Computer Networks 57 (2013) 634–646

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Blog or block: Detecting blog bots through behavioral biometrics

Zi Chu [a,*], Steven Gianvecchio [a], Aaron Koehl [a], Haining Wang [a], Sushil Jajodia [b]

[a] Department of Computer Science, The College of William and Mary, Williamsburg, VA 23187, USA
[b] Center for Secure Information Systems, George Mason University, Fairfax, VA 22030, USA

ARTICLE INFO

Article history:
Received 19 May 2012
Received in revised form 18 August 2012
Accepted 10 October 2012
Available online 17 October 2012

Keywords:
Blog Bot
Behavioral biometrics
Automatic classification
Security
Web

ABSTRACT

Blog bots are automated scripts or programs that post comments to blog sites, often including spam or other malicious links. An effective defense against the automatic form filling and posting from blog bots is to detect and validate the human presence. Conventional detection methods usually require direct participation of human users, such as recognizing a CAPTCHA image, which can be burdensome for users. In this paper, we present a new detection approach by using behavioral biometrics, primarily mouse and keystroke dynamics, to distinguish between human and bot. Based on passive monitoring, the proposed approach does not require any direct user participation. We collect real user input data from a very active online community and blog site, and use this data to characterize behavioral differences between human and bot. The most useful features for classification provide the basis for a detection system consisting of two main components: a webpage-embedded logger and a server-side classifier. The webpage-embedded logger records mouse movement and keystroke data while a user is filling out a form, and provides this data in batches to a server-side detector, which classifies the poster as human or bot. Our experimental results demonstrate an overall detection accuracy greater than 99%, with negligible overhead.

1. Introduction

Blogs (from *weblog*), are a popular application of Web 2.0. Internet users publish articles on blog sites, such as personal online diaries or news on a particular subject. Like normal web pages, blog pages are primarily textual combined with images, videos, and links. The distinctive feature of blog is user interaction, which allows visitors to leave comments to blog articles. A visitor fills in the comment form and submits it, and his comment will display below the article in reverse-chronological order. Unfortunately, the increasing popularity of blogs and the simplicity of posting comments have made it easy for blog bots to automatically post comments with malicious intent. According to the estimation of [1], about 83 percent of blog comments are injected by blog bots, indicating how rampant blog bots are in the blogosphere. Most of these automated comments are associated with spam websites, containing either traceback links to inflate search engine rankings [2], or other content to lure visitors to these sites.

Since the majority of content generated by blog bots is unwanted by blog owners and visitors, blogging software has incorporated a variety of methods to discourage posting from these sources. Fundamentally, detecting human presence is an effective defense against blog bots. Conventional detection methods based on Human Interactive Proofs (HIPs) [3] usually require direct participation from human users, such as CAPTCHA. As a reverse Turing test, it challenges a user with an image carrying alphanumeric text. The user must enter the exact text before the blog site can accept the comment for submission. To cope with an advanced bot's capability for image recognition (namely, De-CAPTCHA), CAPTCHA tools add image noise to the background canvas, and greatly distort characters [4]. However, such a CAPTCHA validation also requires non-trivial effort

* Corresponding author. Tel.: +1 917 698 5015.
E-mail addresses: zichu@cs.wm.edu (Z. Chu), srgian@cs.wm.edu (S. Gianvecchio), amkoeh@cs.wm.edu (A. Koehl), hnw@cs.wm.edu (H. Wang), jajodia@gmu.edu (S. Jajodia).

1389-1286/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2012.10.005

from human users. In some cases, users have to try several times to correctly recognize a CAPTCHA image because it has become more and more difficult even for human to read. Such a validation in place may effect a significant decrease in participation from human visitors.

In this paper, we present a new method based on passive monitoring for blog bot detection, as conventional detection systems have become a nuisance for human users. Our proposed approach employs behavioral biometrics, including mouse and keystroke dynamics, to distinguish between human and bot. It has two major advantages over existing solutions. First, it uses continuous monitoring throughout the entire user session, and eliminates single checkpoints. In contrast, blog sites deployed with the conventional detection face the dilemma of applying one-time test or multiple tests. On one hand, blog bots can pass the one-time test, such as account login, with the help of human. On the other hand, multiple tests, such as recognizing a CAPTCHA image before each comment posting, are too intrusive for human users. Our passive continuous monitoring resolves the above dilemma. Second, our method is non-interactive and completely transparent to users. Moreover, no detection decision needs to be made until the user submits the comment, which in turn saves system resources. We develop a passive, webpage-embedded logger to collect user input activities on a real, active blog site. By measuring and characterizing biometric features of user input data, we discover the fundamental differences between human and blog bot in how they surf web pages and post comments. These results greatly facilitate accurate detection of blog bots.

We build a prototype of an automatic classification system that detects blog bots based on user input data. The system consists of two components, a webpage-embedded logger and a server-side detector. The logger is implemented as a JavaScript snippet that runs in the webpage on the client browser. It records a user's input actions during her stay at the site and streams the data to the server-side detector. The detector processes raw user input (UI) data, and extracts biometrics-related features. The core of the detector is a machine-learning-based classifier which is tuned with training data for the binary classification, namely determining whether the user is human or bot. Informed with the classification result, the server decides whether or not to accept the comment form submission.[1] We evaluate the efficacy of the detection system by conducting a series of experiments over the user input dataset. The experimental results demonstrate that the system can detect 97.9% of current blog bots with extremely low false positive rate of 0.2%.

As defense against bots is a challenging task, we acknowledge that our detection alone cannot eliminate the problem. However, our approach is a significant complement to conventional HIPs. We believe that, with the inherent irregularity and complexity of human behavior, it is extremely difficult if not impossible for a bot to completely mimic human behavior. Our behavior-based detection raises the bar for bot participation during this game of cat-and-mouse.

The remainder of the paper is organized as follows. Section 2 covers related work on blog bot and behavioral biometrics. Section 3 details our measurements and characterization of user inputs from human visitors and blog bots, respectively. Section 4 describes our automatic classification system. Section 5 evaluates the system efficacy for detecting blog bots. Section 6 discusses potential evasion against our detection system. Finally, Section 7 concludes the paper.

2. Background and related work

From the perspective of blog content creation, there are two types of blog bots. The first type is the article posting bot, which automatically publishes blog articles. For example, it pipelines RSS feeds from other sites as articles into the blog site, or posts preset content for a spam blog (also known as a splog). Since the posting of articles usually requires the elevated privilege of the webmaster, article posting bots are not the focus of this study. The second type is the comment posting bot, which posts comments or replies to blog sites. Given a link to a blog site, this bot analyzes the HTML structure of the blog article, especially the "leave a comment" form, fills in input fields, and posts a comment automatically. Most blog sites do not require visitors to register to post comments, and thus give ample space for bots to exploit. The focus of our work is on this bot type, and the term "blog bot" in the remainder of the paper implicitly refers to comment posting bots. Currently, blog bots are mainly created to fulfill two tasks. First, the bot posts a comment with a backlink directing to a specific website (such as that of the bot owner).[2] Posting backlinks to numerous blog sites has the effect of increasing the search engine traffic, in an attempt to boost search rankings for the originating site. The search industry has already employed some mitigation measures, such as Google's no-follow tag to prevent spam from polluting on search rankings. However, bots still massively generate inflation backlinks due to the ease and low cost of posting. Second, bots post comments with spam content (also known as spam comments) aiming to lure visitors to spam-related or other malicious sites.[3] Many blog sites eliminate spam comments based on content filtering, and Akismet [5] is such a distributed anti-spam web service. Each time a new comment is posted to the blog, it is submitted to Akismet, which checks content, runs other tests, and returns the spam detection result to the blog. Our work has the different research direction, and checks posting behavior instead of content posted.

2.1. Existing web bot detection

There have been many previous works on web bot detection. Stassopoulou et al. [6] introduced a probabilistic

[1] For instance, the server can be configured to accept the manual submission from human, and reject the automated form completion from bot.

[2] Here is an example of backlink comment, "*I don't really think this is right, but believe whatever you want. The real story can be found here on my blog: http://myblog.com/blog/*".

[3] Here is an example of spam comment, "*Thousands of cheap replica watches and fashionable designer bands at www.hot-replica.com/*".