FAIR RECOMMENDATION THROUGH CONDITIONING ON SYSTEMATIC PRESENTATION IN RECOMMENDER SYSTEMS

by

Özge Bozal

B.Sc., Electrical and Electronics Engineering, Boğaziçi University, 2018

Submitted to the Institute for Graduate Studies in
Science and Engineering in partial fulfillment of
the requirements for the degree of
Master of Science

Graduate Program in Computational Science and Engineering
Boğaziçi University
2021

ACKNOWLEDGEMENTS

I would like to thank my valuable professors and thesis advisors Prof. Ali Taylan Cemgil and Assist. Prof. Mustafa Gökçe Baydoğan for their expertise and support throughout the writing of this thesis.

I feel privileged to be one of the students of invaluable professors and to befriend a group of friendly and supportive people during my studies.

I am especially indebted to my friends Gökhan Çapan and İlker Gündoğdu whose deep knowledge I referred to during this work. I am grateful for their collaborations and counseling along the way.

Finally, I would like to thank my family for their love and continuous support, which gave me the strength to accomplish my studies.

ABSTRACT

FAIR RECOMMENDATION THROUGH CONDITIONING ON SYSTEMATIC PRESENTATION IN RECOMMENDER SYSTEMS

Recommendation engines are interactive systems aiming to predict users' top preferences according to their choices. A limitation of the recommender systems is that over time the recommendations get narrower in scope. That means the popular items are favored and the less frequently visited ones are censored. Consequently, users are prevented from exploring new alternatives. The so-called filter bubble is the inevitable outcome of the self-reinforcing feedback loops where the preference estimations depend on the user choices.

Self-reinforcing feedback loops are both cause and effect of over- and under-presentation of some content in interactive recommender systems. That results in inaccurate user preference estimates, namely, overestimation of over-presented content and vice versa. The burden is on the recommender system to eliminate the self-reinforcing feedback loops for a more accurate preference estimation. In this regard, we define the "fairness" criteria for an interactive recommender system considering the adverse impacts of these self-reinforcing feedback loops. We also claim that designing an intelligent presentation mechanism is essential to meet those criteria. To prove the claim we address two models that explicitly incorporate or ignore the systematic and limited exposure to alternatives. By simulating real-world biases, we demonstrate that ignoring systematic presentations results in overestimation of promoted options and underestimation of censored alternatives. Simply conditioning on the limited exposure is a remedy for these biases.

ÖZET

ÖNERİ SİSTEMLERİNDE SİSTEMATİK SUNUM ÜZERİNE KOŞULLAMA İLE ADİL ÖNERİLER

Öneri sistemleri, kullanıcılarının tercihlerini doğru şekilde tahminlemeyi amaçlayan etkileşimli sistemlerdir. Bu sistemlerin bir sınırlaması, zamanla tavsiyelerin kapsamının daralmasıdır. Bu ise, popüler seçeneklerin daha çok önerildiği, daha az ziyaret edilenlerin adeta sansürlendiği, böylece kullanıcıların yeni alternatifleri keşfetmesinin engellendiği bir ortam yaratır. Filtre balonu adı verilen bu durum, tercih tahminlemelerinin kullanıcı geri bildirimlerine bağlı olduğu kendi kendini güçlendiren geri bildirim döngülerinin kaçınılmaz sonucudur.

Kendini güçlendiren geri bildirim döngüleri, etkileşimli öneri sistemlerinde bazı içeriklerin gereğinden fazla ve/veya az sunulmasının hem nedeni hem de sonucudur. Bu durum, yanlış kullanıcı tercihi tahminlemelerine, yani fazlaca sunulan içeriklerin fazla tercih edildiği ve aynı şekilde az sunulan içeriğin ise daha az tercih edildiği yanılgısına sebep olur. Daha hassas bir tercih tahminlemesi için kendi kendini güçlendiren geri bildirim döngülerini ortadan kaldırmaktan ise öneri sisteminin kendisi sorumludur. Bu bağlamda, kendi kendini güçlendiren geri bildirim döngülerinin olumsuz etkilerini göz önünde bulundurarak etkileşimli bir öneri sistemi için "adil" olma kriterlerini tanımlıyoruz. Ayrıca, akıllı bir sunum mekanizması tasarlamanın bu kriterleri sağlamak için gerekli olduğunu savunuyoruz. Bu savı kanıtlamak için, sırasıyla alternatiflere sistematik ve sınırlı maruz kalmayı açıkça içeren ve görmezden gelen iki modeli ele alıyoruz. Gerçek dünyadaki yanlılıkları simüle ederek, sistematik sunumları görmezden gelmenin teşvik edilen seçenekleri olduğundan fazla, sansürlenmiş alternatifleri ise olduğundan az tahminlediğini gösteriyoruz. En basit şekilde sınırlı maruz kalma üzerine koşullandırmanın bu yanlılıkları hafifletebileceğini gösteriyoruz.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF SYMBOLS

$B(.)$	Beta function
$Beta(.)$	Beta distribution
$Bin(.)$	Binomial distribution
$C_{1:T}$	Sequence of the variable C from time 1 to T
$Dir(.)$	Dirichlet distribution
$Mult(.)$	Multinomial distribution
$p(.)$	Probability density function
$\Gamma(.)$	Gamma function
θ	Parameters of a probability distribution
$\hat{\theta}$	Estimated parameter
$\phi(.)$	Potential function
$\mathbb{E}[.]$	Expected value
$\mathbb{E}_{p(.)}[.]$	Expectation under a probability density function

LIST OF ACRONYMS/ABBREVIATIONS

MAP	Maximum A-Posteriori
ML	Maximum Likelihood

1. INTRODUCTION

In today's digital world, we are constantly exposed to abundant content of many sorts, which makes it difficult for us to find what we in fact look for. Personalized recommender systems appear on many different platforms, such as music, movie, social media, e-commerce, travel, to provide us with the filtered options among many others that do not interest us. Even though recommender systems have been making our life easier, one should raise the question of whether they are capable of estimating what we truly like. A more important question would be, whether our choices are shaped by what we are exposed to. The answer being yes leaves us with a very serious ethical issue that can harm considerably both individuals and society as a whole.

As in many other machine learning research studies, recommender system researchers also focus on the societal impacts of the algorithms and the models run behind the recommender systems. Are these models capable of correctly modeling human behavior? Are they free of bias? Do they treat each content equally? In other words, do the recommender systems pay regard for fairness concerns, at all?

Personalized recommender systems estimate user preferences based on their behavior (clicks, purchases, ratings, reviews, etc.) and present options to the user. Users, in turn, select the options appealing to them, the information of which is then fed back to the system. This is a chicken-egg problem, in which we cannot differentiate whether user behavior shapes the presented set of options or the presentation mechanism of recommender systems manipulates users' behavior. Obviously, the future choices of users will be affected by the presentation mechanism due to this self-reinforcing feedback loop [1]. Over-estimated, thus over-presented options will be favored, whereas under-estimated, namely, less presented options will be undervalued over time. Societal impacts of bias rooted from self-reinforcing feedback loop are two-fold: Firstly, it prevents users to reach the contents that truly match their interest. Secondly, it violates the right of content providers to be presented under equal terms. Assume an e-commerce platform where the buyers are members of the website and the sellers are both small, medium-sized and large enterprises. If the recommender system of this e-commerce website favors the best selling items, then it most probably will favor the large enterprises which have a larger number of customers in the first place, being unfair to smaller businesses. In fact, smaller businesses might offer lower-priced yet same quality items, which buyers would find more appealing. As in the example, *popular* and *promoted* contents get a lot of exposure leading to a biased estimation of user preferences.

The natural design of personalized recommender systems being solely based on user feedback on presented options raises the question, how we can build a presentation mechanism not over- or under-estimating any option, namely, aiming for a fair and responsible recommendation. The motivation of this thesis is the search for the fairness criteria of personalized recommenders and what kind of a model might meet them.

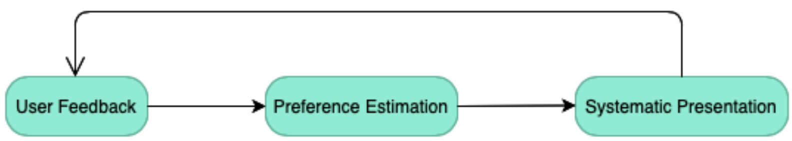

Figure 1.1. Self-reinforcing feedback loop in personalized recommender systems.

1.1. Background on Personalized Recommender Systems

Recommender systems are platforms that provide the users of the platform with the appropriate content. Namely, they serve as a filter for the users to reach what they look for easily. On the other hand, recommendations are especially essential for some businesses where a competent and efficient recommender system can make a huge difference in a company's income. For instance, entertainment platforms such as Netflix, YouTube, social networking websites such as Facebook, Instagram, and any type of e-commerce company highly depend on their recommender systems.

Recommender systems are built upon algorithms based on some assumptions on users and items to produce a personalized recommendation. They utilize both user logs and item features to model user-item relations. We can categorize recommender systems into two in terms of how they leverage data: content-based and collaborative filtering.

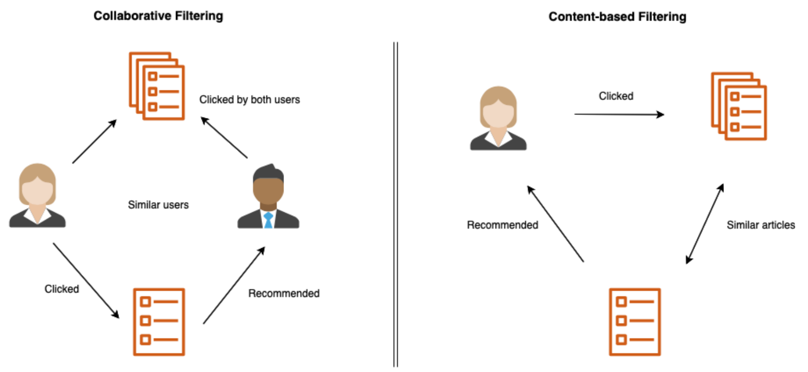

Figure 1.2. Collaborative filtering (left) is based on similar user-item interactions. Content-based filtering (right) is based on user preferences on item features.

The content-based filtering approach reduces recommendation problem in classification or regression, such that for each user a classifier or regressor is trained, provided that item features are known. The idea is straightforward: If user's choice on an item is binary (as *like* or *dislike*), a classifier is trained. If the user's rating on an item is concerned, then the regressor is trained.

In collaborative filtering, the assumption is that a user's next behavior on an item solely depends on his/her previous behavior on other items. That can be of two sorts: memory-based and model-based. Former depends on the similarity among users or items, whereas the latter claims a generative model for user-item interaction.

The memory-based approach in collaborative filtering is the nearest neighbor. There are two options to consider: user-based or item-based. The user-based approach pays regard to actions of similar users to estimate the rating of a user on a specific item. The item-based approach, on the other hand, assumes that similar items get similar ratings from the same user. Both approaches do not consider any user- or item-specific features. There are several limitations of the nearest neighbor method. Firstly, it cannot cope with sparsity when there is an insufficient amount of ratings and secondly, it becomes computationally inefficient for real-time applications when the number of users and items grows rapidly.

Matrix factorization is a method of decomposing the target matrix, in this context the user-item interaction matrix, into low-dimensional latent space matrices, which are assumed to generate the target matrix. When compared to the memory-based models, the matrix factorization model is advantageous in terms of scalability and computational efficiency.

1.2. Related Work on Fair Recommendation

Recommender systems constitute a major part of many different types of platforms such as entertainment applications *Netflix, Spotify, YouTube*, e-commerce platforms like *Amazon, Aliexpress, Hepsiburada, Trendyol*, job matching, and online dating platforms which have more than one stake-holder at hand to be satisfied by the recommendations. Personalized recommender systems facilitate our lives in many ways and in most cases, they save us time.

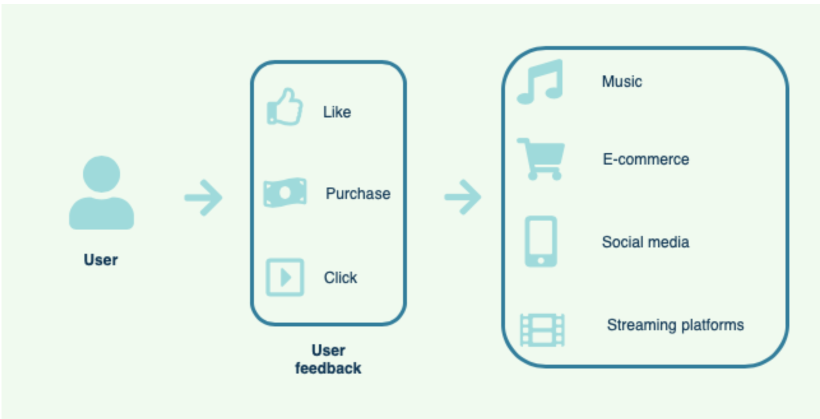

Figure 1.3. Several recommender systems and user-system interaction scenarios.

Lately, fair algorithmic design is a trending topic in research communities [2–8]. Fairness in recommender systems is also of high importance due to the massive societal and economic impacts of recommendation engines on large populations. Recommender systems suffer from statistical biases caused by algorithmic design or inherent in the parties involved in the recommender system platforms. Its interactive nature enables a recommender system to update its user preference estimation, which in turn is expected to provide a better recommendation. Although it seems a legitimate process, it is questionable to what extent they are fair in the estimation of both user preference and item relevance.

Personalized interactive systems present a subset of all the existing alternatives in the platform assuming one would choose the most relevant/appealing option among those alternatives shown to her in the first place. The user feedback can be clicks, purchases, ratings, etc.

When initially *popular* or *promoted* contents are recommended, the rest of the contents are doomed to be under-presented, if at all presented, therefore under-estimated. *Popularity bias* inherent in every recommender system platform gives rise to a vicious cycle such that over-presented items get more attention and/or higher ratings, similarly, under-presented items are obliged to being under-estimated [2–4, 7, 9, 10]. The phenomenon of biased user preference estimation due to users choosing among the options that are presented in the first place is called *observation bias* [6, 11]. If there is no intervention into the system, this *self-reinforcing feedback loop* prevents users of such platforms from reaching a diverse range of alternatives leaving them in a *filter bubble* [1]. One other adverse effect of this self-reinforcing feedback loop is that it censors initially not promoted alternatives violating the right of content providers to be equally presented [12]. In such a setup, initially unpopular or un-promoted contents will continue to appear in recommendations very rarely [13], therefore, users won't have the chance to discover whether they like or dislike such contents.

Censorship of particular contents has economic impacts on the content providers. In his book [14], Chris Anderson argues the efficiency of selling relatively more popular products in a particular market, pointing out, the total of the products having low sales volume, may actually be equal to or exceed the revenue coming from bestselling products. In [15], the long tail problem concerning the recommender systems is discussed. Concordantly, rare books, which do not often available in brick and mortar stores constitute a significant amount of Amazon's book sales [16]. This phenomenon becomes evident in digital platforms, where abundant options are available, most of which, however, ranked only by a few users. Recommending a limited amount of "hit" items will not be the most beneficial thing to practice neither for users nor for product owners [10].

In the evaluation of recommendations, the main concern is usually user satisfaction. However, most of the time user satisfaction is not the only metric to measure the performance of a recommendation. Recommender systems having multiple stakeholders such as e-commerce, online dating, and job matching platforms should consider each stakeholder's interest [17, 18]. Neglecting any party involved in the recommendation would have both obvious and unnoticed social and economic consequences.

Recommender systems with collaborative filtering suffer from a "cold-start problem", namely, introducing a new item or a new user to the platform is a challenge in a fair and accurate recommendation. Collaborative filtering solely relies on past user-item interactions and there are no sufficient data available describing a new user or a new item. Therefore, the initial recommendation is a challenging task whether a new item or a new user is at stake [9]. It is indeed risky to recommend newly introduced items, on the other hand, without being presented, no accurate preference estimation can be made on them. That being said, the recommender systems should be in balance between exploiting already high-ranked options and enabling users to explore their potentially favorite alternatives. On the other hand, the initial choices of a new user will be favored by the system, not only preventing the system from making accurate estimation of her/his preferences, but also censoring other alternatives due to the bias towards the user's initial choices.

One other non-negligible and devastating reason for the filter bubble is in fact a human behavior, which is the tendency to keep one's prior beliefs on subjects [19]. Namely, people's interests are subject to degenerate in the presence or absence of a recommender system. Nevertheless, the recommender systems can accelerate or slow down such degeneracies by offering exploratory recommendations and growing candidate pools [20].

The study [21] lays the theoretical background of the aforementioned bias types induced by the self-reinforcing feedback loops in collaborative filtering recommender systems. To eliminate the feedback loops there are quite an amount of research done so far.

In a collaborative recommendation setting, the authors of [22] focus on recovering the true rating matrix in the absence of feedback loops. To do so, they posed several assumptions on the iterative nature of the recommendation to deconvolve the observed rating matrix to true ratings and the ratings induced by the recommender system via singular value decomposition method.

Based on the findings on the relationship between diversity and recommendation in [23], the users tend to stay in filter bubbles even without recommendation. In this direction, the authors of [24] investigated whether accurate recommendations provide the highest user satisfaction and claimed that a recommender system should discover unexpected yet useful items for the user. For this purpose, they proposed a serendipity metric that evaluates the competency of the recommender system in finding such items for the users. They claimed that traditionally collected user feedbacks are not sufficient and the recommender systems should incorporate additional data on individual user beliefs on the items they haven't seen yet, how preferences change over time, and to what extent the users are risk-avert, *i.e.*, they choose the least uncertain item in parallel to their preferences. They proposed a model based on the expected utility theory that can derive and utilize this information, thereby enables the user to choose a better item than she would without any recommendation.

Another challenge for recommender systems is that some users tend to rate only specific types of items, which may bring bias into the system that those users would not see items belonging to some other types. In [6] the authors addressed fairness concerns in a hybrid (utilizing user-user similarity, item-item similarity, content, and demographic information) recommender system setup that accounts for both observation bias and the bias induced from an imbalance in the data. The authors implemented a set of logical rules accounting for both of the fairness concerns using a probabilistic programming language, that is probabilistic soft logic.

Counterfactual estimation aims to find the real cause of an event discarding the independent factors that might not be noticeable in the first place. It makes use of causal inference, a trending topic in various fields such as social science, economics, politics, etc. The authors of [25] addressed the item-to-item recommendation (recommendation based on item similarity) problem as a counterfactual estimation task and claimed that popularity, the release date of an item or, the time a particular item is in the inventory confound the user feedback to the recommendations leading to biased estimations. The proposed model debiases the listed variables due to the limited exposure of the users to the items, namely, the observation bias, thereby finds the causal parameters to the partial feedback in the real-world cases by using small annotated dataset.

Another remarkable impact of the self-reinforcing feedback loops is the "homogenization" of user behavior due to the collaborative filtering in the recommender systems, which can be defined as similar users being exposed to a very similar set of options over the course of interaction that leaves them with a small chance of discovering new items [26]. The authors' warning to the researchers is algorithmic confounding rooted in self-reinforcing feedback loops increases the homogenization of user behavior without any gain in the utility. To prove that, they compared several interactive recommendation setups in terms of their homogenizing effect across users.

So far, we explain various bias types together with their root causes, and the consequences are discussed concerning the recommender system design. We give several different approaches that address debiasing the algorithmic confounding caused by self-reinforcing feedback loops due to the partial user feedback. A fair recommender system should take the limited exposure of the user into account, otherwise, it penalizes the options the user has never seen. A clever presentation mechanism would provide the user with favorable options and also discover the censored favorites [27].

In this work, we assume an interactive recommendation setup where only partial user feedback is available for user preference estimation. We investigate whether a Bayesian choice model accounting for the limited and systematic exposure to alternatives can eliminate some of the potential biases defined here. We give the details regarding the discrete choice models, Bayesian treatments for preference estimation and, the discovery task, namely, the presentation mechanism in Section 3.

1.3. Scope of the Thesis

In this thesis, we define the fairness criteria of a personalized recommender system given potential biases interfering with the recommendation. We claim that a model which satisfies the independence of irrelevant alternatives and also accounts for the presentation mechanism would break the self-reinforcing feedback loops, therefore, meet those criteria.

For this purpose, we investigate a Bayesian choice model that accounts for the systematic presentation of the recommender system together with the presentation mechanism in [28]. In accordance with the fairness criteria, we prepare simulation scenarios of common challenges for real-world recommender systems.

To prove that accounting for the systematic presentation is crucial in debiasing the recommendation, we compare a Bayesian choice model, the Dirichlet-Luce model, to the naive Dirichlet-Multinomial model which ignores the systematic presentation.

We use the same presentation mechanism with a single-user setup and show that the former model satisfies the fairness criteria. We also investigate whether the presentation mechanism can eliminate the potential bias induced by the model itself.

The thesis is organized as follows: First, we give background information on personalized recommender systems and related work on fairness issues, and the recent studies remedying these concerns in Section 1. Then we gather preliminary background on Bayesian statistics in Section 2. We explain the discrete choice models and corresponding Bayesian treatments defining the methodology in Section 3, and present the simulations and the results in Section 4. Finally, we give the final remarks and possible future work in Section 5.

2. PRELIMINARIES

2.1. Bayesian Statistics

Bayesian statistics is the mathematical procedure which dictates updating the probability density function of an event with the prior belief on this event.

2.1.1. Bayes Theorem

Bayes' theorem constitutes the basis of Bayesian statistics, which formulates the *conditional probability* of an event when given some observed data based on some prior belief on data. The theorem estimates the parameters θ, which are random variables of the underlying distribution generating the observed data D,

$$\begin{aligned} p(\theta|D) &= \frac{p(D|\theta)p(\theta)}{p(D)} \\ &= \frac{p(D|\theta)p(\theta)}{\int p(D|\theta')p(\theta')\,d\theta'}. \end{aligned} \tag{2.1}$$

Bayes' theorem states that when θ are the parameters of the probability distribution of interest, $P(\theta|D)$ is the *posterior probability* which is the probability of θ when D is taken into account, $P(D)$ is called the *evidence*, $P(D|\theta)$ is the *likelihood probability* and $P(\theta)$ is the *prior* which is the prior belief on the the parameters θ. When $p(\theta|D) = p(\theta)$, θ and D are said to be *independent*.

2.1.2. Prediction

To predict the next data point given some observed data, it is usually more efficient to find a single adequate θ rather than computing the entire probability distribution $p(\theta|D)$. This point estimate of θ serves as the best estimate for the point estimator. There are two methods for point estimation.

2.1.2.1. Maximum likelihood (ML) estimation. ML estimates refer to the θ which maximizes the likelihood function, *i.e.*, under which the observed data are most likely to occur. It is common to use log-likelihood functions for the sake of ease in the computation,

$$\begin{aligned} \hat{\theta}_{ML} &= \operatorname*{argmax}_{\theta} p(D|\theta) \\ &= \operatorname*{argmax}_{\theta} \log p(D|\theta). \end{aligned} \tag{2.2}$$

2.1.2.2. Maximum a posteriori (MAP) estimation. MAP estimates refer to the "most likely" θ when given the observed data,

$$\begin{aligned} \hat{\theta}_{MAP} &= \operatorname*{argmax}_{\theta} p(\theta|D) \\ &= \operatorname*{argmax}_{\theta} \frac{p(D|\theta)p(\theta)}{\int p(D|\theta')p(\theta')\,d\theta'} \\ &= \operatorname*{argmax}_{\theta} p(D|\theta)p(\theta) \\ &= \operatorname*{argmax}_{\theta} (\log p(D|\theta) + \log p(\theta)). \end{aligned} \tag{2.3}$$

When the prior distribution is uniform, the ML estimate gives the same result as the MAP estimate, since the $\log p(\theta)$ term will be a constant.

2.1.3. Posterior Predictive Distribution and Conjugate Priors

Bayesian statistics is the tool to update one's belief on a random variable, in the presence of experimental data. Point estimates are usually useful for this purpose, however, they ignore the uncertainty about θ. To better explain the extreme values of unobserved data points this uncertainty should be taken into account, namely, the full posterior $p(\theta|D)$ should be obtained to compute the *posterior predictive distribution.*

When the dimensionality of D is m, the next data point D_{m+1} is conditionally independent of D, given θ. The posterior predictive distribution denoted as $p(D_{m+1}|D)$

can be formulated as follows:

$$\begin{aligned} p(D_{m+1}|D) &= \int p(D_{m+1}|D,\theta)p(\theta|D)\,d\theta \\ &= \int p(D_{m+1}|\theta)p(\theta|D)\,d\theta. \end{aligned} \tag{2.4}$$

Recall that the posterior distribution has an integral in the denominator:

$$p(\theta|D) = \frac{p(D|\theta)p(\theta)}{\int p(D|\theta')p(\theta')\,d\theta'}. \tag{2.5}$$

The denominator in Equation (2.5) is the normalizing constant, which needs to be computed to find the exact posterior distribution. The computation of the posterior probability of a random variable is very expensive, such that we need to compute the normalizing constant *i.e.*, the evidence for each and every variable. However, numerical integration of the evidence is not required, when the prior distribution is *conjugate prior* to the likelihood probability. Conjugate prior is the probability distribution in accordance with the likelihood function, which gives a posterior probability distribution in the same probability distribution family of its own. This property removes the necessity to compute the normalizing constant, and the closed-form posterior distribution can be directly formed with only parameter updates of the prior distribution.

Let's consider a binomial likelihood function and examine the Beta distribution being its conjugate prior distribution to the binomial likelihood function. A binomial distribution is the discrete probability distribution representing the success probability of N independent Bernoulli experiments with a success rate of θ for each trial.

A success/failure event with success probability θ is said to be a Bernoulli event. Therefore, Binomial distribution has two parameters: N and θ

$$x \sim Bin(N,\theta). \tag{2.6}$$

The probability that k successes occur after N Bernoulli trials can be written as:

$$P(x = k) = \binom{N}{k} \theta^k (1 - \theta)^{N-k}. \tag{2.7}$$

The Beta distribution is a continuous probability distribution of a continuous random variable θ defined on the interval $[0, 1]$ with two parameters α and β,

$$\theta \sim Beta(\alpha, \beta) \tag{2.8}$$

$$p(\theta) = \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha, \beta)} \tag{2.9}$$

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)} \tag{2.10}$$

where Γ is the Gamma function.

<u>2.1.3.1. Intuition for the Beta distribution.</u> The Beta distribution is a probability distribution, the random variable of which is the probability of success for a success-failure event. Therefore, its support is the interval $[0, 1]$ $(0 \leq \theta \leq 1)$. The α and β parameters define the shape of the Beta distribution curve. These parameters can be considered as the counts of successes and failures, respectively. If α is much larger than β, the curve shifts to right, meaning that it is more likely to get a higher probability of success than of failure. Similarly, when β is larger, the more mass of the probability favors the failure and the curve shifts left. When α and β parameters are equal, the Beta distribution curve gains a bell-shape equally distributing the probability bulk to both success and failure probabilities. As both α and β simultaneously increase more and more, *i.e.*, we have more data showing the success and failure counts and they turned out to be equal, the bell-shape of the curve gets narrower. In this way, one can be more and more certain on success and failure probabilities being equal. If both α and β are equal to 1, the curve becomes a straight line and one has a uniform prior on the success probability.

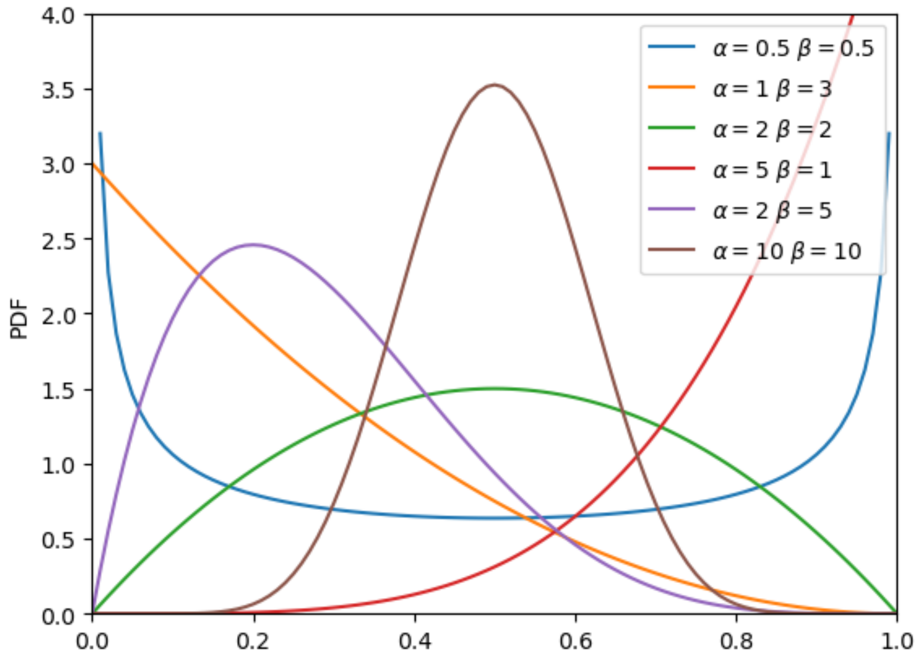

Figure 2.1. Probability density function of the Beta distribution with different α and β parameters.

The Beta distribution might not be the only choice to model the probability of success of an event. However, the essence of Bayesian inference is to update the prior belief on a probability with the additional data at hand. The Beta distribution is the conjugate prior to the binomial distribution, meaning it provides a closed-form solution for the posterior distribution of the probability of success, therefore is the convenient choice.

Assume we have a binomial likelihood function of the form in Equation (2.7) and the conjugate prior distribution of the form in Equation (2.9). Then the posterior distribution $p(\theta|x)$ can be directly computed as:

$$\begin{aligned} p(\theta|x) &= \frac{\binom{N}{k}\theta^k(1-\theta)^{N-k}\frac{1}{B(\alpha,\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}}{\int_0^1 \binom{N}{k}\theta^k(1-\theta)^{N-k}\frac{1}{B(\alpha,\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}\,d\theta} \\ &= \frac{\theta^{k+\alpha-1}(1-\theta)^{N-k+\beta-1}}{\int_0^1 \theta^{k+\alpha-1}(1-\theta)^{N-k+\beta-1}\,d\theta}. \end{aligned} \tag{2.11}$$

The denominator in Equation (2.11) is the normalizing constant $B(k+\alpha, N-k+\beta)$ for the nominator and the resulting probability distribution is again a Beta distribution $Beta(k+\alpha, N-k+\beta)$.

2.2. Dirichlet Distribution

A Dirichlet distribution is a generalized Beta distribution parameterized by a vector $\boldsymbol{\alpha}$ which is often denoted as $Dir(\boldsymbol{\alpha})$. Each elements of $\boldsymbol{\alpha}$ is a real value greater than 0. Suppose

$$\theta \sim Dir(\boldsymbol{\alpha}), \tag{2.12}$$

then the probability density function of Dirichlet distribution is

$$p(\theta) = \frac{1}{B(\boldsymbol{\alpha})} \prod_{k=1}^{K} \theta_k^{\alpha_k - 1} \tag{2.13}$$

where $B(\boldsymbol{\alpha})$ denotes the multivariate Beta function of the form

$$B(\boldsymbol{\alpha}) = \frac{\prod_{k=1}^{K} \Gamma(\alpha_k)}{\Gamma(\sum_{k=1}^{K} \alpha_k)}, \tag{2.14}$$

where $\boldsymbol{\alpha} = \{\alpha_1, \alpha_2, \ldots, \alpha_K\}$.

2.2.1. Intuition for the Dirichlet distribution

Dirichlet distribution is the conjugate prior to the categorical and multinomial distribution, *i.e.*, it is a probability distribution over probabilities of a K-outcome event. The hyperparameters $\alpha_1, \alpha_2 \ldots \alpha_K$ of a Dirichlet distribution define how likely a probability distribution over these K outcomes is. Dirichlet distribution is the generalization of the Beta distribution. The Beta distribution is defined over real numbers in the interval $[0,1]$, whereas the Dirichlet distribution is defined over (K-1)-probability simplex. Each sample drawn from this distribution provides a probability distribution over K outcomes. Therefore, we can think of Dirichlet parameters as the pseudo-counts of each outcome.

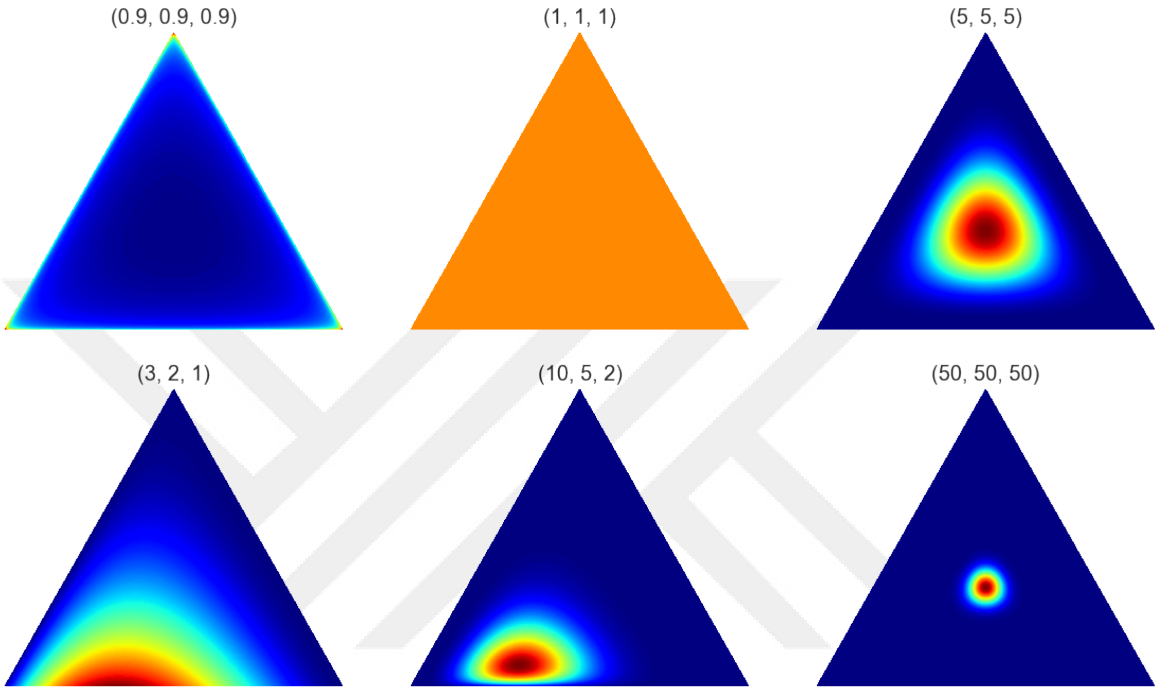

Figure 2.2. Dirichlet distribution with different concentration parameter α on a 2-simplex

Suppose $\boldsymbol{\theta} \sim Dir(\boldsymbol{\alpha})$, expected value of the Dirichlet distribution is

$$\mathbb{E}[\theta_i] = \frac{\alpha_i}{\sum_{k=1}^{K} \alpha_k}, \tag{2.15}$$

$\sum_{k=1}^{K} \alpha_k$ being the normalization constant. Samples of the K-outcome event concentrate around the expected value and the normalization constant determines how densely the samples are located in the simplex. Growing normalization constant results in a more dense concentration of the samples over the simplex. The α parameters favor the probability of their corresponding outcome, therefore the concentration shifts to the highest parameter over the simplex.

2.2.2. Dirichlet-Multinomial model

Dirichlet distribution is the conjugate prior to categorical and multinomial distribution, which are the multivariate Bernoulli and multivariate binomial, respectively. For instance, a coin flip is a Bernoulli experiment, whereas a categorical distribution models a dice roll. In the same direction, N coin flip can be modeled with a binomial distribution where N dice roll will be modeled with a multinomial distribution.

In the Beta-binomial model, the prior probability θ of success rate in each Bernoulli event is parameterized with a Beta distribution. Multinomial distribution, on the other hand, is the distribution over multiple choices, such as dice roll. The conjugate prior Dirichlet distribution models the probabilities over multiple choices which are mutually exclusive and is parameterized with the so-called concentration parameter, which weights each outcome of the event. Assume an event with K outcomes repeated N times,

$$x \sim Mult(N, \theta). \tag{2.16}$$

Therefore the likelihood distribution is

$$\begin{aligned} p(x|\theta, N) &= \frac{N!}{\prod_{k=1}^{K} n_k!} \prod_{k=1}^{K} \theta_k^{n_k} \\ &= \frac{\Gamma(N+1)}{\prod_{k=1}^{K} \Gamma(n_k+1)} \prod_{k=1}^{K} \theta_k^{n_k} \end{aligned} \tag{2.17}$$

n_k's being the count values of each outcome in N trials. Knowing that the Dirichlet distribution is the conjugate prior to the multinomial distribution, the posterior distribution $p(\theta|x)$ will form a Dirichlet distribution, as well. For that reason, we can manipulate the equation without computing the evidence:

$$
\begin{aligned}
p(\theta|x) &\propto p(x|\theta)P(\theta|\boldsymbol{\alpha}) && (2.18)\\
&= \prod_{k=1}^{K} \theta_k^{n_k} \prod_{k=1}^{K} \theta_k^{\alpha_k - 1} \\
&= \prod_{k=1}^{K} \theta_k^{n_k + \alpha_k - 1}. && (2.19)
\end{aligned}
$$

We can directly infer that, $p(\theta|x) = Dir(\boldsymbol{\alpha'})$ where $\alpha'_k = \alpha_k + n_k$.

2.2.3. Dirichlet-Multinomial Distribution

Assume we draw a probability distribution from $Dir(\boldsymbol{\alpha})$ each time before we draw a sample from a categorical distribution. After N such draws, we obtain a sequence $x = [x_1, x_2, \ldots, x_N]$, n_k is the count values for each outcome ($n_k = \sum_i^N \mathbb{1}[x_i = k]$),

$$\theta \sim Dir(\boldsymbol{\alpha}) \tag{2.20}$$

$$x \sim Mult(1, \theta). \tag{2.21}$$

Dirichlet-Categorical distribution (Equation (2.21)) is the compound distribution that represents the marginal distribution over observations provided by the above process. After N trial ($\sum_k n_k = N$) the probability of a particular sequence x is

$$p(x|\theta) = \prod_{k=1}^{K} \theta_k^{n_k}. \tag{2.22}$$

Unlike the multinomial distribution (Equation (2.17)), there is no normalization constant here. This is because multinomial distribution assigns a probability to count values of each outcome in N trials of a K-outcome event, whereas Equation (2.22) represents the probability of N categorical distribution, namely, a particular sequence of outcomes in N trial event.

The marginal distribution is obtained by integrating out θ from the joint distribution $p(x, \theta|\boldsymbol{\alpha}) = p(x|\theta)p(\theta|\alpha)$. From Equation (2.13) and (2.22) the marginal distribution of x is

$$\begin{aligned} p(x|\alpha) &= \int_\theta p(x|\theta)p(\theta|\alpha)\, d\theta \\ &= \int_\theta \frac{1}{B(\boldsymbol{\alpha})} \prod_{k=1}^{K} \theta^{a_k+n_k-1}\, d\theta \\ &= \frac{1}{B(\boldsymbol{\alpha})} \int_\theta \prod_{k=1}^{K} \theta^{a_k+n_k-1}\, d\theta. \end{aligned} \tag{2.23}$$

Provided that $p(x)$ is a probability distribution over x, $\int_x p(x)\, dx = 1$, therefore,

$$\frac{1}{B(\boldsymbol{\alpha})} \int_\theta \prod_k \theta_k^{a_k-1}\, d\theta = 1 \tag{2.24}$$

$$\begin{aligned} \int_\theta \prod_k \theta_k^{a_k-1}\, d\theta &= B(\boldsymbol{\alpha}) \\ &= \frac{\prod_k \Gamma(a_k)}{\Gamma(\sum_k a_k)}. \end{aligned} \tag{2.25}$$

Using this equation, we can rewrite the the marginal distribution in Equation (2.23) such that,

$$\begin{aligned} p(x|\alpha) &= \frac{\Gamma(\sum_k a_k)}{\prod_k \Gamma(a_k)} \frac{\prod_k \Gamma(a_k + n_k)}{\Gamma(\sum_k a_k + n_k)} \\ &= \frac{\Gamma(\sum_k a_k)}{\prod_k \Gamma(a_k)} \frac{\prod_k \Gamma(a_k + n(x_i = k))}{\Gamma(\sum_k a_k + n(x_i = k))}. \end{aligned} \tag{2.26}$$

Another scenario is that

$$\theta \sim Dir(\alpha) \tag{2.27}$$

$$(n_1, n_2, \ldots, n_k) \sim Mult(N, \theta), \tag{2.28}$$

i.e., the samples are drawn from a multinomial distribution parameterized by θ, which is drawn from a Dirichlet distribution, and the total number of draws N. Here the order of the observation sequence itself is not important, and we only care about the count numbers of each outcome. In this case, the total number of possible ordering of observations given a count value vector $n = [n_1, n_2, \ldots, n_K]$ is $\frac{(\sum_k n_k)!}{\prod_k n_k!}$ which is equivalent to $\frac{\Gamma(N+1)}{\prod_k \Gamma(n_k+1)}$.

We can easily modify the Dirichlet-Categorical distribution in Equation (2.26) with this value to obtain the Dirichlet-Multinomial distribution,

$$p(n_1, n_k, \ldots, n_K|\alpha) = \frac{\Gamma(N+1)}{\prod_k \Gamma(n_k+1)} \frac{\Gamma(\sum_k a_k)}{\prod_k \Gamma(a_k)} \frac{\prod_k \Gamma(a_k + n_k)}{\Gamma(\sum_k a_k + n_k)}. \tag{2.29}$$

This distribution is also called Polya distribution, named after George Pólya. Urn model or Polya urn model describes an urn full of balls with K different color, each time we draw a ball from the urn, we return this ball and an additional ball of the same color with the observed ball. The probability over ball counts for each color is a Dirichlet-Multinomial distribution. If there are α_i number of balls in the urn for i'th color and after N draws, the probability of count values of each color is a Dirichlet-Multinomial distribution parameterized by N and α (see Equation (2.29)).

2.2.4. Generalized Dirichlet Distribution

The parameters of Dirichlet distribution share the same variance and also are negatively correlated, *i.e.*, parameters corresponding to entries in the random vector should decrease or remain the same in case another entry increases. Generalized Dirichlet distributions are introduced for coping with the shortcomings of Dirichlet distribution. Instead of giving one degree of freedom, which is the total sample size, generalized Dirichlet distributions enable one to sample from independent Beta distributions for each proportion of random vector. The distribution defined in [29] is

$$p(\theta|\alpha, \beta) = \prod_{i=1}^{K} \frac{\theta_i^{a_i - 1}(1 - \theta_1 - \theta_2 - \cdots - \theta_i)^{\gamma_i}}{B(\alpha_i, \beta_i)} \tag{2.30}$$

where $\gamma_i = \beta_i - \alpha_{i+1} - \beta_{i+1}$ for $1 \leq i \leq K-1$ and $\gamma_K = \beta_K - 1$.

3. METHODOLOGY

The recommendation framework is assumed to be a discrete choice problem in this work. A discrete choice model dictates that the choices are not deterministic, meaning that the utility of choosing item i over j is not a deterministic value, rather it is probabilistic. One's tendency to prefer the former to the latter might depend on her/his mood, the circumstances, etc.

When the pool of alternatives is too large, the decision-maker is supposedly exposed to a subset of the entire set of options. The very same mechanism applies to any recommendation engine. The decision-maker/user would make a selection among the alternatives that he/she is exposed to, which in turn would affect the user preference elicitation of the recommender system. The presentation mechanism, namely, the decision process of what to include in the subset of alternatives plays a crucial role in preference estimation, such that a badly-designed one may bring bias into the process. For instance, a movie recommender would present a set of movies of different genres to the user, the user then would select an action movie. However, when a set of action movies were presented, the user would select a movie other than the one she chose before. To be able to infer the preference ranking, the system should not only account for the presentations but also build its presentation mechanism, wisely.

In this thesis, we dictate and ground the desirable features of a fair recommender system as *free of bias* (*i.e.*, can cope with any kind of biases rooted from the limited and systematic presentations), *novel* (*i.e.*, capable of discovering the censored options and providing the user with his/her favorites even he/she is not aware of), *responsible* (*i.e.*, having a mechanism to eliminate any kind of misinterpretation on the actual preferences due to over-/under-presentation of some alternatives), *representative* (*i.e.*, righteously presenting all alternatives without violating the right of being equally presented) [27].

The task of meeting these fairness criteria that we establish has an *inference task* for the preference estimations and a *discovery task* to establish a presentation mechanism.

In this section, we introduce the discrete choice models in general, then we give the details regarding the Dirichlet-Luce model together with a presentation mechanism which we claim to be a fair alternative for recommendation engines, explaining the corresponding inference and discovery task.

3.1. Discrete Choice Models

Discrete choice models are probabilistic models based on the random utility maximization for social choice which postulates the human choice behavior given a set of options [30]. They are probabilistic based on the hypothesis that human choices may differ depending on the circumstances. In this approach of choice behavior, there is a decision-maker, a set of options, and the features of those options which impose distinctive effects on the decision-maker's choice.

Discrete choice modeling has received attention in recommender system literature [31–33]. It claims a recommendation system where independence of irrelevant alternative holds, *i.e.*, given a universal set of alternatives, relative preferability of an option does not depend on the presence or absence of another. In this scenario, the user of the recommendation engine is the decision-maker who makes choices among the set of all possible alternatives that the recommendation engine can present and the features of the options are the latent factors of the model.

The choices of a decision-maker among a subset of all alternatives can be represented by a restricted multinomial for both pairwise [34] and for L-wise [35, 36] preferences. In this chapter we present Luce's choice axioms [35] and Plackett-Luce [36] model which constitutes a basis to the investigated Dirichlet-Luce model.

3.1.1. Luce's Choice Axioms

Duncan Luce states that the probability of choosing an item over another from a pool of options is independent of the presence or absence of other options in the pool [35].

3.1.1.1. Axiom. Let T be a finite subset of U such that, for every $S \subset T, p_s$ is defined. And $p(x, y)$ is used in replacement of $p_{\{x,y\}}(x)$ when $x \neq y$.

(i) *If* $p(x, y) \neq 0, 1$ *for all* $x, y \in T$, *then for* $R \subset S \subset T$

$$p_T(R) = p_S(R) p_T(S). \tag{3.1}$$

(ii) *If* $p(x, y) = 0$ *for some* $x, y \in T$, *then for every* $S \subset T$

$$p_T(S) = p_{T-\{x\}}(S - \{x\}). \tag{3.2}$$

3.1.1.2. Interpretation. The first part of the axiom can be interpreted as a flow of choices: The set T of all options is categorized into subsets according to some criterion. There is a probability assigned to each subset of options, that is $p_T(S_i)$, S_i being i'th subset of T. Then S_i's are further sub-categorized. The same procedure is applied on each subsequent set until the choice is made.

The second part of the axioms states that if y is always chosen over x from a set of options T, then removing x from T does not change the pairwise probabilities of other options, namely, *if* $p(x, y) = 0$ *for some* $y \in T$ *then* $p_T(x) = 0$. For instance, one might prefer to eat meat over vegetables every time. When included chicken in this option, one might argue whether this person chooses chicken or meat, but definitely not the vegetables. In other words, the vegetable won't be chosen as long as the meat is present in the option set, whether or not there is another option.

3.1.1.3. Independence of irrelevant alternatives. From Luce's choice axiom, one other interpretation is the independence of irrelevant alternatives. Given two options x and y, the relative probability of x being chosen over y is independent of the presence or absence of other options. From axiom i, assuming R is the set of containing a single option x, S is the set of x and y and T is any subset of all options (and $p(x, y)$ is the replacement for $p_{\{x,y\}}(x)$),

$$
\begin{aligned}
p_T(x) &= p_{\{x,y\}}(x) p_T(\{x, y\}) && (3.3)\\
&= p(x, y)[p_T(x) + p_T(y)] \\
p_T(x)[1 - p(x, y)] &= p(x, y) p_T(y) && (3.4)\\
p_T(x) p(y, x) &= p(x, y) p_T(y). && (3.5)
\end{aligned}
$$

By the axiom i, $p(x, y)$ and $p(x, y)$ is different than 0 which leads to,

$$\frac{p(x, y)}{p(y, x)} = \frac{p_T(x)}{p_T(y)}. \tag{3.6}$$

This result indicates that an option being preferred over another is independent of the presence or absence of other options. This property is indeed applicable in many real-time choice scenarios, however, there are also cases this may not present the reality.

3.1.2. Plackett-Luce Model

Plackett explained the same phenomenon with Luce's in his own work [36]. Assume we want to find the probability of particular permutation, *i.e.*, rankings of choices among a set of K alternatives given probability of each alternative ($p_1, p_2, \ldots, p_K$ where $\sum_i p_i = 1$) to be chosen when the entire set of options are presented.

Let's assume the first option has the highest ranking. Then the probability of an alternative having the second ranking among other remaining options are $p_2/(1 - p_1), p_3/(1 - p_1), \ldots, p_K/(1 - p_1)$. After the selection of the alternative having the second ranking, let's say alternative i was chosen, then the resulting probability of an alternative to be chosen is $p_2/(1-p_1-p_i), p_3/(1-p_1-p_i), \ldots, p_{i-1}/(1-p_1-p_i), p_{i+1}/(1-p_1-p_i), \ldots, p_K/(1-p_1-p_i)$. Let's assume r represents the $1, 2, 3, \ldots, K$ ordering of alternatives, then the probability of this ranking $p(r)$ is

$$p(r) = p_1 \frac{p_2}{(1-p_1)} \frac{p_3}{(1-p_1-p_2)} \cdots \frac{p_{N-1}}{1-(\sum_{i=1}^{K-2} p_i)}. \tag{3.7}$$

It is trivial that the last option is automatically ranked after the selection of the alternative before the last one, therefore has the probability of 1.

Let's assume that we want to find the probability of rankings of items given a subset of K options, given p_i's representing the probability of an alternative being chosen among K options ($p_i = p(k_i|K)$). Suppose there are T non-empty subsets of those alternatives $\{C_1, C_2, \ldots, C_T\}$ meaning each alternative existing in a particular subset belongs to the universal set of K options. The generalization of an alternative k_i to be chosen given a subset C_t provided that $k_i \in C_t$ is

$$p(k_i|C_t) = \frac{p_k}{\sum_{\kappa \in C_t} p_\kappa}. \tag{3.8}$$

In this choice modeling, given two items i and j, and two subsets C and C^* of N alternatives,

$$\frac{p(i|C)}{p(j|C)} = \frac{p(i|C^*)}{p(j|C^*)}. \tag{3.9}$$

indicating that the probability of choosing an option i over j is independent of presence or absence of other alternatives. For the pairwise comparisons (where $n(C_t) = 2$) the model is Bradley-Terry model [34] and for L-wise comparisons (where $n(C_t) = L, L > 2$) the model is a Plackett-Luce model [36].

3.2. Dirichlet-Luce Model with a Presentation Mechanism

Suppose among K options, the user is presented with T different presentations and makes the total of T choices and k_t is the selection from the t'th presentation C_t. Also suppose that θ_k is the probability of choosing an alternative when all other alternatives are presented ($p(k|K)$). Assuming a Plackett-Luce model and provided a particular presentation $C \in \mathcal{C}$, $\mathcal{C}$ being the set of all non-empty subsets of K alternatives, the probability that an item is chosen over others is

$$p(k|C) = \frac{\theta_k}{\sum_{\kappa \in C} \theta_\kappa}. \tag{3.10}$$

Conditioned on these T presentations $C_{1:T}$ and the underlying preferences $\theta = \theta_{1:K}$ where θ is a $(K-1)$ probability simplex, the likelihood function of choosing an option set $k_{1:T}$ is

$$p(k_{1:T}|\theta, C_{1:T}) = \frac{\prod_k {\theta_k}^{y_k}}{\prod_{C \in \mathcal{C}} (\sum_{\kappa \in C} \theta_\kappa)^{\mu(C)}} \tag{3.11}$$

where $\mu(C)$ and y_k are

$$\mu(C) = \sum_1^T [C_t = C] \tag{3.12}$$

$$y_k = \sum_1^T [k_t = k], \tag{3.13}$$

$\mu(C)$ is the number of times the subset C is presented and y_k is the number of times the option k is chosen. This is a restricted multinomial likelihood function, due to the restriction in the presentation (subset of all alternatives).

Recall that Dirichlet distribution $\theta \sim Dir(\alpha)$ is the conjugate prior to multinomial likelihood function (see Section (2.2.2)). In fact, the prior distribution to the probability mass function in Equation (3.11) can be formulated as a generalized Dirichlet distribution such that

$$p(\theta|\alpha,\beta) \propto \prod_k {\theta_k}^{\alpha_k - 1} \prod_{C \in \mathcal{C}} \left(\sum_{\kappa \in C} \theta_\kappa \right)^{-\beta(C)}. \tag{3.14}$$

Here $\beta(C)$ and a_k are the pseudo-counts of presentations and choices, respectively. Therefore, $\sum_{\mathcal{C}} \beta(C) = \sum_k a_k$. When $C = [K]$, *i.e.*, all of the alternatives are present to the decision maker, then the prior distribution reduces to Dirichlet distribution.

With these specifications, the posterior distribution is proportional to

$$p(\theta|k_{1:T}, C_{1:T}, \alpha, \beta) \propto \prod_k \theta^{\alpha_k + y_k - 1} \prod_{C \in \mathcal{C}} \left(\sum_{\kappa \in C} \theta_\kappa \right)^{-\mu(C) - \beta(C)}. \tag{3.15}$$

3.3. Preference Learning

An unbiased preference estimator from user feedback for choice model setup is studied by [37] for pairwise rankings and by [38] and [28] for L-wise rankings. In this work, we follow the procedure described in [28] for the inference task. The predicted preference probability of an option k given presentation C_{T+1} is equal to the expectation of preference ratios under the posterior:

$$p(k_{T+1} = k \mid C_{T+1}, k_{1:T}, C_{1:T}, \alpha, \beta) = \mathbb{E}_{p(\theta|k_{1:T}, C_{1:T})} \left[\frac{\theta_k}{\sum_{\kappa \in C_{T+1}} \theta_\kappa} \right]. \tag{3.16}$$

From the posterior potential of the Dirichlet-Luce model (see the posterior proportionality in Equation (3.15)) the log potential $\phi(\theta)$ is defined as:

$$\phi(\theta) = \log \left[\prod_k \theta_k^{y_k + \alpha_k - 1} \prod_{C \in \mathcal{C}} \left(\sum_{\kappa \in C} \theta_\kappa \right)^{-\beta(C) - \mu(C)} \right]. \tag{3.17}$$

3.4. Presentation Mechanism

Dirichlet-Luce model provides a fair preference estimation for a recommender system, such that it does not punish under-presented, yet has not been properly ranked alternatives. However, the interactive system should present those under-presented alternatives eventually to make a correct preference estimation over all items. This calls for a fair presentation mechanism, such that containing a pool of both favorite and also unexplored alternatives.

As [39] explained, self-reinforcing feedback loops can be eliminated by accounting for novelty and serendipity, namely discovering the censored favorites of the users. Combining the exploration of undiscovered items that are yet not presented to the user and exploitation of the similar items to already favorable ones is framed as "bandit problems" [40]. Bandit algorithms are used when each choice's properties are only partially known at the time of allocation and may become better understood as time passes or by allocating resources to the choice, which represents the nature of the user-system interaction in a recommendation setup. In [41] bandit algorithms are presented as the remedy for self-reinforcing feedback loops rooted in the partial user feedback.

The posterior of the Dirichlet-Luce model puts a high probability on the items if they are presented and chosen frequently and also if they are presented less frequently. Using this nature of the posterior, we present top L options which have the highest probability of being chosen to the user, each time after the user chooses among what is presented and the preference estimation is updated. This sampling procedure is an instance of "Thompson sampling" [42] which is proven to be effective for bandit problems and widely explored [43–45]. It enables the system to converge rapidly to the true preference estimates. For the simulations in Section 4, the Thompson sampling procedure described in [28] is used.

4. SIMULATIONS

In this section different simulations are designed to illustrate the behavior of personalized interactive systems under these designed conditions. We demonstrate that explicitly conditioning on systematic exposure guarantees robustness to promotion and unfair comparisons, and discovery of initially censored favorites.

To see the effect of conditioning on the presentations we make a comparison between the Dirichlet-Luce model, a Bayesian choice model which infers preferences based on previous choices from systematic presentations and the naive Dirichlet-Multinomial model which ignores the systematic presentation. For both models, we consider a Thompson sampling-based presentation mechanism.

Contrary to the Dirichlet-Multinomial model, the Dirichlet-Luce model accounts for systematic exposure, *i.e.*, the preference elicitation process is conditioned on what is previously presented to the user. For online learning to present Thompson sampling [42] is used as the bandit algorithm, since it provides a presentation mechanism based on the posterior distribution of preference estimates, thus enables under-presented options to be presented, unlike a greedy algorithm would present only top-L (of K) options, censoring the under-presented options.

These simulations address several biases explained in Section 1.2. Problem setup is as follows: $\theta_k = p(k|[K]) \; \forall k \in [K]$ represents the latent factor for preference k being chosen, *i.e.*, θ_k is the choice probability for item k. Over T presentations $\{C_t|t = 1, \dots, T\}$ and user feedback $k_{1:T}$ to these presentations where $k_t \in C_t$, the system makes a preference elicitation for the particular user. For each case, the simulated user is exposed to 2 items among a total of 5. The choices are made based on the true latent θ^*, namely, the choice probability of item k is $p(k|C) = \frac{\theta^*_k}{\sum_{\kappa \in C} \theta^*_\kappa}$. Pairwise comparisons are then fed to the system ($L = 2$ and $K = 5$). Without loss of generality, it can be assumed that the choice probabilities of items are $\theta^*_1 > \theta^*_2 > \cdots > \theta^*_K$.

Here we should clarify several assumptions regarding the simulations. We conceive a single user-system interaction scenario. However, the setup can be reused in a collaborative recommender system. We ignore the position bias, *i.e.*, the ranking of the alternatives in the presentation is not effective in the choice decision, under the assumption that the user can review the presented alternatives and make a choice based on the latent preferences. We disregard that the repeated and systematic presentation might really alter users' preferences, leading to an echo chamber. We allude to [20] for such an examination, where adequate conditions that lead to interest extremes are given. This is because the Dirichlet-Luce model with the associated presentation mechanism matches at least two essential conditions to maintain a strategic distance from such degeneracy focuses, by normally permitting a "growing pool of alternatives" and randomization inherent in Thompson sampling. We accept that the user picks one of the displayed choices, even though the model can be expanded to incorporate a manufactured 'browse' choice corresponding to preferring not to select. We ground this assumption on the "principle of least effort", namely, the decision-maker makes a selection among the presented options if at least the minimally acceptable option is available.

It should be noted that the emphasis is not on the uniqueness of the model in eliminating feedback loops (see other models [7, 22, 46]) Nevertheless, the Dirichlet-Luce model, when compared to the Dirichlet-Multinomial, is a considerable test bed to assess the effect of eliminating feedback loops in fixing some of the biases prevalent in personalized recommender systems.

The following scenarios represent potentially biased setups that might result in under- or over-estimation of choice probabilities of some items. For each setup, the goal is to observe whether a model which accounts for the independence of irrelevant alternatives coupled with conditioning on the limited exposure to alternatives can cope with the under-/over-presentation phenomenon.

4.1. Robustness to Promotion

One of the main causes of biases in recommender systems is the excessive advertisement of some popular or promoted options. We believe that a fair system would not overestimate the preference of the user towards the promoted options and not underestimate the choice probabilities of other ones that are yet not presented, as well. As mentioned in Section 1, the promotion of popular items is the main cause of the observation bias, which makes the system favor the over-presented items and discredit the less presented ones. The consequences of observation bias are two-fold: biased user preference elicitation which blocks the user from acquiring the most appealing options and, violating the content providers' right to be equally presented.

In this scenario, we assume that some option is promoted to the user such that it exists in each presentation the user encounters. For this purpose, we artificially add a popular option to each presentation that the system would display. We claim that a fair recommender system would correct the over-estimation of the popular item by conditioning on the presentations in the preference elicitation process. In other words, the system is expected not to favor the popular items such that the estimations on preferences converge to the actual user preferences. As illustrated in Figure 4.1, conditioning on presentations as the Dirichlet-Luce model does can cope with the promotion of popular items.

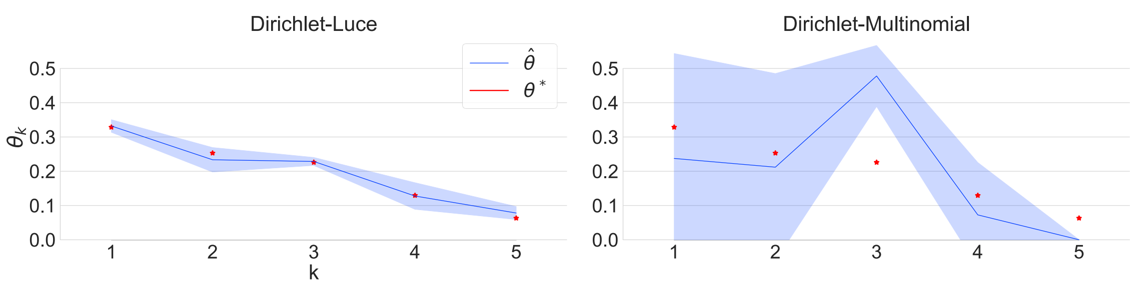

Figure 4.1. Final (after $T = 10000$ interactions) preference estimates ($\hat{\theta}$) averaged over 10 runs when the option 3 is *promoted*, *i.e.*, it is included in every presentation. Over-presented option 3 is over-estimated (right), whereas conditioning on presentations fixes this bias (left). Shaded regions denote the standard deviation.

4.2. Discovery of Censored Favorites

Another problem for the recommender systems is the cold-start problem. Prior to any user-system interaction, the system has no information about a specific user's preferences. The easiest way is to present the user with the most popular options to keep the user online on the platform. However, the initial preferences of the user might not reveal his/her actual preferences, meaning that there might be less popular options that are indeed more appealing to the user.

Discovering the initially censored, maybe less popular items is a challenging task. We claim that a fair recommender system would correct the biases towards initial choices over the course of user-system interactions by giving an equal opportunity to be presented to each option and eventually discover the censored favorites of the user.

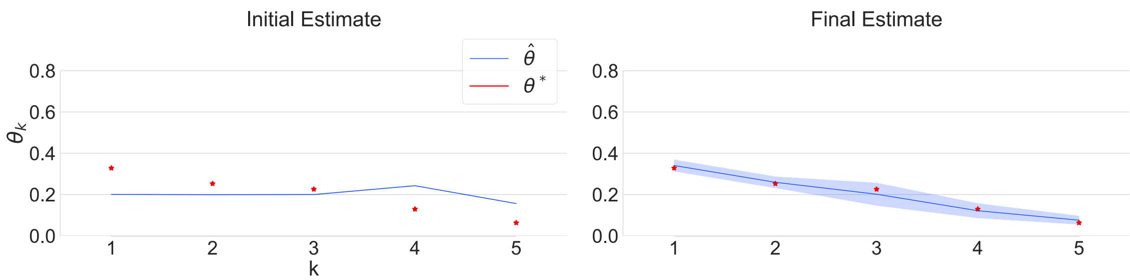

Figure 4.2. Dirichlet-Luce model. Preference estimates after 100 choices are made from originally inferior options $\{4, 5\}$ (left), and final preference estimates after $T = 10000$ interactions (right). Conditioning on presentations does not impose negative bias towards censored options, and the implied personalization model eventually learns to present the best two options.

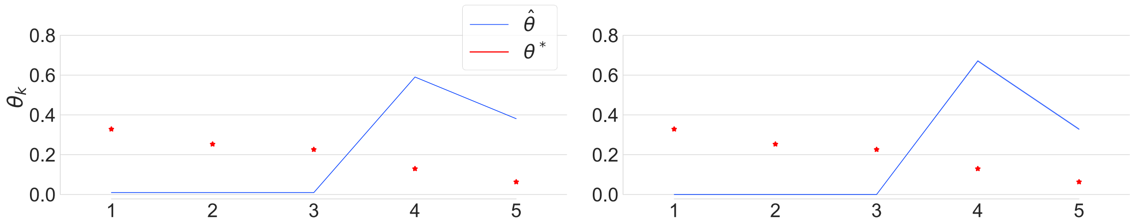

Figure 4.3. Dirichlet-Multinomial model. Preference estimates after 100 choices are made from originally inferior options $\{4, 5\}$ (left), and final preference estimates after $T = 10000$ interactions (right). Over-presented options are overestimated.

In this scenario, the user is presented with initially inferior options over 100 runs and the system is expected to discover the superior preferences of the user. We should note that exploring the favorite options is a computationally costly process. Conditioning on presentations does not impose bias towards the initially censored options, on the contrary, the Dirichlet-Luce model is capable of exploring those options as illustrated in Figure 4.2, whereas the Dirichlet-Multinomial model is incompetent in eliminating negative bias towards initially censored favorites and under-estimate those items favoring the initial choices of the user.

4.3. Robustness to Unfair Comparison

Other than popularity bias and bias towards initial choices, the recommender systems may develop other biases inherited in the user interactions with the presentations. For instance, repetitively presenting a superior option together with a more superior one might lead the system to consider the formerly described option as an overall inferior one. To overcome this misconception, the system should make a balance between the exploitation of the already preferred options and the exploration of not yet presented options. Otherwise, an exploiting system would present the preferred options and lack in capturing the actual pairwise rankings of the options.

In this scenario, the user chooses among two superior alternatives $\{1, 2\}$ over 100 runs, which might mislead the system such that the system considers option 2 as an overall inferior option. The system is expected the overcome this misconception by correcting the preference estimation of option 2 by providing the user with a clever presentation mechanism.

Recalling Luce's choice axioms, the choice probability of an item over another is independent of the presence or absence of other options. Therefore, in those experiments, it is essential for the estimation of actual preferences that the user is forced to make various comparisons among options to eliminate such a comparison bias.

Again, over the course of interactions, the Dirichlet-Luce model's presentation mechanism corrects the misconception introduced by the initial ill-presentation mechanism.

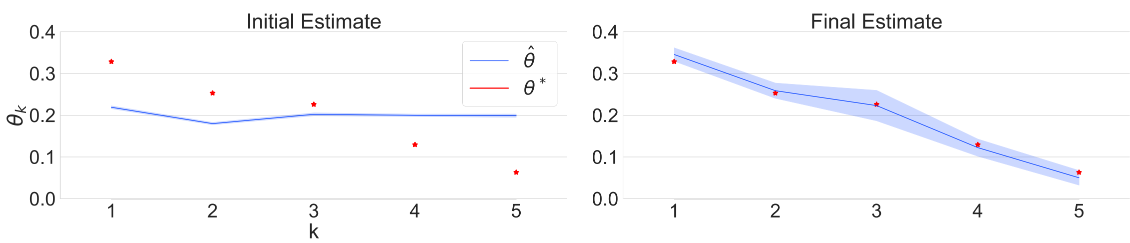

Figure 4.4. When an initial 100 choices are made from the alternatives $\{1, 2\}$, Dirichlet-Luce initially imposes a negative bias towards option 2 (left). Over the course of interactions, the system captures that option 2 is still preferable to other, originally inferior options (right).

5. CONCLUSION

In this thesis, we established fairness criteria for recommendation engines as being free of bias, novel, responsible and representative. Ignoring self-reinforcing feedback loops is *unfair* to both sides (the user and the content provider) of the interactive systems due to overestimation of the over-presented, or underestimation of the under-presented alternatives.

We investigated and show whether a discrete choice model conditioning on presentations, which are indeed cleverly designed, would produce fair personalized recommendations, namely, make unbiased estimations on user preferences.

We prepared simulations imitating the various real-world scenarios which might impose observation bias in recommender systems and demonstrated the difference between *ignoring* and explicitly *including* the systematic presentation to the inference mechanism.

The fact behind the results is parallel to our prior hypothesis, ignoring the systematic and limited exposure to alternatives leads to biased recommendations. This is due to the fact that a personalized recommender system, "learns a mechanism from mechanism induced data [47]." To avoid the inevitable feedback loop, the competent recommender system should account for the presentations themselves.

As future work, these scenarios can be applied to multi-user setups and real-world data to validate the effect of incorporating systematic presentations in preference elicitation for a fair recommendation.

REFERENCES

1. Pariser, E., *The filter bubble: what the internet is hiding from you*, Penguin UK, 2011.

2. Abdollahpouri, H., R. Burke and B. Mobasher, "Managing popularity bias in recommender systems with personalized re-ranking", *The Thirty-Second International Flairs Conference*, 2019.

3. Abdollahpouri, H., "Popularity bias in ranking and recommendation", *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*, p. 529–530, Association for Computing Machinery (ACM), 2019.

4. Abdollahpouri, H., R. Burke and B. Mobasher, "Controlling popularity bias in learning-to-rank recommendation", *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*, p. 42–46, Association for Computing Machinery (ACM), 2017.

5. Mehrotra, R., J. McInerney, H. Bouchard, M. Lalmas and F. Diaz, "Towards a fair marketplace: counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems", *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*, p. 2243–2251, Association for Computing Machinery (ACM), 2018.

6. Farnadi, G., P. Kouki, S. K. Thompson, S. Srinivasan and L. Getoor, "A fairness-aware hybrid recommender system", *arXiv preprint arXiv:1809.09030*, 2018.

7. Sun, W., S. Khenissi, O. Nasraoui and P. Shafto, "Debiasing the human-recommender system feedback loop in collaborative filtering", *Companion Proceedings of the 2019 World Wide Web Conference (WWW '19)*, p. 645–651, Association for Computing Machinery (ACM), 2019.

8. Singh, A. and T. Joachims, “Fairness of exposure in rankings”, *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18)*, p. 2219–2228, Association for Computing Machinery (ACM), 2018.

9. Konstan, J. A. and J. Riedl, “Recommender systems: from algorithms to user experience”, *User Modeling and User-Adapted Interaction*, Vol. 22, No. 1-2, pp. 101–123, 2012.

10. Celma, Ò. and P. Cano, “From hits to niches? or how popular artists can bias music recommendation and discovery”, *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition*, pp. 1–8, 2008.

11. Yao, S. and B. Huang, “Beyond parity: fairness objectives for collaborative filtering”, *Advances in Neural Information Processing Systems*, pp. 2921–2930, 2017.

12. Patro, G. K., A. Chakraborty, N. Ganguly and K. P. Gummadi, “Incremental fairness in two-sided market platforms: on updating recommendations fairly”, *arXiv preprint arXiv:1909.10005*, 2019.

13. Leonhardt, J., A. Anand and M. Khosla, “User fairness in recommender systems”, *Companion Proceedings of the The Web Conference 2018*, p. 101–102, International World Wide Web Conferences Steering Committee, 2018.

14. Anderson, C., *The long tail: why the future of business is selling less of more*, Hachette Books, 2006.

15. Park, Y.-J. and A. Tuzhilin, “The long tail of recommender systems and how to leverage it”, *Proceedings of the 2008 ACM Conference on Recommender Systems (RecSys '08)*, p. 11–18, Association for Computing Machinery (ACM), 2008.

16. Brynjolfsson, E., Y. J. Hu and M. D. Smith, “From niches to riches: anatomy of

the long tail", *Sloan Management Review*, Vol. 47, No. 4, pp. 67–71, 2006.

17. Burke, R. D., H. Abdollahpouri, B. Mobasher and T. Gupta, "Towards multi-stakeholder utility evaluation of recommender systems.", *Extended Proceedings of the 24th ACM Conference on User Modeling, Adaptation and Personalisation (UMAP-ExtProc)*, 2016.

18. Burke, R., "Multisided fairness for recommendation", *arXiv preprint arXiv:1707.00093*, 2017.

19. Bakshy, E., S. Messing and L. A. Adamic, "Exposure to ideologically diverse news and opinion on Facebook", *Science*, Vol. 348, No. 6239, pp. 1130–1132, 2015.

20. Jiang, R., S. Chiappa, T. Lattimore, A. György and P. Kohli, "Degenerate feedback loops in recommender systems", *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*, p. 383–390, Association for Computing Machinery (ACM), 2019.

21. Khenissi, S., B. Mariem and O. Nasraoui, "Theoretical modeling of the iterative properties of user discovery in a collaborative filtering recommender system", *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, p. 348–357, Association for Computing Machinery (ACM), 2020.

22. Sinha, A., D. F. Gleich and K. Ramani, "Deconvolving feedback loops in recommender systems", *Advances in Neural Information Processing Systems*, pp. 3243–3251, 2016.

23. Nguyen, T. T., P.-M. Hui, F. M. Harper, L. Terveen and J. A. Konstan, "Exploring the filter bubble: the effect of using recommender systems on content diversity", *Proceedings of the 23rd International Conference on World Wide Web*, pp. 677–686, 2014.

24. Aridor, G., D. Goncalves and S. Sikdar, "Deconstructing the filter bubble: user

decision-making and recommender systems", *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, p. 82–91, Association for Computing Machinery (ACM), 2020.

25. Schnabel, T. and P. N. Bennett, "Debiasing item-to-item recommendations with small annotated datasets", *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, p. 73–81, 2020.

26. Chaney, A. J. B., B. M. Stewart and B. E. Engelhardt, "How algorithmic confounding in recommendation systems increases homogeneity and decreases utility", *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*, p. 224–232, Association for Computing Machinery (ACM), 2018.

27. Çapan, G., Ö. Bozal, İ. Gündoğdu and A. T. Cemgil, "Towards fair personalization by avoiding feedback loops", *arXiv preprint arXiv:2012.12862*, 2020.

28. Çapan, G., İ. Gündoğdu, A. C. Türkmen, Ç. Sofuoğlu and A. T. Cemgil, "A Bayesian choice model for eliminating feedback loops", *arXiv preprint arXiv:1908.05640*, 2019.

29. Wong, T.-T., "Generalized Dirichlet distribution in Bayesian analysis", *Applied Mathematics and Computation*, Vol. 97, No. 2-3, pp. 165–181, 1998.

30. Train, K. E., *Discrete choice methods with simulation*, Cambridge University Press, 2009.

31. Liang, D., R. G. Krishnan, M. D. Hoffman and T. Jebara, "Variational autoencoders for collaborative filtering", *Proceedings of the World Wide Web Conference*, pp. 689–698, 2018.

32. Covington, P., J. Adams and E. Sargin, "Deep neural networks for YouTube recommendations", *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*, p. 191–198, Association for Computing Machinery (ACM), 2016.

33. Chaptini, B. H., *Use of discrete choice models with recommender systems*, Ph.D. Thesis, Massachusetts Institute of Technology, 2005.

34. Bradley, R. A. and M. E. Terry, "Rank analysis of incomplete block designs: I. the method of paired comparisons", *Biometrika*, Vol. 39, No. 3/4, pp. 324–345, 1952.

35. Luce, R. D., *Individual choice behavior*, Wiley, 1959.

36. Plackett, R. L., "The analysis of permutations", *Applied Statistics*, pp. 193–202, 1975.

37. Caron, F. and A. Doucet, "Efficient Bayesian inference for generalized Bradley–Terry models", *Journal of Computational and Graphical Statistics*, Vol. 21, No. 1, pp. 174–196, 2012.

38. Guiver, J. and E. Snelson, "Bayesian inference for Plackett-Luce ranking models", *Proceedings of the 26th Annual International Conference on Machine learning*, pp. 377–384, 2009.

39. Baeza-Yates, R., "Bias in search and recommender systems", *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, p. 2, Association for Computing Machinery (ACM), 2020.

40. Lattimore, T. and C. Szepesvári, *Bandit algorithms*, Cambridge University Press, 2020.

41. Basilico, J., "Recent trends in personalization: a Netflix perspective", *Thirty-Sixth International Conference on Machine Learning (ICML)*, 2019, `https://slideslive.com/38917692`, accessed in May 2021.

42. Thompson, W. R., "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples", *Biometrika*, Vol. 25, No. 3/4, pp. 285–294, 1933.

43. Sui, Y., V. Zhuang, J. W. Burdick and Y. Yue, "Multi-dueling bandits with dependent arms", *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.

44. Saha, A. and A. Gopalan, "Battle of bandits", *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 06–10, 2018.

45. Sui, Y., M. Zoghi, K. Hofmann and Y. Yue, "Advancements in dueling bandits.", *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 5502–5510, 2018.

46. Liu, W. and R. Burke, "Personalizing fairness-aware re-ranking", *arXiv preprint arXiv:1809.02921*, 2018.

47. Liu, T.-Y., W. Chen and T. Qin, "Mechanism learning with mechanism induced data.", *The Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pp. 4037–4041, the Association for the Advancement of Artificial Intelligence (AAAI), 2015.