Offline Recommender System Evaluation under Unobserved Confounding

OLIVIER JEUNEN, ShareChat, UK

BEN LONDON, Amazon Music, USA

Off-Policy Estimation (OPE) methods allow us to learn and evaluate decision-making policies from logged data. This makes them an attractive choice for the offline evaluation of recommender systems, and several recent works have reported successful adoption of OPE methods to this end. An important assumption that makes this work, is the absence of unobserved confounders: random variables that influence both actions and rewards at data collection time. Because the data collection policy is typically under the practitioner's control, the unconfoundedness assumption is often left implicit, and its violations are rarely dealt with in the existing literature.

This work aims to highlight the problems that arise when performing off-policy estimation in the presence of unobserved confounders, specifically focusing on a recommendation use-case. We focus on policy-based estimators, where the logging propensities are learned from logged data. We characterise the statistical bias that arises due to confounding, and show how existing diagnostics are unable to uncover such cases. Because the bias depends directly on the *true* and unobserved logging propensities, it is non-identifiable. As the unconfoundedness assumption is famously untestable, this becomes especially problematic. This paper emphasises this common, yet often overlooked issue. Through synthetic data, we empirically show how naïve propensity estimation under confounding can lead to severely biased metric estimates that are allowed to fly under the radar. We aim to cultivate an awareness among researchers and practitioners of this important problem, and touch upon potential research directions towards mitigating its effects.

ACM Reference Format:

Olivier Jeunen and Ben London. 2023. Offline Recommender System Evaluation under Unobserved Confounding. In CONSEQUENCES Workshop at ACM RecSys (CONSEQUENCES@RecSys'23), September 19, 2023, Singapore. ACM, New York, NY, USA, 6 pages.

1 INTRODUCTION & MOTIVATION

Inferring cause and effect from observational data is not a straightforward task, due to the presence of "*unobserved confounders*" [20]. Indeed, if an unobserved variable exists that influences both the treatment and its outcome, this can easily lead to biased estimates of the treatment effect. In the extreme case where the estimate changes sign, this is known as "Simpson's Paradox" [12]. Bottou et al. describe how to avoid such unpleasantries in counterfactual learning scenarios, by carefully modelling the data-generating process [3]. Jadidinejad et al. describe how an instance of Simpson's Paradox is prevalent to occur in the offline evaluation of recommender systems—as the test data is often influenced by an unknown recommendation policy that takes the role of the confounder [8]. Their proposed solution to this problem involves *Inverse Propensity Score* (IPS) weighting [19, Ch. 9], with propensities that are estimated from logged data. Estimating propensities is common practice in cases where the true propensities are unknown [1, 26], and empirical results indicate that IPS with estimated propensities can even lead to favourable variance [7]. Whether to simplify experimental setup or to deal with missing information in publicly available datasets, these works often make the limiting assumption that the propensities are independent of user or context (see [8, §6.3] or [26, §3.3]). In the very likely case that these assumptions are violated, this implies the presence of unobserved confounders.¹ Existing diagnostics for validating logged bandit feedback cannot detect these issues [15, 17], as we will show formally.

Unobserved confounding is a well-known issue in general off-policy reinforcement learning, and several methods have been proposed to deal with it in the recent literature. They typically leverage additional data (such as interventions [6, 24]

¹An analogous problem occurs in ranking applications when contextual independence is incorrectly assumed for position bias estimates [5].

Manuscript submitted to ACM

or instrumental variables [25]) or make assumptions about the nature of confounding variables [2, 13, 18]. Analogously, instrumental variables have been leveraged to test for the unconfoundedness assumption [4], and other statistical methods have been proposed to assess the sensitivity of results to potential confounders [16]. Nevertheless, in the absence of additional tools, unconfoundedness is a famously *untestable* assumption, rendering its effects especially troublesome. Focusing on the off-policy bandit setting with a guiding example in recommendation, we aim to answer:

"Can we reliably select the optimal policy from a set of competing policies, under unobserved confounding?"

2 OFF-POLICY ESTIMATION IN THE PRESENCE OF UNOBSERVED CONFOUNDERS

Throughout this work, we denote random variables as *X*, with specific instances as $x \in X$. A contextual bandit problem consists of *contexts X* (e.g. user and context features), *actions A* (e.g. item recommendations), and *rewards R* (e.g. clicks, streams, revenue). Rewards are causally influenced by both contexts and actions, as illustrated by the edges in the causal graph shown in Figure 1. A contextual *policy* π determines which actions are selected (or sampled), thereby inducing a probability distribution over *A*, which is often denoted with the shorthand $\pi(a|x) := P(A = a|X = x; \Pi = \pi)$.



Fig. 1. Probabilistic Graphical Model (PGM) for our setup.

A policy's effectiveness is measured by the expected reward obtained when selecting actions according to that policy: $\mathbb{E}_{a\sim\pi}[R]$. In a recommendation application, this value can be estimated by deploying the policy in an online experiment. However, since such experiments are typically costly, and we may have many policies to evaluate, we would rather obtain reward estimates by other means.

Suppose there is an existing deployed policy, called the *logging policy* π_0 , with which we collect a dataset $\mathcal{D}_0 := \{(a_i, r_i)_{i=1}^N\}$. We will assume, as is often the case, that the logging policy, and the contextual covariates it uses, are unobservable, as indicated by the dashed nodes and edges in Figure 1. Our goal is to leverage data logged under π_0 to estimate the expected reward under π -a problem often referred to as *off-policy* estimation [21].

One simple off-policy estimator is the *Direct Method* (DM). DM computes the expected reward under π (Eq. 1) using a model $\widehat{R}^A_{DM}(a)$ that estimates the reward for every available action. Since we assume that contextual covariates are unavailable, the best we can do for $\widehat{R}^A_{DM}(a)$ is to naïvely count the observed rewards for every action (Eq. 2).

$$\widehat{R}_{\text{DM}}(\pi) = \sum_{a \in \mathcal{A}} \widehat{R}^{A}_{\text{DM}}(a)\pi(a). \qquad (1) \qquad \qquad \widehat{R}^{A}_{\text{DM}}(a) = \frac{\sum_{(a_i, r_i) \in \mathcal{D}_0} \mathbb{1}\{a_i = a\} \cdot r_i}{\sum_{(a_i, r_i) \in \mathcal{D}_0} \mathbb{1}\{a_i = a\}}, \qquad (2)$$

Unfortunately, this estimator is biased, for two reasons: (a) it does not take into account the covariates X (i.e., the model is mis-specified), and (b) it ignores the selection bias from the logging policy π_0 , influencing the estimates in Eq. 2.

In theory, we can bypass both the model mis-specification and selection bias problems by leveraging the *ideal* IPS estimator (Eq. 3), which is provably unbiased. Importantly, ideal IPS requires access to both the contextual covariates and the exact action probabilities (*propensities*) under π_0 —which we assume are unavailable. Accordingly, we will adopt the common practice of using estimated logging propensities $\hat{\pi}_0$ for IPS (Eq. 4). As the estimated propensities cannot properly consider all covariates, this leads to unobserved confounding.

$$\widehat{R}_{\text{ideal}-\text{IPS}}(\pi) = \frac{1}{|\mathcal{D}_0|} \sum_{\substack{(x_i, a_i, r_i) \in \mathcal{D}_0 \\ \pi_0(a_i | x_i)}} r_i \frac{\pi(a_i)}{\pi_0(a_i | x_i)}.$$
 (3)
$$\widehat{R}_{\text{estim}-\text{IPS}}(\pi) = \frac{1}{|\mathcal{D}_0|} \sum_{\substack{(a_i, r_i) \in \mathcal{D}_0 \\ \overline{\pi}_0(a_i)}} r_i \frac{\pi(a_i)}{\overline{\pi}_0(a_i)}.$$
 (4)

Using the fact that the ideal IPS estimator is unbiased, we can quantify the bias of the estimated IPS estimator as:

$$\mathbb{E}[\widehat{R}_{\text{estim-IPS}}(\pi)] - \mathbb{E}_{a \sim \pi}[R] = \mathbb{E}[\widehat{R}_{\text{estim-IPS}}(\pi)] - \mathbb{E}[\widehat{R}_{\text{ideal-IPS}}(\pi)] = \mathbb{E}\left[R\pi(A|X)\left(\frac{1}{\widehat{\pi}_0(A)} - \frac{1}{\pi_0(A|X)}\right)\right].$$
(5)

To further illustrate our point, we resort to Pearl's do-calculus framework [20]. What OPE methods wish to estimate is the expected value of the reward given that a new policy *intervenes* on the action distribution. When unobserved confounders are present, this *interventional* quantity is *not* equal to the *observational* quantity we can estimate from logged data: $\mathbb{E}[R|A = a] \neq \mathbb{E}[R|do(A = a)]$. Instead, we would require the "backdoor adjustment" to obtain:

$$\mathbb{E}\left[R|\operatorname{do}(A=a)\right] = \sum_{x \in \mathcal{X}} \mathbb{E}\left[R|A=a, X=x\right].$$
(6)

It should be clear that without access to X, this estimand is non-identifiable, and this problem is not easily solved.

3 EXISTING DIAGNOSTICS FOR LOGGING PROPENSITIES DO NOT UNCOVER CONFOUNDING BIAS

Several diagnostics have been proposed in the literature to detect data quality issues with logged bandit feedback. In particular, they try to uncover cases where the two classical assumptions of the IPS estimator do not hold [15, 17]: (1) either the empirical action frequencies in the data do not match those implied by the logged propensities, or (2) the logging policy does not have full support over the action space. Note that the presence of unobserved confounders does *not* automatically violate these assumptions. As a result, the diagnostics that were proposed will *not* detect confounding bias. Logging propensities can be estimated by empirically counting logged actions, as shown in Eq. 7. In doing so, we obtain unbiased estimates of the true marginal action probabilities. Indeed, $\lim_{N\to\infty} \widehat{\pi_0}(a) = P(A = a|\Pi = \pi_0)$.

Li et al. propose the use of *arithmetic* and *harmonic* mean tests to compare empirical action frequencies with the logging propensities [15]. As we *define* the logging propensities to be equal to the empirical action frequencies, it should be clear that this test will trivially pass. Alternatively, London and Joachims propose to use the average importance weight as a control variate, whose expectation should equal 1 for any target policy π [17]. Here as well, because the marginal propensities are unbiased (Eq. 7), we can show that the control variate remains unbiased as well (Eq. 8).

$$\widehat{\pi}_{0}(a) = \frac{1}{|\mathcal{D}_{0}|} \sum_{(a_{i},r_{i})\in\mathcal{D}_{0}} \mathbb{1}\{a_{i}=a\} \underset{N\to\infty}{=} \mathsf{P}(A=a|\Pi=\pi_{0}) = \sum_{x\in\mathcal{X}}\mathsf{P}(A=a|X=x,\Pi=\pi_{0})\mathsf{P}(X=x).$$
(7)

THEOREM 3.1. The expectation of importance weights equals 1 for any target policy, when logging propensities are estimated using marginal action probabilities:

$$\mathbb{E}_{\substack{X \sim \mathsf{P}(X) \\ \sim \mathsf{P}(A | X = x, \Pi = \pi_0)}} \left[\frac{\pi(a)}{\widehat{\pi}_0(a)} \right] = 1.$$

а

Proof.

$$\mathbb{E}_{\substack{x \sim \mathsf{P}(X) \\ a \sim \mathsf{P}(A|X=x,\Pi=\pi_0)}} \left[\frac{\pi(a)}{\hat{\pi}_0(a)} \right] = \sum_{a \in \mathcal{A}} \sum_{x \in \mathcal{X}} \frac{\pi(a)}{\hat{\pi}_0(a)} \mathsf{P}(A = a | X = x, \Pi = \pi_0) \mathsf{P}(X = x)
= \sum_{a \in \mathcal{A}} \frac{\pi(a)}{\hat{\pi}_0(a)} \sum_{x \in \mathcal{X}} \mathsf{P}(A = a | X = x, \Pi = \pi_0) \mathsf{P}(X = x)
= \sum_{\substack{n \in \mathcal{A} \\ N \to \infty}} \frac{\pi(a)}{a \in \mathcal{A}} \sum_{x \in \mathcal{X}} \mathsf{P}(A = a | X = x, \Pi = \pi_0) \mathsf{P}(X = x)
= \sum_{\substack{n \in \mathcal{A} \\ X \in \mathcal{X}}} \pi(a) \frac{\sum_{x \in \mathcal{X}} \mathsf{P}(A = a | X = x, \Pi = \pi_0) \mathsf{P}(X = x)}{\sum_{x \in \mathcal{X}} \mathsf{P}(A = a | X = x, \Pi = \pi_0) \mathsf{P}(X = x)} = \sum_{a \in \mathcal{A}} \pi(a) = 1 \quad \Box$$

As such, existing diagnostics are unable to detect issues of unobserved confounding. This implies that the selfnormalised IPS (SNIPS) estimator and its extensions that adopt the same control variate to reduce the variance of the IPS estimator, would suffer equally when unobserved confounders are present [11, 23]. CONSEQUENCES@RecSys '23, September 19, 2023, Singapore

Olivier Jeunen and Ben London

$P(X = x_0)$	$P(X = x_1)$	$\mathbb{E}[\mathbf{R} \mathbf{X},\mathbf{A}]$	a_0	<i>a</i> ₁	$\pi_0(\mathbf{a} \mathbf{x})$	a_0	<i>a</i> ₁		$\widehat{\pi}_{0}(\mathbf{a})$
$1 - \alpha$	α	x_0	1.0	0.7	x_0	$1.0-\epsilon$	ϵ	a_0	$(1-\epsilon)\alpha + \epsilon(1-\alpha)$
		<i>x</i> ₁	0.0	0.7	x_1	ϵ	$1.0 - \epsilon$	a_1	$(1-\epsilon)(1-\alpha)+\epsilon\alpha$
(a) Covariate distribution		(b) Reward distribution			(c) Logging policy distribution			(d) Logging propensity estimate	

Table 1. Data distributions for a guiding example that highlights issues with off-policy estimation under unobserved confounding.



Fig. 2. Estimated differences in reward for various estimators. IPS with estimated propensities suffers from unobserved confounding.

4 EMPIRICAL VALIDATION OF THE EFFECTS OF UNOBSERVED CONFOUNDING ON SYNTHETIC DATA

We now describe a simple guiding example, and provide a Jupyter notebook that implements the methods described earlier, on this example. Consider a setting with two possible actions and a binary covariate $X = \{x_0, x_1\}$, following the distribution in Table 1a (parameterised with $\alpha \in [\frac{1}{2}, 1]$). Rewards are Bernoulli-distributed, following Table 1b. The logging policy is contextual, taking a suboptimal action with probability $\epsilon \in [0, 1]$, as shown in Table 1c. We can map this to an intuitive setting: action a_1 is of general appeal to the entire population (i.e. $R \perp X | A = a_1$); whereas action a_0 , on the other hand, is specifically appealing to a more niche user-base (i.e. $\mathbb{E}[R|X = x_0, A = a_0] > \mathbb{E}[R|X = x_0, A = a_1]$, but $P(X = x_0) < P(X = x_1)$). Estimates for logging propensities can be obtained by empirical counting, as in Eq. 7. The expected value for these estimated context-independent propensities is shown in Table 1d.

Naïve propensity estimation methods suffer from confounding bias. We simulate an off-policy estimation setup where we wish to evaluate deterministic policies $\pi_a(a) \equiv 1$. We obtain $N = 2 \cdot 10^6$ samples from the synthetic distribution described in Table 1, and compute the confounded estimate $\widehat{R}_{estim-IPS}$, as well as the unobservable ideal IPS estimate $\widehat{R}_{ideal-IPS}$. We vary both the level of selection bias ϵ (over the x-axis), and the confounding distribution α (over columns) in Fig. 2, where the y-axis shows the estimated difference in rewards from policies π_{a_1} and π_{a_0} . We shade the positive and negative regions in the plot to clearly visualise when an off-policy estimator allows us to *correctly* identify the optimal policy, or when it does not. We observe that the IPS estimator with estimated propensities fails considerably, in that it will incorrectly identify π_{a_0} as the reward-maximising policy. Only when ϵ is sufficiently high (i.e. approaching a uniform logging policy for $\epsilon = 0.5$, and hence no confounding is present), $\widehat{R}_{estim-IPS}$ is able to correctly identify π_{a_1} . This shows that, even in simplified settings, the estimates we obtain from IPS with confounded estimates lead to misleading conclusions. Furthermore, existing diagnostics cannot detect these problems when they occur.

5 CONCLUSIONS & OUTLOOK

Unobserved confounders lead to biased estimates, both for DM- and IPS-based methods. This problem has received considerable attention in the research literature for general offline reinforcement learning use-cases, but the literature dealing with these issues in recommendation settings remains scarce. Our work highlights that this is problematic—especially in cases where propensities are estimated under simplifying independence assumptions. In doing so, we add to the literature identifying problematic practices that might hamper progress in the field [9, 10, 14, 22].

Offline Recommender System Evaluation under Unobserved Confounding

REFERENCES

- [1] Qingyao Ai, Keping Bi, Cheng Luo, Jiafeng Guo, and W. Bruce Croft. 2018. Unbiased Learning to Rank with Unbiased Propensity Estimation. In The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (Ann Arbor, MI, USA) (SIGIR '18). ACM, 385–394. https://doi.org/10.1145/3209978.3209986
- [2] Andrew Bennett, Nathan Kallus, Lihong Li, and Ali Mousavi. 2021. Off-policy Evaluation in Infinite-Horizon Reinforcement Learning with Latent Confounders. In Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130), Arindam Banerjee and Kenji Fukumizu (Eds.). PMLR, 1999–2007. https://proceedings.mlr.press/v130/bennett21a.html
- [3] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X. Charles, D. Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual Reasoning and Learning Systems: The Example of Computational Advertising. *Journal of Machine Learning Research* 14, 101 (2013), 3207–3260. http://jmlr.org/papers/v14/bottou13a.html
- [4] Xavier de Luna and Per Johansson. 2014. Testing for the Unconfoundedness Assumption Using an Instrumental Assumption. Journal of Causal Inference 2, 2 (2014), 187–199. https://doi.org/doi:10.1515/jci-2013-0011
- [5] Zhichong Fang, Aman Agarwal, and Thorsten Joachims. 2019. Intervention Harvesting for Context-Dependent Examination-Bias Estimation. In Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (Paris, France) (SIGIR'19). ACM, 825–834. https://doi.org/10.1145/3331184.3331238
- [6] Maxime Gasse, Damien Grasset, Guillaume Gaudron, and Pierre-Yves Oudeyer. 2021. Causal Reinforcement Learning using Observational and Interventional Data. arXiv:2106.14421 [cs.LG]
- [7] Josiah Hanna, Scott Niekum, and Peter Stone. 2019. Importance Sampling Policy Evaluation with an Estimated Behavior Policy. In Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97), Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 2605–2613. https://proceedings.mlr.press/v97/hanna19a.html
- [8] Amir H. Jadidinejad, Craig Macdonald, and Iadh Ounis. 2021. The Simpson's Paradox in the Offline Evaluation of Recommendation Systems. ACM Trans. Inf. Syst. 40, 1, Article 4 (sep 2021), 22 pages. https://doi.org/10.1145/3458509
- [9] Olivier Jeunen. 2019. Revisiting Offline Evaluation for Implicit-Feedback Recommender Systems. In Proceedings of the 13th ACM Conference on Recommender Systems (Copenhagen, Denmark) (RecSys '19). ACM, 596–600. https://doi.org/10.1145/3298689.3347069
- [10] Olivier Jeunen, Ivan Potapov, and Aleksei Ustimenko. 2023. On (Normalised) Discounted Cumulative Gain as an Offline Evaluation Metric for Top-n Recommendation. arXiv:2307.15053 [cs.IR]
- [11] Thorsten Joachims, Adith Swaminathan, and Maarten de Rijke. 2018. Deep Learning with Logged Bandit Feedback. In International Conference on Learning Representations. https://openreview.net/forum?id=SJaP_-xAb
- [12] Steven A Julious and Mark A Mullee. 1994. Confounding and Simpson's paradox. Bmj 309, 6967 (1994), 1480-1481.
- [13] Chinmaya Kausik, Yangyi Lu, Kevin Tan, Yixin Wang, and Ambuj Tewari. 2023. Offline Policy Evaluation and Optimization under Confounding. arXiv:2211.16583 [stat.ML]
- [14] Walid Krichene and Steffen Rendle. 2020. On Sampled Metrics for Item Recommendation. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Virtual Event, CA, USA) (KDD '20). ACM, 1748–1757. https://doi.org/10.1145/3394486.3403226
- [15] Lihong Li, Shunbao Chen, Jim Kleban, and Ankur Gupta. 2015. Counterfactual Estimation and Optimization of Click Metrics in Search Engines: A Case Study. In Proceedings of the 24th International Conference on World Wide Web (Florence, Italy) (WWW '15 Companion). ACM, 929–934. https://doi.org/10.1145/2740908.2742562
- [16] Weiwei Liu, S Janet Kuramoto, and Elizabeth A Stuart. 2013. An introduction to sensitivity analysis for unobserved confounding in nonexperimental prevention research. Prev Sci 14, 6 (Dec. 2013), 570–580.
- [17] Ben London and Thorsten Joachims. 2022. Control Variate Diagnostics for Detecting Problems in Logged Bandit Feedback. In CONSE-QUENCES+REVEAL Workshop at ACM RecSys '22 (CONSEQUENCES+REVEAL '22).
- [18] Hongseok Namkoong, Ramtin Keramati, Steve Yadlowsky, and Emma Brunskill. 2020. Off-policy Policy Evaluation For Sequential Decisions Under Unobserved Confounding. In Advances in Neural Information Processing Systems, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 18819–18831. https://proceedings.neurips.cc/paper_files/paper/2020/file/da21bae82c02d1e2b8168d57cd3fbab7-Paper.pdf
- [19] Art B. Owen. 2013. Monte Carlo theory, methods and examples.
- [20] Judea Pearl. 2009. Causality. Cambridge university press.
- [21] Yuta Saito and Thorsten Joachims. 2021. Counterfactual Learning and Evaluation for Recommender Systems: Foundations, Implementations, and Recent Advances. In Proc. of the 15th ACM Conference on Recommender Systems (RecSys '21). ACM, 828–830. https://doi.org/10.1145/3460231.3473320
- [22] Aixin Sun. 2023. Take a Fresh Look at Recommender Systems from an Evaluation Standpoint. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (Taipei, Taiwan) (SIGIR '23). ACM, 2629–2638. https://doi.org/10.1145/3539618. 3591931
- [23] Adith Swaminathan and Thorsten Joachims. 2015. The Self-Normalized Estimator for Counterfactual Learning. In Advances in Neural Information Processing Systems, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Eds.), Vol. 28. Curran Associates, Inc. https://proceedings.neurips. cc/paper_files/paper/2015/file/39027dfad5138c9ca0c474d71db915c3-Paper.pdf

CONSEQUENCES@RecSys '23, September 19, 2023, Singapore

- [24] Lingxiao Wang, Zhuoran Yang, and Zhaoran Wang. 2021. Provably Efficient Causal Reinforcement Learning with Confounded Observational Data. In Advances in Neural Information Processing Systems, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 21164–21175.
- [25] Yang Xu, Jin Zhu, Chengchun Shi, Shikai Luo, and Rui Song. 2023. An Instrumental Variable Approach to Confounded Off-Policy Evaluation. In Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202), Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 38848–38880. https://proceedings.mlr.press/v202/ xu23x.html
- [26] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge Belongie, and Deborah Estrin. 2018. Unbiased Offline Recommender Evaluation for Missing-Not-at-Random Implicit Feedback. In Proceedings of the 12th ACM Conference on Recommender Systems (Vancouver, British Columbia, Canada) (RecSys '18). ACM, 279–287. https://doi.org/10.1145/3240323.3240355

6