



Optimal Baseline Corrections for Off-Policy Contextual Bandits

Shashank Gupta*
University of Amsterdam
Amsterdam, The Netherlands
s.gupta2@uva.nl

Harrie Oosterhuis
Radboud University
Nijmegen, The Netherlands
harrie.oosterhuis@ru.nl

Olivier Jeunen*
ShareChat
Edinburgh, United Kingdom
jeunen@sharechat.co

Maarten de Rijke
University of Amsterdam
Amsterdam, The Netherlands
m.derijke@uva.nl

ABSTRACT

The off-policy learning paradigm allows for recommender systems and general ranking applications to be framed as *decision-making* problems, where we aim to learn decision policies that optimize an unbiased *offline* estimate of an *online* reward metric. With unbiasedness comes potentially high variance, and prevalent methods exist to reduce estimation variance. These methods typically make use of control variates, either additive (i.e., baseline corrections or doubly robust methods) or multiplicative (i.e., self-normalisation).

Our work unifies these approaches by proposing a single framework built on their equivalence in learning scenarios. The foundation of our framework is the derivation of an equivalent baseline correction for all of the existing control variates. Consequently, our framework enables us to characterize the variance-optimal unbiased estimator and provide a closed-form solution for it. This optimal estimator brings significantly improved performance in both evaluation and learning, and minimizes data requirements. Empirical observations corroborate our theoretical findings.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Contextual Bandits; Recommender Systems; Off-policy Learning

ACM Reference Format:

Shashank Gupta, Olivier Jeunen, Harrie Oosterhuis, and Maarten de Rijke. 2024. Optimal Baseline Corrections for Off-Policy Contextual Bandits. In *18th ACM Conference on Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3640457.3688105>

1 INTRODUCTION & MOTIVATION

Recommender systems have undergone a paradigm shift in the last few decades, moving their focus from *rating* prediction in

the days of the Netflix Prize [3], to *item* prediction from implicit feedback [47] and *ranking* applications gaining practical importance [30, 56]. Recently, work that applies ideas from the algorithmic *decision-making* literature to recommendation problems has become more prominent [17, 28, 50, 64]. While this line of research is not inherently new [38, 55], methods based on contextual bandits (or reinforcement learning by extension) have now become widespread in the recommendation field [2, 5, 26, 42, 43, 59, 67]. The *off-policy* setting is particularly attractive for practitioners [63], as it allows models to be trained and evaluated in an offline manner [7–9, 12, 18–21, 25, 27, 31, 39, 41]. Indeed, methods exist to obtain unbiased *offline* estimators of *online* reward metrics, which can then be optimized directly [23].

Research at the forefront of this area typically aims to find Pareto-optimal solutions to the bias-variance trade-off that arises when choosing an estimator: reducing variance by accepting a small bias [22, 57], by introducing control variates [13, 62], or both [58]. Control variates are especially attractive as they (asymptotically) preserve the unbiasedness of the widespread inverse propensity scoring (IPS) estimator. Additive control variates give rise to baseline corrections [16], regression adjustments [15], and doubly robust estimators [13]. Multiplicative control variates lead to self-normalised estimators [35, 62]. Previous work has proven that for off-policy *learning* tasks, the multiplicative control variates can be re-framed using an equivalent additive variate [6, 33], enabling mini-batch optimization methods to be used. We note that the self-normalised estimator is only *asymptotically* unbiased: a clear disadvantage for evaluation with finite samples. The common problem which most existing methods tackle is that of *variance reduction* in offline value estimation, either for learning or for evaluation. The common solution is the application of a control variate, either multiplicative or additive [45]. However, to the best of our knowledge, there is no work that attempts to unify these methods. Our work addresses this gap by presenting these methods in a unifying framework of baseline corrections which, in turn, allows us to find the optimal baseline correction for variance reduction.

In the context of off-policy learning, adding to the well-known equivalence between reward-translation and self-normalisation described by Joachims et al. [33], we demonstrate that the equivalence extends to baseline corrections, regression adjustments, and doubly robust estimators with a constant reward model. Further, we derive a novel baseline correction method for off-policy learning that minimizes the variance of the gradient of the (unbiased) estimator. We further show that the baseline correction can be estimated in a

*Equal contribution.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

RecSys '24, October 14–18, 2024, Bari, Italy
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0505-2/24/10
<https://doi.org/10.1145/3640457.3688105>

closed-form fashion, allowing for easy practical implementation.

In line with recent work on off-policy evaluation/learning for recommendation [25, 26, 29, 48, 52], we adopt an off-policy simulation environment to emulate real-world recommendation scenarios, such as stochastic rewards, large action spaces, and controlled randomisation. This choice also encourages future reproducibility [49]. Our experimental results indicate that our proposed baseline correction for gradient variance reduction enables substantially faster convergence and lower gradient variance during learning.

In addition, we derive a closed-form solution to the optimal baseline correction for off-policy evaluation, i.e., the one that minimizes the variance of the estimator itself. Importantly, since our framework only considers unbiased estimators, the variance-optimality implies overall optimality. Our experimental results show that this leads to lower errors in policy value estimation than widely used doubly-robust and SNIPS estimators [13, 62].

All source code to reproduce our experimental results is available at: https://github.com/shashankg7/recsys2024_optimal_baseline.

2 BACKGROUND AND RELATED WORK

The goal of this section is to introduce common contextual bandit setups for recommendation, both on-policy and off-policy.

2.1 On-policy contextual bandits

We address a general contextual bandit setup [32, 51] with contexts X , actions A , and rewards R . The context typically describes *user* features, actions are the *items* to recommend, and rewards can be any type of *interaction* logged by the platform. A policy π defines a conditional probability distribution over actions x : $P(A = a | X = x, \Pi = \pi) \equiv \pi(a | x)$. Its *value* is the expected reward it yields:

$$V(\pi) = \mathbb{E}_{x \sim P(X)} \left[\mathbb{E}_{a \sim \pi(\cdot | x)} [R] \right]. \quad (1)$$

When the policy π is deployed, we can estimate this quantity by averaging the rewards we observe. We denote the expected reward for action a and context r as $r(a, x) := \mathbb{E}[R | X = x; A = a]$.

In the field of contextual bandits (and reinforcement learning (RL) by extension), one often wants to learn π to maximise $V(\pi)$ [36, 60]. This is typically achieved through gradient ascent. Assuming π_θ is parameterised by θ , we iteratively update with learning rate η :

$$\theta_{t+1} = \theta_t + \eta \nabla_\theta (V(\pi_\theta)). \quad (2)$$

Using the well-known REINFORCE “log-trick” [66], the above gradient can be formulated as an expectation over sampled actions, whereby tractable Monte Carlo estimation is made possible:

$$\begin{aligned} \nabla_\theta (V(\pi_\theta)) &= \nabla_\theta \left(\mathbb{E}_{x \sim P(X)} \left[\mathbb{E}_{a \sim \pi_\theta(\cdot | x)} [R] \right] \right) \\ &= \nabla_\theta \left(\int \sum_{a \in \mathcal{A}} \pi_\theta(a | x) r(a, x) P(X = x) dx \right) \\ &= \int \sum_{a \in \mathcal{A}} \nabla_\theta (\pi_\theta(a | x) r(a, x)) P(X = x) dx \\ &= \int \sum_{a \in \mathcal{A}} \pi_\theta(a | x) \nabla_\theta (\log(\pi_\theta(a | x)) r(a, x)) P(X = x) dx \end{aligned} \quad (3)$$

$$= \mathbb{E}_{x \sim P(X)} \left[\mathbb{E}_{a \sim \pi_\theta(\cdot | x)} [\nabla_\theta (\log(\pi_\theta(a | x)) R)] \right].$$

This provides an unbiased estimate of the gradient of $V(\pi_\theta)$. However, it may be subject to high variance due to the inherent variance of R . Several techniques have been proposed in the literature that aim to alleviate this, mostly using additive *control variates*.

Control variates are random variables with a known expectation [45, §8.9]. If the control variate is correlated with the original estimand — in our case $V(\pi_\theta)$ — they can be used to reduce the estimator’s variance. A natural way to apply control variates to a sample average estimate for Eq. 1 is to estimate a model of the reward $\widehat{r}(a, x) \approx \mathbb{E}[R | X = x; A = a]$ and subtract it from the observed rewards [15]. This is at the heart of key RL techniques (i.a., generalised advantage estimation [54]), and it underpins widely used methods to increase sensitivity in online controlled experiments [1, 6, 11, 46]. As such, it applies to both *evaluation* and *learning* tasks. We note that if the model $\widehat{r}(a, x)$ is biased, this bias propagates to the resulting estimator for $V(\pi_\theta)$.

Alternatively, instead of focusing on reducing the variance of $V(\pi_\theta)$ directly, other often-used approaches tackle the variance of its gradient estimates $\nabla_\theta (V(\pi_\theta))$ instead.

Observe that $\mathbb{E}_{a \sim \pi_\theta(\cdot | x)} [\nabla_\theta (\log(\pi_\theta(a | x)))] = 0$ [44, Eq. 12]. This implies that a translation on the rewards in Eq. 3 does not affect the unbiasedness of the gradient estimate. Nevertheless, as such a translation can be framed as an additive control variate, it will affect its variance. Indeed, “*baseline corrections*” are a well-known variance reduction method for on-policy RL methods [16]. For a dataset consisting of logged contexts, actions and rewards $\mathcal{D} = \{(x_i, a_i, r_i)_{i=1}^N\}$, we apply a *baseline* control variate β to the estimate of the final gradient to obtain:

$$\begin{aligned} \nabla_\theta (V(\pi_\theta)) &\approx \nabla_\theta (\widehat{V}_\beta(\pi_\theta)) \\ &= \frac{1}{|\mathcal{D}|} \sum_{(x, a, r) \in \mathcal{D}} (r - \beta) \nabla_\theta \log \pi_\theta(a | x). \end{aligned} \quad (4)$$

Williams [65] originally proposed to use the average observed reward for β . Subsequent work has derived optimal baselines for general on-policy RL scenarios [10, 16]. However, to the best of our knowledge, *optimal baselines for on-policy contextual bandits have not been considered in previous work*.

Optimal baseline for on-policy bandits. The optimal baseline β for the on-policy gradient estimate in Eq. 4 is the one that minimizes the variance of the gradient estimate. In accordance with earlier work [16], we define the variance of a vector random variable as the sum of the variance of its individual components. Therefore, the optimal baseline is given by:

$$\begin{aligned} &\arg \min_{\beta} \text{Var} \left(\nabla_\theta (\widehat{V}_\beta(\pi_\theta)) \right) \\ &= \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \text{Var} [\nabla_\theta (\log(\pi_\theta(a | x)) (r - \beta))] \end{aligned} \quad (5)$$

$$\begin{aligned} &= \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \mathbb{E} [\nabla_\theta \log(\pi_\theta(a | x))^\top \nabla_\theta \log(\pi_\theta(a | x)) (r - \beta)^2] \\ &\quad - \frac{1}{|\mathcal{D}|} \mathbb{E} [\nabla_\theta \log(\pi_\theta(a | x)) (r - \beta)]^\top \mathbb{E} [\nabla_\theta \log(\pi_\theta(a | x)) (r - \beta)] \end{aligned} \quad (6)$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \mathbb{E} [\|\nabla_{\theta} \log(\pi_{\theta}(a|x))\|_2^2 (r - \beta)^2], \quad (7)$$

where we ignore the second term in Eq. 6, since it is independent of β [44, Eq. 12]. The result from this derivation (Eq. 7) reveals that the optimal baseline can be obtained by solving the following equation:

$$\frac{\partial \text{Var}(\widehat{V}_{\beta}(\pi_{\theta}))}{\partial \beta} = \frac{2}{|\mathcal{D}|} \mathbb{E} [\|\nabla_{\theta} \log(\pi_{\theta}(a|x))\|_2^2 (\beta - r)] = 0, \quad (8)$$

which results in the following optimal baseline correction:

$$\beta^* = \frac{\mathbb{E} [\|\nabla_{\theta} \log(\pi_{\theta}(a|x))\|_2^2 r(a, x)]}{\mathbb{E} [\|\nabla_{\theta} \log(\pi_{\theta}(a|x))\|_2^2]}, \quad (9)$$

and the empirical estimate of the optimal baseline correction:

$$\widehat{\beta}^* = \frac{\sum_{(x,a,r) \in \mathcal{D}} [\|\nabla_{\theta} \log(\pi_{\theta}(a|x))\|_2^2 r(a, x)]}{\sum_{(x,a,r) \in \mathcal{D}} [\|\nabla_{\theta} \log(\pi_{\theta}(a|x))\|_2^2]}. \quad (10)$$

This derivation follows the more general derivation from Green-Smith et al. [16] for partially observable Markov decision processes (POMDPs). We have not encountered its use in the existing bandit literature applied to recommendation problems. In Section 3.2, we show that a similar line of reasoning can be applied to derive a variance-optimal gradient for the off-policy contextual bandit setup.

2.2 Off-policy estimation for general bandits

Deploying π is a costly prerequisite for estimating $V(\pi)$, that comes with the risk of deploying a possible poorly valued π . Therefore, commonly in real-world model validation pipelines, practitioners wish to estimate $V(\pi)$ *before* deployment. Accordingly, we will address this *counterfactual* evaluation scenario that falls inside the field of off-policy estimation (OPE) [52, 64].

The expectation $V(\pi)$ can be unbiasedly estimated using samples from a *different* policy π_0 through *importance sampling*, also known as inverse propensity score weighting (IPS) [45, §9]:

$$\mathbb{E}_{x \sim P(X)} \left[\mathbb{E}_{a \sim \pi(\cdot|x)} [R] \right] = \mathbb{E}_{x \sim P(X)} \left[\mathbb{E}_{a \sim \pi_0(\cdot|x)} \left[\frac{\pi(a|x)}{\pi_0(a|x)} R \right] \right]. \quad (11)$$

To ensure that the so-called *importance weights* $\frac{\pi(a|x)}{\pi_0(a|x)}$ are well-defined, we assume “*common support*” by the logging policy: $\forall a \in \mathcal{A}, x \in \mathcal{X} : \pi(a|x) > 0 \implies \pi_0(a|x) > 0$.

From Eq. 11, we can derive an unbiased estimator for $V(\pi)$ using contexts, actions and rewards logged under π_0 , denoted by \mathcal{D} :

$$\widehat{V}_{\text{IPS}}(\pi, \mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} r. \quad (12)$$

To keep our notation brief, we suppress subscripts when they are clear from the context. In the context of gradient-based optimization methods, we often refer to a minibatch $\mathcal{B} \subset \mathcal{D}$ instead of the whole dataset, as is typical for, e.g., stochastic gradient descent (SGD).

If we wish to learn a policy that maximises this estimator, we need to estimate its gradient for a batch \mathcal{B} . Whilst some previous work has applied a REINFORCE estimator [7, 9, 41], we use a straightforward Monte Carlo estimate for the gradient:

$$\nabla \widehat{V}_{\text{IPS}}(\pi, \mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{(x,a,r) \in \mathcal{B}} \frac{\nabla \pi(a|x)}{\pi_0(a|x)} r. \quad (13)$$

Importance sampling — the bread and butter of unbiased off-policy estimation — often leads to increased variance compared to on-policy estimators. Several variance reduction techniques have been proposed specifically to combat the excessive variance of \widehat{V}_{IPS} [13, 22, 62]. Within the scope of this work, we only consider techniques that reduce variance *without* introducing bias.

Self-normalised importance sampling. The key idea behind *self-normalisation* [45, §9.2] is to use a *multiplicative* control variate to rescale $\widehat{V}_{\text{IPS}}(\pi, \mathcal{D})$. An important observation for this approach is that for any policy π and a dataset \mathcal{D} logged under π_0 , the expected average of importance weights should equal 1 [62, §5]:

$$\mathbb{E}_{\mathcal{D} \sim P(\mathcal{D})} \left[\frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} \right] = 1. \quad (14)$$

Furthermore, as this random variable (Eq. 14) is likely to be correlated with the IPS estimates, we can expect that its use as a control variate will lead to reduced variance (see [e.g., 35]). This gives rise to the asymptotically unbiased and parameter-free self-normalised IPS (SNIPS) estimator, with $S := \frac{1}{D} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)}$ as its normalization term:

$$\widehat{V}_{\text{SNIPS}}(\pi, \mathcal{D}) = \frac{\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} r}{\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)}} = \frac{\widehat{V}_{\text{IPS}}(\pi, \mathcal{D})}{S}. \quad (15)$$

Given the properties of being asymptotically unbiased and parameter-free, this estimator is often a go-to method for off-policy *evaluation* use-cases [52]. An additional advantage is that the SNIPS estimator is invariant to translations in the reward, which cannot be said for \widehat{V}_{IPS} . Whilst the formulation in Eq. 15 is not obvious in this regard, it becomes clear when we consider its gradient:

$$\begin{aligned} \nabla \widehat{V}_{\text{SNIPS}}(\pi, \mathcal{D}) &= \nabla \left(\frac{\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} r}{\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)}} \right) \\ &= \frac{\left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\nabla \pi(a|x)}{\pi_0(a|x)} r \right) \left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} \right)}{\left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} \right)^2} \\ &\quad - \frac{\left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} r \right) \left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\nabla \pi(a|x)}{\pi_0(a|x)} \right)}{\left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} \right)^2} \\ &= \frac{\sum_{(x_i, a_i, r_i) \in \mathcal{D}} \sum_{(x_j, a_j, r_j) \in \mathcal{D}} \frac{\pi(a_i|x_i) \nabla \pi(a_j|x_j)}{\pi_0(a_i|x_i) \pi_0(a_j|x_j)} (r_j - r_i)}{\left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} \right)^2} \\ &= \frac{\sum_{(x_i, a_i, r_i) \in \mathcal{D}} \sum_{(x_j, a_j, r_j) \in \mathcal{D}} \frac{\pi(a_i|x_i) \pi(a_j|x_j)}{\pi_0(a_i|x_i) \pi_0(a_j|x_j)} \nabla \log \pi(a_j|x_j) (r_j - r_i)}{\left(\sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} \right)^2}. \end{aligned} \quad (16)$$

Indeed, as the SNIPS gradient relies on the *relative difference* in observed reward between two samples, a constant correction would not affect it (i.e., if $\bar{r} = r - \beta$, then $r_j - r_i \equiv \bar{r}_j - \bar{r}_i$).

Swaminathan and Joachims [62] effectively apply the SNIPS estimator (with a variance regularisation term [61]) to off-policy *learning* scenarios. Note that while \widehat{V}_{IPS} neatly decomposes into a single sum over samples, $\widehat{V}_{\text{SNIPS}}$ no longer does. Whilst this may be

clear from the gradient formulation in Eq. 16, a formal proof can be found in [33, App. C]. This implies that mini-batch optimization methods (which are often necessary to support learning from large datasets) are no longer directly applicable to $\widehat{V}_{\text{SNIPS}}$.

Joachims et al. [33] solve this by re-framing the task of maximising $\widehat{V}_{\text{SNIPS}}$ as an optimization problem on \widehat{V}_{IPS} with a constraint on the self-normalisation term. That is, if we define:

$$\pi^* = \arg \max_{\pi \in \Pi} \widehat{V}_{\text{SNIPS}}(\pi, \mathcal{D}), \text{ with } S^* = \frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi^*(a|x)}{\pi_0(a|x)}, \quad (17)$$

then, we can equivalently state this as:

$$\pi^* = \arg \max_{\pi \in \Pi} \widehat{V}_{\text{IPS}}(\pi, \mathcal{D}), \text{ s.th. } \frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} = S^*. \quad (18)$$

Joachims et al. [33] show via the Lagrange multiplier method that this optimization problem can be solved by optimising for \widehat{V}_{IPS} with a translation on the reward:

$$\begin{aligned} \pi^* &= \arg \max_{\pi \in \Pi} \widehat{V}_{\lambda^* \text{-IPS}}(\pi, \mathcal{D}), \text{ where} \\ \widehat{V}_{\lambda \text{-IPS}}(\pi, \mathcal{D}) &= \frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} (r - \lambda). \end{aligned} \quad (19)$$

This approach is called BanditNet [33]. Naturally, we do not know λ^* beforehand (because we do not know S^*), but we know that S^* should concentrate around 1 for large datasets (see Eq. 14). Joachims et al. [33] essentially propose to treat λ as a hyper-parameter to be tuned in order to find S^* .

Doubly robust estimation. Another way to reduce the variance of \widehat{V}_{IPS} is to use a model of the reward $\widehat{r}(a, x) \approx \mathbb{E}[R|X = x; A = a]$. Including it as an additive control variate in Eq. 12 gives rise to the doubly robust (DR) estimator, deriving its name from its unbiasedness if *either* the logging propensities π_0 or the reward model \widehat{r} is unbiased [13]:

$$\begin{aligned} \widehat{V}_{\text{DR}}(\pi, \mathcal{D}) &= \\ \frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} &\left(\frac{\pi(a|x)}{\pi_0(a|x)} (r - \widehat{r}(a, x)) + \sum_{a' \in \mathcal{A}} \pi(a'|x) \widehat{r}(a', x) \right). \end{aligned} \quad (20)$$

Several further extensions have been proposed in the literature: one can optimize the reward model $\widehat{r}(a, x)$ to minimize the resulting variance of \widehat{V}_{DR} [14], further parameterise the trade-off relying on \widehat{V}_{IPS} or $\widehat{r}(a, x)$ [58], or shrink the IPS weights to minimize a bound on the MSE of the resulting estimator [57]. One disadvantage of this method, is that practitioners are required to fit the secondary reward model $\widehat{r}(a, x)$, which might be costly and sample inefficient. Furthermore, variance reduction is generally not guaranteed, and stand-alone \widehat{V}_{IPS} can be empirically superior in some scenarios [24].

3 UNIFYING OFF-POLICY ESTIMATORS

Section 2 provides an overview of (asymptotically) unbiased estimators for the value of a policy. We have introduced the contextual bandit setting, detailing often used variance reduction techniques for both on-policy (i.e., regression adjustments and baseline corrections) and off-policy estimation (i.e., self-normalisation and doubly robust estimation). In this section, we demonstrate that they perform equivalent optimization as baseline-corrected estimation.

Subsequently, we characterize the baseline corrections that either minimize the variance of the estimator, or that of its gradient.

3.1 A unified off-policy estimator

Baseline corrections for $\nabla \widehat{V}_{\text{IPS}}(\pi, \mathcal{D})$. Baseline corrections are common in on-policy estimation, but occur less often in the off-policy literature. The estimator is obtained by removing a baseline control variate $\beta \in \mathbb{R}$ from the reward of each action, while also adding it to the estimator:

$$\widehat{V}_{\beta \text{-IPS}} = \beta + \frac{1}{|\mathcal{D}|} \sum_{(x,a,r) \in \mathcal{D}} \frac{\pi(a|x)}{\pi_0(a|x)} (r - \beta). \quad (21)$$

Its unbiasedness is easily verified:

$$\begin{aligned} \mathbb{E}[\widehat{V}_{\beta \text{-IPS}}] &= \mathbb{E}[\beta] + \mathbb{E}\left[\frac{\pi(a|x)}{\pi_0(a|x)} (r - \beta)\right] \\ &= \beta + \mathbb{E}\left[\frac{\pi(a|x)}{\pi_0(a|x)} r\right] - \beta = V(\pi). \end{aligned} \quad (22)$$

From an optimization perspective, we are mainly interested in the gradient of the $\widehat{V}_{\beta \text{-IPS}}$ objective:

$$\nabla \widehat{V}_{\beta \text{-IPS}}(\pi, \mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{(x,a,r) \in \mathcal{B}} \frac{\nabla \pi(a|x)}{\pi_0(a|x)} (r - \beta). \quad (23)$$

Our key insight is that SNIPS and certain doubly-robust estimators have an equivalent gradient to the proposed β -IPS estimator. As a result, optimizing them is equivalent to optimizing $\widehat{V}_{\beta \text{-IPS}}$ for a specific β value.

Self-normalisation through BanditNet and $\widehat{V}_{\lambda \text{-IPS}}(\pi, \mathcal{D})$. If we consider the optimization problem for SNIPS that is solved by BanditNet in Eq. 19 [33], we see that its gradient is given by:

$$\nabla \widehat{V}_{\lambda \text{-IPS}}(\pi, \mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{(x,a,r) \in \mathcal{B}} \frac{\nabla \pi(a|x)}{\pi_0(a|x)} (r - \lambda). \quad (24)$$

Doubly robust estimation via $\widehat{V}_{\text{DR}}(\pi, \mathcal{D})$. As mentioned, a nuisance of doubly robust estimators is the requirement of fitting a regression model $\widehat{r}(a, x)$. Suppose that we instead treat \widehat{r} as a single scalar hyper-parameter, akin to the BanditNet approach. Then, the gradient of such an estimator would be given by:

$$\nabla \widehat{V}_{\widehat{r} \text{-DR}}(\pi, \mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{(x,a,r) \in \mathcal{B}} \frac{\nabla \pi(a|x)}{\pi_0(a|x)} (r - \widehat{r}). \quad (25)$$

Importantly, these three approaches are motivated through entirely different lenses: minimizing gradient variance, applying a multiplicative control variate to reduce estimation variance, and applying an additive control variate to improve robustness. But they result in equivalent gradients, and thus, in equivalent optima. Specifically, for optimization, the estimators are equivalent when $\beta \equiv \lambda \equiv \widehat{r}$.

This equivalence implies that the choice between these three approaches is not important. Since the simple baseline correction estimator $\widehat{V}_{\beta \text{-IPS}}$ (Eq. 21) has an equivalence with all SNIPS estimators and all doubly-robust estimators with a constant reward, we propose that $\widehat{V}_{\beta \text{-IPS}}$ should be seen as an estimator that unifies all three approaches. Accordingly, we argue that the real task is to find the optimal β value for $\widehat{V}_{\beta \text{-IPS}}$, since this results in an estimator that is at least as optimal as any estimator in the underlying families of

estimators, and possibly superior to them.

The remainder of this section describes the optimal β values for minimizing gradient variance and estimation value variance.

3.2 Minimizing gradient variance

Similar to the on-policy variant derived in Eq. 7, we can derive the optimal baseline in the off-policy case as the one which results in the minimum variance for the gradient estimate given by Eq. 13:

$$\arg \min_{\beta} \text{Var} \left(\nabla_{\theta} (\widehat{V}_{\beta\text{-IPS}}(\pi_{\theta}, \mathcal{B})) \right) \quad (26)$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{B}|} \text{Var} \left[\frac{\nabla \pi(a|x)}{\pi_0(a|x)} (r - \beta) \right] \quad (27)$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{B}|} \mathbb{E} \left[\left\| \nabla \pi(a|x) \right\|_2^2 \left(\frac{r - \beta}{\pi_0(a|x)} \right)^2 \right] \quad (28)$$

$$- \frac{1}{|\mathcal{B}|} \mathbb{E} \left[\frac{\nabla \pi(a|x)}{\pi_0(a|x)} (r - \beta) \right]_2^2$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{B}|} \mathbb{E} \left[\frac{\left\| \nabla \pi(a|x) \right\|_2^2}{\pi_0(a|x)^2} (r - \beta)^2 \right], \quad (29)$$

where we can ignore the second term of the variance in Eq. 28, since it is independent of β [44, Eq. 12]. The optimal baseline can be obtained by solving for:

$$\frac{\partial \text{Var} \left(\nabla (\widehat{V}_{\beta\text{-IPS}}(\pi, \mathcal{B})) \right)}{\partial \beta} = \frac{2}{|\mathcal{B}|} \mathbb{E} \left[\frac{\left\| \nabla \pi(a|x) \right\|_2^2}{\pi_0(a|x)^2} (\beta - r) \right] = 0, \quad (30)$$

which results in the following optimal baseline:

$$\beta^* = \frac{\mathbb{E}_{x, a \sim \pi_0, r} \left[\frac{\left\| \nabla \pi(a|x) \right\|_2^2}{\pi_0(a|x)^2} r(a, x) \right]}{\mathbb{E}_{x, a \sim \pi_0, r} \left[\frac{\left\| \nabla \pi(a|x) \right\|_2^2}{\pi_0(a|x)^2} \right]}, \quad (31)$$

with its empirical estimate given by:

$$\widehat{\beta}^* = \frac{\sum_{(x, a, r) \in \mathcal{B}} \left[\frac{\left\| \nabla \pi(a|x) \right\|_2^2}{\pi_0(a|x)^2} r \right]}{\sum_{(x, a, r) \in \mathcal{B}} \left[\frac{\left\| \nabla \pi(a|x) \right\|_2^2}{\pi_0(a|x)^2} \right]}. \quad (32)$$

Note that this expectation is over actions sampled by the *logging* policy. As a result, we can obtain Monte Carlo estimates of the corresponding expectations. The derivation has high similarity with the on-policy case (cf. Section 2.1) [16]. Nevertheless, we are unaware of any work on off-policy learning that uses it. Joachims et al. [33] refer to the on-policy variant with: “we cannot sample new roll-outs from the current policy under consideration, which means we cannot use the standard variance-optimal estimator used in REINFORCE.” Since the expectation is over actions sampled by the *logging* policy and not the *target* policy, we have shown that we do not need new roll-outs. Thereby, our estimation strategy is a novel off-policy approach that estimates the variance-optimal baseline.

THEOREM 1. *Within the family of gradient estimators with a global additive control variate, i.e., β -IPS (Eq. 23), IPS (Eq. 13), BanditNet (Eq. 24), and DR with a constant correction (Eq. 25), β -IPS with our proposed choice of β in Eq. 31 has minimal gradient variance.*

PROOF. Eq. 30 shows that the β value in Eq. 31 attains a minimum. Because the variance of the gradient estimate (Eq. 28) is a quadratic function of β , and hence a convex function (Eq. 29), it must be the global minimum for the gradient variance. \square

3.3 Minimizing estimation variance

Besides minimizing gradient variance, one can also aim to minimize the variance of estimation, i.e., the variance of the estimated value. We note that the β value for minimizing estimation need not be the same value that minimizes gradient variance. Furthermore, since $\widehat{V}_{\beta\text{-IPS}}$ is unbiased, any estimation error will entirely be driven by variance. As a result, the value for β that results in minimal variance will also result in minimal estimation error:

$$\arg \min_{\beta} \text{Var} \left(\widehat{V}_{\beta\text{-IPS}}(\pi, \mathcal{D}) \right) \quad (33)$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \text{Var} \left[\frac{\pi(a|x)}{\pi_0(a|x)} (r - \beta) \right] \quad (34)$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \mathbb{E} \left[\left(\frac{\pi(a|x)}{\pi_0(a|x)} (r - \beta) \right)^2 \right] \quad (35)$$

$$- \frac{1}{|\mathcal{D}|} \left(\mathbb{E} \left[\frac{\pi(a|x)}{\pi_0(a|x)} (r - \beta) \right] \right)^2$$

$$= \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \mathbb{E} \left[\left(\frac{\pi(a|x)}{\pi_0(a|x)} \right)^2 (r - \beta)^2 \right] \quad (36)$$

$$- \frac{1}{|\mathcal{D}|} \left(\mathbb{E} \left[\frac{\pi(a|x)}{\pi_0(a|x)} r \right] - \beta \right)^2.$$

The minimum is obtained by solving for the following equation:

$$\frac{\partial \left(\text{Var} \left(\widehat{V}_{\beta\text{-IPS}}(\pi, \mathcal{D}) \right) \right)}{\partial \beta} \quad (37)$$

$$= \frac{2}{|\mathcal{D}|} \mathbb{E} \left[\left(\frac{\pi(a|x)}{\pi_0(a|x)} \right)^2 (\beta - r) \right] - \frac{2}{|\mathcal{D}|} \left(\beta - \mathbb{E} \left[\frac{\pi(a|x)}{\pi_0(a|x)} r \right] \right) = 0,$$

which results in the following optimal baseline:

$$\beta^* = \frac{\mathbb{E} \left[\left(\left(\frac{\pi(a|x)}{\pi_0(a|x)} \right)^2 - \frac{\pi(a|x)}{\pi_0(a|x)} \right) r(a, x) \right]}{\mathbb{E} \left[\left(\frac{\pi(a|x)}{\pi_0(a|x)} \right)^2 - \left(\frac{\pi(a|x)}{\pi_0(a|x)} \right) \right]}. \quad (38)$$

We can estimate β^* using logged data, resulting in a Monte Carlo estimate of the optimal baseline. Such a sample estimate will not be unbiased (because it is a ratio of expectations), but the bias will vanish asymptotically (similar to the bias of the $\widehat{V}_{\text{SNIPS}}$ estimator).

Next, we formally prove that optimal estimator variance leads to overall optimality (in terms of the MSE of the estimator).

THEOREM 2. *Within the family of offline estimators with a global additive control variate, i.e., β -IPS (Eq. 21), IPS (Eq. 12), and DR with a constant correction (Eq. 25), β -IPS with our proposed β in Eq. 38 has the minimum mean squared error (MSE):*

$$\text{MSE}(\widehat{V}(\pi)) := \mathbb{E}_{\mathcal{D}} \left[\left(\widehat{V}(\pi, \mathcal{D}) - V(\pi) \right)^2 \right]. \quad (39)$$

PROOF. The MSE of any off-policy estimator $\widehat{V}(\pi, \mathcal{D})$ can be

decomposed in terms of the bias and variance of the estimator [57]:

$$\text{MSE}(\hat{V}(\pi)) = \text{Bias}(\hat{V}(\pi), \mathcal{D})^2 + \text{Variance}(\hat{V}(\pi), \mathcal{D}), \quad (40)$$

where the bias of the estimator is defined as:

$$\text{Bias}(\hat{V}(\pi), \mathcal{D}) = |\mathbb{E}_{\mathcal{D}}[\hat{V}(\pi, \mathcal{D}) - V(\pi)]|, \quad (41)$$

and the variance of the estimator is defined previously (see Section 3.3). Eq. 22 proves that β -IPS is unbiased: $\text{Bias}(\hat{V}(\pi), \mathcal{D}) = 0$. Thus, the minimum variance (Eq. 37) implies minimum MSE. \square

We note that SNIPS is not covered by this theorem, as it is only asymptotically unbiased. As a result, the variance reduction brought on by SNIPS might be higher than that by β -IPS, but as it introduces bias, its estimation error (MSE) is not guaranteed to be better. Our experimental results below indicate that our method is always at least as good as SNIPS, and outperforms it in most cases, in both learning and evaluation tasks.

4 EXPERIMENTAL SETUP

In order to evaluate off-policy learning and evaluation methods, we need access to logged data sampled from a stochastic policy involving logging propensities (exact or estimated) along with the corresponding context and action pairs. Recent work that focuses on off-policy learning or evaluation for contextual bandits in recommender systems follows a supervised-to-bandit conversion process to simulate a real-world bandit feedback dataset [25, 26, 29, 52, 57, 58], or conducts a live experiment on actual user traffic to evaluate the policy in an *on-policy* or *online* fashion [7, 9]. In this work, we adopt the Open Bandit Pipeline (OBP) to simulate, in a reproducible manner, real-world recommendation setups with stochastic rewards, large action spaces, and controlled randomization [48]. Although the Open Bandit Pipeline simulates a generic offline contextual bandit setup, there is a strong correspondence to real-world recommendation setups where the environment context vector corresponds to the user context and the actions correspond to the items recommended to the user. Finally, the reward corresponds to the user feedback received on the item (click, purchase, etc.). As an added advantage, the simulator allows us to conduct experiments in a *realistic* setting where the logging policy is sub-optimal to a controlled extent, the logged data size is limited, and the action space is large. In addition, we conduct experiments with real-world recommendation logs from the OBP for off-policy evaluation.¹

The research questions we answer with our experimental results are:

- RQ1** Does the proposed *estimator-variance-minimizing* baseline correction (Eq. 38) improve off-policy learning (OPL) in a full-batch setting?
- RQ2** Does the proposed *gradient-variance-minimizing* baseline correction (Eq. 31) improve OPL in a mini-batch setting?
- RQ3** How does the proposed *gradient-variance-minimizing* baseline correction (Eq. 31) affect gradient variance during OPL?
- RQ4** Does the proposed *estimator-variance-minimizing* baseline correction (Eq. 38) improve OPE performance?

¹<https://research.zozo.com/data.html>

5 RESULTS AND DISCUSSION

5.1 Off-policy learning performance (RQ1–3)

To evaluate the performance of the proposed β -IPS method on an OPL task, we consider two learning setups:

- (1) *Full-batch*. In this setup, we directly optimize the β -IPS policy value estimator (Eq. 21) with the optimal baseline correction, which minimizes the variance of the value (Eq. 38). Given that the optimal baseline correction involves a ratio of two expectations, optimizing the value function directly via a mini-batch stochastic optimization is not possible for the same reason as the SNIPS estimator, i.e., it is not possible to get an unbiased gradient estimate with a ratio function [33]. Therefore, for this particular setting, we use a *full-batch* gradient descent method for the optimization, where the gradient is computed over the entire training dataset.
- (2) *Mini-batch*. In this setup, we focus on optimizing the β -IPS policy value estimator with the baseline correction, which minimizes the gradient estimate (Eq. 31). This setup translates to a traditional machine learning training setup where the model is optimized in a stochastic mini-batch fashion.

Full-batch. The results for the full-batch training in terms of the policy value on the test set are reported in Figure 1, over the number of training epochs. To minimize the impact of external factors, we use a linear model without bias, followed by a softmax to generate a distribution over all actions, given a context vector x (this is a common setup, see [e.g., 24, 31, 53]). We note that the goal of this work is not to get the maximum possible policy value on the test set but rather to evaluate the effect of baseline corrections on gradient and estimation variance. The simple model setup allows us to easily track the empirical gradient variance, given that we have only one parameter vector.

An advantage of the full-batch setup is that we can compute the gradient of the SNIPS estimator directly [62]. SNIPS is a natural baseline method to consider, along with the traditional IPS estimator. Because of practical concerns, we only consider 500 epochs of optimization. Additionally, we use the state-of-the-art and widely used Adam optimizer [34].

The IPS method converges to a lower test policy value in comparison to the SNIPS and the proposed β -IPS methods, even after 500 epochs. A likely reason is the high-variance of the IPS estimator [13], which can cause it to get stuck in bad local minima.

The methods with a control variate, i.e., SNIPS (with multiplicative control variate) and β -IPS (with additive control variate) converge to substantially better test policy values. In terms of the convergence speed, β -IPS converges to the optimal value faster than the SNIPS estimator, most likely because it has lower estimator variance than SNIPS. With this, we can answer RQ1 as follows: in the full-batch setting, our proposed optimal baseline correction enables β -IPS to converge faster than SNIPS at similar performance.

Mini-batch. The results for mini-batch training in terms of the test policy value are reported in Figure 2. Different from the full-batch setup, where the focus is on reducing the variance of the *estimator* value (Section 3.3), in the mini-batch mode, the focus is on reducing the variance of the gradient estimate (Section 3.2). The model and training setup are similar to the full-batch mode, except




Figure 1: Performance of different off-policy learning methods trained in a full-batch gradient descent fashion in terms of the policy value on the test set. x-axis corresponds to the training epoch during the optimization (we use a maximum of 500 epochs for all methods), and y-axis corresponds to the policy value. A decaying learning rate is used. Reported results are averages over 32 independent runs with 95% confidence interval.




Figure 2: Performance of different off-policy learning methods trained in a mini-batch gradient descent fashion in terms of policy value on the test set. The axis labels are similar to Figure 1.




Figure 3: Empirical variance of the gradient of different off-policy learning estimators in a mini-batch optimization setup with varying learning rates (in title). We compute gradient variance for each mini-batch during training and then report the average value across all mini-batches in a training epoch. Results are averaged across 32 independent runs with 95% confidence interval.

that we fixed the batch size to 1024 for the mini-batch experiments. Preliminary results indicated that the batch size hyper-parameter has a limited effect.

Analogous to the full-batch setup, the IPS estimator results in a lower test policy value, most likely because of the high gradient variance which prevents convergence to high performance. In contrast, due to their baseline corrections, BanditNet (Eq. 24) and β -IPS have a lower gradient variance. Accordingly, they also converge to better performance [4], i.e., resulting in superior test policy values.

Amongst these baseline-corrected gradient-based methods (Ban-

ditNet and β -IPS), our proposed β -IPS estimator outperforms BanditNet as it provides a policy with substantially higher value. The differences are observed over different choices of learning rates. Thus we answer RQ2 accordingly: in the mini-batch setting, our proposed gradient-minimizing baseline method results in considerably higher policy value compared to both IPS and BanditNet.

Next, we directly consider the empirical gradient variance of different estimators; Figure 3 reports the average mini-batch gradient variance per epoch. As expected, the IPS estimator has the highest gradient variance by a large margin. For BanditNet, we observe a lower gradient variance, which is the desired result of the additive

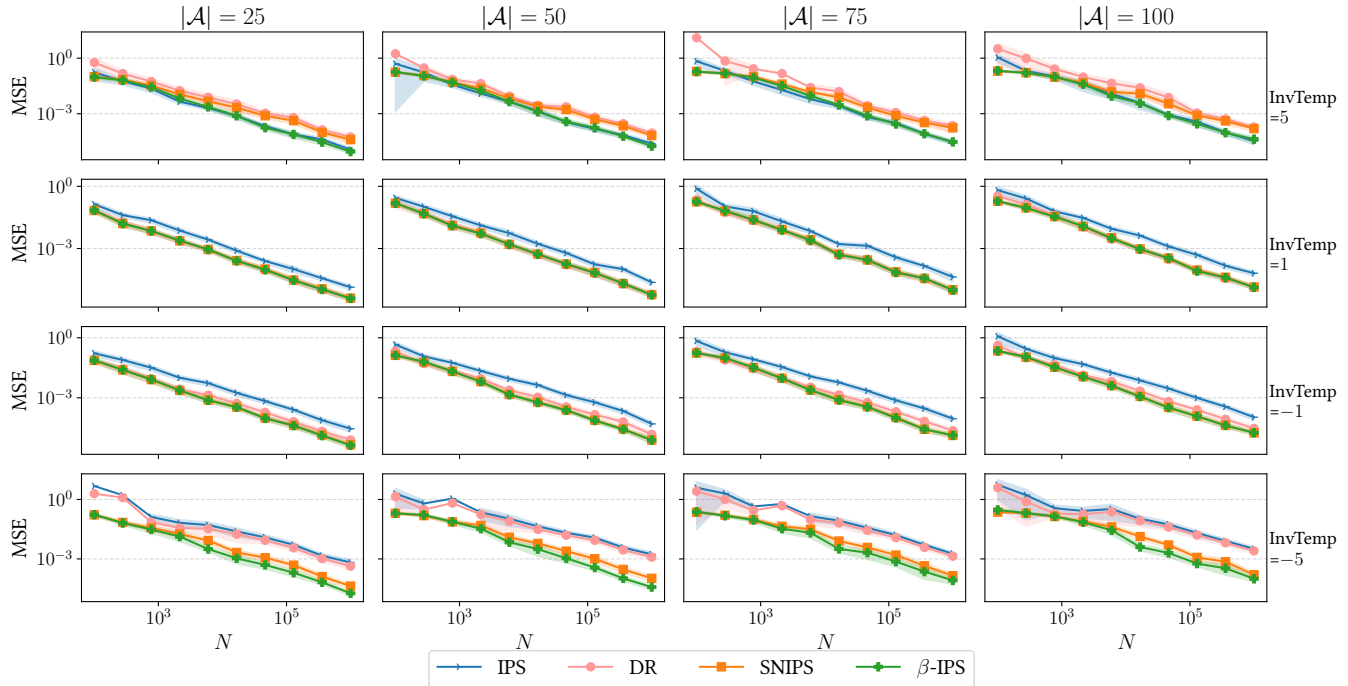


Figure 4: Mean Squared Error (MSE) of different off-policy estimators with varying action space (from left to right), and varying inverse temperature parameter of the softmax logging policy (from top to bottom). X-axis corresponds to the size of the logged data simulated (ranging from 10^2 to 10^6), and the y-axis corresponds to the MSE (evaluated over 100 independent samples of the synthetic data) along with 95% confidence interval. Each row corresponds to a different setting of inverse temperature of the softmax logging policy. We only consider unbiased (asymptotically or otherwise) estimators.

baseline it employs. Finally, we observe that our proposed method β -IPS has the lowest gradient variance. This result corroborates the theoretical claim (Theorem 1) that the β -IPS estimator has the lowest gradient variance amongst all global additive control variates (including IPS and BanditNet). Our answer to RQ3 is thus clear: our proposed β -IPS results in considerably lower gradient variance compared to BanditNet and IPS.

5.2 Off-policy evaluation performance (RQ4)

To evaluate the performance of the proposed β -IPS method, which minimizes the estimated policy value (Eq. 38), in an OPE task, results are presented in Figure 4. The target policy (to be evaluated) is a logistic regression model trained via the IPS objective on logged data and evaluated on a separate full-information test set. We evaluate the MSE of the estimated policy value against the *true* policy value (Eq. 39). To evaluate the MSE of different estimators realistically, we report results with varying degrees of the optimality of the behavior policy (decided by the inverse temperature parameter of the softmax) and with a varying cardinality of the action space. A positive (and higher) inverse softmax temperature results in an increasingly optimal behavior policy (selects action with highest reward probability), and a negative (and lower) inverse softmax temperature parameter results in an increasingly sub-optimal behavior policy (selects actions with lowest reward probability). Our proposed β -IPS method has the lowest MSE in all simulated settings. Interestingly, the proposed β -IPS has a lower MSE than the DR method, which has a regression model-based control variate,

Table 1: Comparison of different OPE methods on real-world recommender system logs of ZOTOTOWN from a campaign targeted towards men with a uniformly random production policy. We report the mean relative absolute error (with std).

OPE estimator	Abs. relative error \downarrow
IPS	0.1277 (0.0142)
SNIPS	0.1113 (0.0372)
DR	0.1144 (0.0366)
β -IPS	0.1078 (0.0383)

arguably more powerful than the constant control variate from the proposed β -IPS method. Similar observations have been made in previous work, e.g., Jeunen and Goethals [24] reported that the DR estimator’s performance heavily depends on the logging policy.

Depending on the setting, we see that β -IPS either has performance comparable to the SNIPS estimator, i.e., when inverse temperature $\in \{-1, 1\}$; or noticeably higher performance than SNIPS, i.e., when inverse temperature $\in \{-5, 5\}$.

Real-world evaluation. To evaluate different estimators in a real-world recommender systems setup, we report the results of OPE from the production logs of a real-world recommender system in Table 1. Similar to the simulation setup, the proposed β -IPS has the lowest absolute relative error amongst all estimators in the comparison. In conclusion, we answer RQ4: our proposed policy-value variance minimizing baseline method results in substantially improved MSE, compared to IPS, SNIPS and DR, in offline evaluation tasks that are typical recommender system use-cases.




Figure 5: Empirical variance of different off-policy estimators with varying action space (from left to right), and varying sub-optimality of a temperature-based softmax behavior policy (from top to bottom). The x-axis corresponds to the size of the logged data simulated (ranging from 10^2 to 10^6), and the y-axis corresponds to the variance of different estimators (evaluated over 100 independent samples of the synthetic data) along with 95% confidence interval. Each row corresponds to a different optimality level of the logging policy, decided by the inverse temperature parameter. We only consider unbiased (asymptotically or otherwise) estimators.

6 CONCLUSION AND FUTURE WORK

In this work, we have proposed to unify different off-policy estimators as equivalent additive baseline corrections. We look at off-policy evaluation and learning settings and propose baseline corrections that minimize the variance in the estimated policy value and the empirical gradient of the off-policy learning objective. Extensive experimental comparisons on a synthetic benchmark with realistic settings show that our proposed methods improve performance in the off-policy estimation (OPE) and off-policy learning (OPL) tasks.

We believe our work represents a significant step forward in the understanding and use of off-policy estimation methods (for both evaluation and learning use-cases), since we show that the prevalent SNIPS estimator can be improved upon with essentially no cost, as our proposed method is parameter-free and – in contrast with SNIPS – it retains the unbiasedness that comes with IPS.

Future work may apply a similar approach to offline reinforcement learning setups [37], or consider extensions of our approach for ranking applications [40].

A APPENDIX: OFF-POLICY ESTIMATOR VARIANCE

In this appendix, we report additional results from the experimental section (Section 4 from the main paper), answering RQ4. Specifically, we look at the empirical variance of various offline estimators for the task of off-policy evaluation. The mean squared error (MSE) of

different offline estimators are reported in Figure 4. In this appendix, we report the empirical variance of various offline estimators in Figure 5.

From the figure, it is clear that our proposed β -IPS estimator with estimator variance minimizing β value (Eq. 38) results in the lowest empirical variance in most of the cases. It is interesting to note that when the logged data is limited ($N < 10^3$), sometimes the SNIPS estimator has lower estimator variance. We suspect that the reason could be a bias in the estimate of the variance-optimal β estimate (Eq. 38), when the dataset size is small, given that it is a ratio estimate of expectations. For practical settings, i.e., when $N > 10^3$, the proposed estimator β -IPS results in a minimum sample variance, thereby empirically validating the effectiveness of our proposed β -IPS estimator for the task of OPE.

ACKNOWLEDGMENTS

The authors would like to thank Adith Swaminathan, Thorsten Joachims, Ben London, and Philipp Hager for valuable discussions and feedback on early drafts of this manuscript.

This research was supported by Huawei Finland, the Dutch Research Council (NWO), under project numbers VI.Veni.222.269, 024.004.022, NWA.1389.20.183, and KICH3.LTP.20.006, and the European Union’s Horizon Europe program under grant agreement No 101070212. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] Shubham Baweja, Neeti Pokharna, Aleksei Ustimenko, and Olivier Jeunen. 2024. Variance Reduction in Ratio Metrics for Efficient Online Experiments. In *Proc. of the 46th European Conference on Information Retrieval (ECIR '24)*. Springer.
- [2] Walid Bendada, Guillaume Salha, and Théo Bontempelli. 2020. Carousel Personalization in Music Streaming Apps with Contextual Bandits. In *Proceedings of the 14th ACM Conference on Recommender Systems (RecSys '20)*. ACM, 420–425. <https://doi.org/10.1145/3383313.3412217>
- [3] James Bennett, Stan Lanning, et al. 2007. The Netflix Prize. In *Proceedings of KDD cup and workshop*, Vol. 2007. 35.
- [4] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. 2018. Optimization Methods for Large-scale Machine Learning. *SIAM review* 60, 2 (2018), 223–311.
- [5] Léa Briand, Théo Bontempelli, Walid Bendada, Mathieu Morlon, François Rigaud, Benjamin Chapus, Thomas Bouabça, and Guillaume Salha-Galvan. 2024. Let's Get It Started: Fostering the Discoverability of New Releases on Deezer. In *Proc. of the 46th European Conference on Information Retrieval (ECIR '24)*. Springer.
- [6] Roman Budylin, Alexey Drutsa, Ilya Katsev, and Valeriya Tsoy. 2018. Consistent Transformation of Ratio Metrics for Efficient Online Controlled Experiments. In *Proc. of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18)*. ACM, 55–63. <https://doi.org/10.1145/3159652.3159699>
- [7] Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H. Chi. 2019. Top-K Off-Policy Correction for a REINFORCE Recommender System. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19)*. ACM, 456–464. <https://doi.org/10.1145/3289600.3290999>
- [8] Minmin Chen, Bo Chang, Can Xu, and Ed H. Chi. 2021. User Response Models to Improve a REINFORCE Recommender System. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM '21)*. ACM, New York, NY, USA, 121–129. <https://doi.org/10.1145/3437963.3441764>
- [9] Minmin Chen, Can Xu, Vince Gatto, Devanshu Jain, Aviral Kumar, and Ed Chi. 2022. Off-Policy Actor-Critic for Recommender Systems. In *Proceedings of the 16th ACM Conference on Recommender Systems (RecSys '22)*. ACM, 338–349. <https://doi.org/10.1145/3523227.3546758>
- [10] Peter Dayan. 1991. Reinforcement Comparison. In *Connectionist Models*, David S. Touretzky, Jeffrey L. Elman, Terrence J. Sejnowski, and Geoffrey E. Hinton (Eds.). Morgan Kaufmann, 45–51. <https://doi.org/10.1016/B978-1-4832-1448-1.50011-1>
- [11] Alex Deng, Ya Xu, Ron Kohavi, and Toby Walker. 2013. Improving the Sensitivity of Online Controlled Experiments by Utilizing Pre-Experiment Data. In *Proc. of the Sixth ACM International Conference on Web Search and Data Mining (WSDM '13)*. ACM, 123–132. <https://doi.org/10.1145/2433396.2433413>
- [12] Zhenhua Dong, Hong Zhu, Pengxiang Cheng, Xinhua Feng, Guohao Cai, Xiquang He, Jun Xu, and Jirong Wen. 2020. Counterfactual learning for recommender system. In *Proceedings of the 14th ACM Conference on Recommender Systems (RecSys '20)*. ACM, 568–569. <https://doi.org/10.1145/3383313.3411552>
- [13] Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. 2014. Doubly Robust Policy Evaluation and Optimization. *Statist. Sci.* 29, 4 (2014), 485 – 511. <https://doi.org/10.1214/14-STSS500>
- [14] Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. 2018. More Robust Doubly Robust Off-policy Evaluation. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 1447–1456. <https://proceedings.mlr.press/v80/farajtabar18a.html>
- [15] David A. Freedman. 2008. On Regression Adjustments to Experimental Data. *Advances in Applied Mathematics* 40, 2 (2008), 180–193. <https://doi.org/10.1016/j.aam.2006.12.003>
- [16] Evan Greensmith, Peter L. Bartlett, and Jonathan Baxter. 2004. Variance Reduction Techniques for Gradient Estimates in Reinforcement Learning. *J. Mach. Learn. Res.* 5 (dec 2004), 1471–1530.
- [17] Shashank Gupta, Philipp Hager, Jin Huang, Ali Vardasbi, and Harrie Oosterhuis. 2024. Unbiased Learning to Rank: On Recent Advances and Practical Applications. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*. ACM, 1118–1121. <https://doi.org/10.1145/3616855.3636451>
- [18] Shashank Gupta, Harrie Oosterhuis, and Maarten de Rijke. 2023. A Deep Generative Recommendation Method for Unbiased Learning from Implicit Feedback. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*. 87–93.
- [19] Shashank Gupta, Harrie Oosterhuis, and Maarten de Rijke. 2023. A First Look at Selection Bias in Preference Elicitation for Recommendation (Abstract). In *CONSEQUENCES Workshop at RecSys '23*. ACM.
- [20] Shashank Gupta, Harrie Oosterhuis, and Maarten de Rijke. 2023. Safe Deployment for Counterfactual Learning to Rank with Exposure-based Risk Minimization. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 249–258.
- [21] Shashank Gupta, Harrie Oosterhuis, and Maarten de Rijke. 2024. Practical and Robust Safety Guarantees for Advanced Counterfactual Learning to Rank. In *CIKM 2024: 33rd ACM International Conference on Information and Knowledge Management*. ACM.
- [22] Edward L. Ionides. 2008. Truncated Importance Sampling. *Journal of Computational and Graphical Statistics* 17, 2 (2008), 295–311. <https://doi.org/10.1198/106186008X320456> arXiv:<https://doi.org/10.1198/106186008X320456>
- [23] Olivier Jeunen. 2021. *Offline Approaches to Recommendation with Online Success*. Ph.D. Dissertation. University of Antwerp.
- [24] Olivier Jeunen and Bart Goethals. 2020. An Empirical Evaluation of Doubly Robust Learning for Recommendation. In *Proc. of the ACM RecSys Workshop on Bandit Learning from User Interactions (REVEAL '20)*.
- [25] Olivier Jeunen and Bart Goethals. 2021. Pessimistic Reward Models for Off-Policy Learning in Recommendation. In *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21)*. ACM, 63–74. <https://doi.org/10.1145/3460231.3474247>
- [26] Olivier Jeunen and Bart Goethals. 2021. Top-K Contextual Bandits with Equity of Exposure. In *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21)*. ACM, 310–320. <https://doi.org/10.1145/3460231.3474248>
- [27] Olivier Jeunen and Bart Goethals. 2023. Pessimistic Decision-Making for Recommender Systems. *ACM Trans. Recomm. Syst.* 1, 1, Article 4 (feb 2023), 27 pages. <https://doi.org/10.1145/3568029>
- [28] Olivier Jeunen, Thorsten Joachims, Harrie Oosterhuis, Yuta Saito, and Flavian Vasile. 2022. CONSEQUENCES—Causality, Counterfactuals and Sequential Decision-Making for Recommender Systems. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 654–657.
- [29] Olivier Jeunen, Sean Murphy, and Ben Allison. 2023. Off-Policy Learning-to-Bid with AuctionGym. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*. ACM, 4219–4228. <https://doi.org/10.1145/3580305.3599877>
- [30] Olivier Jeunen, Ivan Potapov, and Aleksei Ustimenko. 2024. On (Normalised) Discounted Cumulative Gain as an Offline Evaluation Metric for Top-*n* Recommendation. In *Proceedings of the 30th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '24)*. arXiv:2307.15053 [cs.IR]
- [31] Olivier Jeunen, David Rohde, Flavian Vasile, and Martin Bompaire. 2020. Joint Policy-Value Learning for Recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. ACM, 1223–1233. <https://doi.org/10.1145/3394486.3403175>
- [32] Thorsten Joachims and Adith Swaminathan. 2016. Counterfactual Evaluation and Learning for Search, Recommendation and Ad Placement. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 1199–1201.
- [33] 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
- [34] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [35] Augustine Kong. 1992. *A Note on Importance Sampling Using Standardized Weights*. Technical Report 348. University of Chicago, Dept. of Statistics.
- [36] Tor Lattimore and Csaba Szepesvári. 2020. *Bandit Algorithms*. Cambridge University Press.
- [37] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems. *arXiv preprint arXiv:2005.01643* (2020).
- [38] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A Contextual-bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web (Raleigh, North Carolina, USA) (WWW '10)*. ACM, 661–670. <https://doi.org/10.1145/1772690.1772758>
- [39] Yaxu Liu, Jui-Nan Yen, Bowen Yuan, Rundong Shi, Peng Yan, and Chih-Jen Lin. 2022. Practical Counterfactual Policy Learning for Top-K Recommendations. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*. ACM, 1141–1151. <https://doi.org/10.1145/3534678.3539295>
- [40] Ben London, Alexander Buchholz, Giuseppe Di Benedetto, Jan Malte Lichtenberg, Yannik Stein, and Thorsten Joachims. 2023. Self-Normalized Off-Policy Estimators for Ranking. In *CONSEQUENCES Workshop at ACM RecSys '23 (CONSEQUENCES '23)*.
- [41] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Ji Yang, Minmin Chen, Jiayi Tang, Lichan Hong, and Ed H. Chi. 2020. Off-Policy Learning in Two-Stage Recommender Systems. In *Proceedings of The Web Conference 2020 (WWW '20)*. ACM, 463–473. <https://doi.org/10.1145/3366423.3380130>
- [42] James McInerney, Benjamin Lacker, Samantha Hansen, Karl Higley, Hugues Bouchard, Alois Gruson, and Rishabh Mehrotra. 2018. Explore, Exploit, and Explain: Personalizing Explainable Recommendations with Bandits. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. ACM, 31–39. <https://doi.org/10.1145/3240323.3240354>
- [43] Rishabh Mehrotra, Niannan Xue, and Mounia Lalmas. 2020. Bandit based Optimization of Multiple Objectives on a Music Streaming Platform. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. ACM, 3224–3233. <https://doi.org/10.1145/3394486.3403374>
- [44] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. 2020. Monte Carlo Gradient Estimation in Machine Learning. *J. Mach. Learn. Res.* 21, Article 132 (jan 2020), 62 pages.
- [45] Art B. Owen. 2013. *Monte Carlo Theory, Methods and Examples*. <https://artowen.su.domains/mc/>.

- [46] Alexey Poyarkov, Alexey Drutsa, Andrey Khalyavin, Gleb Gusev, and Pavel Serdyukov. 2016. Boosted Decision Tree Regression Adjustment for Variance Reduction in Online Controlled Experiments. In *Proc. of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. ACM, 235–244. <https://doi.org/10.1145/2939672.2939688>
- [47] Steffen Rendle. 2022. *Item Recommendation from Implicit Feedback*. Springer US, New York, NY, 143–171. https://doi.org/10.1007/978-1-0716-2197-4_4
- [48] David Rohde, Stephen Bonner, Travis Dunlop, Flavian Vasile, and Alexandros Karatzoglou. 2018. RecoGym: A Reinforcement Learning Environment for the Problem of Product Recommendation in Online Advertising. *arXiv preprint arXiv:1808.00720* (2018).
- [49] Yuta Saito, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. 2020. Open Bandit Dataset and Pipeline: Towards Realistic and Reproducible Off-policy Evaluation. *arXiv preprint arXiv:2008.07146* (2020).
- [50] 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>
- [51] Yuta Saito and Thorsten Joachims. 2022. Counterfactual Evaluation and Learning for Interactive Systems: Foundations, Implementations, and Recent Advances. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4824–4825.
- [52] Yuta Saito, Takuma Udagawa, Haruka Kiyohara, Kazuki Mogi, Yusuke Narita, and Kei Tateno. 2021. Evaluating the Robustness of Off-Policy Evaluation. In *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21)*. ACM, 114–123. <https://doi.org/10.1145/3460231.3474245>
- [53] Otmame Sakhi, Stephen Bonner, David Rohde, and Flavian Vasile. 2020. BLOB: A Probabilistic Model for Recommendation that Combines Organic and Bandit Signals. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. ACM, 783–793. <https://doi.org/10.1145/3394486.3403121>
- [54] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [55] Guy Shani, Ronen I. Brafman, and David Heckerman. 2002. An MDP-based Recommender System. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (Alberta, Canada) (UAI'02)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 453–460.
- [56] Harald Steck. 2013. Evaluation of Recommendations: Rating-prediction and Ranking. In *Proc. of the 7th ACM Conference on Recommender Systems (RecSys '13)*. ACM, 213–220. <https://doi.org/10.1145/2507157.2507160>
- [57] Yi Su, Maria Dimakopoulou, Akshay Krishnamurthy, and Miroslav Dudík. 2020. Doubly Robust Off-policy Evaluation with Shrinkage. In *International Conference on Machine Learning*. PMLR, 9167–9176.
- [58] Yi Su, Lequn Wang, Michele Santacatterina, and Thorsten Joachims. 2019. CAB: Continuous Adaptive Blending for Policy Evaluation and Learning. In *Proc. of the 36th International Conference on Machine Learning (ICML '19, Vol. 97)*. PMLR, 6005–6014. <https://proceedings.mlr.press/v97/su19a.html>
- [59] Yi Su, Xiangyu Wang, Elaine Ya Le, Liang Liu, Yuening Li, Haokai Lu, Benjamin Lipshitz, Sriraj Badam, Lukasz Heldt, Shuchao Bi, Ed H. Chi, Cristos Goodrow, Su-Lin Wu, Lexi Baugher, and Minmin Chen. 2024. Long-Term Value of Exploration: Measurements, Findings and Algorithms. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*. ACM, 636–644. <https://doi.org/10.1145/3616855.3635833>
- [60] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- [61] Adith Swaminathan and Thorsten Joachims. 2015. Batch Learning from Logged Bandit Feedback through Counterfactual Risk Minimization. *The Journal of Machine Learning Research* 16, 1 (2015), 1731–1755.
- [62] Adith Swaminathan and Thorsten Joachims. 2015. The Self-Normalized Estimator for Counterfactual Learning. In *Advances in Neural Information Processing Systems*, Vol. 28. https://proceedings.neurips.cc/paper_files/paper/2015/file/39027dfad5138c9ca0c474d71db915c3-Paper.pdf
- [63] Bram van den Akker, Olivier Jeunen, Ying Li, Ben London, Zahra Nazari, and Devesh Parekh. 2024. Practical Bandits: An Industry Perspective. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*. ACM, 1132–1135. <https://doi.org/10.1145/3616855.3636449>
- [64] Flavian Vasile, David Rohde, Olivier Jeunen, and Amine Benhalloum. 2020. A Gentle Introduction to Recommendation as Counterfactual Policy Learning. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*. ACM, 392–393. <https://doi.org/10.1145/3340631.3398666>
- [65] Ronald J. Williams. 1988. *Toward a Theory of Reinforcement-learning Connectionist Systems*. Technical Report NU-CCS-88-3. Northeastern University.
- [66] Ronald J. Williams. 1992. Simple Statistical Gradient-following Algorithms for Connectionist Reinforcement Learning. *Machine Learning* 8, 3 (01 May 1992), 229–256. <https://doi.org/10.1007/BF00992696>
- [67] Xinyang Yi, Shao-Chuan Wang, Ruining He, Hariharan Chandrasekaran, Charles Wu, Lukasz Heldt, Lichan Hong, Minmin Chen, and Ed H. Chi. 2023. Online Matching: A Real-time Bandit System for Large-scale Recommendations. In *Proceedings of the 17th ACM Conference on Recommender Systems (RecSys '23)*. ACM, New York, NY, USA, 403–414. <https://doi.org/10.1145/3604915.3608792>