

Distributional Off-Policy Evaluation for Slate Recommendations

Shreyas Chaudhari¹ David Arbour², Georgios Theodorou², Nikos Vlassis²

¹University of Massachusetts Amherst

²Adobe Research

schaudhari@cs.umass.edu, {arbour,theochar,vlassis}@adobe.com

Abstract

Recommendation strategies are typically evaluated by using previously logged data, employing off-policy evaluation methods to estimate their expected performance. However, for strategies that present users with slates of multiple items, the resulting combinatorial action space renders many of these methods impractical. Prior work has developed estimators that leverage the structure in slates to estimate the expected off-policy performance, but the estimation of the entire performance *distribution* remains elusive. Estimating the complete distribution allows for a more comprehensive evaluation of recommendation strategies, particularly along the axes of risk and fairness that employ metrics computable from the distribution. In this paper, we propose an estimator for the complete off-policy performance distribution for slates and establish conditions under which the estimator is unbiased and consistent. This builds upon prior work on off-policy evaluation for slates and off-policy distribution estimation in reinforcement learning. We validate the efficacy of our method empirically on synthetic data as well as on a slate recommendation simulator constructed from real-world data (MovieLens-20M). Our results show a significant reduction in estimation variance and improved sample efficiency over prior work across a range of slate structures.

Introduction

Recommendation services are ubiquitous throughout industry (Bobadilla et al. 2013; Lu et al. 2015). A common variant of recommendation consists of suggesting multiple items to a user simultaneously, often termed recommendation *slates*, where each position, (a.k.a., a *slot*), can take multiple possible values (Sarwar et al. 2000). For example, webpage layouts of news or streaming services have separate slots for each category of content where each slot can display any of the items from the category for that slot. The items are suggested to the user based on a recommendation strategy, called a *policy* and the user response is encoded into a *reward* (Li et al. 2010). A crucial problem for selection and improvement of recommendation strategies is to evaluate the efficacy of a slate policy by estimating the expected reward of that policy.

One of the simplest and most effective approaches to policy evaluation, often employed in industrial settings, is A/B

testing (Gomez-Uribe and Hunt 2015; Kohavi and Longbotham 2017; Feitelson, Frachtenberg, and Beck 2013). This involves randomly assigning users to receive the item recommended by one of two candidate policies, and the relative performance of each policy is directly measured. However, A/B testing involves deploying the new policy online, which may be infeasible in many settings due to practical or ethical considerations. As a result, it is often necessary to employ offline *off-policy evaluation*, in which interaction data collected (offline) from previously deployed policies (off-policy) is used to estimate statistics of the expected performance, risk and other metrics of interest for a new target policy, without actually deploying it online. This ensures that policies with undesirable outcomes are not deployed online.

A large amount of literature addresses the problem of off-policy evaluation in the non-slate setting (Dudík et al. 2014; Wang, Agarwal, and Dudík 2017; Thomas, Theodorou, and Ghavamzadeh 2015), where the majority of methods rely on some version of importance sampling (Horvitz and Thompson 1952). Applied to the slate setting, these methods result in very large importance weights that result in high variance estimates due to the combinatorially large action space on which slate policies operate. Recent work addresses this deficiency, by introducing methods that leverage the structure in slate actions and rewards to address the high variance. Swaminathan et al. (2017); Vlassis et al. (2021) leverage reward additivity across slot actions to propose estimators for estimation of the expected reward of target slate policies. McInerney et al. (2020) propose a Markov structure to the slate rewards for sequential interactions.

All the aforementioned methods provide solutions for estimating the *expected reward* for a target policy. However, in many scenarios where recommendation systems are used, such as those with large financial stakes and in healthcare applications, practitioners are concerned with evaluation metrics such as the behavior of the policy at extreme quantiles and the expected performance at a given risk tolerance (CVaR). These quantities need the full *reward distribution* for estimation, which renders prior work which estimates the expected reward inapplicable. A notable exception is Chandak et al. (2021) who provide a method for off-policy estimation of the target policy’s cumulative reward distribution using ideas similar to importance sampling, allowing for the computation of various metrics of interest, but the

estimator is intractable in the slate setting.

In this work, we propose slate universal off-policy evaluation (SUnO), a method that allows for off-policy estimation of target reward distribution (Theorem 3) for slate recommendation policies. SUnO applies the core ideas from the universal off-policy evaluation (UnO) method (Chandak et al. 2021) to the slate setting, leveraging an additive decomposition of the conditional reward distribution. This makes it possible to perform off-policy estimation in structured high dimensional action spaces without incurring prohibitively high estimation variance. We highlight how the estimator can readily be adapted to other generalized decompositions of reward while continuing to be unbiased. Finally, we provide an empirical evaluation comparing against UnO, where the proposed estimator shows significant variance reduction, improved sample efficiency, and robust performance even when the conditions for unbiasedness of the estimator are not met.

The main contributions of our work are:

- We propose an unbiased estimator for the off-policy reward distribution for slate recommendations under an additively decomposable reward distribution, generalizing prior results for slate off-policy evaluation to the distributional setting.
- We theoretically demonstrate how the estimator readily generalizes to slate rewards that do not decompose additively over slots.
- We empirically demonstrate the efficacy of the proposed estimator on slate simulators using synthetic as well as real world data, on a range of slate reward structures.

Background and Notation

We first formulate the slate recommendation system as a contextual bandit with a combinatorial action space. Each slate action has K dimensions where each dimension is a slot-level action. The user-bandit interaction results in a random tuple (X, A, R) at each step, where $X \sim d_X(\cdot)$ is the user context, A is the slate action generated by the recommendation strategy where $A = [A^k]_{k=1}^K$ is composed of K slot-level actions, and $R \sim d_R(\cdot | A, X)$ is the scalar slate-level reward. Since the rewards are observed only at the slate level and not at a per-slot level, we use reward and slate reward interchangeably. Each slot-level action can take upto N candidate values, leading to a combinatorially large action space of the order $\binom{N}{K}$.

A *logging* policy $\mu(A | X) = \Pr(A | X)$ recommends slate actions conditioned on user context X is deployed online to collect a dataset for offline evaluation. The offline dataset consists of n i.i.d. samples $D_n = \{(X_i, A_i, R_i)\}_{i=1}^n$, generated by the user-bandit interaction. We focus on the case where μ is a factored policy, that is,

$$\mu(A | X) = \prod_{k=1}^K \mu_k(A^k | X)$$

where K is the number of slots. Data collection with factored uniform logging policies is standard in practice (Swaminathan et al. 2017). *Off-policy evaluation* is the task

of utilizing data D_n logged using a policy μ , to evaluate a target policy π by computing evaluation metrics from the target reward under π . Standard methods focus on the estimation of the expected reward under the target policy. In this work, our focus is on the estimation of quantities that go beyond just the expected target reward by estimating the reward *distribution*.

Throughout the paper, the sample estimates of any quantity y are denoted by \hat{y}_n where the subscript indicates the number (n) of data points used for estimation. For instance, the cumulative reward distribution observed for a policy π is denoted by $F^\pi(\nu)$. The sample estimate of the distribution will be denoted by $\hat{F}_n^\pi(\nu)$.

Related Work on Off-Policy Evaluation

Importance sampling (IS) (Horvitz and Thompson 1952; Sutton and Barto 2018), also known as inverse propensity scoring (IPS) (Dudík et al. 2014), provides a technique for unbiased estimation of expected target reward. However, it suffers from high variance in large action spaces. There are numerous extensions to IS for variance reduction (Dudík, Langford, and Li 2011; Thomas 2015; Kallus and Uehara 2019). The IS estimator and all methods derived from it rely on a standard common-support assumption. We use a weaker form of this assumption that requires support at the slot level instead of the entire slate.

Assumption 1 (Common Support). *The set D_n contains i.i.d. tuples generated using μ , such that for some (unknown) $\epsilon > 0$, $\mu_k(A^k | X) < \epsilon \implies \pi(A^k | X) = 0; \forall k, X, A$.*

Applying the methods to slates, (Swaminathan et al. 2017) assume additivity of slate rewards as a way to reduce the estimation variance of the expected reward of the target policy. Further variance reduction is obtained by using control variates (Vlassis et al. 2021). McInerney et al. (2020) assume a Markov structure to the slate rewards during sequential item interaction and propose an estimator for the expected target reward. We refer the reader to these papers for additional references on slate recommendations.

These off-policy estimators, along with most others in the literature, provide estimates of the *expected* reward of the target policy (Li et al. 2018). However, the expected value is usually not sufficient for comprehensive off-policy analysis, particularly in the case of risk assessment that is crucial for recommendation systems (Shani and Gunawardana 2011). Additional metrics of interest, often those computable from the whole *reward distribution* are necessary in practice (Keramati et al. 2020; Altschuler, Brunel, and Malek 2019). For example, metrics like value at risk are used risk analysis of a new recommendation strategy. To that end, work on universal off-policy estimation (UnO) (Chandak et al. 2021) uses ideas motivated by importance sampling to estimate the whole cumulative distribution of the reward under the target policy. However, in combinatorially large action spaces, as with the slate problems we consider here, the UnO estimator can incur prohibitive variance. Our proposed estimator utilizes possible structure in slate rewards to circumvent this issue.

Structure in Slate Rewards

The combinatorial action space of slates becomes a key challenge for most general methods like importance sampling (IS) for off-policy evaluation (Wang, Agarwal, and Dudik 2017). The generality of the approach results in them not fully leveraging the structure present in slate rewards (Sunehag et al. 2015) and thus general IS-based approaches frequently suffer from high variance. Prior work leverages user behavior patterns while interacting with slates, which are encoded into structured rewards (e.g., time spent, items clicked, etc.). Some examples of structure in slate rewards are Markov structure for observed slot-level rewards (McInerney et al. 2020), the dependence of slate reward only on the selected slot (Ie et al. 2019), and unobserved slot-level rewards with an additively decomposable slate reward (Swaminathan et al. 2017; Vlassis et al. 2021).

The last one is of particular interest, where the *additivity of expected reward* (Cesa-Bianchi and Lugosi 2012) posits that the conditional mean slate-level reward decomposes additively as the sum of (arbitrary) slot-level latent functions, i.e., $\mathbb{E}[R | A, X] = \sum_{k=1}^K \phi_k(A_k, X)$. This has been leveraged to obtain a significant reduction in estimation variance for off-policy evaluation (Swaminathan et al. 2017; Vlassis et al. 2021). This decomposition captures the individual effects of each slot. It may readily be generalized to capture non-additive joint effects of more than one slot action, for example, to capture the effects of *pairs* of slot-actions, one may consider the decomposition:

$$\mathbb{E}[R | A, X] = \sum_{k=1}^K \sum_{j=k}^K \phi_{jk}(A_k, A_j, X). \quad (1)$$

It may further be generalized to capture the combined effects of m -slots. Note that in the most general case, for $m = K$, the reward does not permit any decomposition over slots.

Analogous to the above structural conditions, we posit a condition that allows us to perform *consistent and unbiased* estimation of the target off-policy *distribution*.

Assumption 2 (Additive CDF). *There exists an additive decomposition of the conditional cumulative density function (CDF) of the slate reward as the sum of (arbitrary) slot-level latent functions:*

$$F_R(\nu) = \sum_{k=1}^K \psi_k(A_k, X, \nu), \quad \forall \nu$$

where $F_R(\nu) := \Pr(R \leq \nu | A, X)$.

The slot-level rewards, if any, are unobserved. The condition just assumes that an additive decomposition *exists* and does not require knowledge of the constituent slot-level functions. We demonstrate empirically that this condition is often a close approximation for real-world data and that our estimator performs better than more general methods even when this condition happens to be an inexact approximation, i.e., when a perfect decomposition does not exist.

This decomposition may also be generalized to capture the combined effects of m -slots. In line with prior work (Wen, Kveton, and Ashkan 2015; Kale, Reyzin, and

Schapire 2010; Kveton et al. 2015; Swaminathan et al. 2017; Vlassis et al. 2021; Ie et al. 2019), we focus our analysis for the case $m = 1$, which proves to be effective in practice as corroborated by empirical analysis. We additionally provide derivations for how the estimator can readily generalize to cases where $m > 1$, along with a theoretical analysis of its properties.

It is worth noting that an additively decomposable reward CDF always implies an additive expected reward by definition and the former often serves to be a close approximation when the latter holds. This is helpful since a commonly used metric for evaluating the performance of slates, the normalized discounted cumulative gain (nDCG) (Burgess et al. 2005), is an additively decomposable metric, and it has been used in the past for defining the slate reward (Järvelin and Kekäläinen 2017; Swaminathan et al. 2017).

Slate Universal Off-Policy Evaluation

We will now turn to off-policy evaluation in slates as an off-policy reward distribution estimation task. The core idea builds upon the framework of Chandak et al. (2021) who use importance weights ρ in an estimator of the reward CDF of a target policy from logged data $D_n \sim \mu$. In the case of slates, the the importance weight ρ comprises of probability ratios over all slot actions. The most direct approach for defining ρ in the case of a factored logging policy μ is to consider a formulation analogous to importance sampling by taking the **product** of the slot-level probabilities (Equation (2)). This approach will be plagued by high variance when the size of the slate K is large. To remedy this, our proposed algorithm SUnO utilizes the structure in slates provided by Assumption 2 wherein the CDF of the slate level reward admits an additive decomposition. In place of ρ , we define an importance weight G (Equation (2)) that is a **sum** of slot-density ratios.

$$\begin{aligned} \rho &= \frac{\pi(A|X)}{\mu(A|X)} = \prod_{k=1}^K \frac{\pi(A^k | X)}{\mu_k(A^k | X)} \\ G &= \sum_{k=1}^K \left(\frac{\pi(A^k | X)}{\mu_k(A^k | X)} - 1 \right) + 1 \end{aligned} \quad (2)$$

The estimator for the target distribution counts the number of samples for which the reward is less than a threshold ν and reweights that count with importance weights to be reflective of the counts under the target policy π . In expectation, the counts reflect the probabilities of obtaining a reward less than ν , providing the value of the reward CDF at ν , denoted by $F^\pi(\nu)$. Use of the importance weight G in place of ρ results in significantly lower variance in estimation and improved effective sample size, while keeping the estimator unbiased. It is easy to confirm that under a factored logging policy, $\mathbb{E}_\mu[G] = 1$. Below we prove that this importance weight allows for a change of distribution of the expected slate reward and can be used for estimating the target reward CDF. Our main result is the following:

Theorem 3. *Let R be a real-valued random variable that denotes the slate reward and admits an additive decompo-*

Algorithm 1: SUnO(ν)

Input: $\pi, \mu, \nu, \{(X_i, A_i, R_i)\}_{i=1}^n \sim D_n$

Output: $\hat{F}_n^\pi(\nu)$

```

1:  $s_\nu = 0$                                      {Initialize counter and iterate over the logged dataset}
2: for  $i = 1, 2, \dots, n$  do
3:    $G_i \leftarrow 1 - K + \sum_{k=1}^K \frac{\pi(A_i^k | X_i)}{\mu_k(A_i^k | X_i)}$    {Compute the importance weight (Equation (2))}
4:    $s_\nu \leftarrow s_\nu + \mathbf{1}\{R_i \leq \nu\} G_i$    {Add to counter if reward is less than  $\nu$ }
5: end for
6: return  $\hat{F}_n^\pi(\nu) = s_\nu/n$    {Normalize and return counter}

```

sition of its conditional cumulative distribution $F_R(\nu)$ (Assumption 2). Under a factored μ and Assumption 2,

$$F^\pi(\nu) = \mathbb{E}_\mu[G \cdot \mathbf{1}\{R \leq \nu\}], \forall \nu$$

Thus, a weighted expectation of the indicator function, with weights given by G , gives the target CDF. Based on this result, we propose the following sample estimator for $F^\pi(\nu)$ that uses data $D_n \sim \mu$,

$$\hat{F}_n^\pi(\nu) := \frac{1}{n} \sum_{i=1}^n G_i \mathbf{1}\{R_i \leq \nu\}, \forall \nu$$

This estimator, called *slate universal off-policy estimator* (SUnO) is outlined in Algorithm 1. In the following result, we establish that SUnO leverages the additive structure in Assumption 2 to obtain an unbiased and pointwise consistent estimate of the CDF of the target policy. The proofs of both results may be found in the Appendix.

Theorem 4. Under Assumption 2, $\hat{F}_n^\pi(\nu)$ is an unbiased and pointwise consistent estimator of $F^\pi(\nu)$.

It is important to note that analogous to Swaminathan et al. (2017) our estimator does *not* require knowledge of the specific functions (ψ_k 's) in the decomposition of the conditional CDF in Assumption 2; it only assumes the *existence* of a set of such latent functions, and a corresponding additive decomposition of the conditional CDF, to attain unbiased estimation. Even in cases where the assumption is not satisfied, our method (Algorithm 1) performs robustly, as we demonstrate in our experiments.

The estimated target CDF can be used to compute metrics of interest as functions of the CDF (for example, mean, variance, VaR, CVaR, etc.). Some of these metrics are non-linear functions of the CDF (VaR, CVaR) and thus their sample estimates would be biased estimators. This is to be expected (Chandak et al. 2021). Metrics that are linear functions of the CDF however have unbiased sample estimators. Thus, an unbiased target CDF estimator serves to be a “one-shot” solution for most metrics of interest, though unbiasedness holds only for certain metrics. We demonstrate the estimation of some of these metrics in our empirical analysis.

Properties of SUnO

Variance: The estimator enjoys significantly low variance for target estimation. The key factor is that the estimator uses

importance weights that are a *sum* of slot level density ratios as opposed to a *product* as is used UnO (Chandak et al. 2019). Particularly in the slate setting, the latter methods suffer from enormous variance and reduced effective sample size, as we demonstrate empirically.

Consider the worst-case variance of the two estimators. From Assumption 1 we have $0 \leq \frac{\pi(A^k | X)}{\mu_k(A^k | X)} \leq \frac{1}{\epsilon}$, which implies

$$\text{Var}(\text{UnO}) = O\left(\frac{1}{\epsilon^K}\right); \quad \text{Var}(\text{SUnO}) = O\left(\frac{K}{\epsilon}\right)$$

Thus the worst-case variance of SUnO grows linearly with an increase in the size of the slate (K) while that of UnO grows exponentially.

Generalization to m -slot reward decomposition As noted earlier, the structural assumptions in slate rewards may be generalized to account for the joint effects of multiple slots. The proposed estimator readily applies to such generalizations. For instance, consider the case where the conditional reward CDF decomposes into terms composed of m slot-actions.

$$F_R(\nu) := \sum_{1 \leq k_1 < k_2, \dots, < k_m \leq K} \psi_{k_{1:m}}(A^{k_{1:m}}, X, \nu) \quad (3)$$

where $k_{1:m}$ is used as shorthand to denote the indices k_1, k_2, \dots, k_m . It can be seen that this decomposition consists of $\binom{K}{m}$ terms. The importance weight from Equation 2 with a minor change provides an *unbiased* and *consistent* off-policy estimator for this reward structure. Define

$$G_m = \sum_{1 \leq k_1 < k_2, \dots, < k_m \leq K} \left(\prod_{i=1}^m \frac{\pi(A^{k_i} | X)}{\mu_{k_i}(A^{k_i} | X)} - 1 \right) + 1 \quad (4)$$

to be the importance weight. With a derivation similar to Theorem 3 we show the following result,

Corollary 5. Let R be a random variable that denotes the slate reward, and permits a decomposition into (latent) functions of m slots as in Equation 3. Under a factored μ , we have

$$F^\pi(\nu) = \mathbb{E}_\mu[G_m \cdot \mathbf{1}\{R \leq \nu\}], \forall \nu$$

The proof of the result may be found in the Appendix. The derivation highlights how the result may be extended to other forms of decomposition of the slate reward. Note

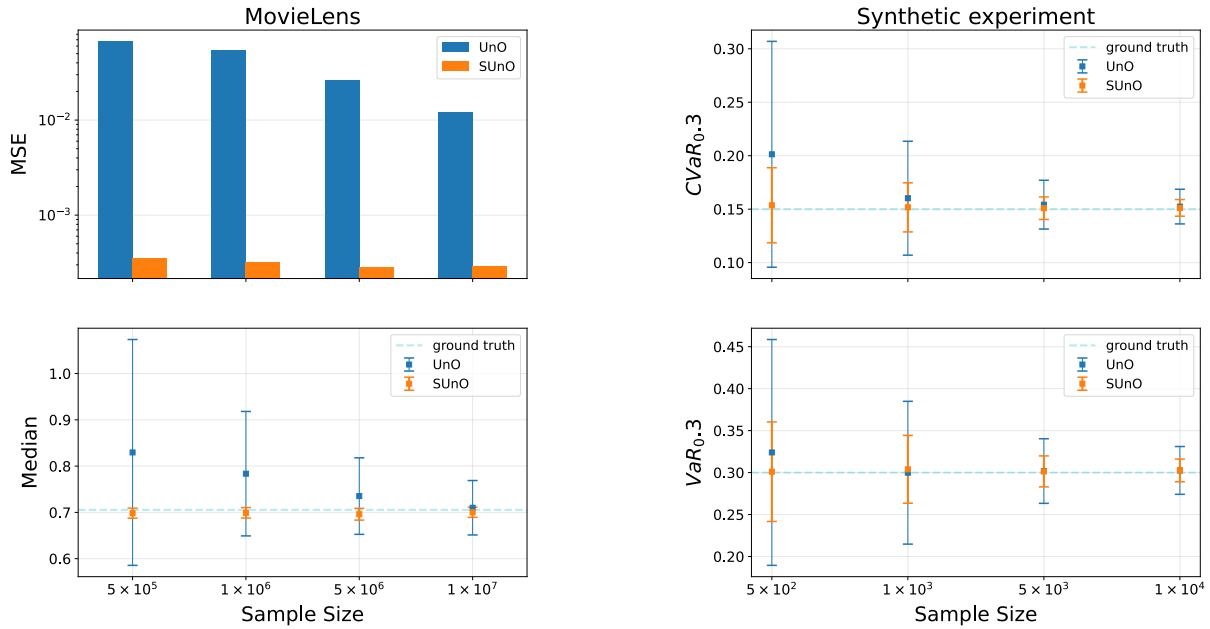


Figure 1: (left) MovieLens-20M: (top) MSE for mean computed from the estimated target CDF for increasing sample sizes. SUnO performs significantly better in terms of bias and variance compared to UnO. The same follows for median estimation (bottom) where it demonstrates much better sample efficiency and lower estimation variance as seen by the error bars. (right) Synthetic Experiment: Estimates of $CVaR_{0.3}$ (top) and $VaR_{0.3}$ (bottom) computed from the estimated target CDF. In this setting where Assumption 2 is satisfied, SUnO performs better than UnO in terms of estimation variance, sample efficiency, and estimation accuracy as expected.

that when $m = K$, the importance weight reduces to $G = \prod_{i=1}^K Y_{k+i}$ which is the same as ρ as used in UnO. This is because at $m = K$ the reward permits no decomposition. The proposed estimator may then be interpreted as a generalization of UnO to various reward decompositions, where UnO forms a special case wherein the reward does not permit any decomposition. We now empirically study the case where $m = 1$, in line with prior work. The reward decomposition at $m = 1$ proves to be a sufficiently accurate approximation to real-world data as we demonstrate in the empirical section that follows.

Empirical Analysis

We investigate the following questions in the empirical analysis:

RQ1: Does the estimator have low estimation variance and high sample efficiency when the slate rewards have an additive structure?

RQ2: Does the method accurately estimate the off-policy distribution, and metrics from it?

RQ3: Does the method apply to settings where conditions for unbiasedness of the estimator do not hold?

To that end, we evaluate the performance of SUnO in a range of slate recommendation settings, comparing its performance against UnO (Chandak et al. 2021), a general estimator that does not make any structural assumptions. Better performance is defined as having lower estimation error and variance, along with improvement in sample efficiency.

We begin by evaluating the estimators on synthetic data that follows the additive CDF structure to corroborate theoretical results (RQ1). We then proceed to real-world data experiments and use the additively decomposable metric nDCG (Swaminathan et al. 2017) as the slate reward (RQ2, RQ3). We test our estimator on a publicly available dataset - MovieLens-20M (Harper and Konstan 2015) - and on a semi-synthetic slate simulator - Open Bandit Pipeline (Saito et al. 2020).

We develop a procedure to *construct a slate simulator* from ratings datasets like MovieLens. We evaluate SUnO in settings where the slate reward does not by construction satisfy either the additive CDF or the additive reward conditions, and see it demonstrate robust performance.

Implementational note: Algorithm 1 outlines the steps for estimating the target CDF at any reward value ν . The reward, in general, takes on continuous real values and implementationally it is not practical to estimate the target CDF at all continuous values of ν . In practice, an empirical estimate of the CDF may be computed at discrete points over the range of rewards. In between those points, the value of the CDF is kept constant. Consequently, we compute the target CDF at evenly spaced points over the range of rewards for both estimators in the experiments that follow. The granularity of this discretization of the domain of the CDF reflects in the granularity of the estimated CDF. To ensure accuracy and relative smoothness in the estimated CDF, we choose a very fine level of discretization relative to the range of reward for

(a) Synthetic experiment

Sample size	0.5×10^3	1×10^3	5×10^3	10×10^3
SUnO	0.131	0.102	0.059	0.049
UnO	0.256	0.191	0.098	0.077

(b) MovieLens

Sample size	0.5×10^6	1×10^6	5×10^6	10×10^6
SUnO	0.169	0.173	0.181	0.184
UnO	0.441	0.410	0.31912	0.252

Table 1: The tables report the average Kolmogorov-Smirnov statistic for the estimated CDFs for (a) Synthetic experiment and (b) MovieLens. The results demonstrate that SUnO estimates the target CDF more accurately with better sample efficiency.

each experiment.

Metrics: We define a few metrics that we utilize in our empirical analysis. Value at Risk (VaR_α) (Rockafellar, Uryasev et al. 2000; Wirch and Hardy 2001) denotes the value of the reward such that the probability of observing rewards greater than that value is $1 - \alpha$. Correspondingly, Conditional Value at Risk (CVaR_α) denotes the expected value of the rewards given that that observed reward is less than VaR_α . These metrics are used for risk assessment of policies and we compute them from the estimated target reward distribution. To evaluate the estimation error in our estimated target distribution, we use the Kolmogorov-Smirnov statistic (Stephens 1974) which measures how well the estimated target distribution matches ground-truth distribution by assessing the largest discrepancy between their cumulative probabilities.

All the experiments have a factored uniform-random logging policy. The error bars denote one standard error. The code is available at: <https://github.com/shreyasc-13/suno>.

Synthetic Experiments

We begin by synthetically generating data where the slate reward permits the additive CDF structure (Assumption 2).

Setup: We consider the non-contextual bandit setting for ease of analysis, and the same may easily be extended to a contextual setting. To construct the data-generating reward distribution, each ψ_k is set to a monotonic non-decreasing function by assigning slices of a sigmoid function to the corresponding ψ_k 's for each (A^k). This manner of construction ensures that the resultant sum of the functions, the CDF, is again a monotonic non-decreasing function. The ψ_k 's are appropriately normalized. For these experiments, we set the number of slots $K = 3$ and the number of actions in each slot to $N = 3$. The target policy is a deterministic policy that chooses one action per slot, where the action for each slot is assigned randomly at the start of the experiment and held constant for all experiments.

Experiment: We compare the performance of the estimators on two fronts:

1. *Goodness-of-fit of CDF:* We report the average Kolmogorov-Smirnov statistic of the estimated target CDFs. The ground truth target CDF is computed by executing the target policy on the simulator.
2. *Tail measures:* We compute the $\text{CVaR}_{0.3}$ and $\text{VaR}_{0.3}$ from the target CDFs estimated by the two estimators.

The experiments are run for different logged data sizes and the results are averaged over 1000 trials.

Results [Table 1, Figure 1]: Since SUnO leverages the additive structure in rewards, it estimates the CDF and the tail

measures with lower variance and estimation error, while being more sample efficient. If a single slot action has a zero probability of occurring under the target policy, the importance weight used by UnO for the entire slate goes to zero since it comprises of the product of slot-level density ratios. This is not the case for SUnO which is thus able to utilize a larger effective sample size and be more sample efficient. Note that with an increase in sample size, both estimators tend to the ground-truth values as they are consistent and unbiased. SUnO has a significantly lower variance in estimation as seen by the error bars in Figure 1.

Real-World Data

In this section, we first introduce a procedure for converting a recommender system ratings dataset (like MovieLens) to a slate recommendation simulator with additive rewards and proceed to evaluate our method on the simulator. The procedure for constructing the simulator follows.

Simulation Setup

1. Learn a user-item preference matrix B along with the user context embedding X . For m users and l item, $B \in \mathbb{R}^{m \times l}$ and $X \in \{0, 1\}^m$. We follow the steps outlined in (Steck 2019) to learn B from rating data. X is a binary vector that encodes user-item interaction history. An alternative method for learning embeddings could be (Elahi and Chandrashekar 2020).
2. To limit our setup to approximately 10k unique users, we trim the set of users to those that have an interaction history of 10 to 15 items.
3. Compute the *ground truth* preference scores for each user by computing the product of a user's context embedding with the preference matrix ($x \cdot B$).
4. To make the simulator tractable, we trim the action set by retaining the top 20 preferred actions per user based on each user's ground truth scores ($N = 20$).
5. For a slate action A , a ranking metric like nDCG can be set as the slate reward R and works well in practice (Vlasis et al. 2021; Swaminathan et al. 2017).

Experiment: First, we set up a slate simulator as described above using the MovieLens-20M dataset to estimate B and X . A uniform random factored logging policy is used for creating the offline dataset for evaluating the estimators. We consider an ϵ -greedy target policy. For each user, it picks the top K preferred actions (one per slot) with probability $1 - N\epsilon$ and a uniform random action from the user's action set with probability ϵ . Here $N = 20$, $K = 5$, $\epsilon = 0.01$ and results are averaged over 50 trials. We analyze:

Sample size	0.5×10^5	1×10^6	5×10^6
SUnO	0.253	0.257	0.269
UnO	0.543	0.541	0.567

Table 2: The table reports the mean squared error for mean computation from the estimated CDFs. Even in settings where the slate reward is not additive, our method continues to perform better than the structure-agnostic estimator.

1. *Goodness-of-fit of CDF*: We report the average Kolmogorov-Smirnov statistic of the estimated CDFs against the ground truth CDF. The ground truth CDF is computed by executing the target policy on the simulator.
2. *Metrics computed from the CDF*: We compute the mean and 0.5-quantile (median) from the estimated CDF.

Results [Table 1, Figure 1]: The experiments demonstrate that although only the additive reward condition is met and not Assumption 2, SUnO estimates the target CDF with fewer samples (Table 1) than UnO. Our estimator has a significantly lower estimation variance for metrics computed from the CDF, as seen by the error bars for median computation and the mean squared error for the expected value computation in Figure 1. Note that the mean squared error (MSE) captures both the bias and variance in estimation.

Non-Additive Reward Structure

Finally, we evaluate the estimators in a setting where neither the additive reward nor the additive CDF conditions are satisfied.

Simulator: We use the Open Bandit Pipeline (OBP) slate bandit simulator (Saito et al. 2020) that uses the synthetic slate reward model described in (Kiyohara et al. 2022). It models higher-order interactions among slot actions and thus does not trivially satisfy Assumption 2 or the additive slate reward structure. We use the *cascade additive reward model* defined in OBP for these experiments.

Experiment: Similar to the MovieLens experiments, we observe the estimation error for the target mean computed from the estimated CDF. A uniform random logging policy is used to generate the offline dataset and the target policy defined here¹ is evaluated. We set $K = 3$, $N = 10$, and the results are averaged over 10 trials.

Results [Table 2]: In this setting, we cannot expect unbiased estimates of the mean from SUnO since the additive CDF condition is required for unbiased estimation of the target CDF. Nonetheless, SUnO continues to perform significantly better in terms of the mean squared error for the mean estimation compared to UnO, which does not make any structural assumptions and is an unbiased estimator in this setting. Here K is set to a relatively small value and a large gap in performance between the two estimators can be expected larger K .

¹https://github.com/st-tech/zr-obp/blob/master/examples/quickstart/synthetic_slate.ipynb

Discussion and Conclusion

We proposed an estimator (SUnO) for off-policy estimation of the target *reward distribution* in slate recommendations modeled as a bandit problem. Under an additively decomposable conditional CDF, the estimator is unbiased and consistent. The proposed estimator leads to significant reduction in estimation variance and an increase in effective sample size as compared to the estimator of Chandak et al. (2021) for the slate setting. We demonstrate estimation gains on synthetic as well as real-world data experiments. The estimator also readily extends to other reward decompositions that capture the joint effects of slot actions.

In future work, variance reduction techniques can be applied for further variance gains. For instance, one may consider a self-normalized version of SUnO that incurs some bias but provides further variance reduction akin to weighted importance sampling (Koller and Friedman 2009). Control variates can also be adopted for variance reduction. For example, by recapitulating the analysis of (Vlassis et al. 2021), one can derive an optimal control variate (w^*) for SUnO at each ν . Further analysis and experiments with such methods are left to future work.

One must note while SUnO provides unbiased estimates for the target CDF and its linear functions under our assumptions, many metrics of interest, such as variance or CVaR, are not linear functions of the CDF. Another direction for future work would be to extend our results to the unbiased estimation of risk functionals like CVaR (Huang et al. 2021). Finally, it would be interesting to develop techniques for discovering the decomposition structure in rewards so that the appropriate unbiased off-policy estimators can be used.

Acknowledgements

The research was supported by and partially conducted at Adobe Research. We are also immensely grateful to the four anonymous reviewers who shared their insights and feedback.

References

- Altschuler, J. M.; Brunel, V.-E.; and Malek, A. 2019. Best Arm Identification for Contaminated Bandits. *J. Mach. Learn. Res.*, 20(91): 1–39.
- Bobadilla, J.; Ortega, F.; Hernando, A.; and Gutiérrez, A. 2013. Recommender systems survey. *Knowledge-based systems*, 46: 109–132.
- Burges, C.; Shaked, T.; Renshaw, E.; Lazier, A.; Deeds, M.; Hamilton, N.; and Hullender, G. 2005. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, 89–96.
- Cesa-Bianchi, N.; and Lugosi, G. 2012. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5): 1404–1422.
- Chandak, Y.; Niekum, S.; da Silva, B.; Learned-Miller, E.; Brunskill, E.; and Thomas, P. S. 2021. Universal off-policy evaluation. *Advances in Neural Information Processing Systems*, 34.

- Chandak, Y.; Theocharous, G.; Kostas, J.; Jordan, S.; and Thomas, P. 2019. Learning action representations for reinforcement learning. In *International conference on machine learning*, 941–950. PMLR.
- Dudík, M.; Erhan, D.; Langford, J.; and Li, L. 2014. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4): 485–511.
- Dudík, M.; Langford, J.; and Li, L. 2011. Doubly robust policy evaluation and learning. *arXiv preprint arXiv:1103.4601*.
- Elahi, E.; and Chandrashekar, A. 2020. Learning representations of hierarchical slates in collaborative filtering. In *Proceedings of the 14th ACM Conference on Recommender Systems*, 703–707.
- Feitelson, D. G.; Frachtenberg, E.; and Beck, K. L. 2013. Development and deployment at facebook. *IEEE Internet Computing*, 17(4): 8–17.
- Gomez-Uribe, C. A.; and Hunt, N. 2015. The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4): 1–19.
- Harper, F. M.; and Konstan, J. A. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4): 1–19.
- Horvitz, D. G.; and Thompson, D. J. 1952. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260): 663–685.
- Huang, A.; Leqi, L.; Lipton, Z.; and Azizzadenesheli, K. 2021. Off-policy risk assessment in contextual bandits. *Advances in Neural Information Processing Systems*, 34: 23714–23726.
- Ie, E.; Jain, V.; Wang, J.; Narvekar, S.; Agarwal, R.; Wu, R.; Cheng, H.-T.; Chandra, T.; and Boutilier, C. 2019. SlateQ: A tractable decomposition for reinforcement learning with recommendation sets.
- Järvelin, K.; and Kekäläinen, J. 2017. IR evaluation methods for retrieving highly relevant documents. In *ACM SIGIR Forum*, volume 51, 243–250. ACM New York, NY, USA.
- Kale, S.; Reyzin, L.; and Schapire, R. E. 2010. Non-stochastic bandit slate problems. *Advances in Neural Information Processing Systems*, 23.
- Kallus, N.; and Uehara, M. 2019. Intrinsically efficient, stable, and bounded off-policy evaluation for reinforcement learning. *Advances in neural information processing systems*, 32.
- Keramati, R.; Dann, C.; Tamkin, A.; and Brunskill, E. 2020. Being optimistic to be conservative: Quickly learning a cvar policy. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 4436–4443.
- Kiyohara, H.; Saito, Y.; Matsui, T.; Narita, Y.; Shimizu, N.; and Yamamoto, Y. 2022. Doubly robust off-policy evaluation for ranking policies under the cascade behavior model. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 487–497.
- Kohavi, R.; and Longbotham, R. 2017. Online Controlled Experiments and A/B Testing. *Encyclopedia of machine learning and data mining*, 7(8): 922–929.
- Koller, D.; and Friedman, N. 2009. *Probabilistic graphical models: principles and techniques*. MIT press.
- Kveton, B.; Wen, Z.; Ashkan, A.; and Szepesvari, C. 2015. Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, 535–543. PMLR.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, 661–670.
- Li, S.; Abbasi-Yadkori, Y.; Kveton, B.; Muthukrishnan, S.; Vinay, V.; and Wen, Z. 2018. Offline evaluation of ranking policies with click models. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1685–1694.
- Lu, J.; Wu, D.; Mao, M.; Wang, W.; and Zhang, G. 2015. Recommender system application developments: a survey. *Decision support systems*, 74: 12–32.
- McInerney, J.; Brost, B.; Chandar, P.; Mehrotra, R.; and Carterette, B. 2020. Counterfactual evaluation of slate recommendations with sequential reward interactions. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1779–1788.
- Rockafellar, R. T.; Uryasev, S.; et al. 2000. Optimization of conditional value-at-risk. *Journal of risk*, 2: 21–42.
- Saito, Y.; Shunsuke, A.; Megumi, M.; and Yusuke, N. 2020. Open Bandit Dataset and Pipeline: Towards Realistic and Reproducible Off-Policy Evaluation. *arXiv preprint arXiv:2008.07146*.
- Sarwar, B.; Karypis, G.; Konstan, J.; and Riedl, J. 2000. Analysis of recommendation algorithms for e-commerce. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 158–167.
- Sen, P. K.; and Singer, J. M. 1994. *Large sample methods in statistics: an introduction with applications*, volume 25. CRC press.
- Shani, G.; and Gunawardana, A. 2011. Evaluating recommendation systems. In *Recommender systems handbook*, 257–297. Springer.
- Steck, H. 2019. Embarrassingly shallow autoencoders for sparse data. In *The World Wide Web Conference*, 3251–3257.
- Stephens, M. A. 1974. EDF statistics for goodness of fit and some comparisons. *Journal of the American statistical Association*, 69(347): 730–737.
- Sunehag, P.; Evans, R.; Dulac-Arnold, G.; Zwols, Y.; Visentin, D.; and Coppin, B. 2015. Deep reinforcement learning with attention for slate markov decision processes with high-dimensional states and actions. *arXiv preprint arXiv:1512.01124*.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.

- Swaminathan, A.; Krishnamurthy, A.; Agarwal, A.; Dudik, M.; Langford, J.; Jose, D.; and Zitouni, I. 2017. Off-policy evaluation for slate recommendation. *Advances in Neural Information Processing Systems*, 30.
- Thomas, P.; Theocharous, G.; and Ghavamzadeh, M. 2015. High-confidence off-policy evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Thomas, P. S. 2015. Safe reinforcement learning.
- Vlassis, N.; Chandrashekar, A.; Amat, F.; and Kallus, N. 2021. Control variates for slate off-policy evaluation. *Advances in Neural Information Processing Systems*, 34.
- Wang, Y.-X.; Agarwal, A.; and Dudik, M. 2017. Optimal and adaptive off-policy evaluation in contextual bandits. In *International Conference on Machine Learning*, 3589–3597. PMLR.
- Wen, Z.; Kveton, B.; and Ashkan, A. 2015. Efficient learning in large-scale combinatorial semi-bandits. In *International Conference on Machine Learning*, 1113–1122. PMLR.
- Wirch, J. L.; and Hardy, M. R. 2001. Distortion risk measures: Coherence and stochastic dominance. In *International congress on insurance: Mathematics and economics*, 15–17.