

Organic Ponies and Sponsored Batteries: A Category-Based CTR Optimization Model

Or Levi
ebay/Marktplaats
olevi@ebay.com

ABSTRACT

A common challenge for E-commerce sites is the allocation of available digital real estate between organic and sponsored results. While methods for optimizing each type of results in isolation have been extensively studied, selective presentation of these two types to optimize overall performance has been largely unexplored.

Our work aims to address this allocation challenge at Marktplaats.nl, one of the largest sites in the ebay classifieds group. To this end, we explore the interplay between organic and sponsored results across a variety of item categories while reflecting on findings by previous works. We hypothesize that in categories of niche items, such as Ponies, organic results perform better than sponsored results, while in categories of commoditized items, such as Batteries, the opposite is true.

Based on our findings, we propose a simple and adaptive allocation model to improve the overall CTR performance. Empirical evaluation attests to the merits of our model, compared to the existing method in production, with a significantly higher click-through rate for both organic and sponsored results.

For future work, we consider the challenges of optimizing the allocation for profitability, rather than clicks, and taking into account additional factors beyond category, such as personal user preferences.

KEYWORDS

E-commerce, Sponsored Advertising, Click Prediction

ACM Reference format:

Or Levi. 2019. Organic Ponies and Sponsored Batteries: A Category-Based CTR Optimization Model. In *Proceedings of ACM SIGIR Workshop on eCommerce, Paris, France, July 2019 (SIGIR 2019 eCom)*, 5 pages.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR 2019 eCom, July 2019, Paris, France

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

E-commerce sites often present users with two types of results: organic and sponsored. Both serve important business functions; Organic results represent consumer-to-consumer listings that help to maintain an active user base, whereas sponsored results represent business-to-consumer ads which allow for monetization. This gives rise to the challenge of allocating available digital real estate between these two types of results. Previous works [2, 5, 6, 8] have addressed the challenge of optimizing organic and sponsored results in isolation. However, selective presentation of these two types to optimize overall performance has been largely unexplored.

Our work aims to address this allocation challenge at Marktplaats.nl, one of the largest sites in the ebay classifieds group. The Marktplaats homepage feed, presented in figure 1, the largest placement on the site in terms of traffic and revenues, employs a paradigm that allocates equal amounts of organic and sponsored impressions on a per category basis. The homepage feed holds two desirable traits for our study on the relevancy of results. First, unlike the search result pages, where the presentation order of organic and sponsored results can affect the performance, all results in each page of the feed are shuffled together, producing a random order and eliminating position bias towards one type of results. Second, while the sponsored results on search result pages are marked with a badge, there is no similar mark for sponsored results on the feed, removing the disclosure effect on user behavior.

To address the allocation challenge, we study the relationship between the item category and the relative performance of the two types of results. Our hypothesis is that in some categories organic results perform better than sponsored results while in others the opposite is true, due to the different nature of these two types. Organic results usually reflect more second hand stuff or niche items, while sponsored results are geared more towards new products and commoditized items. For example, users looking for Ponies are more likely to be interested in the organic results, while users looking for Batteries are likely to find the sponsored results more relevant.

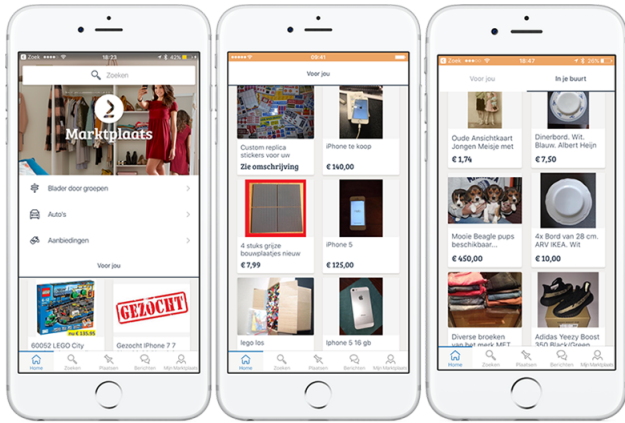


Figure 1: The Markplaats Homepage Feed

Our main contribution is a framework for selective presentation of organic and sponsored results to optimize the overall performance of an E-commerce site operator. We show through empirical evaluation that our method outperforms the existing method in production, with a significantly higher click-through rate for both organic and sponsored results.

2 RELATED WORK

Previous works [1, 4, 7] have studied the interplay between organic and sponsored results on the search results page. Yang et al. [7] studied whether the presence of organic listings on a search engine is associated with a positive, a negative, or no effect on the click-through rates of paid search advertisements. Their findings suggest that clicks on organic listings have a positive interdependence with clicks on paid listings, and vice versa, and that this positive interdependence is asymmetric such that the impact of organic clicks on increases in utility from paid clicks is much stronger.

Danescu et al. [4] investigated the perceived relative usefulness of the results with respect to the nature of the query. They found that when both sources focus on the same intent, for navigational queries there is a clear competition between ads and organic results, while for non-navigational queries this competition turns into synergy. Similarly, Agarwal et al. [1] found that an increase in organic competition leads to a decrease in the click performance of sponsored advertisements. However, organic competition helps the conversion performance of sponsored ads and leads to higher revenue.

3 DATA EXPLORATION

To test our hypothesis, we collect a dataset based on the responses of users to several hundred millions of impressions, across more than one thousand categories, over a one-month period, from the logs at Marktplats.nl. We use the data to




Figure 2: Ratio of organic CTR to sponsored CTR per category, sorted in descending order. For instance, organic results perform better for 'Ponies', but sponsored results are relatively better for 'Batteries'. Since organic results outperform across the majority of categories, our method employs a normalization by the median CTR ratio, such that in half of the categories we show more organic and in the other half more sponsored.

Rank	Category
1	Animals and Accessories Ponies
2	Animals and accessories Dogs
3	Mopeds Honda
4	Animals and Accessories Cats
5	Computer Games Nintendo Game Boy

Table 1: The top 5 item categories with highest organic CTR compared to the sponsored CTR

Rank	Category
1	Cell Phones Chargers and car chargers
2	Audio, TV and Photography Batteries
3	Holiday homes Italy
4	Car miscellaneous Stickers
5	Services and Professionals Movers

Table 2: The top 5 item categories with highest sponsored CTR compared to the organic CTR

explore the relative performance of organic and sponsored results across the different categories while also reflecting on findings by previous works.

Our study reveals, in accordance with prior findings [3], that organic results generally attain higher click-through rate than sponsored results, as shown in figure 2. However, the ratio of organic CTR to sponsored CTR varies to a large degree across the different categories.

Tables 1 and 2 show the top 5 item categories (translated from Dutch) with the highest organic CTR compared to the

Organic Ponies and Sponsored Batteries:
A Category-Based CTR Optimization Model

sponsored CTR, and vice versa. Among the categories with the highest relative organic CTR, 'Animals and Accessories' is very dominant. This might be a result of users generally looking for animals while the sponsored ads are selling accessories, such as dog harnesses and horse food. The other categories with relatively high organic CTR are 'Mopeds | Honda' and 'Computer Games | Nintendo Game Boy'. A quick examination of the inventory of ads in these categories reveals that there is no business-to-consumer seller of 'Honda mopeds' or 'Nintendo Game Boy', but only ads for moped parts and console games, respectively.

Among the categories with the highest relative sponsored CTR, it is not surprising to see 'Batteries', 'Phone Chargers' and 'Car Stickers', given that users normally do not buy these items second hand. Further examination of more categories with relatively high sponsored CTR reveals multiple examples in 'Holiday homes' and 'Services and Professionals'. It could be that users value the reputation and expertise of a business-to-consumer seller in these categories in particular. Overall, this confirms our hypothesis with regard to the different nature of organic and sponsored results, and the potential to adjust the allocation on a per category basis.

4 METHOD

Our work aims to allocate impressions between organic and sponsored results to improve the overall performance. While the profitability of clicks on sponsored results is straightforward to measure, it is much more difficult to evaluate how much clicks on organic results are worth, given that organic results do not generate revenue directly, but help to maintain an active user base.

Consequently, we make two simplifying assumptions. First, we focus on optimizing for clicks, rather than profitability, as a common denominator for organic and sponsored results. Our assumption is that more clicks would translate to more leads with organic results, and more revenues with sponsored results. Second, we bypass the question of how much an organic click is worth compared to a sponsored click, by keeping the preexisting overall balance of organic and sponsored impressions while showing, on a per category basis, more of the type that is expected to perform better. In other words, we maintain the same total numbers of organic and sponsored impressions, but only allocate them in a smarter way between the categories, such that the click-through rate, and the clicks, for both organic and sponsored results increase.

Given that organic results generally attain higher click-through rate than sponsored results, as discussed in section 3, a straightforward allocation model based on historical CTR performance is likely to impair the overall balance of impressions, resulting in significantly more organic results.

SIGIR 2019 eCom, July 2019, Paris, France




Figure 3: The proposed allocation per category based on our method. Contrary to a naive equal amount allocation, the new allocation is highly consistent with the actual CTR ratio (shown overlaid)

Our model, on the contrary, uses historical CTR performance to allocate the impressions between the two types of results, proportionally to their expected relative performance, while normalizing with the expected relative performance of the median category. This helps to account for the a priori preference of users towards organic results and maintain the preexisting overall balance of organic and sponsored impressions, given that the correlation between the relative performance and the category size is not significant (0.03 Pearson correlation). We also apply a multiplier of 0.5 such that the impressions of the median category are divided equally.

$$Allocation(x) = 0.5 \cdot \frac{CTR Ratio(x)}{Median CTR Ratio} \quad (1)$$

where $CTR Ratio(x) = Organic CTR(x) / Sponsored CTR(x)$ for a category x and such that $80\% \geq Allocation(x) \geq 20\%$.

We limit the proposed allocation, such that we never show less than 20% of the impressions from one type. This guarantees that we will have sufficient data regarding the performance of both the organic and sponsored results in each category, to continue updating the model. Contrary to the naive equal amount method, the proposed allocation per category is highly consistent with the actual CTR ratio, as presented in figure 3.

To employ the proposed allocation, we produce a table with the calculated ratio of organic to sponsored results per each category. This table is then loaded into an ElasticSearch index in production. In query time, we look-up the ratio per the relevant category and the impressions are allocated between organic and sponsored results accordingly. We use Apache Spark to build a pipeline for collecting the data and

calculating the optimal allocation. This process runs end-to-end offline, which allows for a simple and scalable solution. The Spark job runs weekly to support dynamic allocation that adapts based on changes in performance.

5 EVALUATION

The allocation challenge can be seen as classification task of predicting on a per category basis, whether sponsored results will perform better or worse than organic results. We use the data collected in section 3 to evaluate the predictions of our model in an offline setting. We split the data by weeks and use each consecutive pair of weeks as the train and test sets, predicting based on the historical CTR of the prior week and evaluating using the next one. For this classification task, the baseline with an equal amount of organic and sponsored impressions has no predictive ability, meaning that it does not provide any insight regarding the relative performance of organic and sponsored results per category. On the contrary, our model is able to predict between sponsored and organic results with precision of 0.82 and recall of 0.81, as shown in table 3.

We further evaluate our model through an online A/B test over a two-week period. Each group is assigned with an equal size of the traffic divided randomly by user ID. The evaluation demonstrates the superiority of our model, compared to the existing method in production of equal amount allocation, with a significantly higher click-through rate for both organic and sponsored results, as shown in table 4. The two-tailed paired t-test with a 0.05 significance level was used for testing statistical significance of performance differences. Further examination confirms that the overall balance of organic and sponsored impressions remains unchanged as planned.

To illustrate why the CTR increases for both organic and sponsored results, consider the following 'toy' example. Assuming we have two categories: A and B, and in each we show 100 impressions, of which 50 organic and 50 sponsored. Moreover, if we assume that in the initial state, users clicked on all the sponsored results in category A, and only those, and vice versa with the organic results in category B, then the initial CTRs for both organic and sponsored, across both categories, are 50%. With our method, we allocate 80% of the 100 impressions in category A to sponsored results and 80% of the 100 impressions in category B to organic results (respecting the 20% lower bound). If the user behavior would remain 100% consistent, we would expect the CTR for both organic and sponsored results to increase to 80%. In practice, the behavior is not fully consistent due to temporal changes in the ad inventory and user preferences, however this approximation allows to shift the allocation in a desirable direction, as demonstrated in our results. The increase

Precision	Recall	F1
0.82	0.81	0.81

Table 3: The allocation challenge as a classification task. While a naive equal amount allocation has no predictive ability, our model is able to predict between sponsored and organic results with f1-score of 0.81

Organic Results	Sponsored Results	Overall
5.98%*	8.31%*	7.10%*

Table 4: Main Results. Increase in click-through rate based on our method compared to the existing method in production. Statistically significant differences are marked with ''**

in clicks reflects that the results are generally more relevant to the users and is translated, as assumed, to an increase in leads of 0.9% with the organic results and an increase in revenues of 1.1% with the sponsored results.

6 CONCLUSION AND FUTURE WORK

Our work addressed the challenge of allocating digital real estate between organic and sponsored results. We studied the interplay between these two types of results across different categories, and found that organic results generally attain higher CTR, in accordance with prior findings, but this varies to a large degree across the different categories, confirming our hypothesis with regard to the different nature of organic and sponsored results. Based on these findings, we proposed a simple and adaptive impression allocation model that accounts for the a-priory preference of users towards organic results and is highly consistent with the actual CTR ratio per category. Empirical evaluation demonstrated the superiority of our model, compared to the existing method in production, with a significant increase in click-through rate for both organic and sponsored results, that has made a great impact on the relevancy of the results and revenues at Marktplaats.nl.

As avenues for future work, we plan to extend this work to further placements on the site. Specifically, this work has focused on the homepage feed. Next, we plan to experiment with the impression allocation method on the search result pages.

Furthermore, in this work we have made a couple of simplifying assumptions due to the difficulty of estimating the worth of clicks on organic results. Consequently, we employed a constrain to keep the overall balance of organic and sponsored impressions. This leaves room for future work to propose models for estimating the monetary value of organic

clicks, and remove this constrain, to optimize for overall profitability directly.

Lastly, a generalization of our approach could employ a confidence-based classifier to predict how good are the organic or sponsored results in a category. Note that this would still require a normalization scheme, perhaps using the a priori class probabilities. The features for this method can be based on historical performance as in our work. We also plan to study the effect of factors, such as user preferences with regard to price, buying new versus second hand, and more, on the interplay between organic and sponsored results. We envision that these features could be utilized in a contextual bandit setting to learn a personalized optimal allocation, per user and category.

7 ACKNOWLEDGMENT

We thank our colleagues at Marktplaats.nl and especially the Finding team for their support in implementation and set up of the experiment.

REFERENCES

- [1] K. Hosanagar A. Agarwal and M. Smith. 2015. Do Organic Results Help or Hurt Sponsored Search Performance?. In *Information Systems Research*. 291–300.
- [2] I. Markov L. Stout F. Xumara A. Grotov, A. Chuklin and M. de Rijke. 2015. A Comparative Study of Click Models for Web Search. In *CLEF*. 78–90.
- [3] M. Resnick B. Jansen. 2006. An examination of searcher’s perceptions of nonsponsored and sponsored links during ecommerce Web searching. In *J. Assoc. Inf. Sci. Technol.*
- [4] E. Gabrilovich V. Josifovski C. Danescu-Niculescu-Mizil, A.Z. Broder and B. Pang. 2010. Competing for users’ attention: on the interplay between organic and sponsored search results. In *WWW*. 291–300.
- [5] H. Cheng and E. Cantu-Paz. 2010. Personalized click prediction in sponsored search. In *Proceedings of the third ACM international conference on Web search and data mining, WSDM*.
- [6] T. Joachims. 2002. Optimizing Search Engines using Clickthrough Data. In *KDD*. 133–142.
- [7] A. Ghose S. Yang. 2010. Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?. In *Journal Marketing Science*.
- [8] T. Borchert T. Graepel, J. Q. Candela and R. Herbrich. 2010. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th International Conference on Machine Learning*.