

Technical Disclosure Commons

Defensive Publications Series

October 01, 2019

Detecting Spam Publishers By Serving Honeypot Ads

Walter Bogorad

Will Nelson

Daniel Summerhays

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Bogorad, Walter; Nelson, Will; and Summerhays, Daniel, "Detecting Spam Publishers By Serving Honeypot Ads", Technical Disclosure Commons, (October 01, 2019)
https://www.tdcommons.org/dpubs_series/2538



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

Detecting Spam Publishers By Serving Honeypot Ads

ABSTRACT

Click fraud, wherein bots or other unauthorized users click on ads to falsely inflate click-through rates, is a major problem in the online ad industry. This disclosure describes a type of ad, known as a honeypot ad, that is not particularly attractive to humans, but for bots is indistinguishable from a genuine ad. Publishers who employ bots to fraudulently inflate click-through rates, or to misrepresent the popularity of their app or website, are detected when the number of clicks on such honeypot ads are substantially larger than the number of clicks on genuine ads.

KEYWORDS

- Click fraud
- Click spam
- Spam publisher
- Honeypot ad
- Online advertising
- Click-through rate (CTR)

BACKGROUND

Click fraud, where bots or other unauthorized users click on ads to falsely inflate click-through rates (CTR), is a major problem in the online ad industry. Since revenue flows from advertiser to publisher on a cost-per-click basis, ill-intentioned publishers habitually employ bots to falsely increase the popularity of their web properties or apps.

At present, spam publishers are detected using automated filters. However, these filters are known to be unreliable; hence ad exchanges also conduct manual, offline analysis and

remove any ad clicks that are deemed invalid before advertisers are charged. Ad exchanges also launch investigations based on advertisers' reports of suspicious activity. Any time malicious clicks are detected, such clicks are labeled as "invalid" and credits are issued to the advertiser's account. The present approaches are generally reactive, unreliable, and rely on manual sifting of data.

DESCRIPTION

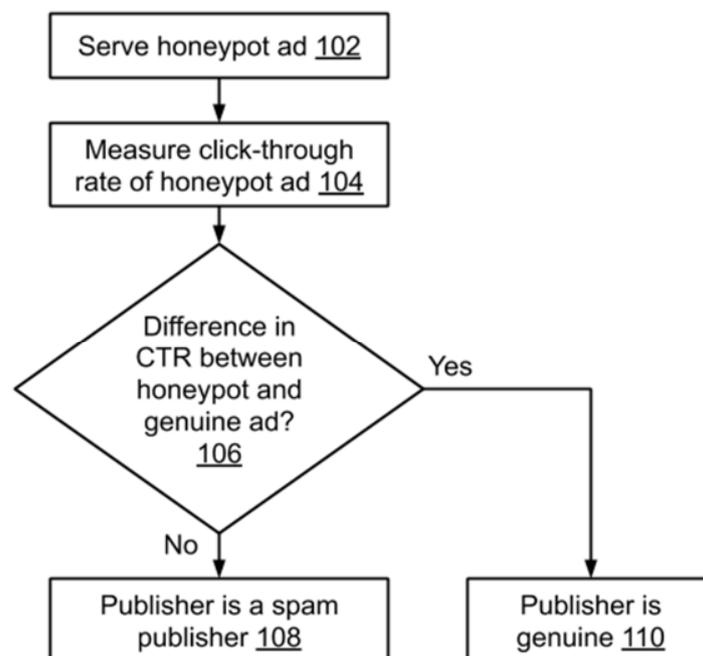


Fig. 1: Detecting spam publishers using honeypot ads

Fig. 1 illustrates detecting spam publishers using honeypot ads, per techniques of this disclosure. A honeypot ad is randomly served (102), e.g., by an ad exchange or network. Honeypot ads are ads that are not especially attractive to human users, but for bots are indistinguishable from genuine ads. Honeypot ads are also referred to herein as unattractive ads. A honeypot ad can be an ad with a low predicted CTR. The predicted CTR of an ad is available from models employed by ad exchanges or networks during the bidding process. Alternatively,

an ad exchange can select an ad randomly from an ad repository to serve as a honeypot ad, since a randomly-selected ad will have a lower CTR than an ad predicted to have a high CTR. A honeypot ad can also be an empty ad, an ad that carries a service error, e.g., “cannot complete request at this time,” etc. Empty ads or ads with a service message are to be used as honeypot ads with caution, as these may confuse genuine publishers or users into thinking that there is a problem with the ad exchange. Also, spam publishers can quickly learn to avoid empty ads or ads with service messages.

The click-through rate of the honeypot ad is measured (104). If there is a substantial difference between the CTR of the honeypot ad and a genuine ad (106), or if there is a substantial difference between the actual CTR of the honeypot ad and the predicted CTR of the honeypot ad, then the publisher is flagged as being a genuine publisher (110). If there is no substantial difference between the CTR of the honeypot ad and a genuine ad, or if there is no substantial difference between the CTR of the honeypot ad and the predicted CTR of the honeypot ad, then the publisher is flagged as spam (108).

In this manner, the techniques of this disclosure can be used to proactively identify and block spam publishers.

CONCLUSION

This disclosure describes a type of ad, known as a honeypot ad, that is not particularly attractive to humans, but for bots is indistinguishable from a genuine ad. Publishers who employ bots to fraudulently inflate click-through rates, or to misrepresent the popularity of their app or website, are detected when the number of clicks on such honeypot ads are substantially larger than the number of clicks on genuine ads.

REFERENCES

- [1] Haddadi, Hamed. "Fighting online click-fraud using bluff ads." *ACM SIGCOMM Computer Communication Review* 40, no. 2 (2010): 21-25.
- [2] Xu, Haitao, Daiping Liu, Aaron Koehl, Haining Wang, and Angelos Stavrou. "Click fraud detection on the advertiser side." In *European Symposium on Research in Computer Security*, pp. 419-438. Springer, Cham, 2014.