

A Non-Parametric Choice Model That Learns How Users Choose Between Recommended Options

Thorsten Krause

thorsten.krause@ru.nl

Radboud University

Nijmegen, The Netherlands

Harrie Oosterhuis

harrie.oosterhuis@ru.nl

Radboud University

Nijmegen, The Netherlands

Abstract

Choice models predict which items users choose from presented options. In recommendation settings, they can infer user preferences while countering exposure bias. In contrast with traditional univariate recommendation models, choice models consider which competitors appeared with the chosen item. This ability allows them to distinguish whether a user chose an item due to *preference*, i.e., they liked it; or *competition*, i.e., it was the best available option. Each choice model assumes specific user behavior, e.g., the multinomial logit model. However, it is currently unclear how accurately these assumptions capture actual user behavior, how wrong assumptions impact inference, and whether better models exist.

In this work, we propose the *learned choice model for recommendation (LCM4Rec)*, a non-parametric method for estimating the choice model. By applying kernel density estimation, LCM4Rec infers the most likely error distribution that describes the effect of inter-item cannibalization and thereby characterizes the users' choice model. Thus, it simultaneously infers *what* users prefer and *how* they make choices. Our experimental results indicate that our method (i) can accurately recover the choice model underlying a dataset; (ii) provides robust user preference inference, in contrast with existing choice models that are only effective when their assumptions match user behavior; and (iii) is more resistant against exposure bias than existing choice models. Thereby, we show that learning choice models, instead of assuming them, can produce more robust predictions. We believe this work provides an important step towards better understanding users' choice behavior.

CCS Concepts

• Information systems → Recommender systems.

Keywords

Recommender Systems, User Choice Modeling, Exposure Bias

ACM Reference Format:

Thorsten Krause and Harrie Oosterhuis. 2025. A Non-Parametric Choice Model That Learns How Users Choose Between Recommended Options. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems (RecSys '25)*, September 22–26, 2025, Prague, Czech Republic. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3705328.3748090>



This work is licensed under a Creative Commons Attribution 4.0 International License. RecSys '25, Prague, Czech Republic

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1364-4/2025/09

<https://doi.org/10.1145/3705328.3748090>

1 Introduction

Traditional implicit collaborative filtering methods assume that all items are chosen independently of each other [35, 36]. Accordingly, they estimate an item's relevance score purely from information about which users interacted with it [24]. In reality, users must often decide for one of multiple items that the recommender system exposes together, for example, when ordering lunch [30], booking a flight [27], or choosing a movie [37]. Consequently, the choices for one item and against others are dependent.

One prominent class of models that accounts for dependencies between choices are so-called choice models. Choice models predict the choice probabilities according to the preference towards each option, while also considering the competitive effect between options. They model this competitive effect through a probabilistic noise distribution, which is their key-characteristic. For instance, the multinomial logit model (MNL) [40], which is equivalent to the softmax, is characterized by its Gumbel noise distribution. Many sequential models, such as BERT4Rec [38], employ the softmax loss to predict the next user interaction. Furthermore, Krause et al. [29] showed that choice models are significantly more robust against exposure bias, particularly against exposure bias from competition among items. By providing choice probabilities, choice models can also contribute to related tasks such as balancing trade-offs in responsible recommendation [20] or, potentially, informing *whether* to make a recommendation and *how many* items to recommend.

However, existing approaches are only effective when the assumed choice model matches the true user behavior. Employing the wrong choice model can bias choice probabilities [40]. This poses a significant limitation, since the most accurate choice model for a given use case is generally unknown. Identifying it through benchmarking is practically infeasible since infinitely many possible choice models exist and some choice models, such as the popular probit, lack a closed-form analytical solution [9].

In this work, we directly address this limitation by proposing the *learned choice model for recommendation (LCM4Rec)*, a novel user-modeling method that simultaneously learns the most likely choice model and user preferences. Instead of assuming the probabilistic noise distribution underlying the true users' choice model, we utilize a non-parametric kernel density estimator to approximate it. Specifically, LCM4Rec maximizes the log-likelihood of observed interaction data, and thereby, infers both the most likely utility distribution (i.e., the user preferences) and noise distributions (i.e., the choice model) underlying their interactions. Because kernel density estimators can be used to approximate any arbitrary probability distribution, LCM4Rec is applicable to any possible choice model.¹ Moreover, because the choice model is inferred from interaction

¹Any choice model for which the error terms are IID.

data, we do not rely on a priori assumptions about what choice model users follow,² in contrast with previous work [29]. Additionally, we investigate the negative implications from assuming an inaccurate choice model with respect to ranking, choice probability estimation and exposure bias robustness.

In summary, LCM4Rec resolves the need for benchmarking or assuming choice models and is applicable to any choice model.² It simultaneously learns user preferences towards items as well as the competitive effects between presented items. Our experimental results on synthetic datasets indicate LCM4Rec is robust to what choice model underlies interaction data. Whilst existing methods can only accurately infer preferences when their assumed choice model matches the true model, LCM4Rec provides reliable and robust performance regardless of the underlying model. Furthermore, our results show that LCM4Rec is better at mitigating exposure bias. We also found that employing inaccurate choice models skews choice probabilities and increases exposure bias vulnerability.

To the best of our knowledge, we introduce the first method that both learns users preferences as how users choose between recommended options. Besides the improved effectiveness and robustness of LCM4Rec, we believe our work contributes to the better modeling and understanding of user behavior whilst avoiding assumptions.

2 Related work and Background

2.1 Discrete choice models for recommendation

Discrete choice models estimate with which probability $P_{ij|C}$ a user $i \in I$ would choose an item $j \in J$ from a discrete choice set $C \in \mathcal{C}$ by utilizing the transitive information on which items are preferred over another. These models assume that each user would always chose the choice alternative $j \in C$ from which they expect the highest utility U_{ij} . Econometricians use discrete choice models for determining how a variable affects preferences [21] and operations researchers use them for solving optimization problems based on observed choice data [14]. Discrete choice models have also been used for recommendation, most prominently under the term *collaborative competitive filtering* [42]. Their key difference from classical implicit feedback collaborative filtering models is that they require information about displayed choice alternatives. In recent years, more such data has become available, for example in the *finn.no* dataset [18] or the 2024 RecSys Challenge dataset³.

To account for the fact that humans do not strictly chose transitively [10], discrete choice models are probabilistic. They model the total utility U_{ij} as the sum of a deterministic utility V_{ij} and a random error term ϵ_{ij} s.t. $U_{ij} = V_{ij} + \epsilon_{ij}$. In the recommendation context, the deterministic utility V_{ij} corresponds to the relevance scores. The random error is interpreted as the share of utility that analysts cannot observe [40]. Thereby, it enables discrete choice models to capture that identical circumstances can result in different choices, or that succeeding choices can appear contradicting.

Our study focuses on the shape of the distribution that underlies ϵ_{ij} . The most commonly assumed error distribution is the IID Gumbel distribution of the multinomial logit model which results

in the following softmax-like choice probabilities [40]:

$$P_{ij|C} = \frac{\exp(V_{ij})}{\sum_{j' \in C} \exp(V_{ij'})}. \quad (1)$$

The closed analytical form and easy interpretability of its choice probabilities have contributed to the popularity of the multinomial logit model [40]. Its key property is its assumption of the irrelevance of independent alternatives (IIA), which states that the ratio of choice alternatives between two items does not depend on any third item's utility [4]. However, this key property is also a source of criticism [5]. Empirical studies indicate that the IIA assumption is inaccurate in many settings [6].

2.2 Non-parametric discrete choice models

Instead of having the analyst guess the correct choice model, non-parametric estimators learn the error distribution. This approach builds on the idea that the choice probabilities can be re-written as:

$$\begin{aligned} P_{ij|C} &= P(\forall j' \in C, j' \neq j \rightarrow U_{ij} > U_{ij'}) \\ &= P(\forall j' \in C, j' \neq j \rightarrow V_{ij} + \epsilon_{ij} - V_{ij'} > \epsilon_{ij'}). \end{aligned} \quad (2)$$

Under the assumption that all error terms ϵ_{ij} are IID distributed to some random variable ϵ with cdf F_ϵ and pdf ρ_ϵ , Manski [32] points out that the choice probabilities can be further reformulated as:

$$P_{ij|C} = \int \rho_\epsilon(\epsilon_{ij}) \prod_{j' \in C, j' \neq j} F_\epsilon(V_{ij} + \epsilon_{ij} - V_{ij'}) d\epsilon_{ij}, \quad (3)$$

and proposes a *maximum score estimator* that estimates a ranking function and that enables computing the effect of observed variables on the utility V_{ij} consistently in the binary case with two choice options. Cosslett [15] criticizes the maximum score estimator for being inapplicable to inferring choice probabilities and instead provides an estimator for the cdf F_ϵ in the binary case. As outlined for the multivariate case by Briesch et al. [8], the core idea is that given the deterministic utilities V_{ij} , the choice distribution exclusively depends on F_ϵ and ρ_ϵ . Thus, for N observations $\{(i_n, j_n, C_n)\}_{n \leq N} \subseteq I \times J \times \mathcal{C}$, one can estimate ϵ by defining a broad family of distributions \mathcal{F} , and subsequently select the member which minimizes the negative log-likelihood:

$$\hat{F} := \operatorname{argmin}_{F \in \mathcal{F}} - \sum_{n=1}^N \log(P_{i_n j_n | C_n}). \quad (4)$$

However, Cosslett [15] uses sorting and linear interpolation for estimating F , which is intractable for gradient descent. They also only consider the binary case. To the best of our knowledge, no other existing approach is directly applicable to the recommender systems setting either. Matzkin [34] drops any assumptions on the shape of V , but derives a discrete optimization problem that returns a non-differentiable estimator. Horowitz [23] smoothes the maximum score estimator but still cannot infer choice probabilities. For the binary case, Klein and Spady [28] estimate the choice probabilities $P_{ij|C}$ directly. However, generalizing their method to the multivariate case requires density estimation with as many dimensions as choice alternatives, which is unfeasible for our setting where choice sets can contain 20 items or more [18]. More recent works focus increasingly on challenges that are beyond current concerns in recommender systems research. For example, Hirano [22]

²To be precise, we only make the assumption that choice model error terms are IID.

³<https://recsys.eb.dk/dataset/>

Table 1: Cannibalization patterns according to two different choice models. For items j_1, j_2, j_3 with utilities $V_{ij_1} = 3, V_{ij_2} = 1, V_{ij_3} = 2$, the models agree on the choice probabilities between j_1 and j_2 , but disagree when the option j_3 is also available.

Model	Choice set	Choice probability		
		j_1	j_2	j_3
Multinomial logit	$\{j_1, j_2\}$	88.1%	11.9%	—
	$\{j_1, j_2, j_3\}$	66.5%	9.0%	24.5%
Exponential	$\{j_1, j_2\}$	88.1%	11.9%	—
	$\{j_1, j_2, j_3\}$	70.2%	5.3%	24.5%

develops a solution for random effect autoregressive models with non-parametric idiosyncratic shocks. Others use techniques that we deem incompatible with gradient descent-based optimization such as discrete estimators [8, 25], or linear programming [12, 39].

2.3 Competitive effects and cannibalization

The choice model also describes the competitive effect between items, i.e., how the availability and utility of other item options affects choice probabilities [3, 12]. Specifically, the concept of cannibalization refers to how the choice probabilities of items decrease due to a substitute's presence [3, 31]. Each choice model predicts a different effect in such cases. For example, under the multinomial logit model all choice probabilities decrease by the same proportion, whereas under the exponential choice model, low-utility alternatives lose a larger share of their choice probabilities than high-utility alternatives. Table 1 displays the differences between these models in an example scenario based on the one in Alptekinoglu and Semple [3]. Whilst both models agree in the ordering of the most probable choices, their predicted choice probabilities differ. Consequently, in order to accurately infer user preferences, the choice model should match the true user behavior. Alptekinoglu and Semple [3] argue that the exponential model better represents real-world cannibalization patterns than the multinomial logit model, while referring to the empirical observations by Blattberg and Wisniewski [7].

2.4 Exposure bias in recommender systems

Exposure bias occurs whenever a model's previous recommendation affect user behavior and subsequently the logged feedback data. Whilst different formal definitions of exposure bias exist, we follow previous work [29, 41] that considers a model *exposure bias resistant* if the exposure distribution does not affect the model's asymptotic rankings. This definition does not imply convergence towards the true ranking because other biases could still occur.

Mitigating exposure bias should be a priority for all stakeholders. It manipulates future model generations to recommend similar items as in the past, resulting in user-model feedback loops [11, 26]. Through exposure bias, the model over-proportionally prefers popular items [19] and can fail to cater to the special interests of user subgroups [2]. In the long-term, its outputs to different user groups can homogenize [33] and in multi-stakeholder environments, the unfair exposure can translate into intra-provider unfairness [1, 16].

Krause et al. [29] demonstrated that, in addition to *overexposure*, which occurs when some items are exposed more often than others, exposure bias can result from *competition*, which occurs when some items tend to be exposed with popular or unpopular alternatives. They also demonstrated that discrete choice models can effectively reduce both forms of exposure bias. Their formal proof for unbiasedness, however, assumed that the model matches a ground-truth choice model. While using an inaccurate model could impair unbiasedness, Krause et al. did not further investigate implications for exposure bias.

3 Method: The Learned Choice Model for Recommendation (LCM4Rec)

Our goal is to propose a choice model approach for modeling user choices in recommender settings while avoiding assumptions about user behavior. The non-parametric choice models discussed in Section 2.2 prevent the need to assume users' true choice model, and thus appear to provide a great alternative to their parametric counterparts. However, each is incompatible with state of the art recommender systems due to at least one of the following reasons: (i) they utilize non-differentiable estimators and optimization methods; (ii) they are only applicable to the binary choice cases where users are only presented with two options; or (iii) they are computationally complex for application on large choice sets. As a result, the existing non-parametric methods do not provide any approach that is relevant to current recommender systems.

In this section, we address these shortcomings directly and propose our novel non-parametric choice model: the *learned choice model for recommendation (LCM4Rec)*. Its core feature is a smooth kernel density estimator that approximates the cdf F_ϵ and the pdf ρ_ϵ of the error distribution (Equation 3). Importantly, the estimator can approximate any arbitrary error distribution and provide a differentiable likelihood function for them; thereby, LCM4Rec can approximate any possible choice model with IID error terms. Moreover, LCM4Rec can be optimized to predict observed interaction data, which enables learning what choice model best explains user behavior, instead of assuming this a priori. To the best of our knowledge, this makes LCM4Rec the first non-parametric choice model that is compatible with state-of-the-art recommender systems.

3.1 Learned choice model architecture

In order to approximate any arbitrary error distribution underlying a choice model, we use kernel density estimation to construct a family of functions \mathcal{F} that contains approximately close members to any possible cdf F_ϵ . Subsequently, we search for a member of \mathcal{F} that minimizes the negative log-likelihood of observed interaction data (equation 4). We start our description of this procedure by defining and discussing our choice for \mathcal{F} . Consider the family \mathcal{F} of sigmoid kernel-based functions with the following shape:

$$\hat{F}_K(x) := \sum_{k=1}^K w_k \sigma\left(\frac{x - x_k}{h_k}\right) \in C^\infty(\mathbb{R}), \quad (5)$$

where σ is the sigmoid function, $K \in \mathbb{N}$, $x_k \in \mathbb{R}$, $w_k \in \mathbb{R}_{>0}$, $h_k \in \mathbb{R}_{>0}$ and $\sum_{k=1}^K w_k = 1$. In other words, each \hat{F}_K is a weighted average of sigmoid functions scaled and translated by the h_k and x_k

parameters. Importantly, the derivative of \hat{F}_K is numerically stable:

$$\hat{\rho}_K(x) = \sum_{k=1}^K \frac{w_k}{h_k} \sigma\left(\frac{x - x_k}{h_k}\right) \left(1 - \sigma\left(\frac{x - x_k}{h_k}\right)\right) \in C^\infty(\mathbb{R}). \quad (6)$$

Any function $\hat{F}_K \in \mathcal{F}$ is strictly monotonically increasing and bounded in the interval $[0, 1]$. Moreover, \hat{F}_K and $\hat{\rho}_K(x)$ are continuously differentiable with respect to their parameters and to the utility estimates V_{ij} so that we can search for optimal representatives with regard to the empirical log-likelihood function (Equation 4) via gradient descent.

While directly learning the parameters x_k is possible, not restricting them leads to local minima w.r.t. the loss function. For this reason, we set x_k to be uniformly spread out over a closed interval:

$$x_1 = -l, x_2 = l\left(\frac{2}{K-1} - 1\right), \dots, x_K = l, \quad (7)$$

for a learnable scale parameter $l > 0$. This makes optimization much easier without restricting the number of learnable functions when N is large enough. To enforce w_k, h_k , and ω to be positive, we introduce $\alpha_k \in \mathbb{R}, \beta_k \in \mathbb{R}$, and $\lambda \in \mathbb{R}$, and define:

$$w_k := \text{softmax}(\alpha_k), h_k := \frac{l}{K} \text{softplus}(\beta_k), l := \text{softplus}(\lambda). \quad (8)$$

For numerical stability, we bound $-0.1 < \beta_i < 5$ and $0.1 < \lambda < 10$. These bounding constraints ensure that the estimator does not collapse during the first iterations when the estimated utilities are very small. With our model of the error distribution, we still require a model of the utilities V_{ij} ; our definition of these applies commonly-used learned-embedding interactions:

$$V_{ij} := u_i \cdot v_j + c_j, \quad (9)$$

where $u_i \in \mathbb{R}^m$ and $v_j \in \mathbb{R}^m$ with $m \in \mathbb{N}$ are user- and item-embeddings, and c_j is an item-specific constant as is common in discrete choice modeling.

This completes the model architecture underlying LCM4Rec: the error distribution of ϵ is modelled using a kernel density estimator that applies a weighted average of sigmoid functions (Equation 5) where an error term ϵ_{ij} is IID sampled per user-item combination; the user utilities are modelled by the dot-product between two learned embeddings and an item-specific constant (Equation 9); and the final utility is simply the sum of these: $U_{ij} = V_{ij} + \epsilon_{ij}$. Recall that choice models assume the item with the highest utility is chosen (see Section 2). Whilst our LCM4Rec is extremely expressive and can capture any arbitrary error distribution, as we show below, we have restricted it by assuming the error terms ϵ_{ij} are IID independent. We note that this restriction rules out some choice models such as the nested logit model [40]; we leave the extension of our model to the more general non-IID setting as future work.

Finally, we show that LCM4Rec can approximate any choice model with IID error terms. Specifically, we show that for any cdf F , there always exists a member $\hat{F} \in \mathcal{F}$ that is arbitrarily close to F .

THEOREM 1. *Let F be a continuous, strictly monotone cdf with finite support on a closed interval S . Let \mathcal{F} be defined as above. Then, for all $\epsilon > 0$, there exists a member $\hat{F} \in \mathcal{F}$ such that*

$$|\hat{F}(x) - F(x)| < \epsilon, \quad \forall x \in S. \quad (10)$$

PROOF. Because F operates on a closed interval, it is uniformly continuous. We can therefore find $\delta > 0$ such that

$$|F(x) - F(y)| < \frac{\epsilon}{4} \quad \forall x, y \in S : |x - y| < \delta. \quad (11)$$

Based on equation 5, set the number of kernels K and the width λ large enough so that any $x \in S$ is within close range of two design points, i.e., $\exists k_1, k_2 \leq K : x_{k_1} < x < x_{k_2}$ with $0 < x_{k_2} - x_{k_1} < \delta$. Set the weights w_k so that $\sum_{l=1}^K w_l = F(x_k)$, $\forall k \leq K$. Then,

$$\hat{F}_K(x_k) = F(x_{k-1}) + \underbrace{\frac{F(x_k) - F(x_{k-1})}{2}}_{>0} + \xi \leq F(x_k) \quad \forall 2 \leq k \leq K. \quad (12)$$

where ξ is a small error term that results from the sigmoids that are placed around the other design points $x_{k'}$ with $k' \neq k$. We can ignore ξ as $\lim_{h \rightarrow 0} \xi = 0$. Hence,

$$F(x_{k-1}) \leq \hat{F}_K(x_k) \leq F(x_k) \quad \forall 2 \leq k \leq K \text{ and } h \ll 1. \quad (13)$$

For $x \in S$ and its nearest point $k^* := \operatorname{argmax}_{k \leq K} (x_k \leq x)$ we get:

$$\begin{aligned} |\hat{F}_K(x) - F(x)| & \leq |\hat{F}_K(x) - \hat{F}_K(x_{k^*})| + |\hat{F}_K(x_{k^*}) - F(x_{k^*})| + |F(x_{k^*}) - F(x)| \\ & \stackrel{\Delta\text{-ineq.}}{\leq} |F(x_{k^*+1}) - F(x_{k^*-1})| + |F(x_{k^*+1}) - F(x_{k^*})| + |F(x_{k^*}) - F(x_{k^*-1})| \\ & \stackrel{(13)}{\leq} 2|F(x_{k^*}) - F(x_{k^*-1})| + 2|F(x_{k^*+1}) - F(x_{k^*})| < \epsilon. \\ & \stackrel{\Delta\text{-ineq.}}{\leq} \end{aligned} \quad (14)$$

□

Therefore, for any continuous, strictly monotone cdf F with bounded support, the function family \mathcal{F} contains an arbitrarily good approximator \hat{F} .

3.2 Optimizing for the most likely choice model

Having constructed the family of functions \mathcal{F} , we now introduce a method for finding a good candidate $\hat{F}_K \in \mathcal{F}$ for approximating F_ϵ , based on N observed choices $\{(i_n, j_n, C_n)\}_{n \leq N} \subseteq I \times J \times C$ of users, chosen items, and choice sets. Importantly, any \hat{F}_K is differentiable with respect to its parameters $\alpha_k, \beta_k, \lambda$, and the user- and item-parameters u_i, v_j , and c_j . Accordingly, we propose to learn all model parameters by optimizing the negative log-likelihood (NLL) function in Equation 4 via gradient descent.

However, the main challenge to this approach is that we lack a closed form for the integrals in Equation 4, and thus, we cannot evaluate them directly. As a solution, we propose to approximate its gradient via Monte Carlo integration. Unfortunately, there is no closed form solution for the inverse cdf of F_K , which means we cannot sample from it directly. Instead, we use the fact that F_K is a weighted average of sigmoid cdfs from which samples can be drawn directly. Specifically, for each observation n and kernel $k \in \{1, \dots, K\}$ (cf. Equation 5), we draw S samples; for observation n and kernel k , sample number s is generated through the inverse cdf of the sigmoid kernel:

$$\tilde{\epsilon}^{n,k,s} := x_k + h_k \log\left(\frac{u^{n,k,s}}{1 - u^{n,k,s}}\right), \quad u^{n,k,s} \sim \mathcal{U}(0, 1). \quad (15)$$

With these samples, the choice probability can be approximated by:

$$\begin{aligned}
& \int \hat{\rho}_K(e_{ijn}) \prod_{j' \in C_n, j' \neq j_n} \hat{F}_K(V_{ijn} + e_{ijn} - V_{ijn'}) de_{ijn} \\
&= \int \sum_{k=1}^K \frac{w_k}{h_k} \sigma\left(\frac{e_{ijn} - x_k}{h_k}\right) \left(1 - \sigma\left(\frac{e_{ijn} - x_k}{h_k}\right)\right) \\
&\quad \cdot \prod_{j' \in C_n, j' \neq j_n} \hat{F}_K(V_{ijn} + e_{ijn} - V_{ijn'}) de_{ijn} \quad (16) \\
&\approx \frac{1}{S} \sum_{s=1}^S \sum_{k=1}^K w_k \underbrace{\prod_{j' \in C_n, j' \neq j_n} \hat{F}_K(V_{ijn} + \tilde{e}^{n,k,s} - V_{ijn'})}_{=: \hat{P}_{ns}} =: \hat{P}_n.
\end{aligned}$$

Note that due to Jensen's inequality, taking the logarithm in Equation 16 introduces bias which we correct for through third-order Taylor expansion as in [17]. This leads to our final approximated negative log-likelihood function:

$$\hat{\mathcal{L}} := - \sum_n \log(\hat{P}_n) + \frac{\sum_{s=1}^S (\hat{P}_{ns} - \hat{P}_n)^2}{2S(S-1)\hat{P}_n^2} - \frac{\sum_{s=1}^S (\hat{P}_{ns} - \hat{P}_n)^3}{3S(S-1)(S-2)\hat{P}_n^3} \quad (17)$$

Finally, we can minimize the approximated loss $\hat{\mathcal{L}}$ through gradient descent to obtain an estimate for F_ϵ . Thereby, LCM4Rec can potentially recover any underlying choice model with IID error terms from observation data, without further a priori assumptions.

3.3 Computational complexity

We apply a straightforward computation of the loss in equation 17 which has a computational complexity of $\mathcal{O}(K^2NS|C|)$. The quadratic scaling in K results from our sampling strategy; as described above, we must sample from the pdf $\hat{\rho}_K$ in a way that is differentiable with respect to the weights w_k . As a solution, we conveniently generate samples per kernel, resulting in the sums over kernels and samples-per-kernel in equation 17. In practice, we see no benefit in increasing the resolution beyond $K = 10$, because very fine grained peaks in the error function—if existent—barely affect choice probabilities. Thus, since K remains small, the quadratic complexity is not a significant problem in practice. Alternatively, linear scaling in K can be achieved by sampling uniformly over the support $\hat{\rho}_K$ as in vanilla Monte-Carlo. We initially implemented this solution but found it less efficient for small values of K than our straightforward computation.

3.4 Identification and regularization

Finally, we tackle a challenge in the prevention of overfitting our method: The variability of l and c_j renders L2-regularization on the user- and item-embeddings \hat{u}_i and \hat{v}_j ineffective, because for any decrease in \hat{u}_i or \hat{v}_j , the unregularized state can be reconstructed by shrinking l , h_k and c_j accordingly. Consequently, L2-regularization would lead to a collapse of $\hat{\rho}_K$. Thus, in order to prevent overfitting, we scale the item-specific constants c_j into the

interval $[0, 1]$ after every gradient update. This forces the width parameter l and the utilities to stay within the same scale as c_j .⁴

4 Experimental Setup

We perform two experiments to answer three research questions. A key property of LCM4Rec is that it does not assume what error distribution underlies user choices, thus its accuracy should not depend on them. Accordingly, our first research question is:

RQ1 Does LCM4Rec return accurate and robust choice probabilities under different true choice models, i.e., different error distributions?

However, accurate choice probabilities do not necessarily indicate the correct choice model is identified, since an incorrect model can still overfit, and thereby, (partially) compensate for an incorrect error distribution. To address this possibility, we also directly validate whether LCM4Rec finds the correct choice model:

RQ2 Can LCM4Rec accurately recover the true choice model from interaction data?

Accurate choice models are more robust to exposure bias, especially from item-co-exposure [29], due to better modeling competition between items. Thus, our third research question concerns:

RQ3 Is LCM4Rec more robust to exposure bias than parametric choice models?

To pursue our research questions and understand the importance of assumptions about user behavior, our experiments compare our LCM4Rec to parametric alternatives, in cases where the parametric models and the correct error distribution match, and in cases where they do not match. Our data and implementation are publicly available at https://github.com/krauthor/LCM4Rec_RecSys2025.

4.1 Experiment 1: Accuracy and robustness

To assess **RQ1** and **RQ2**, we apply a setup similar to that of Krause et al. [29], with the important difference that we vary the choice model underlying interactions.

Dataset. Our first generated dataset consists of the choices of 500 users on 500 items, where every user performed up to 500 choices from uniformly random subsets of four items. Users could interact with the same item multiple times. We chose the large number of interactions per users to investigate systematic, asymptotic errors that result from learning with inaccurate choice models. Each user and item are represented by an embedding, $u_i, v_j \in \mathbb{R}^3$ respectively, sampled uniformly from the unit sphere of radius $\sqrt{2}$. Item-specific constants $c_j \in \mathbb{R}$ are sampled from the uniform distribution: $\mathcal{U}(0, 1)$. Final utilities are $U_{ij} := u_i \cdot v_j + c_j + \epsilon_{ij}$ (cf. equation 9).

We use three different error distributions, and thereby different choice models, to generate the ϵ error terms: (i) **Gumbel** with $\epsilon \sim \text{Gumbel}(0, 0.75)$; (ii) **Signed exponential** with $\epsilon \sim -\text{Exponential}(0, 0.75)$; (iii) **Gaussian Mixture** with $\epsilon = \frac{1}{3}N_1 + \frac{2}{3}N_2$, where $N_1 \sim \mathcal{N}(-0.75, 0.25)$ and $N_2 \sim \mathcal{N}(0.75, 0.25)$. The Gumbel distribution was chosen to represent non-conservative choice behavior, whereas the signed exponential distribution represents conservative choice behavior. Lastly, the Gaussian mixture distribution is

⁴Fixing the scale parameter l does not suffice as the model could still make the estimated distribution infinitely thin by placing the weight on only a few design points.

included because it lacks a corresponding parametric model. Thus, its output can only be fit by non-parametric models such as ours.

To properly evaluate choice models, it is crucial that the training set contains no user-item interactions that are later used for evaluation, including negative interactions. We follow Krause et al. [29] and construct the training and validation datasets based on the interactions of a subset of users $U_{\text{Train}} \subset U$ on all items together with the interactions of the remaining users $U_{\text{Eval}} = U \setminus U_{\text{Train}}$ with half of all items $J_{\text{Train}} \subset J$. The test set was constructed based on the interactions of U_{Eval} with the remaining half of all items $J_{\text{Eval}} = J \setminus J_{\text{Train}}$.

Models. The following models are included in our comparison:

- The **multinomial logit model (MNL)** [40] is the de-facto standard discrete choice model [40]. It assumes **Gumbel**-distributed errors. Its choice probabilities equal a softmax over the deterministic utilities with non-conservative cannibalisation patterns.
- The **exponential model (ENL)** [3] assumes that the random errors are generated by a **signed exponential** distribution with very conservative cannibalization patterns.
- The **binary logit model (BL)** [40] is the univariate variant of the MNL model. It minimizes a logistic loss over all observed positive and negative interactions individually. Thus, it ignores any possible competitive effects of alternative options.
- The **binary cross-entropy (BCE) with negative sampling** is a basic matrix factorization recommendation model that is equivalent to the BL model except that it samples the negative labels from the item corpus.
- The **generalized binary cross-entropy loss (gBCE)** [35] is designed to capture confidence scores more accurately than the BCE loss and the MNL with negative sampling.
- Our **learned choice model for recommendation (LCM4Rec)** features non-parametric kernel density estimation to infer the correct choice probabilities and cannibalization pattern from the data, uniquely, without assuming an error distribution a priori.

Hyper-parameters. We set the embedding size for each model to match the size used in the generation process. Learning rates and optimizer choices are determined via grid-search on the Gumbel distributed data and applied to all distributions. For LCM4Rec we set $K = 5$, $S = 5$, and for gBCE we set $t = 1$. We tried many values of K without running into any computational constraints, optimization instability, or notable overfitting. Values larger than $K = 5$ did not result in better predictive accuracy, thus we chose this value to show that a small number of kernels suffice and improvements can be achieved at low computational costs.

Evaluation. To measure the accuracy of the choice probabilities, we use the Kullback-Leibler divergence (KLD) on the distributions over the entire evaluation item set J_{Eval} . These distributions represent the users' preferences if they could choose from the entire evaluation corpus and would be relevant during inference [20]. We measure predictive performance in the forms of the nDCG, negative log-likelihood (NLL), and accuracy (Acc) for the choices among all alternatives in the respective choice sets on the test set. For the univariate models we compute the probability of choosing an item from a choice set as the probability of choosing only that item and rejecting the others. All results are means over 100 simulation

repetitions. Bootstrapped one-sided t-tests determine statistical significance of performance differences to the next worse value.

4.2 Experiment 2: Correcting for exposure bias

To analyze how the models react to exposure bias for **RQ3**, we evaluate their learned behavior under different exposure distributions. **Dataset.** Our approach to measuring the effect of exposure is to take a small item subset $J^{\text{Bias}} \subset J$, vary its exposure, and subsequently see whether the items are treated differently as a result. Again, we prevent information leakage by avoiding interactions in the training set with user-item pairs that are also used during evaluation (see Section 4.1). The three error distributions are reused: Gumbel, signed exponential and Gaussian Mixture, for each we generate two pairs of two datasets with different exposure distributions: O' and O'' . Subsequently, by comparing the pairs of learned models on the test set, we can see whether the change in exposure results in a difference in the learned choice models.

Our pairs of exposure distributions follow those proposed by Krause et al. [29], to capture different kinds of exposure bias. The first pair considers the effect of *non-uniform exposure frequencies*, i.e., some items getting more exposure than others. To generate O' , all items are presented equally often; whereas for O'' , we first randomly select 25 items and then include at least one of them in every second choice set, resulting in 1.9 times as much exposure per item. The second pair considers bias due to *non-uniform competition*, i.e., the effect of being presented together with more or less popular alternatives. For this pair, 25 items are randomly selected again, to generate O' these are only presented together with the top 20 percent most-popular items, for O'' they are only presented with the bottom 20 percent most popular items (thus the least-popular items). In both cases, if the choice model is robust to exposure bias, the learned preferences for the selection of 25 items should be the same for O' and O'' in expectation. Therefore, large differences indicate a substantial effect from exposure, and thus, susceptibility to exposure bias, and vice-versa, small differences indicate robustness.

The number of users, items, and choices per user, the choice set size, the embeddings, and the error distributions are the same as for the first experiment (Section 4.1). The training, validation and test set were also constructed in the same way.

Models in comparison. For the second experiment, we include the MNL, ENL, and LCM4Rec models. The univariate models are excluded because previous work has already shown them to be more vulnerable to exposure bias than MNL and ENL [29].

Hyper-parameters. We employ the same hyper-parameters as the previous experiment except that we apply the SGD optimizer to ENL. Preliminary experiments showed that, for ENL, the Adam optimizer is more sensitive to exposure bias from non-uniform competition without significantly improving accuracy.

Evaluation. Our evaluation procedure also follows Krause et al. [29]. To measure the effect of the simulated exposure bias we consider the average rank of the selection of items as predicted by a model trained on O' or O'' . Our metric is the difference in predicted ranks for each pair of equivalent models where one was trained on O' and the other on O'' . Thereby, this difference can show whether

the exposure bias, i.e., non-uniform exposure or non-uniform competition, can affect the preferences inferred by the choice models. All reported results are averages over 100 simulation repetitions.

4.3 Assumptions and possible extensions

While we keep modeling assumptions general, several specifications common in real world data can be made: First, we generated 500 choices per user with 4 options per choice, which is realistic in only some scenarios. Non-parametric estimators typically converge slower than parametric ones so that an evaluation on more sparse data sets would be of interest. Second, we sampled choice sets uniformly. Real-world data often exhibits a long-tail item distribution and how such a distribution could affect our model’s robustness remains open. Likewise, we sampled embeddings uniformly while a long-tailed utility distribution could be more realistic. Third, we assumed knowledge over users’ choice sets. The sets’ compositions could be partially unobserved in real-world applications as fully tracking observed options can be complex and as users can discover and interact with items outside of the observable domain. To account for such behavior one can introduce a *no-choice alternative* that represents the decision *not* to choose any presented option [18, 40]. Moreover, padding and cropping allows processing choice sets of varying sizes.

5 Results and Discussion

5.1 Accuracy and robustness

Our discussion starts by considering **RQ1**: whether LCM4Rec returns accurate and robust choice probabilities under different true choice models. Table 2 displays the performance of all models in the comparison under three different true choice models.

We see that two of the univariate models BCE and gBCE consistently achieve the worst performance across all error distributions and metrics (except ENL has worse KLD under Gumbel). This is unsurprising, since univariate models are not designed for the choice model setting and ignore the effect of competition.

Surprisingly, the simplest univariate model, BL, consistently outperforms BCE and gBCE. In particular, under the Gumbel distribution, BL reaches the highest performance and even outperforms the multivariate ENL on all metrics except nDCG. BL is the univariate variant of MNL which assumes a Gumbel distribution, which may explain why BL performs so well in this setting. Under the other distributions, BL no longer reaches the highest performance for KLD, NLL and Acc. It seems that, whilst overall the multivariate models are more accurate in our setting, univariate models can still benefit from matching the true error distribution.

Next we consider the parametric choice models: MNL and ENL. Clearly, we observe that each model has the best performance across all metrics when their assumed error distribution matches the true distribution, i.e., Gumbel for MNL and signed exponential for ENL. This is expected, since the structure of the fitted model matches the true model structure. Surprisingly, MNL and ENL always reach the highest observed nDCG score and are close the highest accuracy, regardless of the true error distribution. The NLL is more affected but differences are still quite marginal: always below 0.08 of the best NLL. However, for KLD this is not the case: MNL and ENL have substantially worse KLD when their assumed error distribution

Table 2: Inference performance and robustness. The correctly specified models are highlighted with \dagger and were expected to perform best. The KLD scores refer to the choice probabilities from the entire item corpus J_B . Best results in bold font, second best underlined. Significantly better scores than the next worse models’ follow $^*p < 0.1$; $^{}p < 0.05$; $^{***}p < 0.01$.**

		KLD \downarrow	NLL \downarrow	nDCG \uparrow	Acc \uparrow
Gumbel	BL	0.028^{***}	1.184^{***}	0.996^{***}	0.477^{***}
	BCE	0.122 ^{***}	1.242 ^{***}	<u>0.983^{***}</u>	0.451 ^{***}
	gBCE	0.163 ^{***}	1.272	0.967	0.415
	MNL †	0.028^{***}	1.184^{***}	0.996^{***}	0.477^{***}
	ENL	7.952	<u>1.188^{***}</u>	0.996^{***}	<u>0.475^{***}</u>
	LCM4Rec	<u>0.085^{***}</u>	1.184^{***}	0.996^{***}	0.477^{***}
Signed Exponential	BL	1.574 ^{***}	1.063	0.997[*]	<u>0.538^{***}</u>
	BCE	2.310 ^{***}	1.146 ^{***}	<u>0.986^{***}</u>	0.510 ^{***}
	gBCE	2.386	1.183	0.974	0.474
	MNL	1.491 ^{***}	<u>1.062[*]</u>	0.997^{***}	0.539[*]
	ENL †	0.345^{***}	1.056^{***}	0.997^{***}	0.539^{**}
	LCM4Rec	<u>0.459^{***}</u>	1.056^{***}	0.997^{***}	0.539[*]
Gaussian Mixture	BL	1.166 ^{***}	1.127 ^{***}	0.997^{***}	<u>0.515^{***}</u>
	BCE	1.815 ^{***}	1.199 ^{***}	<u>0.985^{***}</u>	0.482 ^{***}
	gBCE	1.881	1.231	0.971	0.446
	MNL	1.108 ^{***}	<u>1.126</u>	0.997^{***}	<u>0.515^{***}</u>
	ENL	<u>0.625^{***}</u>	1.130 ^{***}	0.997^{***}	0.514 ^{***}
	LCM4Rec	0.337^{***}	1.123^{***}	0.997^{***}	0.516

Table 3: Mean KLDs and standard deviation between the true and modeled error distributions. † Correctly specified model.

	Gumbel	Sign. Exp.	Gauss. Mix.
MNL	0.00 †	1.25	0.43
ENL	1.19	0.00 †	0.87
LCM4Rec (ours)	0.13 \pm 0.07	0.26 \pm 0.05	0.31 \pm 0.18

is wrong. It thus appears that the MNL and ENL are only able to compensate for their incorrect choice model assumption for the NLL, nDCG, and accuracy metrics. But this results in substantial degraded match between the learned and true preferences in terms of KLD. Importantly, we note that our dataset generation is idealized as choice sets are uniform random samples, potentially, this compensation does not occur in less ideal settings.

Finally, we discuss our non-parametric multivariate LCM4Rec model which is the only model that estimates the error distribution. Strikingly, LCM4Rec reaches the best observed performance for the NLL, nDCG and Acc metrics across all error distributions. Moreover, it has the best KLD under the Gaussian mixture distribution and second best for the signed exponential and Gumbel distributions. Therefore, it appears LCM4Rec is only outperformed by parametric models when they correctly assume the true choice model, where

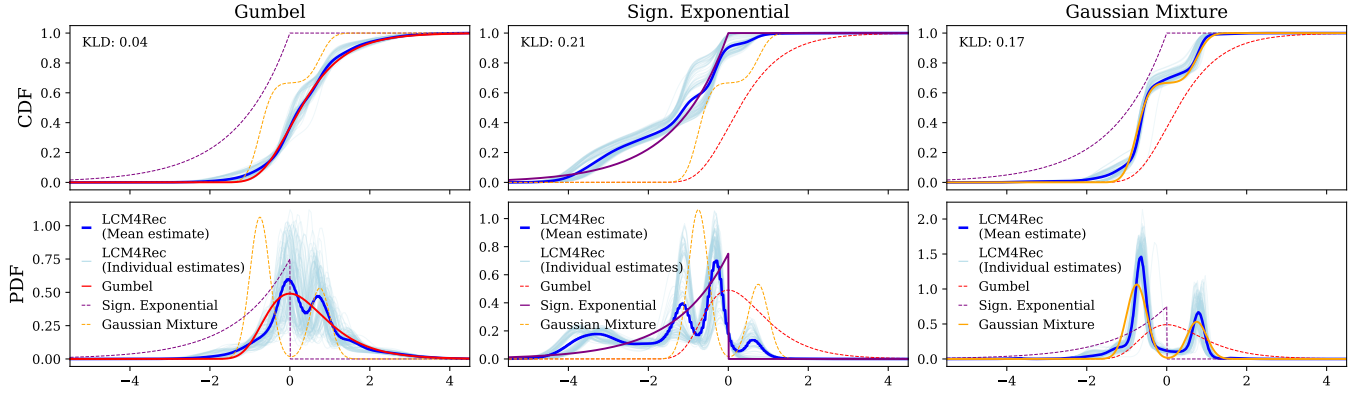


Figure 1: Estimated versus true cdfs and pdfs and KLD score between the mean estimate and the true distribution, smoothed to avoid null sets. Estimates are shifted to minimize the KLD to the true distribution (translation invariance in choice models).

the difference in KLD is still limited. In the case of the Gaussian mixture, where there is no matching choice model, LCM4Rec has significantly better performance in terms of KLD and NLL than all other models. As a result, LCM4Rec is the only model that has consistent good performance across all distributions and metrics. We attribute this to the fact that LCM4Rec infers the most likely error distribution, avoiding assumptions about the true choice model.

Therefore, we can answer **RQ1** as follows: LCM4Rec is the only model that can accurately predict choice probabilities regardless of the true error distribution. Thereby it is robust to the true user choice model and the safest choice when the exact model of user behavior is unknown. Additionally, our results indicate that the robustness of LCM4Rec does not come with a trade-off in predictive performance, as it reaches competitive performance in all metrics.

Lastly, we note that all our discrete choice models scored comparably well in terms of nDCG. Thus, it appears the underlying choice model does not seem to critically affect ranking performance in our first experiment setting. Moreover, performance of BL shows that, under uniform exposure, exact information on item co-exposure and choice sets is not even needed for ranking purposes. We believe that this is an artifact of our setup where item exposure is uniform (in expectation), hence our second experiment concerns a setting where this is not the case (see Section 5.3).

5.2 Learning the choice model

For **RQ2**, we evaluate whether LCM4Rec can accurately recover the true error distribution. Figure 1 displays the error distributions' cdfs and pdfs in our experimental setting and LCM4Rec's estimates.

Upon visual inspection, the estimates accurately approximate the error distributions. The estimator, which learns the cdf, consistently matches the strongest cdf. Its derivatives also follow the true pdfs' shapes. We observe some oscillation, especially when the error is signed exponential, a likely artifact from the resolution ($K = 5$) of our kernel density estimator. Nevertheless, LCM4Rec correctly identifies Gumbel's right skew, signed exponential's left skew, and the Gaussian mixture's two columns. The results in Table 2 also reveal that the estimates produce competitive predictive performance (see Section 5.1). To quantify the match, Table 3 shows the

Table 4: Exposure bias. Values show how many ranks items are ranked higher due to overexposure/competition, including 95%-CIs. Best (lowest abs.) bold, second best underlined.

ϵ	Model	Overexposure	Competition
Gumbel	MNL	$-0.427 \pm .401$	$-1.314 \pm .427$
	ENL	<u>$-0.418 \pm .461$</u>	$1.701 \pm .486$
	LCM4Rec (ours)	$-0.258 \pm .425$	$-0.000 \pm .448$
Sign. Exp.	MNL	$-0.784 \pm .311$	$-6.423 \pm .329$
	ENL	<u>$-0.581 \pm .321$</u>	<u>$-2.021 \pm .337$</u>
	LCM4Rec (ours)	$-0.212 \pm .351$	$-1.571 \pm .359$
Gauss. Mix.	MNL	$-0.247 \pm .345$	$-2.856 \pm .381$
	ENL	<u>$-0.100 \pm .378$</u>	<u>$0.294 \pm .407$</u>
	LCM4Rec (ours)	$-0.026 \pm .363$	$0.118 \pm .389$

mean KLD scores between the true, the assumed, and the estimated distributions. LCM4Rec scores low on KLD and much lower than MNL and ENL when their assumed distribution is incorrect. Hence, we affirm **RQ2**: LCM4Rec can accurately recover the true error distribution's shape and learn the true user choice model.

5.3 Exposure bias resistance

Finally, we turn to **RQ3** and compare the robustness to exposure bias of LCM4Rec with that of parametric choice models. Table 4 shows how many additional ranks items obtain on average in the model predicted ranking of items when they are (i) overexposed versus uniformly exposed; or (ii) presented with popular alternatives versus unpopular alternatives. Additionally, Figure 2 displays the difference in observed and true item ranks.

Our results indicate that the effect of (i) overexposure on the behavior of all three models is limited: Figure 2 shows differences are close to zero for all true item ranks. This is also visible in Table 4 which also reveals that when considering the confidence intervals of the results, the differences between the models do not appear

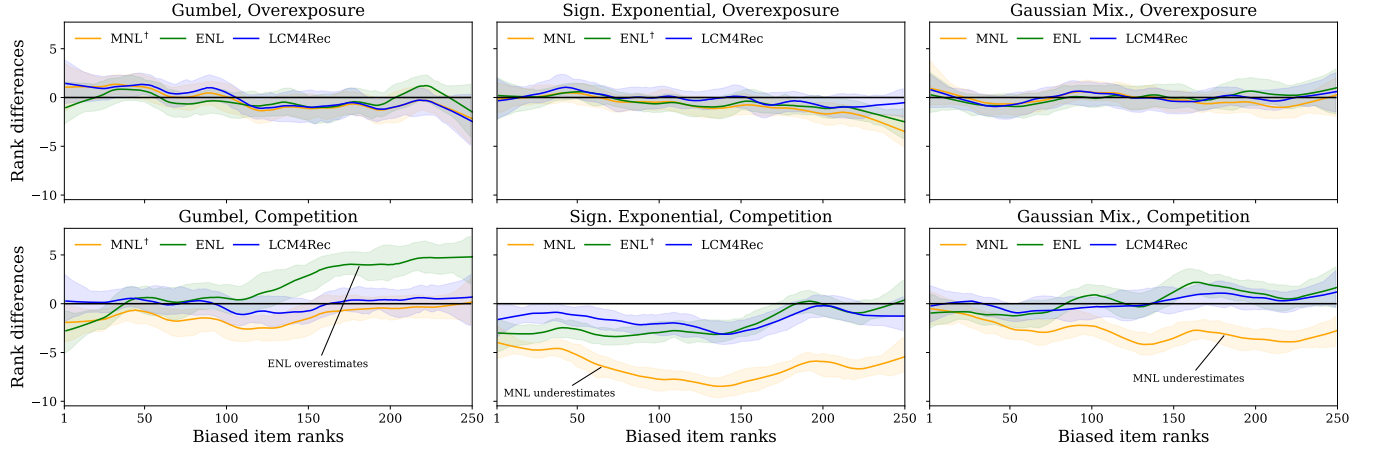


Figure 2: Exposure bias w.r.t. items' mean true ranks. Regression curves and CIs based on LOWESS [13] and bootstrapping.

meaningful. Surprisingly, LCM4Rec is less affected by bias than both other models, even when their assumed error distribution is correct. Nevertheless, due to the variance in these results, we conclude that all models appear similarly robust to overexposure.

The effect of competition results in much larger differences in Table 4. The ENL model under Gumbel and the MNL model under signed exponential and Gaussian mixture are heavily affected by bias. Figure 2 reveals that more preferred items are affected the most, the bias of the conservative ENL places these items lower in the ranking, while the non-conservative MNL places them higher. Again, LCM4Rec is the least affected model across all distributions, surprisingly, even when the parametric models match the true distribution. We speculate that inferring the error distribution and user preferences simultaneously leads to more robustness.

To conclude, we answer RQ3 as follows: Whilst all three choice models appear robust to exposure bias from overexposure, the parametric models can be affected exposure bias from competition depending on what the users' choice model is. In contrast, our results indicate LCM4Rec is the only model that is robust to both types of bias regardless of what the true users' choice model is.

6 Conclusion

Our work concerns the optimization of multivariate choice models that predict how users choose from a set of recommended items. Specifically, we addressed the limitation that existing parametric choice models require a priori assumptions about the error distribution that characterizes how users make choices; existing parametric choice models learn user preferences but not *how* choices are made. Moreover, our experimental results reveal that when their assumptions do not match the true user behavior, they return inaccurate choice probabilities and are susceptible to exposure bias.

In response, we propose LCM4Rec, the first non-parametric multivariate model for recommendation. In contrast with parametric models, LCM4Rec simultaneously infers the most likely error distribution *and* user preferences. Thereby, it both learns *what* users prefer and *how* they choose. It optimizes a kernel density estimator that can approximate any distribution to find the most likely error distribution underlying the users' choice model. As a result, LCM4Rec

alleviates the need to make assumptions about the correct users' choice model. Our experimental results show that LCM4Rec successfully recovers the correct choice model from observed user interactions. Furthermore, our results indicate that—across all tested users' choice models—only LCM4Rec has consistent competitive predictive performance and is robust to exposure bias coming from competition. Therefore, we conclude that LCM4Rec provides the most robust way to learn user preferences from their observed choices and model users' choice behavior, while avoiding a priori assumptions about user behavior.

Our work presents promising directions for future work: LCM4Rec assumes errors are IID, a future extension could tackle the non-IID case. Finally, our experiments were limited to synthetic data but LCM4Rec has the potential to find the most likely real-world user choice model from real-world choice set data.

Acknowledgments

This work is supported by the Dutch Research Council (NWO) under grant VI.Veni.222.269. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- [1] Himan Abdollahpouri and Masoud Mansoury. 2020. Multi-sided Exposure Bias in Recommendation. arXiv:2006.15772 [cs.LR] <https://arxiv.org/abs/2006.15772>
- [2] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2020. The Connection Between Popularity Bias, Calibration, and Fairness in Recommendation. In *Proceedings of the 14th ACM Conference on Recommender Systems (Virtual Event, Brazil) (RecSys '20)*. Association for Computing Machinery, New York, NY, USA, 726–731. doi:10.1145/3383313.3418487
- [3] Aydın Alptekinoglu and John H Semple. 2016. The exponential choice model: A new alternative for assortment and price optimization. *Operations Research* 64, 1 (2016), 79–93.
- [4] Moshe Ben-Akiva and Michel Bierlaire. 1999. *Discrete Choice Methods and their Applications to Short Term Travel Decisions*. Springer US, Boston, MA, 5–33. doi:10.1007/978-1-4615-5203-1_2
- [5] Moshe E Ben-Akiva and Steven R Lerman. 1985. *Discrete choice analysis: theory and application to travel demand*. Vol. 9. MIT press, Boston, USA.
- [6] Austin R. Benson, Ravi Kumar, and Andrew Tomkins. 2016. On the Relevance of Irrelevant Alternatives. In *Proceedings of the 25th International Conference on World Wide Web (Montréal, Québec, Canada) (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 963–973. doi:10.1145/2872427.2883025

- [7] Robert C Blattberg and Kenneth J Wisniewski. 1989. Price-induced patterns of competition. *Marketing science* 8, 4 (1989), 291–309.
- [8] Richard A Briesch, Pradeep K Chintagunta, and Rosa L Matzkin. 2010. Nonparametric discrete choice models with unobserved heterogeneity. *Journal of Business & Economic Statistics* 28, 2 (2010), 291–307.
- [9] Sebastian Büscher and Dietmar Bauer. 2024. Weighting strategies for pairwise composite marginal likelihood estimation in case of unbalanced panels and unaccounted autoregressive structure of the errors. *Transportation research part B: methodological* 181 (2024), 102890.
- [10] Jerome R. Busemeyer and Jörg Rieskamp. 2014. Chapter 3: Psychological research and theories on preferential choice. In *Handbook of Choice Modelling*. Edward Elgar Publishing, Cheltenham, UK. doi:10.4337/9781781003152.00008
- [11] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems* 41, 3 (2023), 1–39.
- [12] Khai Chiong, Yu-Wei Hsieh, and Matthew Shum. 2017. Counterfactual Evaluation in Semiparametric Multinomial Choice Models.
- [13] William S Cleveland. 1979. Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association* 74, 368 (1979), 829–836.
- [14] James L Corner and Craig W Kirkwood. 1991. Decision analysis applications in the operations research literature, 1970–1989. *Operations Research* 39, 2 (1991), 206–219.
- [15] Stephen R. Cosslett. 1983. Distribution-Free Maximum Likelihood Estimator of the Binary Choice Model. *Econometrica* 51, 3 (1983), 765–782. <http://www.jstor.org/stable/1912157>
- [16] Abhisek Dash, Abhijnan Chakraborty, Saptarshi Ghosh, Animesh Mukherjee, and Krishna P. Gummadi. 2021. When the Umpire is also a Player: Bias in Private Label Product Recommendations on E-commerce Marketplaces. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 873–884. doi:10.1145/3442188.3445944
- [17] James Durbin and Siem Jan Koopman. 1997. Monte Carlo maximum likelihood estimation for non-Gaussian state space models. *Biometrika* 84, 3 (1997), 669–684.
- [18] Simen Eide, David S Leslie, Arnoldo Frigessi, Joakim Rishaug, Helge Jenssen, and Sofie Verrewaere. 2021. FINN.no Slates Dataset: A new Sequential Dataset Logging Interactions, all Viewed Items and Click Responses/No-Click for Recommender Systems Research. In *Proceedings of the 15th ACM Conference on Recommender Systems* (Amsterdam, Netherlands) (RecSys '21). Association for Computing Machinery, New York, NY, USA, 556–558. doi:10.1145/3460231.3474607
- [19] Andres Ferraro, Dmitry Bogdanov, Xavier Serra, and Jason Yoon. 2019. Artist and style exposure bias in collaborative filtering based music recommendations. arXiv:1911.04827 [cs.LG]. <https://arxiv.org/abs/1911.04827>
- [20] Paula G. Duran, Pere Gilabert, Santi Seguí, and Jordi Vitrià. 2024. Overcoming diverse undesired effects in recommender systems: A deontological approach. *ACM Transactions on Intelligent Systems and Technology* 15, 4 (2024), 1–23.
- [21] Milad Haghani, Michiel CJ Bliemer, and David A Hensher. 2021. The landscape of econometric discrete choice modelling research. *Journal of choice modelling* 40 (2021), 100303.
- [22] Keisuke Hirano. 2002. Semiparametric Bayesian inference in autoregressive panel data models. *Econometrica* 70, 2 (2002), 781–799.
- [23] Joel L. Horowitz. 1992. A Smoothed Maximum Score Estimator for the Binary Response Model. *Econometrica* 60, 3 (1992), 505–531. <http://www.jstor.org/stable/2951582>
- [24] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *2008 Eighth IEEE International Conference on Data Mining*. IEEE, Piscataway, NJ, 263–272. doi:10.1109/ICDM.2008.22
- [25] Hidehiko Ichimura and T Scott Thompson. 1998. Maximum likelihood estimation of a binary choice model with random coefficients of unknown distribution. *Journal of Econometrics* 86, 2 (1998), 269–295.
- [26] Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. 2019. Degenerate Feedback Loops in Recommender Systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (Honolulu, HI, USA) (AIES '19). Association for Computing Machinery, New York, NY, USA, 383–390. doi:10.1145/3306618.3314288
- [27] Ho Yin Kan, Dennis Wong, and Keith Chau. 2024. A Personalized Flight Recommender System Based on Link Prediction in Aviation Data. *IEEE Access* 12 (2024), 36961–36973. doi:10.1109/ACCESS.2024.3369487
- [28] Roger W. Klein and Richard H. Spady. 1993. An Efficient Semiparametric Estimator for Binary Response Models. *Econometrica* 61, 2 (1993), 387–421. <http://www.jstor.org/stable/2951556>
- [29] Thorsten Krause, Alina Deriyeva, Jan H. Beinke, Gerrit Y. Bartels, and Oliver Thomas. 2024. Mitigating Exposure Bias in Recommender Systems—A Comparative Analysis of Discrete Choice Models. *ACM Trans. Recomm. Syst.* 3, 2, Article 19 (Nov. 2024), 37 pages. doi:10.1145/3641291
- [30] Jiayu Li, Aixun Sun, Weizhi Ma, Peijie Sun, and Min Zhang. 2024. Right Tool, Right Job: Recommendation for Repeat and Exploration Consumption in Food Delivery. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (RecSys '24). Association for Computing Machinery, New York, NY, USA, 643–653. doi:10.1145/3640457.3688119
- [31] Stan Lipovetsky and Michael Conklin. 2014. Finding items cannibalization and synergy by BWS data. *Journal of choice modelling* 12 (2014), 1–9.
- [32] Charles F Manski. 1975. Maximum score estimation of the stochastic utility model of choice. *Journal of econometrics* 3, 3 (1975), 205–228.
- [33] Masoud Mansoury, Himan Abdollahpour, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback Loop and Bias Amplification in Recommender Systems. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (Virtual Event, Ireland) (CIKM '20). Association for Computing Machinery, New York, NY, USA, 2145–2148. doi:10.1145/3340531.3412152
- [34] Rosa L. Matzkin. 1992. Nonparametric and Distribution-Free Estimation of the Binary Threshold Crossing and the Binary Choice Models. *Econometrica* 60, 2 (1992), 239–270. <http://www.jstor.org/stable/2951596>
- [35] Aleksandr Vladimirovich Petrov and Craig Macdonald. 2023. gSASRec: Reducing Overconfidence in Sequential Recommendation Trained with Negative Sampling. In *Proceedings of the 17th ACM Conference on Recommender Systems* (Singapore, Singapore) (RecSys '23). Association for Computing Machinery, New York, NY, USA, 116–128. doi:10.1145/3604915.3608783
- [36] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) (UAI '09). AUAI Press, Arlington, Virginia, USA, 452–461.
- [37] Harald Steck, Linas Baltrunas, Ehtsham Elahi, Dawen Liang, Yves Raimond, and Justin Basilico. 2021. Deep learning for recommender systems: A Netflix case study. *AI Magazine* 42, 3 (2021), 7–18.
- [38] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (Beijing, China) (CIKM '19). Association for Computing Machinery, New York, NY, USA, 1441–1450. doi:10.1145/3357384.3357895
- [39] Pietro Tebaldi, Alexander Torgovitsky, and Hanbin Yang. 2023. Nonparametric estimates of demand in the california health insurance exchange. *Econometrica* 91, 1 (2023), 107–146.
- [40] Kenneth E. Train. 2009. *Discrete Choice Methods with Simulation* (2nd ed.). Number 9780521766555 in Cambridge Books. Cambridge University Press, Cambridge, UK.
- [41] Qi Wan, Xiangnan He, Xiang Wang, Jiancan Wu, Wei Guo, and Ruiming Tang. 2022. Cross Pairwise Ranking for Unbiased Item Recommendation. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France) (WWW '22). Association for Computing Machinery, New York, NY, USA, 2370–2378. doi:10.1145/3485447.3512010
- [42] Shuang-Hong Yang, Bo Long, Alexander J. Smola, Hongyuan Zha, and Zhaohui Zheng. 2011. Collaborative competitive filtering: learning recommender using context of user choice. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Beijing, China) (SIGIR '11). Association for Computing Machinery, New York, NY, USA, 295–304. doi:10.1145/2009916.2009959