

Counterfactual Evaluation of Ads Ranking Models through Domain Adaptation

Mohamed A. Radwan
Meta
USA

Himaghna Bhattacharjee
Meta
USA

Quinn Lanners
Duke University
USA

Jiasheng Zhang
Meta
USA

Serkan Karakulak
Meta
USA

Houssam Nassif
Meta
USA

Murat Ali Bayir
Meta
USA

ABSTRACT

We propose a domain-adapted reward model that works alongside an Offline A/B testing system for evaluating ranking models. This approach effectively measures reward for ranking model changes in large-scale Ads recommender systems, where model-free methods like IPS are not feasible. Our experiments demonstrate that the proposed technique outperforms both the vanilla IPS method and approaches using non-generalized reward models.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Online Advertising**.

KEYWORDS

Recommender Systems, Counterfactual Evaluation, Selection Bias

ACM Reference Format:

Mohamed A. Radwan, Himaghna Bhattacharjee, Quinn Lanners, Jiasheng Zhang, Serkan Karakulak, Houssam Nassif, and Murat Ali Bayir. 2024. Counterfactual Evaluation of Ads Ranking Models through Domain Adaptation. In *ACM Conference on Recommender Systems Workshops, 2024, Bari, Italy*. ACM, New York, NY, USA, 4 pages.

1 INTRODUCTION

Over the past decade, online advertising has undergone significant transformations, particularly with the integration of Artificial Intelligence and Deep Learning technologies. Leveraging the vast amount of traffic managed by popular recommender services, even minor adjustments to such recommendation models can significantly affect business objectives. Understanding and predicting the impact of such model changes in advance highlights the importance of Offline Evaluation in large scale recommender systems [8].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys'24, 2024, Bari, Italy

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

While offline evaluation of ranking models can be a trivial task in traditional machine learning setups, it is a challenge in the context of large-scale ad recommendation systems. In such setups multiple models and various system components such as auction mechanisms [7], content selection for organic feed and ads [2], publisher logic, and client-side app logic together to determine which ads are shown to users. The complexity of these systems significantly affects model-free approaches like Inverse Propensity Scoring (IPS) [3], as the propensity computation must account not only for the model's output, but also for all other factors that contribute to ad display decisions.

Selection bias is another challenge for offline evaluation of large-scale recommender systems [4, 6]. Offline evaluation uses labeled historical data from the current production model to estimate the performance of a new model that does not have any labeled data from online activity. However, with a large inventory of ads, the new offline model will often rank the ads differently than the current online production model. The distribution difference between the ads presented to users by the current production model and the ads that the new model would have presented can result in biased offline reward estimates, making it challenging to accurately predict online performance for a new model [5, 9].

Motivated by the challenges outlined above, we propose a new domain-adapted reward model that works on top of an offline A/B testing system to facilitate more accurate ranking model evaluation. Unlike IPS, our approach eliminates the need to account for the complexity of the ads system since the Offline A/B Testing simulation layer handles this aspect. Our experiments demonstrate that the proposed technique outperforms both the vanilla IPS method and basic direct-method reward models.

2 METHODOLOGY

2.1 Overview

The high level setup for an offline counterfactual evaluation system is shown in Figure 1. It employs an offline serving simulator that operates a separate instance for each target domain (i.e ads recommended by a given ranking model). The simulator runs recommendation request under each domain. Following simulation, the recommended ads for each corresponding target domain are collected. Subsequently, we leverage the reward model to estimate




Figure 1: Counterfactual evaluation setup

the expected reward for each domain. In the following sections, we detail how to train such a generalized reward model to effectively handle multiple domains.

2.2 Notation and set-up

Let $X \subset \mathcal{R}^d$ be a d -dimensional context space of user and request features, $A \subset \mathcal{R}^q$ be a q -dimensional ad space, and $Y \subset \mathcal{R}$ be a one-dimensional reward space (which may be a discrete set or continuous interval depending on the use case).

Each recommendation model defines a policy which induces a domain in the $X \times A$ space. We use S and $T_k \in \{T_1, \dots, T_K\}$ to refer to source and target policies. Finally, we let $p(y | x, a)$ be the conditional distribution of the reward given a particular context and ad combination, which is assumed to be independent of the choice of T_k . The objective is to evaluate a set of K target recommendation models denoted by T_1, \dots, T_K in an offline setting, using data generated by source policy S .

2.3 Metric

Our goal is to correctly rank target policies, and therefore we need the reward model to perform fairly across different domains. To achieve this, we choose the coefficient of variance of recovery (Rec_{cv}) as the main performance indicator for the reward model:


$$Rec_{cv} = \frac{Rec_{dev}}{Rec_{avg}}, \quad Rec(T_k, S) = \frac{\widehat{Lift}(T_k, S)}{Lift(T_k, S)}, \quad (1)$$

where Rec_{dev} is the average absolute deviation of $Rec(\cdot)$, and Rec_{avg} is the average of $Rec(\cdot)$ across all target domains. See appendix A for lift definition and additional details.

2.4 Estimating Lifts Between Domains

We introduce a model based approach to estimate the lift between two domains. Our approach focuses on modeling the non-overlapping regions between source and target domains and then using the trained reward model to calculate the expected lift of the target domain T_k over the source domain S (for each target domain k). Related domain adaptation methods were explored in [1, 10] within the context of the auction optimization problem. The details of the reward model training are discussed in section 2.5.

We use $D_S = \{(x_i, a_i, y_i)\}_{i=1}^n$, a set of n labeled source domain samples, and $D_{T_k} = \{(x_i, a_i^k)\}_{i=1}^n$, unlabeled target domain data for

Figure 2: Rec_{CV} of each method used with synthetic data

the same n samples, to estimate the lift as

$$Lift(T_k, S) = \frac{1}{n} \left(\sum_{(x_i, a_i) \in D_{T_k, n}} h(x_i, a_i) - \sum_{(x_i, a_i) \in D_{S, n}} h(x_i, a_i) \right), \quad (2)$$

where $h(\cdot)$ is the reward model discussed in the preceding paragraph. Because our offline evaluation setup sends the same requests to all domains, as mentioned in section 2.1 and shown in Figure 1, these datasets have the same sample size and sample contexts x_i . However, they differ in the ads they recommend, a_i , and we only have labels, y_i , for the online source recommend model S .

2.5 Reward Model Training

Let $p_c(a | x)$ be the probability of observing ad a under context x with policy c . We define a weight

$$w_a^k = \frac{p_{T_k}(a | x)}{p_S(a | x)}. \quad (3)$$

w_a^k is used in the per sample weights for training a reward model for all K target domains by minimizing the following loss function on the labeled D_S :

$$\sum_{(x_i, a_i, y_i) \in D_S} L[h(x_i, a_i, \theta), y_i] \times \left[\sum_{k=1}^K |w_{a_i}^k - 1| + \beta \sum_{\substack{k=1 \\ k' > k}}^K |w_{a_i}^k - w_{a_i}^{k'}| \right], \quad (4)$$

where θ is the model parameters and β is a hyper-parameter controlling the contribution of the term $|w_{a_i}^k - w_{a_i}^{k'}|$. This term controls the deviation term of Recovery CV by ensuring that the reward model performs equally across all domains, as detailed in Appendix B. The $|w_{a_i}^k - 1|$ term emphasizes the non-overlapping regions between target and source domains.

3 EXPERIMENTAL RESULTS

We report our findings on both synthetic and online experiment results for a CTR prediction model. For the synthetic environment, we generated domain policies, where each target domain or test variant represents an incremental improvement over the source domain or control variant. The reward function is a linear model of the context and ad covariates, where the coefficients are drawn

from a normal distribution based on the ad context vector values. We train the baseline solely on the source domain data. We train our proposed reward model using Section 2.5 weighted target domain information. Figure 2 plots results from three target domains, and shows that our proposed reward outperforms Baseline and IPS.

To evaluate reward models on real experiments, we explored the utilization of a completed A/B test for a CTR prediction model. The test included seven suggested variants that improved upon the control. We used the actual lifts reported per variant as ground truth for evaluating the reward models. Given the intractability of propensity score weight in a complex recommendation systems, we train an impression probability estimator per target domain to estimate Eq 3 weights. The estimated weights are then used to weight the source data D_S , as in Eq 4, to train the proposed reward model. Baseline is trained on D_S without any weighting similar to the synthetic experiment. Our proposed reward model showed a 17.6% improvement on the Recovery CV metric.

REFERENCES

- [1] Murat Ali Bayir, Mingsen Xu, Yaojia Zhu, and Yifan Shi. 2019. Genie: An Open Box Counterfactual Policy Estimator for Optimizing Sponsored Search Marketplace. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (Melbourne VIC, Australia) (WSDM '19). Association for Computing Machinery, New York, NY, USA, 465–473.
- [2] Ari Biswas, Thai T Pham, Michael Vogelsong, Benjamin Snyder, and Houssam Nassif. 2019. Seeker: Real-Time Interactive Search. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. 2867–2875.
- [3] Léon Bottou, Jonas Peters, Joaquín Quiñero Candela, Denis Xavier Charles, Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Y. Simard, and Ed Snelson. 2013. Counterfactual Reasoning and Learning Systems: The Example of Computational Advertising. *Journal of Machine Learning Research* 14 (2013), 3207–3260.
- [4] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Trans. Inf. Syst.* 41, 3, Article 67 (2023).
- [5] Tanner Fiez, Sergio Gamez, Arick Chen, Houssam Nassif, and Lalit Jain. 2022. Adaptive Experimental Design and Counterfactual Inference. In *Workshops of Conference on Recommender Systems (RecSys)*.
- [6] 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 (2021).
- [7] Geng Ji, Wentao Jiang, Jiang Li, Fahmid Morshed Fahid, Zhengxing Chen, Yinghua Li, Jun Xiao, Chongxi Bao, and Zheqing Zhu. 2024. Learning to bid and rank together in recommendation systems. *Machine Learning* 113, 5 (2024), 2559–2573.
- [8] Fanjie Kong, Yuan Li, Houssam Nassif, Tanner Fiez, Shreya Chakrabarti, and Ricardo Henao. 2023. Neural Insights for Digital Marketing Content Design. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. 4320–4332.
- [9] Justin Weltz, Tanner Fiez, Alexander Volfovsky, Eric Laber, Blake Mason, Houssam Nassif, and Lalit Jain. 2023. Experimental Designs for Heteroskedastic Variance. In *Conference on Neural Information Processing Systems (NeurIPS)*. 65967–66005.
- [10] Shuxi Zeng, Murat Ali Bayir, Joseph J. Pfeiffer, Denis Charles, and Emre Kiciman. 2021. Causal Transfer Random Forest: Combining Logged Data and Randomized Experiments for Robust Prediction. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (Virtual Event, Israel) (WSDM '21). Association for Computing Machinery, New York, NY, USA, 211–219.

Appendix A METRIC DETAILS

The metric of interest in our off policy evaluation for ad recommendation models is the performance difference of the offline model with respect to the online source model. We define a lift metric for a target policy T_k :

$$Lift(T_k, S) = E_{p_{T_k}}[y] - E_{p_S}[y]. \quad (5)$$

The expected reward for target policies cannot be directly estimated from the logged data, which brings us to the core problem: for off-policy evaluation, we want to accurately estimate the lift of a target policy in an offline setting. To reflect this use case, we define a *recovery* metric to compare the performance of the reward model for each target policy K , $Rec(T, T_S) = \frac{\widehat{Lift}(T_k, S)}{Lift(T_k, S)}$, where $\widehat{Lift}(T_k, S)$ is the reward lift estimated by the reward model, and $Lift(T_k, S)$ is the realised lift. The recovery metric for a domain where the reward model estimates the lift perfectly is 1.

Moreover, for accurate estimation of policy ranks, we need a metric to perform fairly across all target domains. Therefore, we define an aggregate metric across all target domains using the coefficient of variance:

$$Rec_{cv} = \frac{Rec_{dev}}{Rec_{avg}}, \quad (6)$$

where

$$Rec_{avg}(T_1, T_2, \dots, T_K, S) = \frac{1}{K} \sum_{T \in \{T_1, \dots, T_K\}} Rec(T, S), \quad (7)$$

$$Rec_{dev} = \frac{1}{K} \sum_{T \in \{T_1, \dots, T_K\}} |Rec(T, S) - Rec_{avg}|. \quad (8)$$

Here, Rec_{dev} and Rec_{avg} are the standard deviation and average of Rec across all target domains.

Appendix B DERIVATION OF RECOVERY LOSS

B.1 Single-Domain Recovery Optimization

Let us look at the Recovery metric when there is one target domain T_k . Using the definition of recovery from Eq 1 and the definition for lift from Eq 5, we have $Rec(T_k, S) = 1 - \frac{r_{diff} - \widehat{r_{diff}}}{r_{diff}}$, where r_{diff} and $\widehat{r_{diff}}$ are the true and estimated difference in rewards between the target and source domains, respectively. Hence the optimization goal is equivalent to minimizing $|r_{diff} - \widehat{r_{diff}}|$. With domain distributions, the distance can be further expanded as:

$$\begin{aligned} r_{diff} - \widehat{r_{diff}} &= (E_{p_{T_k}}[y] - E_{p_S}[y]) - (\widehat{E_{p_{T_k}}[y]} - \widehat{E_{p_S}[y]}) \\ &= E_{p_{T_k}}[y(a, x) - \widehat{y(a, x)}] - E_{p_S}[y(a, x) - \widehat{y(a, x)}] \\ &= E_{p_S} \left[\left(\frac{p_{T_k}(a | x)}{p_S(a | x)} - 1 \right) (y(a, x) - \widehat{y(a, x)}) \right] \\ &\approx \frac{1}{N} \times \sum_{(x_i, a_i, y_i) \in D_S} \left(\frac{p_{T_k}(a | x)}{p_S(a | x)} - 1 \right) (y(a, x) - \widehat{y(a, x)}). \end{aligned} \quad (9)$$

It now becomes obvious that the prediction error $|y(a_i, x_i) - \widehat{y(a_i, x_i)}|$ is weighted by $\left(\frac{p_{T_k}(a | x)}{p_S(a | x)} - 1 \right)$, contributing to the overall distance. In other words, samples (a_i, x_i) with higher $\left(\frac{p_{T_k}(a | x)}{p_S(a | x)} - 1 \right)$ are more important for optimizing $Rec(\cdot)$.

B.2 Multi-Domain Optimization

For the multi-domain use-case, the metric to be optimized is $Rec_{cv}(\cdot)$ defined in Eq 6. The Rec_{dev} defined in Eq 8 can be minimized by

minimizing Rec between two target domains T_k and $T_{k'}$:

$$\begin{aligned} Rec_{dev} &= \frac{1}{K} \sum_k \left| Rec_{T_k} - \sum_{k'} \frac{Rec_{T_{k'}}}{K} \right| \\ &= \frac{1}{K} \sum_k \left| \sum_{k'} \frac{Rec_{T_k} - Rec_{T_{k'}}}{K} \right| \\ &\leq \frac{1}{K^2} \sum_k \sum_{k'} |Rec_{T_k} - Rec_{T_{k'}}|. \end{aligned} \quad (10)$$

Furthermore,

$$\begin{aligned} |Rec_{T_k} - Rec_{T_{k'}}| &= \left| \frac{rdiff_{f_{T_k}} - \widehat{rdiff}_{f_{T_k}}}{rdiff_{f_{T_k}}} - \frac{rdiff_{f_{T_{k'}}} - \widehat{rdiff}_{f_{T_{k'}}}}{rdiff_{f_{T_{k'}}}} \right| \\ &\approx \left| \frac{1}{N} \sum_{(x_i, a_i, y_i) \in D_S} \left[\frac{1}{rdiff_{f_{T_k}}} \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} - 1 \right) \right. \right. \\ &\quad \left. \left. - \frac{1}{rdiff_{f_{T_{k'}}} } \left(\frac{p_{T_{k'}}(a|x)}{p_S(a|x)} - 1 \right) \right] (y(a, x) - \widehat{y(a, x)}) \right| \\ &= \left| \frac{1}{N} \sum_{(x_i, a_i, y_i) \in D_S} \left[\right. \right. \\ &\quad \frac{1}{2} \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} - 1 + \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} - 1 \right) \left(\frac{1}{rdiff_{f_{T_k}}} - \frac{1}{rdiff_{f_{T_{k'}}} } \right) \\ &\quad \left. \left. + \frac{1}{2} \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} - \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} \right) \left(\frac{1}{rdiff_{f_{T_k}}} + \frac{1}{rdiff_{f_{T_{k'}}} } \right) \right] \cdot (y(a, x) - \widehat{y(a, x)}) \right| \end{aligned} \quad (11)$$

$$\begin{aligned} &\leq \left| \frac{1}{N} \sum_{(x_i, a_i, y_i) \in D_S} \frac{1}{2} \left[\right. \right. \\ &\quad \left| \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} + \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} - 2 \right) \left(\frac{1}{rdiff_{f_{T_k}}} - \frac{1}{rdiff_{f_{T_{k'}}} } \right) (y(a, x) - \widehat{y(a, x)}) \right| \\ &\quad \left. \left. + \left| \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} - \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} \right) \left(\frac{1}{rdiff_{f_{T_k}}} + \frac{1}{rdiff_{f_{T_{k'}}} } \right) (y(a, x) - \widehat{y(a, x)}) \right| \right] \right| \end{aligned} \quad (12)$$

Eq 12 follows from the triangle inequality. From it we can deduce that $|Rec_{T_k} - Rec_{T_{k'}}|$ can be minimized by minimizing two separate terms. The first is:

$$\begin{aligned} &\frac{1}{N} \sum_{(x_i, a_i, y_i) \in D_S} \frac{1}{2} \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} - 1 + \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} - 1 \right) \\ &\quad \left(\frac{1}{rdiff_{f_{T_k}}} - \frac{1}{rdiff_{f_{T_{k'}}} } \right) (y(a, x) - \widehat{y(a, x)}), \end{aligned} \quad (13)$$

which is similar to reducing the recovery term for each domain as discussed in Appendix B.1. The second term,

$$\begin{aligned} &\frac{1}{N} \sum_{(x_i, a_i, y_i) \in D_S} \frac{1}{2} \left(\frac{p_{T_k}(a|x)}{p_S(a|x)} - \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} \right) \\ &\quad \left(\frac{1}{rdiff_{f_{T_k}}} + \frac{1}{rdiff_{f_{T_{k'}}} } \right) (y(a, x) - \widehat{y(a, x)}), \end{aligned} \quad (14)$$

emphasizes the influence of the error from samples whose ad impression probability is different between target domains. We use these two weights to derive the sample weights for our Eq 4 loss function:

$$\begin{aligned} &\sum_{(x_i, a_i, y_i) \in D_S} L[h(x_i, a_i, \theta), y_i] \times \left[\sum_{k=1}^K \left| \frac{p_{T_k}(a|x)}{p_S(a|x)} - 1 \right| + \right. \\ &\quad \left. \beta \sum_{\substack{k=1 \\ k' > k}}^K \left| \frac{p_{T_k}(a|x)}{p_S(a|x)} - \frac{p_{T_{k'}}(a|x)}{p_S(a|x)} \right| \right]. \end{aligned} \quad (15)$$