LogoMotive: detecting logos on websites to identify online scams

Thijs van den Hout*, Thymen Wabeke*, Giovane C. M. Moura*†, Cristian Hesselman*◇
*SIDN Labs, Arnhem NL †TU Delft, Delft NL ◇University of Twente, Enschede NL
{firstname}.{lastname}@sidn.nl | +31 26 352 55 55

1 INTRODUCTION

Logos give a website a familiar feel and promote trust. Scammers take advantage of that by using well-known organizations' logos on malicious websites. Unsuspecting Internet users see the logos and think that they are looking at a legitimate webshop or government website, when it's actually a phishing site, a fake webshop or a site set up to spread misinformation.

This work focuses on identifying various types of online abuse and scams that rely on logo misuse. To that end, we present *LogoMotive* (§2), an application that employs deep-learning to detect logos on websites.

We present two cases studies, in which we apply *LogoMotive* to the 6.2M domains present in the `.nl` DNS zone – the country-code top-level domain (ccTLD) of The Netherlands, operated by SIDN. As such, it is the largest study on logo detection on websites to date (the largest before us analyzed 350k websites [2]).

In the first case study, we partner with the Dutch national government to detect government impersonation scams (§3). We detect 168 instances of government logo misuse, including phishing, spear phishing, dormant phishing attacks. These phishing websites were removed from the `.nl` zone after the usual legal due diligence. In the second case study, we team up with *Thuiswinkel Waarborg*, a widely recognized trust mark certificate issuer for webshops in The Netherlands (§4). We detect 208 domain names leading to webshops that falsely claimed to be certified by the trust mark organization by displaying their logo, thereby misleading consumers. The trust mark organization requested these websites to remove the logo.

LogoMotive is operational and has been active in the `.nl` zone for 8 months for both use cases here presented. *LogoMotive* can be applied to any DNS zone and easily trained to support different logos. Hence, we make *LogoMotive*'s source code available upon request for academic purposes and actively promote deployment by peer registries such that *LogoMotive* can be used to find abuse in other DNS zones. A full version of this paper will appear in the forthcoming PAM2022 conference [6].

2 LOGOMOTIVE

We developed *LogoMotive* to perform logo recognition on websites. It has three main modules: *Crawler*, which takes a list of domain names as input and generates screenshots from their homepage, *Logo Detector*, which applies a deep learning algorithm to detect logos on those screenshots, and the *Dashboard*, which is used by abuse analysts who are responsible for labeling the results.

The detector is built upon YOLO [5], a supervised machine learning (deep learning) algorithm designed to perform object detection. We generated a synthetic training dataset to train YOLO without having to manually label data by (i) crawling 25k `.nl` random domain names which resulted in 64k screenshots, and (ii) overlaying the logos we are interested in at random locations on these screenshots. We randomly augment the logos such that the model becomes robust against the various appearances of logos on websites. We obtained 90% precision and 98% recall in offline experiments.

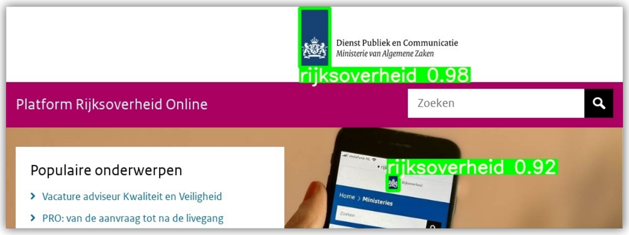
Dienst Publiek en Communicatie
Ministerie van Algemene Zaken
rijksoverheid 0.98
Platform Rijksoverheid Online
Zoeken
Populaire onderwerpen
Vacature adviseur Kwaliteit en Veiligheid
PRO: van de aanvraag tot na de livegang
rijksoverheid 0.92

Figure 1: YOLO detects the government logo on a website.

Label	Full-Zone	Newly-Registered
Total	12862 (100%)	53 (100%)
Without gov. logo (FP)	1164 (9.05%)	0 (0%)
With gov. logo (TP)	11698 (90.95%)	53 (100%)
Benign	10595 (82.37%)	32 (60.38%)
Government impersonation	151 (1.17%)	17 (32.09%)
Phishing	3 (0.02%)	3 (5.66%)
Potential threat	73 (0.57%)	9 (16.98%)
Other	75 (0.58%)	5 (9.43%)
Government domains	952 (7.40%)	4 (7.55%)
In portfolio	636 (4.94%)	2 (0.00%)
Not in portfolio	316 (2.46%)	2 (3.77%)
Added	109 (0.85%)	1 (1.89%)
Pending	207 (1.61%)	1 (1.89%)

Table 1: Manual validation results for government study.

We follow a human-in-the-loop principle, because we do not want our system to make autonomous decisions about domain names – ultimately protecting domains from being misclassified and its potential consequences. This means that *LogoMotive* detects the presence of logos on websites, but relies on abuse analysts to manually evaluate each result.

3 GOVERNMENT IMPERSONATION STUDY

After training *LogoMotive*, we apply it to detect Dutch national government impersonation scams in the `.nl` zone, which is the primary TLD used by the national government. We apply our system in two modes: we evaluate five monthly snapshots of the entire zone (March-July 2021) and evaluate every newly registered domain name to detect short-lived scams for a period of two months.

In total, *LogoMotive* detected 12.8K domain names, 11.7K of which indeed displayed the government logo (91% precision). Abuse analysts at the Dutch national government manually went through the 12.8K results and categorized all of them. Table 1 shows the results. We found 168 domains that hosted an impersonation attack and discovered 318 new government websites. We briefly highlight four insights

Phishing: We found 6 phishing websites. Their target group comprised all citizens of The Netherlands. Only 2 were found in the full zone, because phishing attacks tend to be short-lived The remaining 4 were found in newly registered domain names of which 3 were not present in Netcraft's This suggests that *LogoMotive* complements existing techniques. These malicious domain names were taken down.

Potential threats: Further, the analysts classified 82 domain names as potential threats. This includes domains that return an HTTP redirect to a legitimate government domain but are registered by a third party who has no connection with the government. Among these there are many typo-squatted and other suspicious domain names that are a dormant risk because third parties may, at any time, direct users to a scam or send e-mails that appear to be from the government.

Dormant spear phishing: 2 out of the 9 potential threats detected in newly registered domains are likely dormant spear phishing attacks. One redirected to a specialized branch of the government that is likely not known by the general public. This domain also published MX records which pointed to a mail server that is often used in shady activities, according to our abuse analysts. The other redirected to a service that is only intended for government employees. These domain names could become a serious threat because compromising a national-level agency could have severe implications. Given that spear phishing is difficult to detect, they tend to not appear on lists like Netcraft. These malicious domain names were taken down after the usual legal due diligence.

Discovered government domains: Finally, we discovered 318 legitimate government domain names that were not listed in the website portfolio of the Dutch national government. This is against the government policy for registering domains and has several risks. For instance, these domains can expire and be re-registered by a third party. This has happened before and led to data breaches at the police and a health organization in The Netherlands [3, 4]. We also observed that domains *not* in the government portfolio have lower adoption rates for both DNSSEC and DMARC, meaning that for those domain names users are not protected against DNS and e-mail spoofing.

4 TRUSTMARK ABUSE STUDY

Thuiswinkel Waarborg is a widely recognized trust mark certificate issuer for webshops in The Netherlands [1]. Hence, consumers are more likely to trust a webshop having the *Thuiswinkel Waarborg* logo on it, and therefore, be more likely to shop on these certified webshops. As a consequence, online shops have an incentive to obtain the trust mark legitimately, or to *abuse* it.

As in §3, we apply *LogoMotive* to the entire `.nl` zone to detect pages that contain *Thuiswinkel Waarborg*'s logo (June-Sept 2021). In total, we detected 10,669 domain names, 10,586 of which indeed displayed the trust mark (99.22% precision). Abuse analysts at *Thuiswinkel Waarborg* manually went through these domains and categorized all of them. Table 2 shows the results.

The majority falls in the *benign* category, *i.e.,* certified webshops. We also found 208 trust mark abuse cases. These are domains that link to webshops that display the *Thuiswinkel Waarborg* trust mark, while they are not a member. These shops are unlikely to meet the requirements that members must meet and thus pose a risk to consumers who are likely not aware of this deception. *Thuiswinkel Waarborg* contacted the companies behind the abusive domains with the request to remove the trust mark. At the time of writing, 104 abusive domains removed the trust mark from their website.

Label	Domains	Unique-URLs
Total	10669	3890
Without trust mark	83 (0.78%)	64 (1.65%)
With trust mark	10586 (99.22%)	3826 (98.35%)
Benign	10324 (96.77%)	3691 (94.88%)
Trustmark abuse	208 (1.95%)	106 (2.72%)
Discovered	54 (0.51%)	29 (0.75%)

Table 2: Manual validation results for trust mark abuse study.

We manually analyzed a sample of the 208 domains, and most of them seem to be legitimate shops, with rich and well-designed websites, and some even mention a valid Chamber of Commerce number, which indicates that it is an existing business. Next we look into the average age of the domains: half of the benign ones are at least 11 years old and half of the trust mark abuse domains are least 6 years old. That is very different from phishing, in which domains tend to be short-lived.

5 CONCLUSIONS

Logos are widely used on websites, with both benign and malicious intentions. We proposed *LogoMotive*, a system that detects logos on `.nl`-websites and provides analysts with insights into their logo's (mis)use.

Our work tangibly contributed to increasing the operational security of the Internet: the Dutch national government acted upon 168 embodiments of impersonation and 104 websites that misused the *Thuiswinkel Waarborg* trust mark removed the logos. It also allowed to discover government websites that were registered outside the official regulations

LogoMotive has proven a useful system and will be developed further for a production version at SIDN. In future research, we plan to use labels obtained during the two use cases and external data sources to automatically prioritize websites that are likely to use logos maliciously. We currently focus on the home page of websites. In future work, we will explore to what extend it is feasible to traverse internal pages as well.

REFERENCES

[1] ACM. 2016. Onderzoek naar de kennis, houding en gedrag van consumenten ten aanzien van keurmerken. https://web.archive.org/web/20180420203000/https://www.thuiswinkel.org/data/uploads/publication/ACM_en_GfK_onderzoek_keurmerken_2016.pdf. Accessed: 2021-10-20.

[2] Yun Lin et al. 2021. Phishpedia: A Hybrid Deep Learning Based Approach to Visually Identify Phishing Webpages. In *30th USENIX Security Symposium (USENIX Security 21)*.

[3] RTL Nieuws. 2017. Politiegeheimen op straat door verlopen mailadressen. https://www.rtlnieuws.nl/nieuws/nederland/artikel/240411/politiegeheimen-op-straat-door-verlopen-mailadressen. Accessed: 2021-10-15.

[4] RTL Nieuws. 2019. Groot datalek bij Jeugdzorg: dossiers duizenden kwetsbare kinderen gelekt. https://www.rtlnieuws.nl/tech/artikel/4672826/jeugdzorg-datalek-dossiers-kinderen-utrecht-email. Accessed: 2021-10-15.

[5] Joseph Redmon et al. 2016. You Only Look Once: Unified, Real-Time Object Detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.

[6] Thijs van den Hout et al. 2022. LogoMotive: detecting logos on websites to identify online scams - a TLD case study. In *Passive and Active Measurement*.