preserving recommendation quality.

Evaluation Metrics for Multisided Fairness

- Current focus on individual stakeholders (e.g., accuracy for users, exposure for providers).
- No unified framework for jointly evaluating fairness across all stakeholders.
- Developing comprehensive evaluation methodologies remains an open research problem.

Dynamic and Longitudinal Fairness

- Fairness is often assessed in single sessions, ignoring long-term effects.
- Past recommendations influence future opportunities, especially under scarcity.
- Ensuring fair allocation over time, rather than at a single point, is still largely unresolved.

Open Challenges

Explainability and Transparency

- Fairness-aware recommendations require explanations that clarify decisions and trade-offs.
- Integrating explainability with multisided fairness is still underexplored.
- Transparent systems are crucial for trust among users, providers and platforms.

Real-World Deployment and User Studies

- Most methods are evaluated offline, with limited real-world validation.
- The impact of fairness interventions on user behavior, provider engagement and system sustainability remains unclear.
- More online experiments and user studies are needed.

Thank You!

Part of the ALTER project:

Fairness, Explanations, and Interactivity in Automated Decision Systems

Funded by the **Research Council of Finland**

References

- 1.. Azzalini, D., Azzalini, F., Criscuolo, C., Dolci, T., Martinenghi, D., Amer-Yahia, S.: Loss compensation in multi-session recommendation under limited availability. In: EDBT. pp. 522–533 (2024)
2. Bauer, C., Ferraro, A.: Strategies for mitigating artist gender bias in music recommendation: a simulation study. In: Music Recommender Systems Workshop (2023)
3. Dalla Vecchia, A., Migliorini, S., Quintarelli, E., Stefanidis, K.: Multi-sided fairness in sequential task assignment. Information Systems 139(3), 102700 (7 2026)
4. Gao, H., Lu, Y., Zhu, N.: A Fairness-Aware Recommender System Based on Genetic Algorithms for User, Provider and Platform Optimization. ICAACE (2025)
5. Gómez, E., Contreras, D., Boratto, L., Salamó, M.: Enhancing recommender systems with provider fairness through preference distribution-awareness. Information Management Data Insights 5(1), 100311 (2025)
6. Kaya, M., Bogers, T.: Mapping stakeholder needs to multi-sided fairness in candidate recommendation for algorithmic hiring. In: RecSys. p. 257–267 (2025)
7. Liu, S., Sun, J., Deng, X., Wang, H., Liu, W., Zhu, C., Yang, L.T., Wu, C.: Towards platform profit-aware fairness in personalized recommendation. International Journal of Machine Learning and Cybernetics 15(10), 4341–4356 (2024)
8. Naghiaei, M., Rahmani, H.A., Deldjoo, Y.: Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. In: ACM SIGIR (2022)
9. Patro, G.K., Biswas, A., Ganguly, N., Gummadi, K.P., Chakraborty, A.: FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. The Web Conference (2020)
10. Qi, T., Wu, F., Wu, C., Sun, P., Wu, L., Wang, X., Huang, Y., Xie, X.: ProFairRec: Provider Fairness-aware News Recommendation. SIGIR pp. 1164–1173 (7 2022)

References

11. Schellingerhout, R., Barile, F., Tintarev, N.: Okra: An explainable, heterogeneous, multi-stakeholder job recommender system. In: Advances in IR (2025)
12. Shafiloo, R., Stratigi, M., Peltonen, J., Olsson, T., Stefanidis, K.: The many facets of fairness in recommender systems: Consumers, providers and items. Information Systems 137, 102643 (4 2026)
13. Shafiloo, R., Stratigi, M., Peltonen, J., Stefanidis, K.: Multisided fairness under limited item availability in recommender systems. Information Sciences 731 (2026)
14. Tang, J., Shen, S., Wang, Z., Gong, Z., Zhang, J., Chen, X.: When fairness meets bias: a debiased framework for fairness aware top-n recommendation. In: RecSys (2023)
15. Xu, C., Chen, S., Xu, J., Shen, W., Zhang, X., Wang, G., Dong, Z.: P-mmfr: Provider max-min fairness re-ranking in recommender system. In: The Web Conf. (2023)
16. Xu, L., Lin, Z., Wang, J., Chen, S., Zhao, W.X., Wen, J.R.: Promoting Two-sided Fairness with Adaptive Weights for Providers and Customers in Recommendation. RecSys (2024)
17. Ye, X., Xu, C., Xu, J., Xie, X., Noah', H., Lab, A., Dong, Z., Wang, G.: Regret aware Re-ranking for Guaranteeing Two-sided Fairness and Accuracy in Recommender Systems. ACM Transactions on Information Systems 1(1) (4 2025)
- 18.. Boratto, L., Fenu, G., Marras, M.: Interplay between upsampling and regularization for provider fairness in recommender systems. User Modeling and User-Adapted Interaction 31(3), 421–455 (2021)
19. Pitoura, E., Stefanidis, K., Koutrika, G.: Fairness in rankings and recommendations: an overview. VLDB J. 31(3), 431–458 (2022)
20. Stratigi, M., Nummenmaa, J., Pitoura, E., Stefanidis, K.: Fair sequential group recommendations. In: SAC. pp. 1443–1452 (2020)