Improving Effectiveness and Security in Federated Online Learning to Rank

Shuyi Wang[1,*,†]

[1]The University of Queensland, Australia

Abstract

Online Learning to Rank (OLTR) optimizes ranking models using implicit user's feedback, such as clicks, to directly manipulate search engine results in production. Compared to traditional Learning to Rank (LTR) approaches, OLTR overcomes a number of drawbacks such as the high cost and time required by the editorial annotation effort, the fact that the intent identified by human annotators may be different from that of the user, or the issues associated with rapid changes of intents underlying queries. However, this process requires OLTR methods to collect searchable data, user queries and clicks; current methods are not suited to situations in which users want to maintain their privacy, i.e. not sharing data, queries and clicks. Federated Online Learning to Rank (FOLTR), which implements OLTR under a Federated Learning (FL) scenario, has been proposed to provide a solution that can address the privacy issue in OLTR. Existing work has shown promise; however, FOLTR methods lag behind traditional OLTR, which has been studied in the centralised setting. In particular, challenges currently faced by FOLTR methods include low effectiveness, unclear robustness with respect to biased data distribution across clients, and unclear susceptibility to attacks on the learning process and the ranking model. This research aims to comprehensively investigate the aforementioned challenges and build effective, secure federated online learning to rank methods.

Keywords

Online Learning to Rank, Federated Learning, Privacy Aware Information Retrieval

1. Introduction

1.1. Background and Motivation

Online Learning to Rank (OLTR) methods allow to learn a ranker from observing users queries and interactions with search engine result pages (SERPs), e.g., clicks. This is achieved by iteratively train and update a ranker in production. By exploiting user interactions rather than explicit relevance labels, OLTR overcomes a number of drawbacks associated with traditional learning to rank approaches such as the high cost and time required by the editorial annotation effort, the fact that the intent identified by editors may not be the one the user had in mind [1], and the issues associated with rapid changes of intents underlying queries [2].

In traditional OLTR, search results are produced by a centralised search service, which also

FDIA'22: Proceedings of the 10th Symposium on Future Directions in Information Access, July, 2022, Lisbon, Portugal

*Corresponding author.

shuyi.wang@uq.edu.au (S. Wang)

https://karlywang.github.io/ (S. Wang)

0000-0002-4467-5574 (S. Wang)

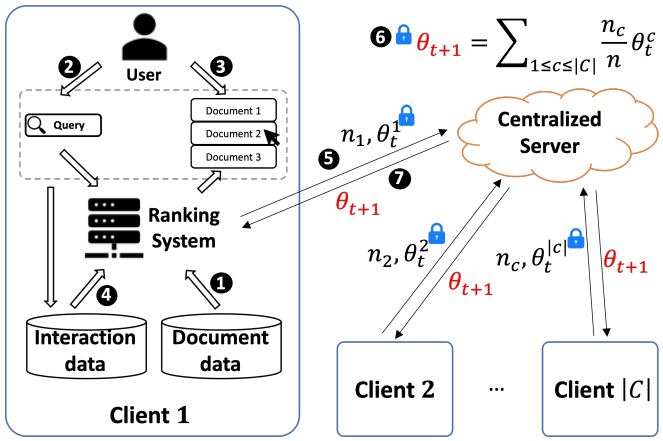

Figure 1: Schematic representation of Federated Online Learning to Rank (FOLTR) setting.

keeps track of user queries and interactions that are used as signal for learning an effective ranker [3, 4, 5, 6, 7, 8]. This centralised search service has also a complete index of the data to be searched (though likely this index is distributed and replicated across servers and data centres). A drawback of this solution is that the search service has access to the indexed data, and also actively collects users queries and interactions such as clicks: thus *user privacy is limited* in that all this user data is collected and exploited by the search service. So, for example, if users were interested for the search service to offer search functionalities on their private data (e.g. emails, desktop files) they need to surrender this to the search service. Similarly, the OLTR-based search service would also not meet the expectations of those users that wish for their search behaviour not to be collected to protect their privacy. To address this issue, Federated OLTR (FOLTR), which conducts OLTR by leveraging Federated Learning (FL), has been proposed as a solution. Federated Learning [19] considers a machine learning setting where several data owners (clients) collaboratively train a model without sharing the data. To this aim, each participating client only sends what is essential for the training (for example, the local gradient) to the central server. The Federated Averaging (FedAvg) algorithm [19] relies on local stochastic gradient descent (SGD), which is computed by each client, while the centralised server performs the global model update by weighted averaging the client's updates.

In a typical FOLTR setting, pictured in Figure 1, data on which to search, along with user queries and interactions are withheld from the central server, and so is also the responsibility of producing search results. Search is instead performed within the user device ❷, which also indexes the users data ❶ collects users interactions ❸ and performs the online updates required by the OLTR method ❹. Updates are then shared by several users devices with a central server ❺. It is the responsibility of the central server to combine the ranker updates from different clients to produce a new version of the ranker ❻. This new version is then distributed to the users devices ❼. The advantage of a federated OLTR solution over a traditional OLTR approach is that neither user data nor users queries and interactions are seen by or shared with anyone, thus potentially *preserving user privacy*.

FOLtR-ES is the first and only FOLTR method proposed by Kharitonov [9], which uses

evolutionary strategies similar to those in genetic algorithms to make client rankers explore the feature space, and a parametric privacy preserving mechanism to further anonymise the feedback signal that is shared by clients to the central server. However, the effectiveness of previous approaches for federated OLTR (i.e., FOLtR-ES) exhibits large gaps in performance compared to current state-of-the-art, not federated, OLTR methods [10]. Data heterogeneity is another factor that has been neglected by FOLTR and many other search and recommendation tasks that use FL. Data heterogeneity or non-IID data means the training data involved in FL is non independent and identically distributed across clients, which violates frequently-used independent and identically distributed (IID) assumptions in common machine learning optimization. Previous works mainly conduct on synthetic homogeneous datasets, either for evaluating the proposed methods, or simulating the process of federation. Meanwhile, real-world applications often display an inevitable phenomenon when FL methods are applied to heterogeneous data, in which the FL performance are forced to deteriorate [11, 12]. Several existing works have proposed solutions to address this [13, 14, 15]. However, while FOLTR systems are on their own rights a type of federated learning system, the presence and effect of non-IID data in FOLTR has not been studied. Meanwhile, the privacy-preserving mechanism leveraging FL is required to protect the gradient updates from privacy attacks [16]. In fact, despite the federate learning process not involving the sharing of user data, a malicious attacker that is able to observe the ranker model and gradient updates can potentially reconstruct the data that produced such an update [17], thus finally exposing user data or poisoning the global model performance. Thus, approaches for designing an effective, secure and applicable FOLTR system are worth investigating. It is those research gaps that we want to fill in this project.

1.2. Research Questions

To achieve our objectives, we plan to explore the following research questions.

RQ1: Effectiveness and Privacy-preservation. How to design an effective and generalisable FOLTR method which can achieve competitive performance compared to the state-of-the-art OLTR methods?
- a) Does the existing FOLTR method (FOLtR-ES) generalize to other commonly-used datasets and metrics in OLTR?
- b) How to integrate the state-of-the-art OLTR method with the widely-used, easy-to-implement FedAvg framework and meanwhile, further protect user's privacy?

RQ2: Robustness under non-IID Data. Are FOLTR methods robust to non-IID data among clients? If not, how can this challenge be addressed?
- a) What are the potential types of non-IID data in FOLTR and their impact on model performance?
- b) How to address the impact of non-IID data in the context of FOLTR?

RQ3: Security and Defense against Malicious Attacks. Are FOLTR methods secure with respect to potential risks brought by malicious participants? If not, how to address this problem?
- a) How to design a novel poisoning attack strategy targeted on model integrity of FOLTR?
- b) How to defend against the model integrity attack for FOLTR?

2. Methodologies

2.1. RQ1: Effectiveness and Privacy-Preservation

2.1.1. Re-evaluation of current FOLTR methods

To date, aside from our proposed method, FOLtR-ES is the only FOLTR algorithm [9]. However, the original research study that introduced this method only evaluated it on a small Learning to Rank (LTR) dataset and with no conformity with respect to current OLTR evaluation practice like metrics and experimental settins. It further did not explore specific parameters of the method, such as the number of clients involved in the federated learning process, and did not compare FOLtR-ES with the current state-of-the-art OLTR method. To address these limitations, we have reproduced this method and conduct empirical investigation based on this [10]. Our findings question whether FOLtR-ES is a mature method that can be considered in practice: its effectiveness largely varies across datasets, click types, ranker types and settings. We have found that its performance, for both search effectiveness and user experience, is also far from that of current state-of-the-art OLTR, raising the need for more effective, stable and generalisable FOLTR methods.

2.1.2. New SOTA for FOLTR

For this task, we propose a novel Federated OLTR method, called FPDGD [18], which leverages the state-of-the-art Pairwise Differentiable Gradient Descent (PDGD) [6] and adapts it to the Federated Averaging [19] framework. For a strong privacy guarantee, we further introduce a noise-adding clipping technique based on the theory of differential privacy to be used in combination with FPDGD.

Compared to traditional OLTR methods, our method considers a situation where user's privacy is required. Our method also aims to improve FOLTR in effectiveness compared to FOLtR-ES [9]. The method we proposed, FPDGD, includes three main parts: (1) we use PDGD as the core OLTR optimisation algorithm - this method is shown to perform well and converge faster than previous OLTR methods and is unbiased with respect to users' preferences; (2) we adopt Federated Averaging as the federated learning framework to implement our adaptation of PDGD to the federated setting; (3) we secure the communication between clients and server, by leveraging differential privacy, so as to protect user privacy.

We further demonstrate the effectiveness of FPDGD through empirical experiments comparing the method to its non-federated counterpart (PDGD) and the other only federated OLTR method, FOLtR-ES. Compared to the non-federated PDGD, our method converges slower when the same total interaction budget (i.e. number of queries) is considered: in particular, the difference is significant when differential privacy is considered – this is the price to pay for not sharing the clients' data with the central sever (federated setup) and for protecting the local gradient updates (differential privacy). When compared with FOLtR-ES, the proposed method showcases higher performance. This resonates with previous work that showed the performance of FOLtR-ES to be highly variable across datasets and settings, making the method unstable and of risky use in practice [10]. FPDGD instead represents a reliable, stable, and secured alternative to non-federated state-of-the-art OLTR methods.

2.2. RQ2: Robustness under Non-IID Data

2.2.1. Impact of non-IID data in FOLTR

A well-known factor that affects the performance of federated learning systems, and that poses serious challenges to these approaches, is the fact that there may be some type of bias in the way training data is distributed across the clients [11]. This is also called data heterogeneity or non independent and identically distributed (non-IID) data, which violates frequently-used assumptions in common machine learning optimization. While FOLTR systems are on their own rights a type of federated learning system, the presence and effect of non-IID data in FOLTR has not been studied. To this aim, in our recent publication [20], we first enumerate possible data distribution settings that may showcase data bias across clients and thus give rise to the non-IID problem. Then, we study the impact of each settings on the performance of the current state-of-the-art FOLTR approach, the Federated Pairwise Differentiable Gradient Descent (FPDGD), and we highlight which data distributions may pose a problem for FOLTR methods. We also explore how common approaches proposed in the federated learning literature address non-IID issues in FOLTR. This allows us to unveil new research gaps that, we argue, future research in FOLTR should consider.

In total, we put forward four different ways data specific to FOLTR could be distributed in a non-IID manner across clients. Specifically this data may not be IID because of biases across clients due to: document label distributions, document preferences, click preferences, data quantity. We found that the presence of non-IID characteristics in the distribution of document preferences and specific cases of document labels have severe effects on the effectiveness of FPDGD. Conversely, if data is distributed across clients in a non-IID manner concerning click preferences or data quantity, no significant effects on the quality of FPDGD are observed. We also show that existing solutions employed in general federated learning to mitigate the non-IID data problem do not apply to the FOLTR setting, despite some of these non-IID cases (and especially for bias in document preferences) being likely to occur across many FOLTR systems. These findings contribute an understanding of under which data distributions it is safe to use FOLTR and when it is not, further showing that researching how to address non-IID data in FOLTR is a worthwhile area of investigation.

2.2.2. Handling non-IID data in FOLTR

Based on the results of Sec 2.2.1, we focus on addressing the impact of non-IID in the distribution of document labels and document preferences.

To eliminate the impact of non-IID data, current approaches in horizontal FL [1] can be classified into (i) data based approaches, (ii) algorithm based approaches, and (iii) system based approaches [11]. Data based approaches focus on augmenting data by creating a small dataset shared globally [13, 22, 14]. We consider this type of approach as a straightforward and effective way to deal with non-IID in the distribution of document labels. One challenge for this approach is the difficulty of collecting the uniformly distributed and globally shared data. Another problem relates to the possible violation of privacy preserving requirements, which is instead a

[1] The scenario in which training data of different clients share the same feature space but have different sample space [21]: this is the scenario we consider in FOLTR.

key motivation for FL. We are then interested in investigating: (1) which amount of globally shared data is required to achieve noteworthy improvements; (2) the necessity, feasibility and frequency of updating the shared data across clients during the online training.

To dealing with non-IID in document preferences, as each client has personal preferences thus leading to different judgement criteria on documents, global data sharing seems not feasible. We will then investigate algorithm based methods which allow each client to have a customized local model by implementing local fine-tuning for both linear and neural rankers. This aims to fine-tune the local models after receiving the global model from the server using the client's local data [23]. The challenges of using local fine-tuning can be roughly divided into two types: (1) how to find a suitable initial shared model, and (2) how to combine local and global information thus obtaining better user experience through the fine-tuned local model.

System based approaches provide a solution for dealing with both document labels skewness and document preferences skewness. For example, client clustering is proposed to construct a multi-center framework by grouping clients into different clusters and clients with similar local training data are allocate to the same cluster. The main challenges of this approach is to estimate the data distribution of each client as local data needs to remain private and be kept secret from the central server and other clients.

2.3. RQ3: Security and Defense against Malicious Attacks.

For this task, we want to study a scenario where one or several malicious attackers participate in the FOLTR process. The intent of a malicious attacker is to destroy the global model integrity through compromising the local data or the local model during the training phase. This type of attack is refered to as *poisoning attack*. Through this study, we want to further understand the potential risks of *poisoning attacks* and the security of FOLTR systems against this kind of attacks. Generally, poisoning attacks during the training phase can be sourced from two types of poisoned updates: (1) *data poisoning attacks* during local data collection; and (2) *model poisoning attacks* during the local model training process [31].

One common example of data poisoning attacks is the label-flipping attack [32, 33], in which the labels of honest training examples of one class are flipped to another class while the features of the data are kept unchanged. In the context of FOLTR, the impact of such an attack on model performance is unknown. With this respect, we plan to investigate: (1) to which extent participants engaging in these attacks affect the FOLTR model; (2) to which extent the amount of training data being poisoned affects the FOLTR model; (3) whether the central server can detect data poisoning attacks by one or several malicious clients.

Model poisoning attacks aim to poison local model updates before sending them to the central server or insert hidden backdoors into the global model [34]. An early study by Bhagoji et al. [35] demonstrated that model poisoning attacks are much more effective than data poisoning in FL settings for classification tasks. While the majority of these attacks are only focused on perturbing the results in general classification tasks (for example, image classification and next word prediction) [34, 36], in the context of FOLTR, it is difficult to determine whether it is more cost-effective for an adversary to poison the local model with manipulated gradients. In fact, FOLTR methods are designed for learning from biased and noisy implicit user feedbacks, and thus are optimized towards completely different learning objectives, leaving the impact of

poisoning attacks of this category largely unpredictable. This brings some unique challenges: (1) as the federated environment significantly limits the adversary's prior knowledge to its local information (i.e., malicious user's data and local models) and the global model, the randomness of user's queries and clicks effects the success of attacking; (2) the poisoned global model should also maintain high ranking performance thus avoiding easy detection.

Taking one step further, we wonder how to design against potential poisoning attacks could be devised. We will start from deploying the existing defensive strategies, such as Bulyan [37] and Trimmed Mean [38]. Then, we plan to develop strategies targeted to our proposed attacking strategies, specific to FOLTR. We believe this work is an important contribution to the community as: while the security and vulnerability of machine learning have been widely studied in recent years, the potential risks that take place in the context of FOLTR are yet to be thoroughly understood.

3. Conclusion

This research tackles a range of largely unexplored areas for Federated Online Learning to Rank methods. Our research will help practitioners to gain in-depth insights into FOLTR in terms of unbiased and effective optimization, robustness to non-IID data and poisoning attacks. Both users and search service providers will benefit from the research outcomes of this project: (1) from the perspective of users, effective rankings will be provided under the respect of user's data privacy; (2) from the perspective of search service providers, high costs for annotation will be reduced and effective, robust, and secure rankers will be trained under the Federated Learning framework.

References

[1] M. Sanderson, Test collection based evaluation of information retrieval systems, Now Publishers Inc, 2010.

[2] S. Zhuang, G. Zuccon, 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, 2021.

[3] Y. Yue, T. Joachims, Interactively optimizing information retrieval systems as a dueling bandits problem, in: Proceedings of the 26th Annual International Conference on Machine Learning, 2009, pp. 1201–1208.

[4] A. Schuth, H. Oosterhuis, S. Whiteson, M. de Rijke, Multileave gradient descent for fast online learning to rank, in: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, 2016, pp. 457–466.

[5] H. Wang, S. Kim, E. McCord-Snook, Q. Wu, H. Wang, 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, 2019, pp. 835–844.

[6] H. Oosterhuis, M. de Rijke, Differentiable unbiased online learning to rank, in: Proceedings of the 2018 ACM on Conference on Information and Knowledge Management, ACM, 2018.

[7] S. Zhuang, G. Zuccon, Counterfactual online learning to rank, in: European Conference on Information Retrieval, Springer, 2020, pp. 415–430.
[8] K. Hofmann, S. Whiteson, M. De Rijke, A probabilistic method for inferring preferences from clicks, in: Proceedings of the 20th ACM international conference on Information and knowledge management, 2011, pp. 249–258.
[9] E. Kharitonov, Federated online learning to rank with evolution strategies, in: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, 2019, pp. 249–257.
[10] S. Wang, S. Zhuang, G. Zuccon, Federated online learning to rank with evolution strategies: A reproducibility study, in: European Conference on Information Retrieval, 2021.
[11] H. Zhu, J. Xu, S. Liu, Y. Jin, Federated learning on non-iid data: A survey, arXiv preprint arXiv:2106.06843 (2021).
[12] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al., Advances and open problems in federated learning, arXiv preprint arXiv:1912.04977 (2019).
[13] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, V. Chandra, Federated learning with non-iid data, arXiv preprint arXiv:1806.00582 (2018).
[14] N. Yoshida, T. Nishio, M. Morikura, K. Yamamoto, R. Yonetani, Hybrid-fl for wireless networks: Cooperative learning mechanism using non-iid data, in: ICC 2020-2020 IEEE International Conference on Communications (ICC), IEEE, 2020, pp. 1–7.
[15] C. Briggs, Z. Fan, P. Andras, Federated learning with hierarchical clustering of local updates to improve training on non-iid data, in: 2020 International Joint Conference on Neural Networks (IJCNN), IEEE, 2020, pp. 1–9.
[16] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, R. Rogers, Protection against reconstruction and its applications in private federated learning, arXiv preprint arXiv:1812.00984 (2018).
[17] J. Geiping, H. Bauermeister, H. Dröge, M. Moeller, Inverting gradients–how easy is it to break privacy in federated learning?, arXiv preprint arXiv:2003.14053 (2020).
[18] S. Wang, B. Liu, S. Zhuang, G. Zuccon, Effective and privacy-preserving federated online learning to rank, in: Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval, 2021, pp. 3–12.
[19] B. McMahan, E. Moore, D. Ramage, S. Hampson, B. A. y Arcas, Communication-efficient learning of deep networks from decentralized data, in: Artificial intelligence and statistics, PMLR, 2017, pp. 1273–1282.
[20] S. Wang, G. Zuccon, Is non-iid data a threat in federated online learning to rank?, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022.
[21] Q. Yang, Y. Liu, T. Chen, Y. Tong, Federated machine learning: Concept and applications, ACM Transactions on Intelligent Systems and Technology (TIST) 10 (2019) 1–19.
[22] T. Tuor, S. Wang, B. J. Ko, C. Liu, K. K. Leung, Overcoming noisy and irrelevant data in federated learning, in: 2020 25th International Conference on Pattern Recognition (ICPR), IEEE, 2021, pp. 5020–5027.
[23] K. Wang, R. Mathews, C. Kiddon, H. Eichner, F. Beaufays, D. Ramage, Federated evaluation of on-device personalization, arXiv preprint arXiv:1910.10252 (2019).
[24] A. Fallah, A. Mokhtari, A. E. Ozdaglar, Personalized federated learning with theoretical

guarantees: A model-agnostic meta-learning approach., in: NeurIPS, 2020.
[25] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in: International Conference on Machine Learning, PMLR, 2017, pp. 1126–1135.
[26] F. Hanzely, P. Richtárik, Federated learning of a mixture of global and local models, arXiv preprint arXiv:2002.05516 (2020).
[27] C. T Dinh, N. Tran, T. D. Nguyen, Personalized federated learning with moreau envelopes, Advances in Neural Information Processing Systems 33 (2020).
[28] Y. Mansour, M. Mohri, J. Ro, A. T. Suresh, Three approaches for personalization with applications to federated learning, arXiv preprint arXiv:2002.10619 (2020).
[29] A. Ghosh, J. Chung, D. Yin, K. Ramchandran, An efficient framework for clustered federated learning, Advances in Neural Information Processing Systems 33 (2020).
[30] F. Sattler, K.-R. Müller, W. Samek, Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints, IEEE transactions on neural networks and learning systems (2020).
[31] L. Lyu, H. Yu, J. Zhao, Q. Yang, Threats to federated learning, in: Federated Learning, Springer, 2020, pp. 3–16.
[32] B. Biggio, B. Nelson, P. Laskov, Poisoning attacks against support vector machines, arXiv preprint arXiv:1206.6389 (2012).
[33] C. Fung, C. J. Yoon, I. Beschastnikh, Mitigating sybils in federated learning poisoning, arXiv preprint arXiv:1808.04866 (2018).
[34] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, V. Shmatikov, How to backdoor federated learning, in: International Conference on Artificial Intelligence and Statistics, PMLR, 2020, pp. 2938–2948.
[35] A. N. Bhagoji, S. Chakraborty, P. Mittal, S. Calo, Analyzing federated learning through an adversarial lens, in: International Conference on Machine Learning, PMLR, 2019, pp. 634–643.
[36] M. Fang, X. Cao, J. Jia, N. Gong, Local model poisoning attacks to byzantine-robust federated learning, in: 29th {USENIX} Security Symposium ({USENIX} Security 20), 2020, pp. 1605–1622.
[37] R. Guerraoui, S. Rouault, et al., The hidden vulnerability of distributed learning in byzantium, in: International Conference on Machine Learning, PMLR, 2018, pp. 3521–3530.
[38] D. Yin, Y. Chen, R. Kannan, P. Bartlett, Byzantine-robust distributed learning: Towards optimal statistical rates, in: International Conference on Machine Learning, PMLR, 2018, pp. 5650–5659.