



Reinforcement online learning to rank with unbiased reward shaping

Shengyao Zhuang¹ · Zhihao Qiao¹ · Guido Zuccon¹

Received: 16 June 2021 / Accepted: 30 May 2022 / Published online: 4 August 2022
© The Author(s) 2022

Abstract

Online learning to rank (OLTR) aims to learn a ranker directly from implicit feedback derived from users' interactions, such as clicks. Clicks however are a biased signal: specifically, top-ranked documents are likely to attract more clicks than documents down the ranking (position bias). In this paper, we propose a novel learning algorithm for OLTR that uses reinforcement learning to optimize rankers: Reinforcement Online Learning to Rank (ROLTR). In ROLTR, the gradients of the ranker are estimated based on the rewards assigned to clicked and unclicked documents. In order to de-bias the users' position bias contained in the reward signals, we introduce unbiased reward shaping functions that exploit inverse propensity scoring for clicked and unclicked documents. The fact that our method can also model unclicked documents provides a further advantage in that less users interactions are required to effectively train a ranker, thus providing gains in efficiency. Empirical evaluation on standard OLTR datasets shows that ROLTR achieves state-of-the-art performance, and provides significantly better user experience than other OLTR approaches. To facilitate the reproducibility of our experiments, we make all experiment code available at <https://github.com/ielab/OLTR>.

Keywords Online learning to rank · Unbiased reward shaping · Reinforcement learning

1 Introduction

Learning to rank (LTR) is a supervised machine learning technique that has been widely used in modern search engines to learn rankers. *Explicit feedback*, consisting of assessors manually judging the relevance of query-document pairs, is required for LTR (Liu, 2011; Li, 2011). This labelled dataset requirement poses obvious limitations: these datasets take substantial effort and cost to compile (Qin et al., 2010; Qin & Liu, 2013;

✉ Guido Zuccon
g.zuccon@uq.edu.au

Shengyao Zhuang
s.zhuang@uq.edu.au

Zhihao Qiao
zhihao.qiao@uq.edu.au

¹ The University of Queensland, St Lucia, Australia

Chapelle & Chang, 2011), labelling personal documents is unethical and often impossible (Wang et al., 2016), static datasets cannot model user intent change over time (Lefortier et al., 2014; Zhuang & Zuccon, 2021), and user preferences may not agree with that of annotators (Sanderson, 2010).

In order to overcome these limitations, *implicit feedback*, such as clicks, has been leveraged. This type of training signal is not affected by the above limitations and has been an attractive alternative to annotated datasets (Joachims, 2002). However, training rankers with implicit feedback has its own drawbacks and challenges. For example, clicks are a weak relevance signal because they often are affected by a number of biases and noise. One of the most prominent bias in web search is the *position bias*, where higher ranked documents have a higher chance to be observed and thus gain more clicks, even if they may be not relevant (Guan & Cutrell, 2007; Pan et al., 2007; Hofmann et al., 2016; Joachims et al., 2017). Therefore, it is important to consider the influence of such biases.

Two main families of approaches have emerged that attempt to learn effective rankers from users' implicit feedback (Jagerman et al., 2019): *counterfactual learning to rank* (CLTR) (Joachims et al., 2017) and *online learning to rank* (OLTR) (Yue & Joachims, 2009). In CLTR, given a historical click through log, clicks are treated as pure binary relevance labels and *inverse propensity scoring* (IPS) is used to re-weight clicks in order to discount the effect of biases. Rankers are trained in an offline manner and deployed online after training. This offline batch updates pipeline can avoid the risk of exposing users to low-quality results since it only displays the best search engine results that are possible for a given CLTR algorithm and training data.

On the other hand, OLTR algorithms interactively update rankers after each user interaction has taken place, thus being more responsive to a non-stationary user environment (Zhuang & Zuccon, 2021). In contrast to CLTR, current OLTR methods do not directly model position or selection bias, and only assume relevant documents are more likely to be clicked than non-relevant documents (Joachims, 2002). The biases and noise of users clicks are handled by *online interventions* (Jagerman et al., 2019), i.e., slightly perturbed result lists are displayed and preferences towards rankers are informed by users clicks. This is one of the key aspects of OLTR, but also one of its biggest disadvantages that limits OLTR's uptake in practice: such online interventions carry the risk of displaying a "sub-optimal" ranking list directly to the user, thus hurting user experience (as measured by online evaluation metrics). Hence, it is a requirement for OLTR methods to efficiently leverage the click feedback so that a good ranker is learnt as fast as possible to avoid displaying low quality search results to a large number of users.

To move beyond the limitations of existing OLTR methods, in this article, we propose a novel OLTR algorithm called *Reinforcement Online Learning to Rank* (ROLTR), which exploits the reinforcement learning (RL) approach adapting it to OLTR. Our motivation for using RL in OLTR is based on the following observations:

1. RL suits OLTR setting very well: RL is powerful for modeling interactive environments and maximizing the long-term rewards yield from the environment. In the OLTR setting, the interactive environment is composed by the users and the search system, and the rewards are the users' satisfaction.
2. RL with carefully designed reward functions allows OLTR algorithms to directly remove the biases present in the user clicks: this is currently hard for other OLTR algorithms to achieve.

As we show in this article, our proposed ROLTR can directly remove position bias and thus effectively and, importantly, efficiently (i.e. within less impressions) update the ranker. To achieve this, we formalize OLTR as a *Markov Decision Process* (MDP) problem and use *policy gradient* with rewards assigned on clicked and unclicked documents to estimate the update gradients. In order to de-bias users' clicks, we further introduce unbiased reward shaping functions that re-weight the rewards for both clicked and unclicked documents. We mathematically prove that the gradient estimation of ROLTR is unbiased with respect to position bias, and it can directly optimize IR metrics such as discounted cumulative gain (DCG). The idea of leveraging unclicked data has been recently explored in offline counterfactual learning studies (Hu et al., 2019; Wang et al., 2021a), however, our method is currently the only OLTR approach that can gain unbiased learning signals from unclicked documents, thus speeding up convergence. Empirical results further show that ROLTR significantly outperforms traditional OLTR methods and is at par with current state-of-the-art methods (offline performance), although requiring fewer user interactions. As a result, our method delivers considerably better user experience (online performance).

2 Related work

2.1 Counterfactual learning to rank

Unlike traditional LTR where rankers are learned from explicitly labelled datasets (Liu, 2011), counterfactual LTR (Agarwal et al., 2018; Ai et al., 2018; Joachims et al., 2017) uses historical interaction data, typically click logs, to learn a ranker. However, clicks are a biased signal. The most prominent bias in the click signal is the *position bias*: assuming that users examine search engine result pages (SERPs) from top to bottom, then the results that are ranked higher are more likely to be observed by the users (Joachims et al., 2017). Joachims et al. refer to the probability of a search result at a rank i to be observed as its *propensity* (Joachims et al., 2017). They then define the *inverse propensity scoring* (IPS) method to re-weight user clicks: when IPS is used, the estimated ranking score is unbiased with respect to position bias. One crucial requirement of IPS-based CLTR methods is the prior knowledge of user propensity. This is usually estimated by conducting online result randomization which can negatively affect the user experience. To address this issue, Ai et al. (2018) proposed the Dual Learning Algorithm (DLA) to jointly learn an unbiased ranker and an unbiased propensity model, thus avoiding the preprocessing of propensity estimations. Unlike these CLTR works, our method also uses IPS but in the context of reinforcement learning for OLTR and we further propose a new IPS method for unclicked documents so as to gain an unbiased training signal from unclicked documents. This accelerates the learning process because both click and unclick information is used for training, thus increasing the number of implicit signals obtained from each search and used for training.

In addition to position bias, recent work in counterfactual LTR has considered correcting the *selection bias* (Ovaisi et al., 2020; Oosterhuis & de Rijke, 2020a), in which some documents have zero probability of being observed by the users (thus never having a chance to be clicked and contribute a training signal). We note that the selection bias is fundamentally different from the position bias in terms of that this bias is introduced by the system itself: some documents will never be included by the system in the top SERPs. In contrast, the position bias instead comes from the users. This often occurs in web search




Fig. 1 A schematic view of the online learning to rank process. Users pose queries to the ranking system. This uses an online LTR algorithm to identify the candidate documents to display to the user. Depending on the OLTR algorithm used, candidate documents will be used to form a displayed list. The user examines the displayed search results (SERP) and clicks on items of interest. The click feedback is used by the OLTR algorithm to perform ranker updates

where SERPs only show a small subset of documents in the first page (e.g., 10) and users do not proceed beyond the first SERP. In this circumstance, documents ranked beyond the first SERP have no chance of being observed, hence never get identified as positive training examples. To debias the selection bias in the click signals, Oosterhuis and de Rijke (2020a) proposed a policy-aware counterfactual estimator for CLTR to directly account for the selection bias introduced by a stochastic logging policy. On the other hand, Ovaisi et al. (2020) adapted Heckman’s two-stage method to account for selection and position bias in LTR systems. In our work, we follow the standard OLTR and CLTR experimental setup and simulate typical, real-world circumstances; we also include selection bias in our experiments by only placing 10 documents in the SERPs. Although in this paper we directly focus on correcting position bias, selection bias is also partially corrected by the proposed method via online intervention. Nevertheless, recently proposed methods (Ovaisi et al., 2020; Oosterhuis & de Rijke, 2020a) could be used to further modify our reward shaping functions: this would have the effect of further reducing the selection bias of the gradient estimations.

2.2 Online learning to rank

Similarly to CLTR, online LTR (OLTR) also considers implicit user feedback to learn a ranker. Unlike CLTR, however, this is done online, by directly interacting with users. This online training process allows to control data acquisition and handle biases and noise through online interventions with regards to which documents to display. Figure 1 provides a schematic representation of the OLTR process.

The *Dueling Bandit Gradient Descent* (DBGD) (Yue & Joachims, 2009) models OLTR as a dueling bandit problem and uses *online evaluation* with users’ clicks on interleaved or multileaved SERPs to indicate user preference among a pool of candidate rankers (Schuth et al., 2015; Oosterhuis et al., 2016; Hofmann et al., 2011a, 2013; Schuth et al., 2016; Hofmann et al., 2011b). More recent work has studied methods to reduce the variance of the gradients estimated by online evaluation (Wang et al., 2019, 2018a). However, because gradients are updated towards the winning rankers in the candidate pool, the quality of the gradient estimation is influenced by the number of candidate rankers. When the number of candidate rankers is large, the performance of DBGD is limited in terms of both effectiveness and efficiency (Schuth et al., 2015; Li et al., 2020). The recently proposed *Counterfactual Online Learning to Rank* (COLTR) (Zhuang & Zuccon, 2020) attempts to overcome

the issues associated with online evaluation. This is done by replacing online evaluation with counterfactual evaluation for DBGD. Instead of interleaving, COLTR uses clicks collected by the current ranker to evaluate candidate rankers thus providing high efficiency. However, empirical results have shown COLTR requires more exploration, thus possibly hurting user experience (online evaluation metrics).

Unlike DBGD and COLTR, the *Pairwise Differentiable Gradient Descent* (PDGD) (Oosterhuis & de Rijke, 2018) does not require to sample candidate rankers for online evaluation. Instead, PDGD directly estimates gradients based on pairwise preferences between documents in the SERP, inferred by users' clicks; then, stochastic gradient descent is used to update the ranker. The gradient estimation of PDGD is unbiased with respect to user document pair preferences (Oosterhuis & de Rijke, 2018). PDGD is empirically found to be significantly better than DBGD in terms of final convergence, learning speed and user experience during optimization, making PDGD the current state-of-the-art method for OLTR (Jagerman et al., 2019; Zhuang & Zuccon, 2020; Oosterhuis & de Rijke, 2019; Wang et al., 2021c). PDGD has also been adapted to the federated OLTR context (Wang et al., 2021b), exhibiting again state-of-the-art performance. However, both PDGD and DBGD based methods do not directly use IPS to handle position bias, which has been proven to be very important in previous CLTR results (Joachims et al., 2017).

As mentioned above, CLTR and OLTR differ from the setting considered for learning the ranker: use an historical log of implicit interactions for counterfactual, versus use direct observation and interaction with rankings in the case of online LTR. We note that Jagerman et al. (2019) have specifically studied the similarity and differences between counterfactual and online LTR. They suggested that, from a theoretical standpoint, counterfactual LTR exploits position bias better, however they also indicated that empirical results have shown that OLTR (and in particular PDGD) is more reliable. On the other hand, recent works have been focusing on adapting the offline counterfactual learning to the online setting (Ai et al., 2021; Zhuang & Zuccon, 2020; Oosterhuis & de Rijke, 2020b, 2021). These works have suggested that OLTR algorithms can benefit from the counterfactual learning framework. Based on this context, our proposed method can also be thought of as embracing this direction.

2.3 Reinforcement learning to rank

Reinforcement learning (RL) has previously been applied to offline LTR problems, but not to online LTR. Wei et al. (2017) have formalized ranking as a Markov Decision Process (MDP) problem and introduced the MDPRank algorithm to optimize a linear ranker. Specifically, the document ranking is modelled as a sequential decision making process where each time step corresponds to a position in the ranking and the action taken at each time step corresponds to the selection of a document for the corresponding position. The rewards given by the environment are generated according to the relevance label of the documents, and the classic policy gradient algorithm of REINFORCE (Sutton et al., 2000) is used to maximize the expectation of cumulative rewards received by the ranker. An attractive property of MDPRank is that it can use gradient descent to directly optimize non-differentiable ranking metrics such as DCG. However, Jun et al. (2020) have recently shown that the gradient estimation of the original MDPRank exhibits high variance, and thus this method requires more training episodes to learn an effective ranker. A similar reinforcement learning framework has also been used for search result diversification (Feng

et al., 2018; Xia et al., 2017), multi-page search (Zeng et al., 2018) and recommendation system (Zhao et al., 2018a, b, 2021).

In this paper, we also apply a MDPRank-like algorithm, although in the context of OLTR and with two important adaptations. First, the original MDPRank is designed for traditional offline LTR, i.e., the relevance labels of query-document pairs are provided. In the OLTR setting, however, only biased user interaction data such as clicks is available: thus the rewards given by the online environment are biased. In order to obtain unbiased rewards, we introduce unbiased reward shaping functions for MDPRank to discount position bias. Secondly, the high variance of the gradient estimation of MDPRank makes it converge slowly. However, for OLTR to be viable, we need the algorithm to converge fast so that it does not hurt the user experience too much. To reduce variance, we simplify the objective function of the policy gradient used in the original MDPRank to achieve an objective function with much lower variance of gradient estimation, without changing its optimization target.

3 Method

3.1 Ranking as a markov decision process for online learning to rank

Ranking can be formalized as an MDP problem, where the search engine (agent) has to decide which document to place (action) at rank i , given the current candidate document set (state) (Wei et al., 2017). In the offline LTR setting, i.e., when relevance labels are provided, any IR evaluation measure can be used as a reward function for the ranking created by the search engine. In the online LTR setting, however, rewards are typically inferred via users' implicit feedback, such as clicks. Following Wei et al. (2017), who have limited their attention to offline LTR only, we define each component of the MDP ranking for the OLTR as follows:

States S indicates the set of states that the agent (search engine) will observe from the environment. In the MDP ranking, at a time step t , a state $s_t \in S$ is the candidate document set D_t which contains the documents that need to be ranked.

Actions A is the set of possible actions that the agent can take when at a state s_t . An action at a time step t , denoted as $a_t \in A(s_t)$, consists of selecting a document $d_{m(a_t)} \in D_t$ to place at rank i , where rank $k = t + 1$,¹ $m(a_t)$ is the function that maps the action to the document index, and $d_{m(a_t)}$ is represented by a feature vector. We note the feature elements in $d_{m(a_t)}$ are not only obtained based on the document but also on the query.

Transition $T(s, a)$ is the function $T(s_t, a_t) = s_{t+1}$, which maps a state s_t to its next state s_{t+1} based on the action a_t , where s_{t+1} is the candidate document set D_t without the selected document $d_{m(a_t)}$:

$$s_{t+1} = s_t \setminus d_{m(a_t)} = D_t \setminus d_{m(a_t)} = D_{t+1} \quad (1)$$

Policy $\pi(a|s, \theta)$ is a probability distribution over all possible actions that the agent can take when in the state s , given the current model parameters θ (where θ is a vector containing the model's parameters). We compute this distribution using the softmax function over document scores:

¹ Note that the time step t starts from 0 and the rank position k starts from 1.

$$\pi(a_t | s_t, \theta) = \frac{\exp(f_\theta(d_{m(a_t)}))}{\sum_{a \in A(s_t)} \exp(f_\theta(d_{m(a)}))} \quad (2)$$

$f_\theta(d_{m(a_t)})$ is the document relevance score estimated by the ranking model. In our experiments, we use a linear ranking model, this means the dimensionality of θ , i.e., the number of parameters in the ranking model, is equal to the dimensionality of the feature vector. However, we note that the ranking model can be extended to any dimensionality, e.g., a neural ranking model. For simplicity, we use $a_t \sim \pi(s_t)$ to denote the action a_t selected according to the policy π at state s_t .

Reward $R(s, a)$ is the reward function that models the immediate reward given by the environment and its value represents the reward assigned to the action a that has been selected by the agent in state s .

When full information is provided (i.e., the true relevance labels are given), it is straightforward to define the reward function on the basis of IR evaluation measures such as DCG. However, in the OLTR setting, relevance labels are not provided – in place of these, the ranker observes implicit, noisy and biased user feedback such as clicks. Thus, in such a partial information setting, the reward function is defined based on the implicit feedback, e.g., on click labels. We define the value of the reward function next.

3.2 Unbiased reward shaping

In RL, reward shaping is used to reshape the original reward function to better guide the direction of the gradient update (Ng et al., 1999). Prior knowledge about the environment is needed to formalize a reliable reward shaping function to avoid otherwise to bias learning (Popov et al., 2017).

3.2.1 Naïve reward function

In the (offline) LTR approach by Wei et al. (2017), DCG scores have been used as the original reward function to instruct the search engine to learn a ranker, when true relevance labels are supplied. This reward function $R_{DCG}(s_t, a_t)$ is defined as:

$$R_{DCG}(s_t, a_t) = \lambda(t) \cdot y_{m(a_t)} \quad (3)$$

where $\lambda(t) = \frac{1}{\log_2(t+2)}$ is the DCG weight term, and $y_{m(a_t)} \in \{0, 1\}$ is the binary relevance label of the document at rank $t + 1$. This is a reasonable reward function, provided that relevance labels are known. This reward function directly corresponds to the DCG evaluation metric (because of the DCG weight term), and thus the agent will attempt to learn a policy π that directly maximizes the cumulative rewards Δ (DCG scores) of the ranking episode, for any given initial state s_0 :

$$\Delta_{R_{DCG}}(s_0, \pi, y) = \sum_{t=0} R_{DCG}(s_t, \pi(s_t)) = \sum_{y_{m(a_t)}=1} \lambda(t) \quad (4)$$

However, in the *online* LTR setting, relevance labels are unknown (i.e., this is a partial information setting); instead, users' clicks are used as learning signal. A naïve reward function in this context is then to treat clicks as binary relevance labels and thus use the DCG reward function of equation 3:

$$R_{NAIVE^+}(s_t, a_t) = \lambda(t) \cdot c_{m(a_t)} \tag{5}$$

where $c_{m(a_t)} = 1$ represents that a user has clicked on the document at rank $t + 1$, while $c_{m(a_t)} = 0$ represents no click on the document. If we assume that the probability of a document $d_{m(a_t)}$ to be observed by a user (known as propensity) only depends on the rank position $t + 1$ (note that t starts from 0) and no click noise is present, i.e., $P(o_{t+1} = 1 | t + 1)$, then the expectation of the final ranking rewards is biased to the users' propensity (Joachims et al., 2017; Jagerman et al., 2019):

$$\begin{aligned} \mathbb{E}_o[\Delta_{R_{NAIVE^+}}(s_0, \pi, c)] &= \sum_{t=0} R_{NAIVE^+}(s_t, \pi(s_t)) \\ &= \sum_{y_{m(a_t)}=1} P(o_{t+1} = 1 | t + 1) \cdot \lambda(t) \end{aligned} \tag{6}$$

This is the so-called position bias effect: top-ranked documents usually have a larger chance of being observed, and thus the rewards assigned to them are biasedly higher.

3.2.2 IPS reshaping

Recent work in CLTR has attempted to account for this position bias effect (Ai et al., 2018; Joachims et al., 2017). One of the most commonly used approaches to mitigate position bias is *Inverse Propensity Scoring* (IPS). We follow this direction to unbiased the click signal in our RL framework, and define the IPS reward shaping function as F_{IPS^+} and the reshaped reward function as R_{IPS^+} :

$$\begin{aligned} R_{IPS^+}(s_t, a_t) &= F_{IPS}(R_{NAIVE^+}(s_t, a_t)) \\ &= \frac{R_{NAIVE^+}(s_t, a_t)}{P(o_{t+1} = 1 | t + 1)} = \frac{\lambda(t)}{P(o_{t+1} = 1 | t + 1)} \cdot c_{m(a_t)} \end{aligned} \tag{7}$$

It can be proven that the cumulative naïve rewards reshaped by F_{IPS^+} are an unbiased estimate of the ranking rewards with true relevance labels (Joachims et al., 2017): $\mathbb{E}_o[\Delta_{R_{IPS^+}}(s_0, \pi, c)] = \Delta_{R_{DCG}}(s_0, \pi, y)$. Thus, F_{IPS} can be used to obtain unbiased cumulative rewards with respect to position bias, likely providing more reliable gradient estimations than the naïve reward function.

3.2.3 Negative rewards

The reward functions above only provide positive rewards to the clicked/relevant documents, while they assign a zero reward to the unclicked/irrelevant documents. However, in other RL problems, it is often found that negative rewards help the agent to avoid selecting poor actions (Sutton & Barto, 2018): for a ranker, negative rewards can help avoid selecting irrelevant documents from the document set, thus leading to better user experience. Based on this, we introduce a negative DCG reward function for online LTR (a similar one could be defined for offline LTR):

$$R_{NAIVE^-}(s_t, a_t) = \lambda(t) \cdot (c_{m(a_t)} - 1) \tag{8}$$

Hence, the negative cumulative ranking score is calculated by:

$$\Delta_{R_{NAIVE^-}}(s_0, \pi, c) = \sum_{t=0} R_{NAIVE^-}(s_t, \pi(s_t)) = \sum_{c_{m(a_t)}=0} -\lambda(t) \tag{9}$$

This means highly ranked unclicked documents will be penalised more and thus drag down the final ranking reward. It can be proven that this naïve negative reward function is a biased estimate of the negative DCG score of the ranking ($\Delta_{R_{DCG^-}}(s_0, \pi, y)$):

$$\begin{aligned} & \mathbb{E}_o[\Delta_{R_{NAIVE^-}}(s_0, \pi, c)] \\ &= \mathbb{E}_o\left[\sum_{t=0} R_{NAIVE^-}(s_t, \pi(s_t))\right] = \mathbb{E}_o\left[\sum_{c_{m(a_t)}=0} -\lambda(t)\right] \\ &= \mathbb{E}_o\left[\sum_{y_{m(a_t)}=0} -\lambda(t)\right] + \mathbb{E}_o\left[\sum_{o_{t+1}=0 \wedge y_{m(a_t)}=1} -\lambda(t)\right] \tag{10} \\ &= \sum_{y_{m(a_t)}=0} -\lambda(t) - \sum_{y_{m(a_t)}=1} (1 - P(o_{t+1} = 1|t + 1)) \cdot \lambda(t) \\ &= \Delta_{R_{DCG^-}}(s_0, \pi, y) - \sum_{y_{m(a_t)}=1} (1 - P(o_{t+1} = 1|t + 1)) \cdot \lambda(t) \end{aligned}$$

That is, the naïve negative reward overestimates the true negative DCG reward by counting the relevant but not observed documents as irrelevant.

3.2.4 Negative IPS reward shaping function

In order to get unbiased negative DCG rewards, we propose the negative IPS reward shaping function F_{IPS^-} and denote the reshaped negative DCG reward as R_{IPS^-} :

$$\begin{aligned} R_{IPS^-}(s_t, a_t) &= F_{IPS^-}(R_{NAIVE^-}(s_t, a_t)) \\ &= \lambda(t) \cdot (c_{m(a_t)} - 1) + \left(\frac{1 - P(o_{t+1} = 1|t + 1)}{P(o_{t+1} = 1|t + 1)}\right) \lambda(t) \cdot c_{m(a_t)} \tag{11} \end{aligned}$$

We mathematically prove that $\mathbb{E}_o[\Delta_{R_{IPS^-}}(s_0, \pi, c)]$ is an unbiased estimate of $\Delta_{R_{DCG^-}}(s_0, \pi, y)$:

$$\begin{aligned}
 \mathbb{E}_o[\Delta_{R_{IPS^-}}(s_0, \pi, c)] &= \mathbb{E}_o \left[\sum_{t=0} F_{IPS^-}(R_{NAIVE^-}(s_t, \pi(s_t))) \right] \\
 &= \mathbb{E}_o \left[\sum_{t=0} \lambda(t) \cdot (c_{m(a_t)} - 1) + \left(\frac{1 - P(o_{t+1} = 1|t+1)}{P(o_{t+1} = 1|t+1)} \right) \lambda(t) \cdot c_{m(a_t)} \right] \\
 &= \mathbb{E}_o \left[\sum_{t=0} \lambda(t) \cdot (c_{m(a_t)} - 1) \right] + \\
 &\quad + \mathbb{E}_o \left[\sum_{t=0} \left(\frac{1 - P(o_{t+1} = 1|t+1)}{P(o_{t+1} = 1|t+1)} \right) \lambda(t) \cdot c_{m(a_t)} \right] \\
 &= \mathbb{E}_o[\Delta_{R_{NAIVE^-}}(s_0, \pi, c)] + \\
 &\quad + \mathbb{E}_o \left[\sum_{o_{t+1}=1 \wedge y_{m(a_t)}=1} \left(\frac{1 - P(o_{t+1} = 1|t+1)}{P(o_{t+1} = 1|t+1)} \right) \lambda(t) \right] \\
 &= \Delta_{R_{DCG^-}}(s_0, \pi, y) - \sum_{y_{m(a_t)}=1} (1 - P(o_{t+1} = 1|t+1)) \cdot \lambda(t) \\
 &\quad + \sum_{y_{m(a_t)}=1} (1 - P(o_{t+1} = 1|t+1)) \cdot \lambda(t) \\
 &= \Delta_{R_{DCG^-}}(s_0, \pi, y)
 \end{aligned} \tag{12}$$

This reward shaping function allows to assign unbiased negative rewards to unclicked documents so as to obtain an unbiased cumulative negative DCG ranking score.

3.2.5 Prior knowledge of propensity

All the unbiased reward shaping functions above require to know a priori the users' propensities $P(o_{t+1} = 1|t+1)$. Many recent works have considered estimating such propensities from historical click-logs (Agarwal et al., 2017; Fang et al., 2019; Agarwal et al., 2019) and during training (Ai et al., 2018; Wang et al., 2018b). We regard propensity estimation as being beyond the scope of this article and in our experiments we assume the propensities to be known. Nevertheless, for completeness, we also test how sensitive our method is to propensity mismatch (Sect. 5.5).

3.3 Learning with policy gradient

Following previous work (Wei et al., 2017; Yao et al., 2020), we learn the policy model parameters θ with REINFORCE (Williams, 1992; Sutton & Barto, 2018; Sutton et al., 2000), a widely used policy gradient algorithm. In REINFORCE, the objective is to find an optimal policy that can maximize the expectation of cumulative reward from the beginning of each episode, $J(\theta) = \mathbb{E}_\pi[G_t]$, where $t = 0$. Here G_t is the discounted future cumulative reward starting from time step t , $G_t = \sum_{m=t}^M \gamma^{m-t} \cdot R(s_m, a_m)$, and M is the maximum depth of the ranking episode, and $\gamma \in [0, 1]$ is the discount factor. In previous offline LTR work (Wei et al., 2017; Jun et al., 2020), γ has been set to 1 (maximum value), which results in $G_0 = \Delta_{R_{DCG}}$. This value results in the learning algorithm directly maximizing DCG. However, it is well known in RL that larger γ values will lead the agent to care more about

future rewards but at the same time to produce gradient estimations with significantly high variance, thus slowing down the learning speed (Jun et al., 2020; Sutton & Barto, 2018). In OLTR settings, however, learning speed is very important as it is highly entangled with user experience, as measured by the online performance of the ranker.

Hence, in order to reduce variance and speed up learning, we simplify MDP to Contextual Bandits (Jagerman et al., 2020; Adomavicius et al., 2005; Hofmann et al., 2011c) by setting $\gamma = 0$. This setting makes REINFORCE to choose a_t so as to maximize only the expectation of immediate reward $R(s_t, a_t)$:

$$J(\theta) = \mathbb{E}_{\pi} [R(s_t, a_t)] \quad (13)$$

This setting actually does not change the fact that the objective is to optimize DCG. This is because maximizing the expected immediate reward at a time step t is equivalent to selecting the most likely relevant document from the candidate set D_t for the state s_t : this is guaranteed to have maximum expected cumulative reward for that ranking episode. Thus, it is safe to ignore future rewards in this case. In Sect. 5.1, we empirically show that $\gamma = 0$ enjoys a much faster learning speed than other settings, without any loss in final convergence. It is important to note that the above is not true for tasks such as search result diversification, since greedily choosing the most relevant document at each rank position may cause lower final ranking scores (Feng et al., 2018). Thus, we use the full MDP algorithm for this special situation.

Following the standard policy gradient practice (Wei et al., 2017; Sutton & Barto, 2018), we estimate gradients with *Monte Carlo sampling* and the gradient $\nabla_{\theta} J(\theta)$ can be calculated as:

$$\nabla_{\theta} J(\theta) = R(s_t, a_t) \nabla_{\theta} \log(\pi(a_t | s_t, \theta)) \quad (14)$$

Intuitively, the gradients will update the ranker parameters towards the actions that yield the highest immediate reward $R(s_t, a_t)$.

The complete procedure of ROLTR is described in Algorithm 1. At each ranking episode i , i.e., at each round of user interaction (line 2 for-loop), the search engine receives a query q_i and the initial candidate document set D_0 is generated (line 3). Then the algorithm first draws an action (document) from the distribution created by the current policy (line 6). Then, the selected action, the state information, and the corresponding document are recorded (line 7) and the environment moves to the next state (line 8). The same procedure is repeated for the next rank positions, until the algorithm reaches the maximum depth for the ranking (line 5 for-loop). After finishing ranking, the final result list is shown to the user, who provides feedback to the search engine in the form of click labels (line 10). Next, for each recorded state-action pair (line 11 for-loop), the gradient is calculated (line 13) and the current policy is updated at the end of each ranking episode (line 15).

Algorithm 1 Reinforcement Online Learning to Rank (ROLTR).

```

1: Input: Initial weights  $\theta$ , learning rate  $\alpha$ , reshaped reward function  $R$ , transition function
    $T$ , number of docs on SERPs  $M$ .
2: for  $i \leftarrow 0 \dots \infty$  do
3:    $q_i \leftarrow \text{receive\_query}(i)$ ,  $D_0 \leftarrow \text{candidate\_set}(q_i)$ 
4:    $s_0 \leftarrow D_0$ ,  $L_i \leftarrow [ ]$ ,  $\Delta\theta \leftarrow 0$ 
5:   for  $t \leftarrow 0 \dots M$  do
6:      $a_t \leftarrow \pi(s_t, \theta)$  // Eq.2
7:      $L_i.append(s_t, a_t, d_{m(a_t)})$ 
8:      $s_{t+1} \leftarrow T(s_t, a_t)$  // Eq.1
9:   end for
10:   $C_i \leftarrow \text{receive\_clicks}(L_i)$ 
11:  for  $t \leftarrow 0 \dots M$  do
12:     $s_t, a_t \leftarrow L_i.get(t)$ ,  $c_{m(a_t)} \leftarrow C_i.get(t)$ 
13:     $\Delta\theta \leftarrow \Delta\theta + R(s_t, a_t) \nabla_{\theta} \log(\pi(a_t | s_t, \theta))$  // Eq.14
14:  end for
15:   $\theta \leftarrow \theta + \alpha \Delta\theta$ 
16: end for

```

3.4 Guarantee of Unbiasedness for ROLTR

ROLTR has two main parts: the learning part (Sect. 3.3), which uses the REINFORCE algorithm, and the unbiased reward shaping part (Sect. 3.2). The theoretical guarantees of REINFORCE are well studied in reinforcement learning, including the guarantee that the gradient estimate is unbiased with respect to maximizing the expected rewards (Sutton & Barto, 2018). However, if the rewards are biased (which are when the naive reward function is used), then the learning algorithm is biased. This is where the second part of ROLTR comes into play: the unbiased rewards shaping functions guarantee the reward signals assigned to REINFORCE are unbiased with respect to position bias (mathematical proof in Eq. 12). Thus, the update gradients estimated by ROLTR are guaranteed to maximize unbiased expected rewards, and thus the gradient estimation is unbiased with respect to position bias.

4 Empirical evaluation

To study the effectiveness of ROLTR, we designed a number of empirical experiments aimed to answer the following research questions:

- RQ1 How does the reward discount factor in ROLTR affect gradient variance and final convergence (i.e., the offline nDCG score on the test dataset)?
- RQ2 How do the unbiased reward shaping functions of ROLTR impact performance?
- RQ3 How does ROLTR compare in terms of convergence and learning speed against current representative OLTR methods?
- RQ4 Does ROLTR deliver better user experience than current OLTR methods, i.e., higher online nDCG?
- RQ5 How sensitive is ROLTR to propensity mismatch?

Table 1 Click probabilities for different user behaviours

| $rel(d)$ | 0 | $P(c = 1 o = 1, rel(d))$ | | | |
|----------------|-----|----------------------------|-----|-----|-----|
| | | 1 | 2 | 3 | 4 |
| <i>perfect</i> | 0.0 | 0.2 | 0.4 | 0.8 | 1.0 |
| <i>noisy</i> | 0.4 | 0.6 | 0.7 | 0.8 | 0.9 |

4.1 Datasets and synthetic data generation

We consider three benchmark datasets that are commonly used to evaluate OLTR (Zhuang & Zuccon, 2020; Schuth et al., 2015; Oosterhuis et al., 2016; Oosterhuis & de Rijke, 2018; Wang et al., 2019) and CLTR (Jagerman et al., 2019; Ai et al., 2018; Joachims et al., 2017): MSLR-WEB10K (Qin & Liu, 2013), Yahoo! Webscope (Chapelle & Chang, 2011), and Istella (Dato et al., 2016). MSLR-WEB10K contains 10,000 queries and 125 retrieved documents on average; documents are represented by 136 features. Yahoo! is a bigger dataset, with 29,921 queries and an average of 23.7 documents per query, represented using 700 features. Istella is the largest dataset we consider, with 33,118 queries and an average of 315 documents per query, represented by 220 features. Query-document pair relevance labels for all datasets are recorded on a five-point scale from not relevant (0) to perfectly relevant (4) and have been split into training, validation and test sets (according to the standard splits in the datasets). Queries in the three sets are disjoint. We use the training set to train the rankers, the validation set to tune the hyper-parameters, and the test set to evaluate the rankers' performance.

To avoid hurting user experience, it is common for research in OLTR and CLTR to simulate users' clicks by relying on the relevance labels recorded in the datasets (Zhuang & Zuccon, 2020; Schuth et al., 2015; Oosterhuis et al., 2016; Oosterhuis & de Rijke, 2018; Jagerman et al., 2019; Joachims et al., 2017; Vardasbi et al., 2020). This also allows to fully control users' biases and noise so that algorithms can be tested under different, controllable conditions. Queries are uniformly sampled from the dataset (sampling query IDs). The candidate document set associated with the query ID is provided by each dataset. Then OLTR algorithms generate a result list of documents to display. Clicks are simulated based on two fixed variables: the click probability and the position bias.

The click probability is the probability of a user clicking on a document after observing it. This probability is conditioned on the document's relevance label. Following previous OLTR work (Oosterhuis & de Rijke, 2018; Zhuang & Zuccon, 2020; Hofmann et al., 2011a, 2013; Oosterhuis et al., 2016), we set two types of click behaviour: *perfect* and *noisy*. The click probability of the *perfect* click behaviour is proportional to the relevance level of the documents, and has 0 probability for non-relevant documents. This simulates an ideal user that is able to always determine the relevance of a document in the SERP. The *noisy* click behaviour mimics instead a realist behaviour on SERPs by assigning a small click probability to non-relevant documents and a small skip probability to relevant documents. Table 1 provides the click probabilities for the two user models.

Position bias is modelled by the document observation probabilities; we assume the observation probabilities only depend on the rank position of the document and set these probabilities to:

$$P(o_k = 1 | k) = \left(\frac{1}{k}\right)^\eta \quad (15)$$

where k is the rank position and η is a parameter that determines the level of position bias. Following Joachims et al. (2017) and Jagerman et al. (2019), we set $\eta = 1$. Thus, the probability of a click occurring on a document at rank k in the result list is:

$$P(c_k = 1) = P(c_k = 1 | o_k = 1, rel(d_k))P(o_k = 1 | k) \quad (16)$$

4.2 Evaluation measures

We measure effectiveness using standard OLTR evaluation practice, which considers two aspects: *offline* and *online* performance.

Offline performance is the final convergence of the learned ranker. We evaluate this using the average nDCG@10 of the ranker over the queries in the held-out test-set across 100, 000 impressions, as all rankers would have reached convergence at this point.

Online performance measures user experience during training. This is quantified by the nDCG@10 obtained by the rank list L_i that the user observes in the training episode i , times a discount rate $\tau < 1$:

$$online_performance = \sum_i \tau^i \cdot NDCG(L_i) \quad (17)$$

The discount rate assigns less weight to the later impressions to reward OLTR algorithms that learn an effective ranker fast, so to limit the amount of low-quality user experience. As in previous work (Oosterhuis & de Rijke, 2018; Wang et al., 2019; Zhuang & Zuccon, 2020; Hofmann et al., 2013), we set $\tau = 0.9995$: this means that impressions that occur after 10, 000 iterations have less than a 1% impact. We further note that, while this online performance measure aims to quantify the user experience during training, it only does so partially: the relevance of the results in SERP is in fact only one of the many aspects influencing user experience (Al-Maskari & Sanderson, 2010; Maxwell et al., 2017). These other aspects however are not measurable in the context of the typical simulated experiments performed in OLTR research.

4.3 Experimental runs

We compare ROLTR to four OLTR baselines. First, the *Dual Bandit Gradient Descent* (DBGD) (Yue & Joachims, 2009) method is used as it is one of the standard and most influential algorithms for OLTR. This method uses interleaving for online evaluation, where only one candidate rank is compared to the production ranker at each update step.

The second baseline we consider is the *Probabilistic Multileaving Gradient Descent* (PMGD) (Oosterhuis et al., 2016), which has been reported to be the best traditional OLTR method that uses online evaluation with multileaving comparison (Oosterhuis & de Rijke, 2018; Zhuang & Zuccon, 2020). For this baseline, we use the same hyper-parameters settings reported in previous work (Zhuang & Zuccon, 2020), where the number of candidate rankers is $n = 49$, the step size is $\delta = 1$ and learning rate is $\alpha = 0.01$.

The third baseline we consider is *Counterfactual Online Learning to Rank* (COLTR) (Zhuang & Zuccon, 2020), which uses DBGD to update candidate rankers but is combined with counterfactual evaluation. This method has been reported to be more effective than online evaluation based methods in terms of final convergence, but it does present an

overall deterioration of user experience during the learning cycle. For this baseline, we use the best hyper-parameters from the original paper (Zhuang & Zuccon, 2020).

The last baseline we consider is the *Pairwise Differentiable Gradient Descent (PDGD)* (Oosterhuis & de Rijke, 2018), which represents the current state-of-the-art in OLTR. This method does not require online evaluation; instead, it directly optimizes a ranker using gradient descent, and its gradient estimations are unbiased with respect to document pair preferences inferred from user clicks. The comparison between ROLTR and PDGD is interesting because both methods use gradient descent but they differ in the way the unbiased gradient estimation is computed. Following the original PDGD paper (Oosterhuis & de Rijke, 2018), we set its learning rate to $\alpha = 0.1$.

ROLTR has two hyper-parameters: the reward function and the learning rate α . We study six reward functions in total: R_{IPS^+} , R_{IPS^-} , $R_{IPS^+} + R_{IPS^-}$ (i.e. using both clicked and unclicked signal) and their corresponding unshaped naive reward functions. We set the learning rate $\alpha = 0.01$ for MSLR-WEB10K and Istella, and $\alpha = 0.005$ for Yahoo! (values tuned based on validation set).

To simulate selection bias, we set $M = 10$ to only display 10 documents in the result lists for all experimental runs. For fair comparison, all methods are used to optimize a linear ranker. In order to measure statistically significant differences between methods, all runs are repeated 15 times spread evenly over the available dataset folds with different random seeds. The results are reported and compared using averaged performance and the two-tailed t-test.

In addition to comparing ROLTR with the other OLTR baselines mentioned above, we also compare our method against the MDPRank trained using full information, i.e., the offline LTR settings with actual relevance labels. This is to be thought of as the skyline for ROLTR as this method is built on MDPRank but training in ROLTR occurs with partial, noisy information (user clicks). For MDPRank, we use the experiment settings from the original paper (Wei et al., 2017).

5 Results

5.1 RQ1: Impact of the reward discount factor

To answer RQ1, we study the reward discount factor γ (set to 1 in previous work (Wei et al., 2017)). As discussed in Sect. 3.3, we expect that setting $\gamma = 0$ will lead to smaller variance in gradient estimation and a consequent speed up of training, without loss in terms of final convergence.

To confirm this assumption, in this section we empirically investigate the convergence of the learned rankers with different values of γ . In this study, we use the reward function R_{NAIVE^+} and R_{IPS^+} , and simulate the user propensity with $\eta = 1$. We report results on MSLR10k. Figure 2 illustrates the offline nDCG@10 learning curves obtained throughout the training process. The plots clearly show that the learning curves obtained when $\gamma = 0$ (blue curves) is always above that associated with other γ values, for both reward functions and both click settings. For naive reward and noisy click setting, rankers trained with lower value of γ achieved higher offline nDCG@10 score at the point of 100,000 impressions. This indicates that ROLTR learns fastest when $\gamma = 0$. On the other hand, ROLTR converges slower when γ increases, and when $\gamma = 1$ (red curves) the convergence speed is the slowest. Interestingly, unbiased reward function R_{IPS^+} seems more sensitive to γ than the

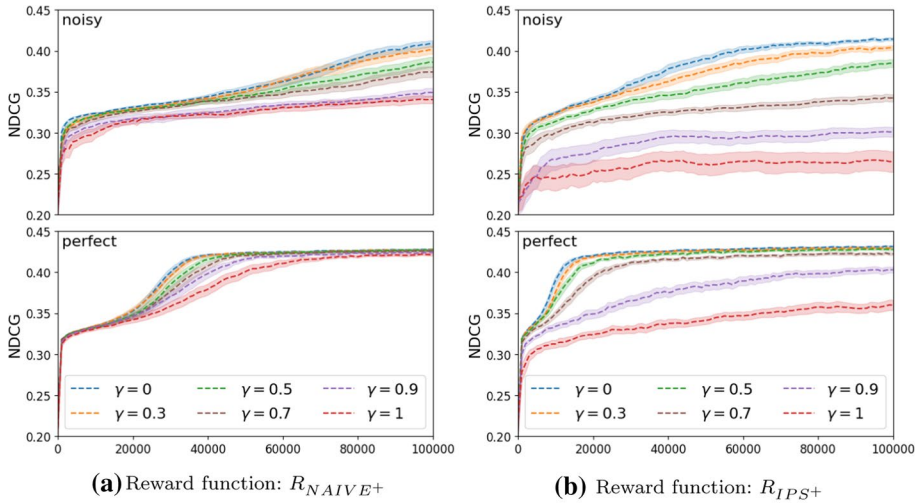


Fig. 2 Offline nDCG@10 score of ROLTR with different reward discount factor γ . (MSLR10K dataset)

naive reward function R_{NAIVE+} . Although R_{IPS+} has faster learning speed than R_{NAIVE+} when γ is small (e.g., $\gamma = 0, 0.3, 0.5$), the learning speed decreases dramatically when γ is large, and R_{IPS+} even fails to converge at 100,000 impressions under perfect clicks. We also studied the convergences of ROLTR with the R_{IPS+} reward function and $\gamma = 0, 0.5, 1$ for long-term impressions (1,000,000 impressions). These results are presented in Fig. 3. Similar to the trends in Fig. 2, larger γ values exhibit slower convergence and fail to converge to the same nDCG value as when $\gamma = 0$, even after performing one million impressions. This is due to the IPS significantly enlarging the click signal at lower ranks and this effect being cumulated to the rewards assigned to the clicks at higher ranks with large γ , thus resulting in considerably larger variance.

Following Jun et al. (2020), we also directly compare the variance of the estimated gradients when $\gamma = 0, \gamma = 0.5$ and $\gamma = 1$ for both naive and unbiased reward functions. For this, we calculate the variance of the gradient vectors at each training episode using the *trace of the covariance matrix*. We report the results for all runs on the MSLR10K in Fig. 4 (along to those for $\gamma = 0.5$ to show the overall trend for varying values between the two extremes). The variance from the noisy click runs is higher than that from the perfect click runs for all values of γ and both reward functions. This indicates that strong click noise causes high variance in gradient estimations. Moreover, the variance recorded for $\gamma = 0$ is almost half that recorded for $\gamma = 1$, for both noisy and perfect clicks and both reward functions. This finding agrees with the reinforcement learning practice: smaller discount factors always result in lower gradient variance. When comparing R_{NAIVE+} and R_{IPS+} , we notice that R_{IPS+} yields almost 10 times larger variance than that of the same setting for R_{NAIVE+} . This explains why R_{IPS+} is more sensitive to γ and larger γ significantly slows down the learning speed, as shown in Fig. 2.

In reinforcement learning, it is well known that γ controls the trade-off between variance and final convergence. Larger values of γ result in higher final convergence at the expense of higher variance (thus slower learning speed), and vice versa: smaller values result in lower variance (faster learning speed) but suboptimal convergence. However, in OLTR, our experiments show that optimizing the expected immediate reward (i.e. $\gamma = 0$)




Fig. 3 Long-term offline nDCG@10 score of ROLTR with different reward discount factor γ . (MSLR10K dataset, reward function: R_{IPSt})

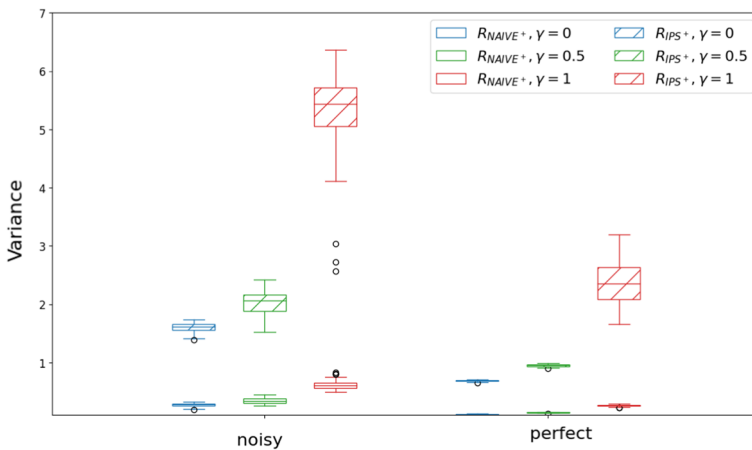


Fig. 4 Average variance of the gradient vector for each sampled episode. Error bars correspond to the 95% confidence intervals. (MSLR10K dataset)

will not hurt the final convergence at least after 100,000 impressions: the optimal policy learned with $\gamma = 0$ is equal to the optimal policy learned with $\gamma = 1$, which is maximizing the DCG score of the ranking. We further note that previous work that formalizing LTR as an MDP shows high variance in gradient estimation (Wei et al., 2017) and that recent work

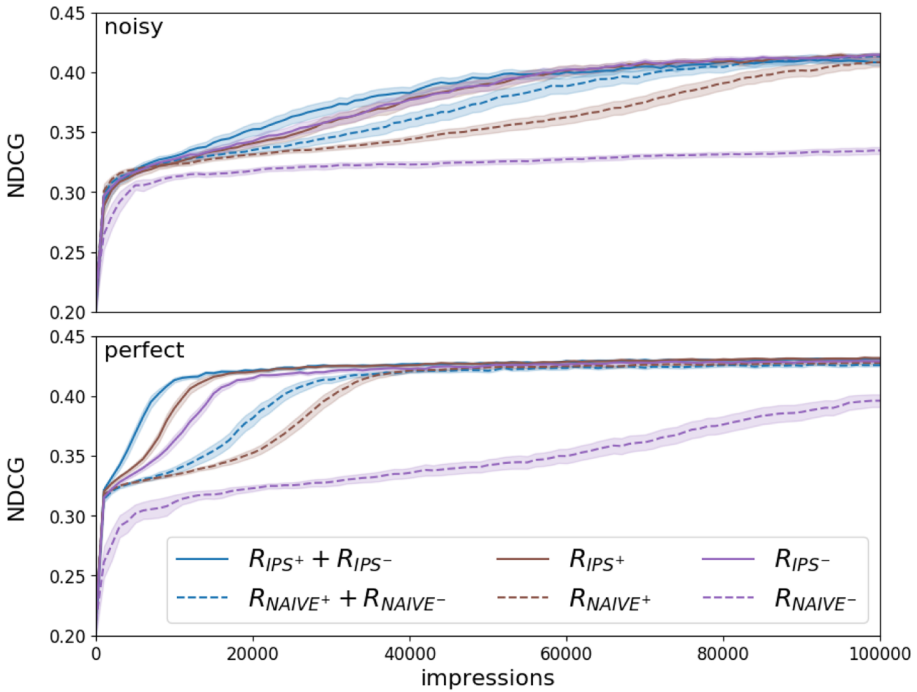


Fig. 5 Offline nDCG for ROLTR when different reward shaping functions are used (MSLR10K dataset)

has attempted to fix this issue for offline LTR (Jun et al., 2020): setting $\gamma = 0$ can be considered a simpler way to fix this problem. We believe this insight may have positive uptake for offline LTR.

With respect to RQ1, then, setting $\gamma = 0$ would not hurt the offline performance (convergence) of the learnt ranker, while delivering lower gradient variance than when $\gamma = 1$. Hence, we set $\gamma = 0$ in the remaining experiments.

5.2 RQ2: Impact of reward shaping functions

To answer RQ2, we consider the impact of different reward shaping in ROLTR. Figure 5 reports the *offline* performance of ROLTR with different reward functions over 100,000 impressions on the MSLR-WEB10K dataset. For both noisy and perfect click settings, all IPS reward functions outperform their naive versions in terms of learning speed. This means that it is important to de-bias the otherwise biased reward shaping function. In addition, we observe that ROLTR with $R_{IPS+} + R_{IPS-}$ and $R_{NAIVE+} + R_{NAIVE-}$ converges faster than when only using positive or negative reward functions. This is under both click settings, and it suggests that leveraging both clicked and unclicked signals is beneficial. This phenomenon is however less obvious with noisy clicks due to noisy clicks decreasing all the rankers' performance, thus resulting in worse and closer nDCG@10 scores. On the other hand, when noisy clicks are considered, R_{IPS+} and R_{IPS-} have a very similar learning curve; while, in the perfect click settings, R_{IPS+} converges much faster than R_{IPS-} . We explain this difference of behaviour between perfect and noisy click settings as follows.

In the perfect click setting, R_{IPS+} only uses signal from the clicked documents, which are relevant documents only in this click setting. R_{IPS-} instead also exploits signal from not-clicked documents (it down-weights this signal), which are non-relevant documents only in this click settings – this however is a worse signal than that from relevant documents only (R_{IPS+}). Thus, it is logical to see R_{IPS+} performing better (faster learning curve) than R_{IPS-} . In the noisy click setting, however, the signals exploited by R_{IPS+} and R_{IPS-} are similar: while still R_{IPS+} only uses clicked documents (while R_{IPS-} also uses not-clicked), in the noisy click setting both relevant and non-relevant documents are clicked. This therefore reduces the differences in performance observed between R_{IPS+} and R_{IPS-} .

When studying the final convergence at 100, 000 impressions, we do not record significant differences (i.e., $p > 0.05$) except for R_{NAIVE-} . This exception can be explained as R_{NAIVE-} only assigns biased negative rewards to unclicked documents and as result, in this case ROLTR just aims to avoid selecting documents that give low rewards, but it ignores valuable documents.

With respect to RQ2, then, we found that using unbiased reward shaping functions to reshape naive rewards make ROLTR converge faster (i.e. it requires less impressions). Specifically, the use of both unbiased positive and negative reward functions make ROLTR converge the fastest. Hence, in the remaining experiments, we set $R_{IPS+} + R_{IPS-}$ as reward function.

The fact that biased rewards do not change the final performance but slow down convergence may seem counter-intuitive at first. In fact, in CLTR studies, the biased learning objective results in a local optimal for the final convergence (Joachims et al., 2017). However, this may not be the case for OLTR. This is because CLTR experiments use a logging ranker to collect clicks and then train a new ranker with the click log: thus the rank of a document is decided by the logging ranker and has been fixed in the log. In OLTR settings like those used here, instead, the deployed ranker is interactively updated after each session: this means every document has a chance to be ranked at top of the displayed rankings hence be observed: these online interventions can eliminate position bias. However, this does not come for free, as it requires a large number of interactions to eventually obtain a good convergence. This is exactly what ROLTR attempts to solve: reduce the number of interactions to reach convergence. ROLTR in fact is the first OLTR algorithm that directly de-biases user position bias, with the results above showing that it can significantly speed up learning (this in turn will translate in better user experience, see Sect. 5.4).

5.3 RQ3: Final convergence (offline nDCG)

To answer RQ3, we compare the *offline* performance of the learned rankers over 100, 000 impressions. Figure 6 reports the learning curves for different methods. COLTR is clearly better than PMGD and DBGD in terms of learning speed and final convergence across almost all datasets: this is similar to previous findings (Oosterhuis et al., 2016; Zhuang & Zuccon, 2020). The gradient descent based methods (ROLTR and PDGD) significantly outperform candidate ranker sampling based methods (DBDG, PMGD and COLTR), rendering DBGD based methods outdated.

On the other hand, ROLTR and PDGD have similar learning curves across all three datasets. This is especially the case when perfect clicks are used. The final convergence of the two methods across all datasets and click settings is not statistically significantly different ($p > 0.05$). However, when noisy clicks are considered, we find that ROLTR has a faster learning speed at the beginning of training. Figure 7 considers the learning curves for

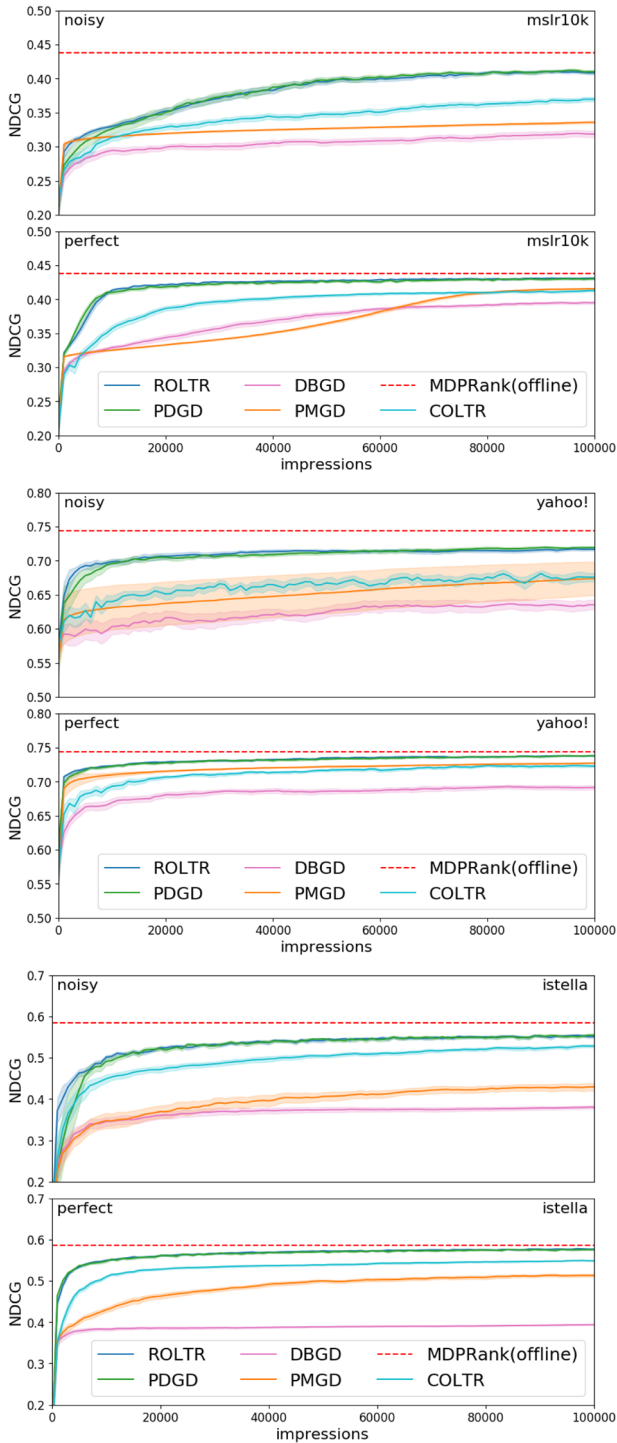


Fig. 6 Offline nDCG@10 for OLTR methods and MDRank under noisy and perfect clicks

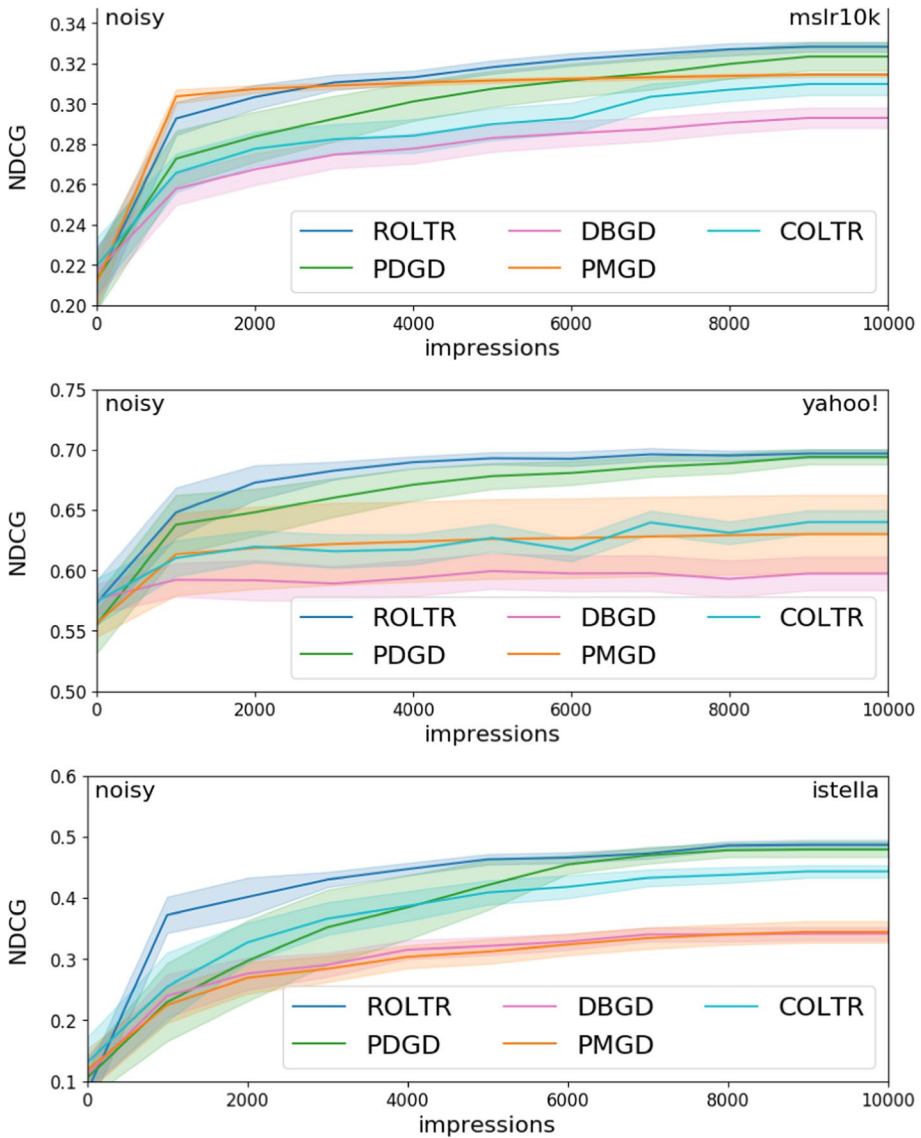


Fig. 7 Offline nDCG@10 for the first 10,000 impressions (noisy clicks)

the early impressions (first 10, 000) when noisy clicks are used. This figure clearly shows that the learning curves of ROLTR are almost always above other OLTR baselines at early impressions (except for the first 3, 000 impressions on MSLR10k, for which PMGD is best).

With respect to RQ3, then, we conclude that, when clicks are noisy, the learning speed of ROLTR is faster than PDGD in the early stage of training.

Finally, comparing the final convergence of ROLTR with the effectiveness of the offline MDPRank (trained with relevance labels and thus indicating skyline effectiveness), we can

Table 2 Online nDCG@10. Bold values indicate the highest average performance.

| Clicks | Algorithm | MSLR10K | Yahoo! | Istella |
|----------------|-----------|--------------------------------|----------------|--------------------------------|
| <i>Perfect</i> | DBGD | 519.99 | 1165.11 | 510.34 |
| | PMGD | 545.22 | 1191.26 | 564.58 |
| | COLTR | 448.65 | 1121.92 | 348.15 |
| | PDGD | 579.22 | 1310.99 | 741.45 |
| | ROLTR | 587.46 Δ | 1302.11 | 808.66 \blacktriangle |
| <i>Noisy</i> | DBGD | 477.29 | 1116.89 | 408.09 |
| | PMGD | 535.42 | 1137.39 | 426.77 |
| | COLTR | 431.27 | 1105.05 | 238.30 |
| | PDGD | 516.77 | 1227.28 | 481.85 |
| | ROLTR | 543.28 \blacktriangle | 1238.10 | 654.04 \blacktriangle |

Bold values represent the best value for each dataset/click setting

Significant gains and losses of ROLTR over PDGD are marked by Δ ($p < 0.05$) and \blacktriangle ($p < 0.01$)

clearly observe that ROLTR's effectiveness is at par to that of MDPRank, when the perfect click setting is used (Fig. 6). Recall that the perfect click setting still exhibits position bias. This findings empirically demonstrates the unbiasedness of ROLTR. However, when noisy clicks are used for ROLTR, then it's performance is sensibly lower to that of the offline MDPRank, showing how much noise in the click signal hurts performance, compared to perfect relevance labels.

5.4 RQ4: User experience during training (online nDCG)

To answer RQ4, we consider the quality of the experience users withstand while rankers are trained (online nDCG). Table 2 reports the *online* performance of each method. As expected, online NDCG@10 scores obtained when learning from perfect clicks are higher than for noisy clicks, suggesting the latter hurt user experience. Although COLTR has better offline performance than online evaluation based methods, it does lead to a worse user experience during the learning phase. This indicates that using counterfactual evaluation for candidate ranker comparison requires a lot more exploration. These findings are in agreement with our early findings (Zhuang & Zuccon, 2020). We also find that online evaluation based methods are significantly worse than ROLTR and PDGD ($p < 0.01$) for both perfect and noisy clicks, agreeing with the offline nDCG learning curves (Fig. 6). This is because ROLTR and PDGD assemble result lists by sampling documents from a softmax probability distribution, while online evaluation based methods use interleaving or multi-leaving to create these lists. This suggests that online evaluation based methods perform more exploration, thus hurting user experience more.

When comparing ROLTR and PDGD, we find that ROLTR obtains the best average online performance across all datasets under noisy clicks (agreeing with the learning curves in Fig. 7 for offline nDCG). Even for perfect clicks, for which the offline learning curves of ROLTR and PDGD are similar, ROLTR does statistically significantly outperform PDGD on MSLR10K and Istella ($p < 0.05, p < 0.01$ respectively). There are no statistically significant differences between ROLTR and PDGD on Yahoo! ($p > 0.05$). This may be because the average proportion of irrelevant documents for each query is small

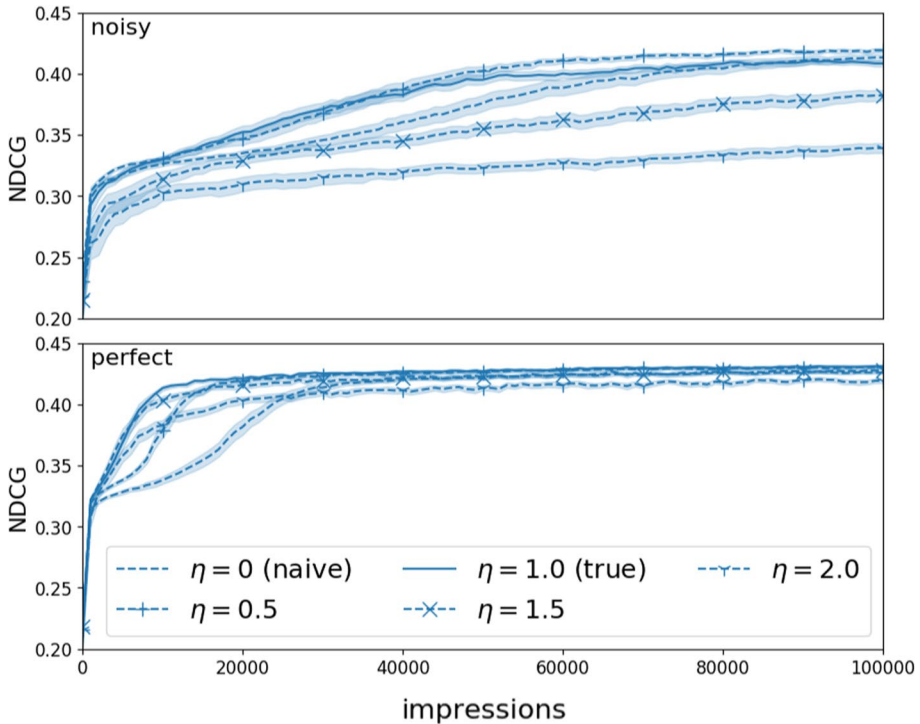


Fig. 8 Offline nDCG@10 for ROLTR when different prior propensities are used (MSLR10K dataset)

(21.92%), and thus it may be less likely the methods perform a ranking error, making them hard to distinguish. To conclude, with respect to RQ4, we found that ROLTR delivers the best user experience among the investigated methods.

5.5 RQ5: Sensitivity to propensity mismatch

To answer RQ5, we investigate the sensitivity of ROLTR to mismatched user propensity. Our previous experiments, in fact, assumed that the user propensity is known a priori and used the true propensity to obtain the unbiased rewards. However, this assumption does not always hold true in practice, as the user propensity could be overestimated or underestimated for various reasons (Joachims et al., 2017).

Figure 8 shows the *offline* nDCG@10 of ROLTR under different propensity values η and for different click settings. When $\eta = 0$, ROLTR assumes no user position bias and the unbiased reward function R_{IPS} reduces to R_{NAIVE} . On the other hand, $\eta = 1$ is the same propensity used to simulate clicks, and thus the propensity assumed by the ranker and that seen in the clicks are the same. Note that when simulating the user, we always assume the user has propensity $\eta = 1$ (while instead we change ROLTR's propensity value).

From Fig. 8 we can observe that, in the noisy click setting, an underestimation of propensity ($\eta < 1$) leads to a similar final convergence as that obtained when the true propensity is considered. In fact, when $\eta = 0.5$ we record a higher final convergence. However,

when the propensity is overestimated ($\eta > 1$), both learning speed and final convergence become much worse. A similar behaviour has been found in CLTR and relates to the large variance introduced by the extreme IPS weights (Jagerman & de Rijke, 2020): Because of the overestimated propensities, the IPS weights will be much bigger than those of the true propensity, resulting in a large amount of variance in the gradient estimation, thus hurting the final performance. There are many ways to prevent this from happening, e.g., ‘propensity clipping’ which trades-off bias against variance.

On the other hand, under perfect clicks, the final convergence of all propensity priors is similar, except for the extremely overestimated propensity ($\eta = 2.0$). In fact, in the perfect setting the main difference is with respect to the learning speed: overestimated propensities converge faster than underestimated propensities.

With respect to RQ5, then, we conclude that, when clicks are perfect, ROLTR is not sensitive to propensity mismatch in terms of final convergence; but when clicks are noisy, only underestimated propensities are not sensitive while overestimated propensities will hurt both the final convergence and learning speed.

6 Conclusion

We proposed a novel OLTR algorithm, ROLTR, which formalises OLTR as an MDP ranking problem and uses reinforcement learning with unbiased reward shaping functions to directly optimize an IR metric. In contrast to traditional online evaluation based methods and the current state-of-the-art PDGD method, ROLTR does not fully rely on online interventions to overcome user position bias. Instead, ROLTR directly uses the IPS reward shaping functions to de-bias rewards given by the environment which can be further used to guide gradient estimation. As a result, the gradient calculated by ROLTR is unbiased with respect to position bias. In order to accelerate the learning speed and obtain better user experience, we first simplified the MDP ranking by setting the reward discount factor γ to 0, which we empirically confirmed it provides a lower gradient variance without hurting final convergence. Furthermore, to fully leverage user click feedback in each training episode, we introduced a negative IPS reward shaping function for ROLTR which provides additional unbiased reward learning signal from unclicked documents. We have proven that the rewards reshaped by the negative IPS reward shaping function are also unbiased with respect to position bias. Our experimental results show that ROLTR achieves state-of-the-art offline performance requiring less user interactions, which results in considerably better user experience (online performance) over other OLTR methods.

Future work will be directed towards considering estimating user observation propensity during training since our method requires propensity to be known a priori. Furthermore, other biases such as selection bias and other type of user behaviour signals, such as dwell time, mouse move, etc, could be modelled within ROLTR, potentially improving its performance. This could be achieved by simply changing the reward function while maintaining other parts of the method unchanged. This makes ROLTR flexible and with plenty of room for further extension.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long

as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adomavicius, G., Sankaranarayanan, R., Sen, S., & Tuzhilin, A. (2005). Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Transactions on Information Systems (TOIS)*, 23(1), 103–145.
- Agarwal, A., Basu, S., Schnabel, T., & Joachims, T. (2017). Effective evaluation using logged bandit feedback from multiple loggers. In Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining, (pp. 687–696).
- Agarwal, A., Zaitsev, I., & Joachims, T. (2018). Counterfactual learning-to-rank for additive metrics and deep models. Preprint at [arXiv:1805.00065](https://arxiv.org/abs/1805.00065).
- Agarwal, A., Zaitsev, I., Wang, X., Li, C., Najork, M., & Joachims, T. (2019). Estimating position bias without intrusive interventions. In Proceedings of the twelfth ACM international conference on web search and data mining, (pp. 474–482).
- Ai, Q., Bi, K., Luo, C., Guo, J., & Croft, W.B. (2018). Unbiased learning to rank with unbiased propensity estimation. In The 41st international ACM SIGIR conference on research & development in information retrieval, (pp. 385–394).
- Ai, Q., Yang, T., Wang, H., & Mao, J. (2021). Unbiased learning to rank: online or offline? *ACM Transactions on Information Systems (TOIS)*, 39(2), 1–29.
- Al-Maskari, A., & Sanderson, M. (2010). A review of factors influencing user satisfaction in information retrieval. *Journal of the American Society for Information Science and Technology*, 61(5), 859–868.
- Chapelle, O., & Chang, Y. (2011). Yahoo! learning to rank challenge overview. In Proceedings of the learning to rank challenge, (pp. 1–24).
- Dato, D., Lucchese, C., Nardini, F. M., Orlando, S., Perego, R., Tonello, N., & Venturini, R. (2016). Fast ranking with additive ensembles of oblivious and non-oblivious regression trees. *ACM Transactions on Information Systems (TOIS)*, 35(2), 1–31.
- Fang, Z., Agarwal, A., & Joachims, T. (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, (pp. 825–834).
- Feng, Y., Xu, J., Lan, Y., Guo, J., Zeng, W., & Cheng, X. (2018). From greedy selection to exploratory decision-making: Diverse ranking with policy-value networks. In The 41st international ACM SIGIR conference on research & development in information retrieval, (pp. 125–134).
- Guan, Z., & Cutrell, E. (2007). An eye tracking study of the effect of target rank on web search. In Proceedings of the SIGCHI conference on human factors in computing systems, (pp. 417–420).
- Hofmann, K., Li, L., & Radlinski, F. (2016). Online evaluation for information retrieval. *Foundations and trends in information retrieval*, 10(1), 1–117.
- Hofmann, K., Schuth, A., Whiteson, S., & de Rijke, M. (2013). Reusing historical interaction data for faster online learning to rank for ir. In Proceedings of the sixth ACM international conference on Web search and data mining, (pp. 183–192), ACM.
- Hofmann, K., Whiteson, S., & De Rijke, M. (2011). Balancing exploration and exploitation in learning to rank online. In European conference on information retrieval, (pp. 251–263), Springer.
- Hofmann, K., Whiteson, S., & De Rijke, M. (2011). A probabilistic method for inferring preferences from clicks. In Proceedings of the 20th ACM international conference on information and knowledge management, (pp. 249–258).
- Hofmann, K., Whiteson, S., de Rijke, M., et al. (2011). Contextual bandits for information retrieval. In NIPS 2011 Workshop on Bayesian optimization, experimental design, and bandits, granada, (vol. 12, p. 2011).
- Hu, Z., Wang, Y., Peng, Q., & Li, H. (2019). Unbiased lambdamart: an unbiased pairwise learning-to-rank algorithm. In The World Wide Web conference, (pp. 2830–2836).
- Jagerman, R., Markov, I., & Rijke, M. D. (2020). Safe exploration for optimizing contextual bandits. *ACM Transactions on Information Systems (TOIS)*, 38(3), 1–23.

- Jagerman, R., Oosterhuis, H., & de Rijke, M. (2019). To model or to intervene: a comparison of counterfactual and online learning to rank from user interactions. In Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval, SIGIR'19, (pp. 15–24). Association for Computing Machinery.
- Jagerman, R., & de Rijke, M. (2020). Accelerated convergence for counterfactual learning to rank. In Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval, (pp. 469–478).
- Joachims, T. (2002). Optimizing search engines using clickthrough data. In Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, (pp. 133–142).
- Joachims, T., Swaminathan, A., & Schnabel, T. (2017). Unbiased learning-to-rank with biased feedback. In Proceedings of the tenth ACM international conference on web search and data mining, (pp. 781–789).
- Jun, X., Zeng, W., Long, X., Yanyan, L., Dawei, Y., Xueqi, C., & Ji-Rong, W. (2020). Reinforcement learning to rank with pairwise policy gradient. In Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval.
- Lefortier, D., Serdyukov, P., & De Rijke, M. (2014). Online exploration for detecting shifts in fresh intent. In: Proceedings of the 23rd ACM international conference on conference on information and knowledge management, (pp. 589–598).
- Li, C., Markov, I., Rijke, M. D., & Zoghi, M. (2020). Mergedts: a method for effective large-scale online ranker evaluation. *ACM Transactions on Information Systems (TOIS)*, 38(4), 1–28.
- Li, H. (2011). Learning to rank for information retrieval and natural language processing. *Foundations and Trends® in Information Retrieval*, 3(3), 225–331.
- Liu, T. Y. (2011). *Learning to rank for information retrieval*. Germany: Springer Science & Business Media.
- Maxwell, D., Azzopardi, L., & Moshfeghi, Y. (2017). A study of snippet length and informativeness: behaviour, performance and user experience. In Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval, (pp. 135–144).
- Ng, A. Y., Harada, D., & Russell, S. (1999). Policy invariance under reward transformations: theory and application to reward shaping. In Proceedings of the sixteenth international conference on machine learning, (vol. 99, pp. 278–287).
- Oosterhuis, H., & de Rijke, M. (2018). Differentiable unbiased online learning to rank. In Proceedings of the 27th ACM international conference on information and knowledge management, (pp. 1293–1302), ACM.
- Oosterhuis, H., & de Rijke, M. (2019). Optimizing ranking models in an online setting. In European conference on information retrieval, (pp. 382–396), Springer.
- Oosterhuis, H., & de Rijke, M. (2020). Policy-aware unbiased learning to rank for top-k rankings. In Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval.
- Oosterhuis, H., & de Rijke, M. (2020). Taking the counterfactual online: Efficient and unbiased online evaluation for ranking. In Proceedings of the 2020 ACM SIGIR on international conference on theory of information retrieval, (pp. 137–144).
- Oosterhuis, H., & de Rijke, M. (2021). Unifying online and counterfactual learning to rank: a novel counterfactual estimator that effectively utilizes online interventions. In Proceedings of the 14th ACM international conference on web search and data mining, (pp. 463–471).
- Oosterhuis, H., Schuth, A., & de Rijke, M. (2016). Probabilistic multileave gradient descent. In European conference on information retrieval, (pp. 661–668), Springer.
- Ovaisi, Z., Ahsan, R., Zhang, Y., Vasilaky, K., & Zheleva, E. (2020). Correcting for selection bias in learning-to-rank systems. In Proceedings of the web conference 2020, (pp. 1863–1873).
- Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., & Granka, L. (2007). In google we trust: Users' decisions on rank, position, and relevance. *Journal of Computer-Mediated Communication*, 12(3), 801–823.
- Popov, I., Heess, N., Lillicrap, T., Hafner, R., Barth-Maron, G., Vecerik, M., Lampe, T., Tassa, Y., Erez, T., & Riedmiller, M. (2017). Data-efficient deep reinforcement learning for dexterous manipulation. Preprint at [arXiv:1704.03073](https://arxiv.org/abs/1704.03073).
- Qin, T., & Liu, T. Y. (2013). Introducing letor 4.0 datasets. Preprint at [arXiv:1306.2597](https://arxiv.org/abs/1306.2597).
- Qin, T., Liu, T. Y., Xu, J., & Li, H. (2010). Letor: A benchmark collection for research on learning to rank for information retrieval. *Information Retrieval*, 13(4), 346–374.
- Sanderson, M. (2010). Test collection based evaluation of information retrieval systems. *Foundations and Trends® in Information Retrieval*, 4(4), 247–375.
- Schuth, A., Brintjes, R. J., Büttner, F., van Doorn, J., Groenland, C., Oosterhuis, H., Tran, C. N., Veeling, B., van der Velde, J., Wechsler, R., et al. (2015). Probabilistic multileave for online retrieval

- evaluation. In Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval, (pp. 955–958), ACM.
- Schuth, A., Oosterhuis, H., Whiteson, S., & de Rijke, M. (2016). Multileave gradient descent for fast online learning to rank. In proceedings of the ninth ACM international conference on web search and data mining, (pp. 457–466).
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: an introduction*. Cambridge, MA: MIT Press.
- Sutton, R.S., McAllester, D.A., Singh, S.P., & Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In Advances in neural information processing systems, (pp. 1057–1063).
- Vardasbi, A., de Rijke, M., & Markov, I. (2020). Cascade model-based propensity estimation for counterfactual learning to rank. Preprint at [arXiv:2005.11938](https://arxiv.org/abs/2005.11938).
- Wang, H., Kim, S., McCord-Snook, E., Wu, Q., & Wang, H. (2019). Variance reduction in gradient exploration for online learning to rank. In Proceedings of the 42Nd international ACM SIGIR conference on research and development in information retrieval, SIGIR'19.
- Wang, H., Langley, R., Kim, S., McCord-Snook, E., & Wang, H. (2018). Efficient exploration of gradient space for online learning to rank. In The 41st international ACM SIGIR conference on research & development in information retrieval, (pp. 145–154).
- Wang, N., Qin, Z., Wang, X., & Wang, H. (2021). Non-clicks mean irrelevant? propensity ratio scoring as a correction. In Proceedings of the 14th ACM international conference on web search and data mining, (pp. 481–489).
- Wang, S., Liu, B., Zhuang, S., & Zuccon, G. (2021). Effective and privacy-preserving federated online learning to rank. In Proceedings of the 2021 ACM SIGIR international conference on theory of information retrieval, (pp. 3–12).
- Wang, S., Zhuang, S., & Zuccon, G. (2021). Federated online learning to rank with evolution strategies: a reproducibility study. In The 43rd european conference on information retrieval (ECIR).
- Wang, X., Bendersky, M., Metzler, D., & Najork, M. (2016). Learning to rank with selection bias in personal search. In Proceedings of the 39th international ACM SIGIR conference on research and development in information retrieval, (pp. 115–124).
- Wang, X., Golbandi, N., Bendersky, M., Metzler, D., & Najork, M. (2018). Position bias estimation for unbiased learning to rank in personal search. In Proceedings of the eleventh ACM international conference on web search and data mining, (pp. 610–618).
- Wei, Z., Xu, J., Lan, Y., Guo, J., & Cheng, X. (2017). Reinforcement learning to rank with markov decision process. In Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval, (pp. 945–948).
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3–4), 229–256.
- Xia, L., Xu, J., Lan, Y., Guo, J., Zeng, W., & Cheng, X. (2017). Adapting markov decision process for search result diversification. In Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval, (pp. 535–544).
- Yao, J., Dou, Z., Xu, J., & Wen, J.R. (2020). Rlper: A reinforcement learning model for personalized search. In Proceedings of the web conference 2020, (pp. 2298–2308).
- Yue, Y., & Joachims, T. (2009). Interactively optimizing information retrieval systems as a dueling bandits problem. In Proceedings of the 26th annual international conference on machine learning, (pp. 1201–1208).
- Zeng, W., Xu, J., Lan, Y., Guo, J., & Cheng, X. (2018). Multi page search with reinforcement learning to rank. In Proceedings of the 2018 ACM SIGIR international conference on theory of information retrieval, (pp. 175–178).
- Zhao, X., Gu, C., Zhang, H., Yang, X., Liu, X., Liu, H., & Tang, J. (2021). Dear: Deep reinforcement learning for online advertising impression in recommender systems. In Proceedings of the AAAI conference on artificial intelligence, (vol. 35, pp. 750–758).
- Zhao, X., Xia, L., Zhang, L., Ding, Z., Yin, D., & Tang, J. (2018). Deep reinforcement learning for page-wise recommendations. In Proceedings of the 12th ACM conference on recommender systems, (pp. 95–103).
- Zhao, X., Zhang, L., Ding, Z., Xia, L., Tang, J., & Yin, D. (2018). Recommendations with negative feedback via pairwise deep reinforcement learning. In Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining, (pp. 1040–1048).
- Zhuang, S., & Zuccon, G. (2020). Counterfactual online learning to rank. In European conference on information retrieval, (pp. 415–430), Springer.

Zhuang, S., & Zuccon, G. (2021). How do online learning to rank methods adapt to changes of intent? In Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.