Learning Product Rankings Robust to Fake Users

Learning Product Rankings Robust to Fake Users

NEGIN GOLREZAEI, MIT Sloan School of Management, Massachusetts Institute of Technology
VAHIDEH MANSHADI, Yale School of Management, Yale University
JON SCHNEIDER, Google Research
SHREYAS SEKAR, Department of Management, University of Toronto Scarborough and Rotman School of Management

In many online platforms, customers' decisions are substantially influenced by product rankings as most customers only examine a few top-ranked products. Concurrently, such platforms also use the same data corresponding to customers' actions to learn how these products must be ranked or ordered. These interactions in the underlying learning process, however, may incentivize sellers to artificially inflate their position by employing fake users, as exemplified by the emergence of click farms. Motivated by such fraudulent behavior, we study the ranking problem of a platform that faces a mixture of real and fake users who are indistinguishable from one another. We first show that existing learning algorithms—that are optimal in the absence of fake users—may converge to highly sub-optimal rankings under manipulation by fake users. To overcome this deficiency, we develop efficient learning algorithms under two informational environments: in the first setting, the platform is aware of the number of fake users, and in the second setting, it is agnostic to the number of fake users. For both these environments, we prove that our algorithms converge to the optimal ranking, while being robust to the aforementioned fraudulent behavior; we also present worst-case performance guarantees for our methods, and show that they significantly outperform existing algorithms. At a high level, our work employs several novel approaches to guarantee robustness such as: (*i*) constructing product-ordering graphs that encode the pairwise relationships between products inferred from the customers' actions; and (*ii*) implementing multiple levels of learning with a judicious amount of bi-directional cross-learning between levels. Overall, our results indicate that online platforms can effectively combat fraudulent users without incurring large costs by designing new learning algorithms that guarantee efficient convergence even when the platform is completely oblivious to the number and identity of the fake users.
A full version of this paper is available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3685465.

CCS Concepts: • **Applied computing** → **Online shopping**; • **Theory of computation** → **Regret bounds**.

Additional Key Words and Phrases: product ranking, sequential search, robust learning, fake users, online platforms

ACM Reference Format:
Negin Golrezaei, Vahideh Manshadi, Jon Schneider, and Shreyas Sekar. 2021. Learning Product Rankings Robust to Fake Users. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3465456.3467580

1 EXTENDED ABSTRACT

The abundance of substitutable products on online shopping platforms combined with consumers' limited attention has resulted in a new form of competition among products: the race for visibility. For example, an Amazon user is typically presented with a ranking of thousands of search results—displayed in a sequence of web-pages each containing a few dozen products—even though she is

EC '21, July 18–23, 2021, Budapest, Hungary

ACM ISBN 978-1-4503-8554-1/21/07.
https://doi.org/10.1145/3465456.3467580

unlikely to go beyond the first page.Consequently, the success of a product crucially depends on its position in the ranking. Cognizant of such position effects, online platforms tend to rank more popular products higher (i.e., make them more visible). However, because the popularity of products is a priori unknown, the platform seeks to learn them through the same process, i.e., by presenting a ranked assortment of products to users and getting feedback from them. Such an online, real-time learning process opens the possibility of manipulations in the race for visibility: "click farms" have emerged in which firms employ *fake users* who would click on designated products in the hope of boosting their popularity and thus misleading the platform to rank them in top positions.The emergence and prevalence of such fraudulent behavior raises the following key question: *can an online platform efficiently learn the optimal product ranking in the presence of fake users?*

We pursue this question in the context of an online platform that presents each arriving customer with a fixed set of products, displayed in a particular order (a ranking). Customers examine the products sequentially until they identify and click on the desired product, exhibiting position bias as they are more likely to only view products in top ranks. The platform then seeks to learn product preferences from click feedback in order to refine its ranking for future customers. However, it faces the threat of manipulation from fake users. In particular, F out of the T customers who visit the platform constitute 'fake users'; such users may strategically click on certain products in order to boost their position or withhold clicks to achieve the opposite effect. Crucially, the platform is not aware of the identity of these fake users and cannot simply ignore their feedback. Therefore, their actions can distort the platform's perception of product popularity, and lead to downstream consequences for real customers who may see undesirable products at top positions. In the face of these challenges, developing learning algorithms that are robust to fake users is clearly a priority.

Summary of Contributions. In this work, we follow a regret analysis framework and assess the performance of learning algorithms by proving worst-case guarantees parameterized by the number of fake users F, which we refer to as the *fakeness budget*. Given the above model, we show the following results.

(1) We show that commonly used learning algorithms for product ranking are vulnerable to fake users in that their regret can be $\Omega(T)$, even when the number of fake users is small.
(2) For the setting where the *fakeness budget* F is known to the platform, we design a deterministic online algorithm called *Fake-Aware Ranking (FAR)* whose worst-case regret is $O(\log(T) + F)$.
(3) For a more challenging setting where the fakeness budget is unknown to the platform, we design a randomized online algorithm called *Fake-Oblivious Ranking with Cross-Learning (FORC)* whose worst-case regret is $O(F \log(T))$.
(4) Finally, we carry out a numerical study using synthetic data that illustrates the superior performance of FORC even though the algorithm is unaware of the fakeness budget.

All together, our results show that *an online platform can effectively combat fake users without incurring too much cost by employing learning algorithms that are robust to such fraudulent behavior.*[1]

ACKNOWLEDGMENTS

Negin Golrezaei would like to acknowledge the MIT Junior Faculty Research Assistance Program. Shreyas Sekar would like to acknowledge the funding and support of the TD Management Data and Analytics Lab at the Rotman School of Management for this research.

[1] A full version of this paper is available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3685465.