Fighting abusive registrations in .eu

The problem

- Registration of DNs with fraudulous/criminal intentions
- Types of abuse
  - To attract traffic to websites (use of reputation of somebody else)
  - To distribute malware
  - To send SPAM
  - To sell illegal products (drugs, counterfeited goods, fake medication, …)
  - To sell products and not deliver
  - Often (very) short term use (hours or days)
- Key issue : fake identity of the registrant

Prevention

Homoglyph bundling

ana.eu = ANA.eu = AnA.eu
Latin

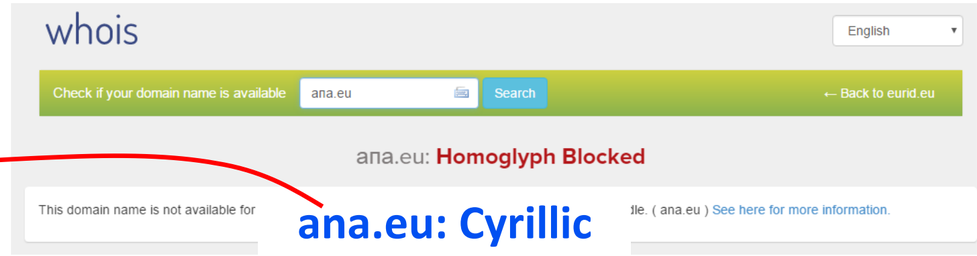

ana.eu: Cyrillic

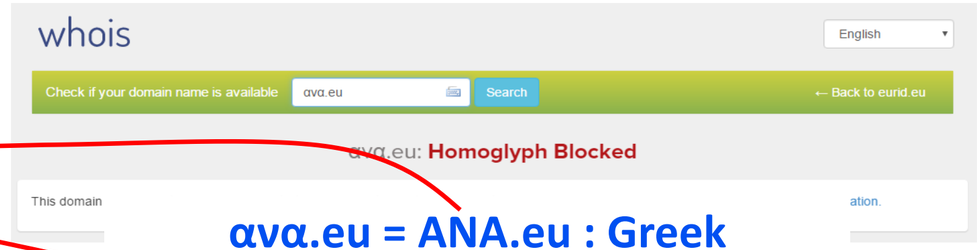

αvα.eu = ANA.eu : Greek

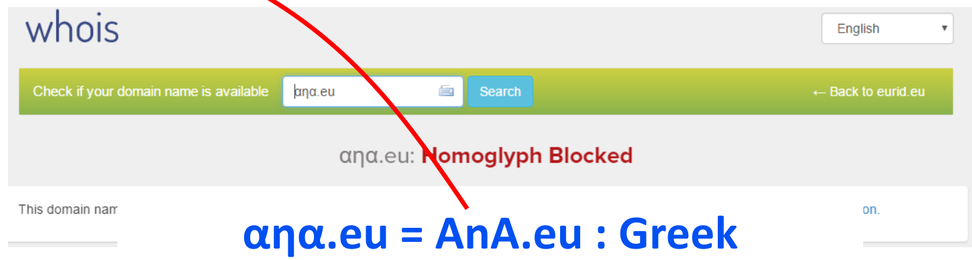

αɳα.eu = AnA.eu : Greek

Predictive Model

Objective : Predict at time of registration whether a DN will be used abusively

Previous registrations for which the results (abuse/no abuse) is known

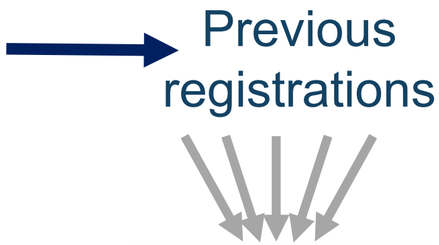

Different models are trained :
- Similarity-based agglomerative clustering
- Reputation Based Classification

For each new registration, the system predicts if the domain will be used for malicious activity

Prediction Model

New registration

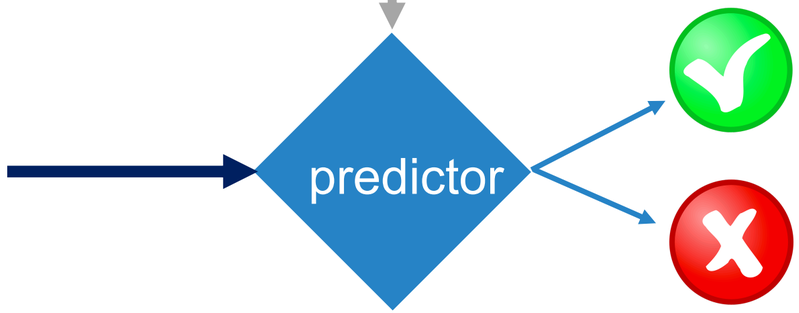

Domains with malicious intent can be
- Early detected
- Delayed
- Prevented from being registered

Similarity Based Clustering

- Rationale : Domains belonging to the same campaign have very similar registration data
- For all malicious registrations in the past period, the similarity with other malicious registrations is calculated and expressed as a metric
- Based on the inter-registration similarity, registrations are clustered into clusters of ‘very similar’ registrations,
  i.e. ‘campaigns’
- For each new registration, the distance to the malicious clusters is calculated

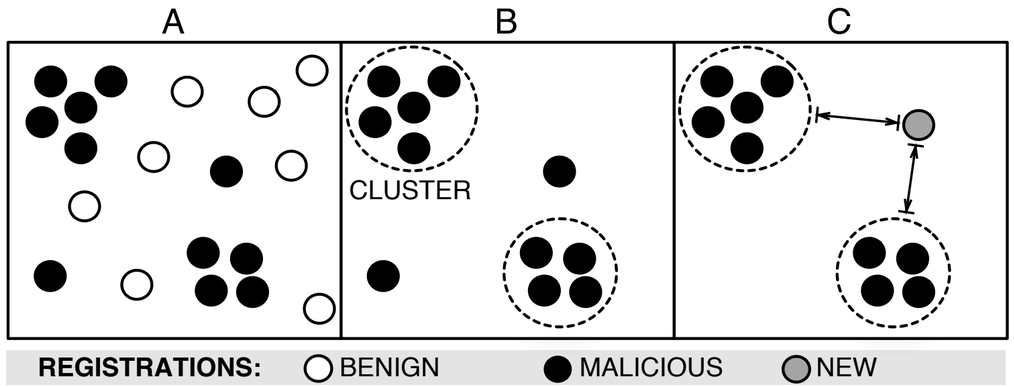

Results test phase

Prediction

Reality	Abuse	No Abuse
Abuse	True Positives (TP)	False Negatives (FN)
No Abuse	False Positives (FP)	True Negatives (TN)

Results

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

$$Recall = \frac{TP}{TP + FN}$$

How many did we find ?
(of the category we were looking for)

$$Precision = \frac{TP}{TP + FP}$$

How many were correct ?
(of those we predicted as a hit)

$$False\ Positive\ Rate = \frac{FP}{FP + TN}$$

How many were incorrectly classified as a hit ?
(of those that were not abusive)

Optimization

What is most important ?
- Find all the cases (recall ↑) with low precision ?
- Predict correctly (precision ↑) and miss a lot of cases ?
- As accurate as possible ?

Results test phase

	TP	FP->TP	FP	TN	FN	Recall	Prec.	FPR
10/01/2019 - 02/01/2019	64	254	248	28045	60	84.13%	56.18%	0.88%
02/06/2018 - 10/01/2019	1575	3919	1311	334821	1759	75.75%	80.73%	0.39%
02/04/2018 - 20/06/2018	1996	1301	488	93023	378	89.71%	87.11%	0.52%
28/03/2018 - 24/04/2018	643	1085	222	37504	140	92.51%	88.62%	0.59%
10/01/2018 - 28/03/2018	4055	24	1089	80551	867	82.47%	78.93%	1.33%

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$ How accurate is our prediction ?

$$Recall = \frac{TP}{TP + FN}$$ How many did we find ?
(of the category we were looking for)

$$Precision = \frac{TP}{TP + FP}$$ How many were correct ?
(of those we predicted as a hit)

$$False\ Positive\ Rate = \frac{FP}{FP + TN}$$ How many were wrong ?
(on total benign)

Average

TPR : 82.32%
(pct reported abuses found)

Precision: 81.62%
(pct correct on predicted abuses)

FPR : 0.58%
(abuses predicted on total benign)

Production phase (no delay)

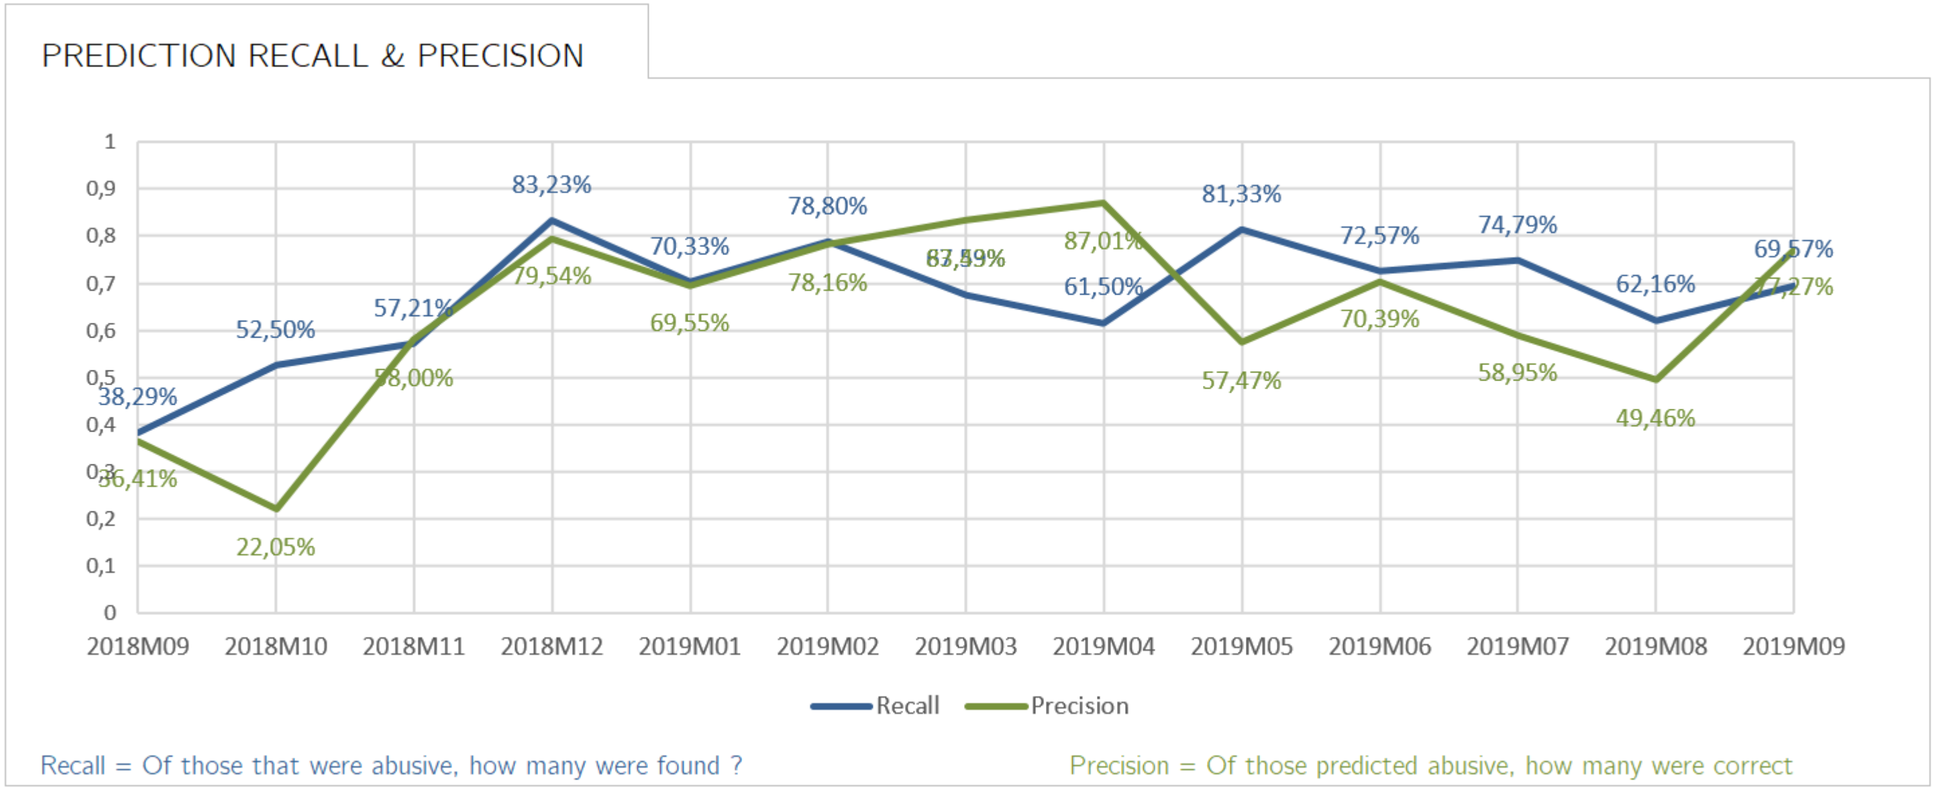

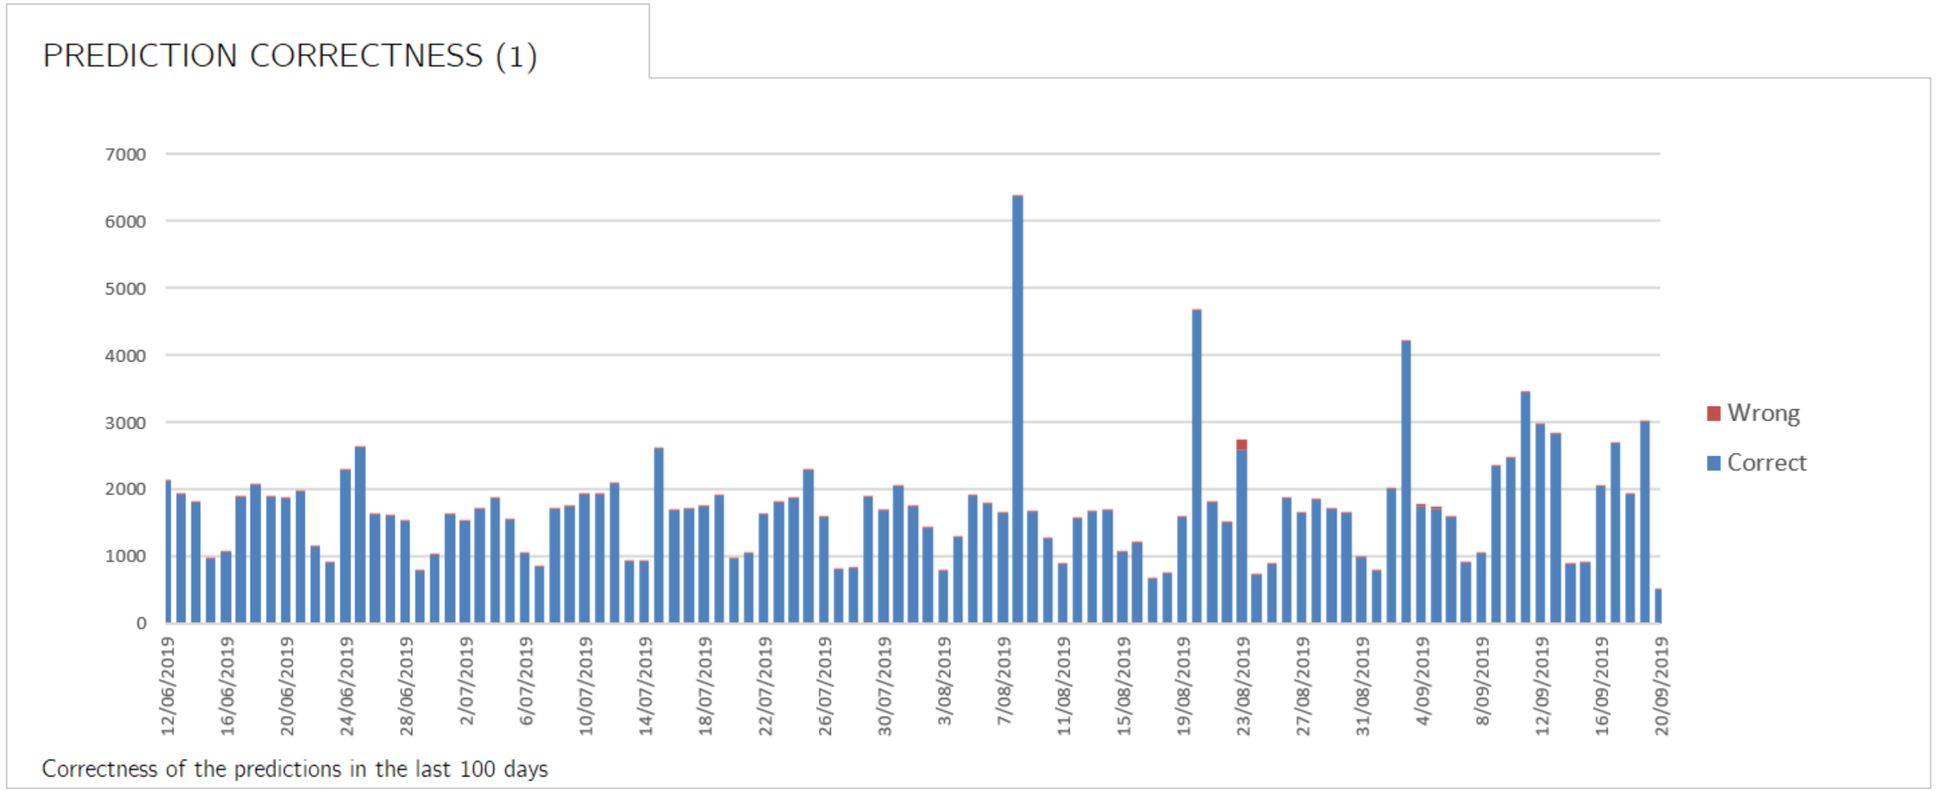
PREDICTION CORRECTNESS (1)
7000
6000
5000
4000
3000
2000
1000
0
12/06/2019
16/06/2019
20/06/2019
24/06/2019
28/06/2019
2/07/2019
6/07/2019
10/07/2019
14/07/2019
18/07/2019
22/07/2019
26/07/2019
30/07/2019
3/08/2019
7/08/2019
11/08/2019
15/08/2019
19/08/2019
23/08/2019
27/08/2019
31/08/2019
4/09/2019
8/09/2019
12/09/2019
16/09/2019
20/09/2019
Wrong
Correct
Correctness of the predictions in the last 100 days

PREDICTION CORRECTNESS (2)

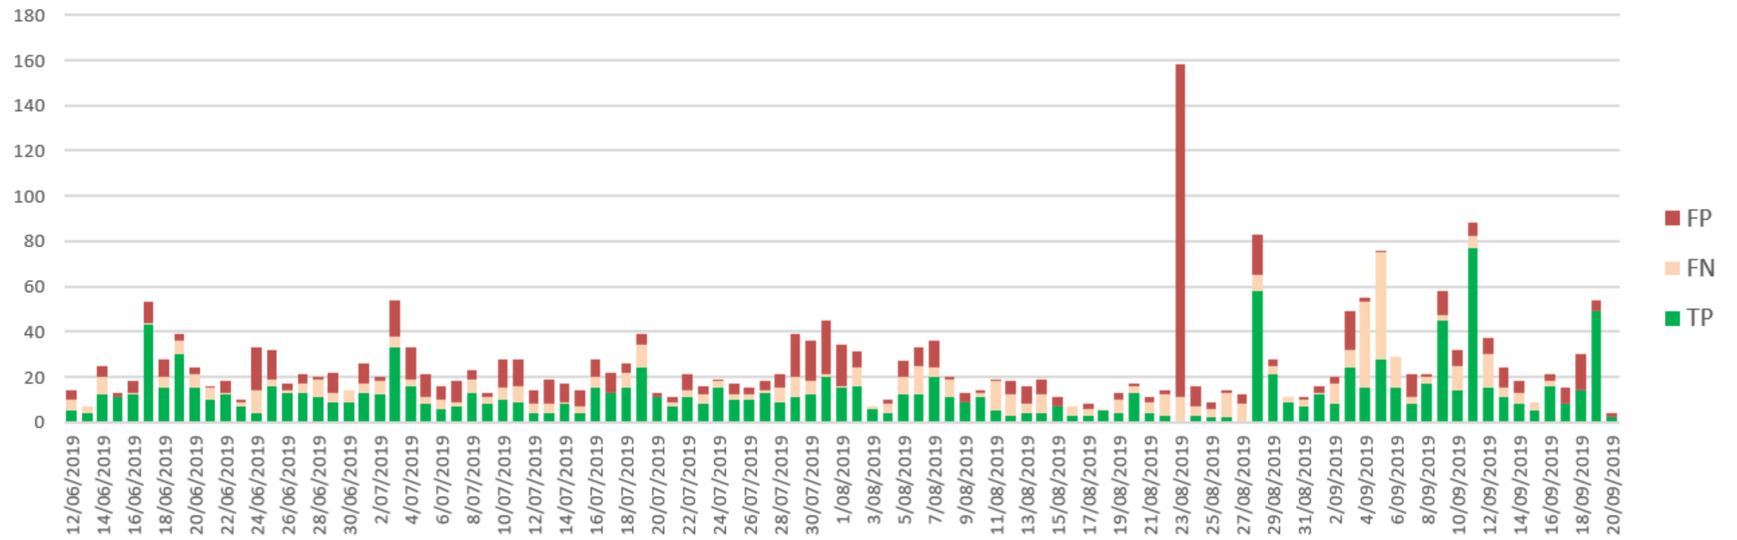

TP : Nbr of DNs that were correctly predicted as abusive in the last 100 days

FN : Nbr of DNs that were incorrectly predicted as not abusive in the last 100 days (= missed cases)

FP : Nbr of DNs that were incorrectly predicted as abusive in the last 100 days (= wrongly delayed)

Note that the FP may still turn out to be TP in the future. It just means that at the time of the report, they were not yet captured as abusive by the monitoring systems.

Effectiveness

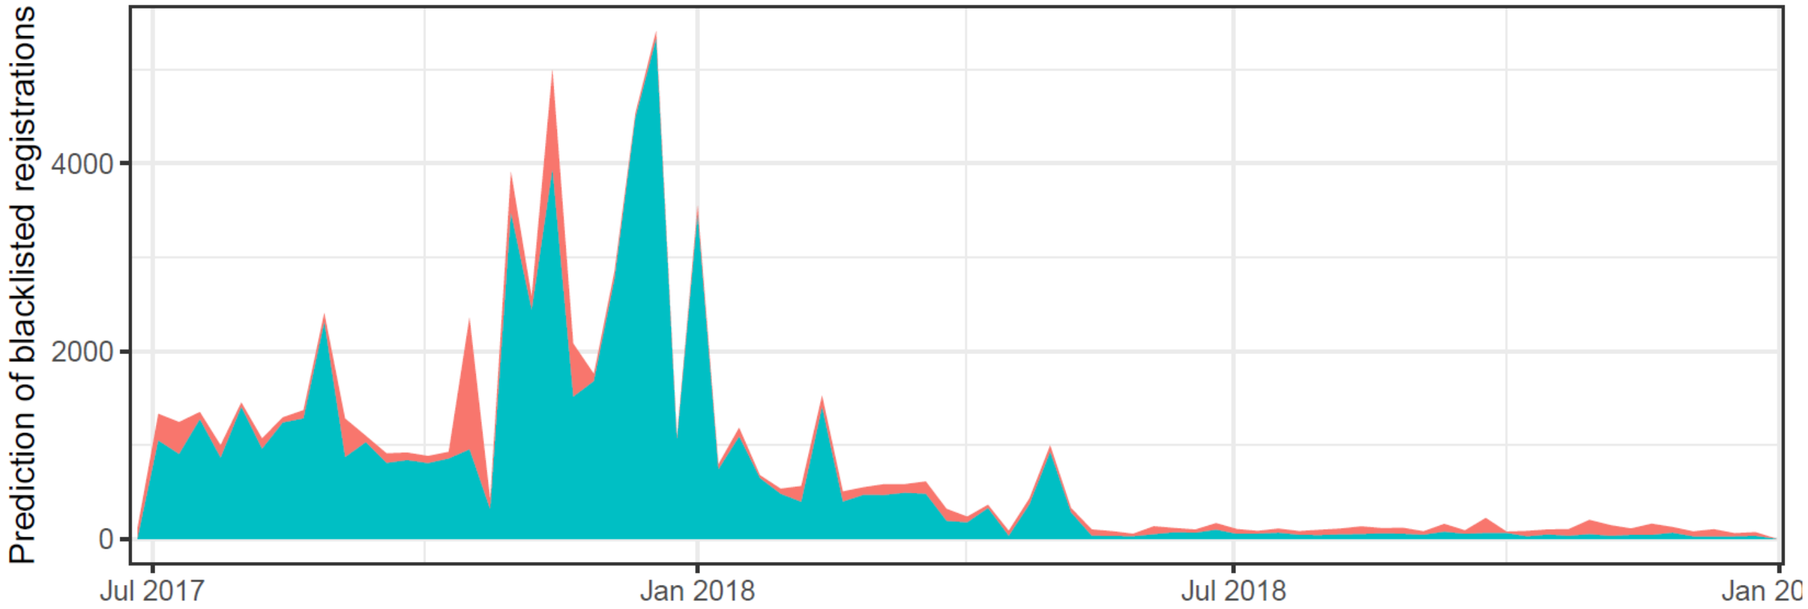

Figure 8: The weekly prediction of blacklisted registrations for the selected ensemble predictor during operations. The red area plots the total number of blacklisted registrations on that week, whereas the green area represents the predictions.

Delayed Delegation

Predict at time of registration whether a DN will be used abusively

Status :

- Running in production without delayed delegation
- Currently 80% Recall and 80% Precision

Next Steps :

- Improve algorithms (add categorisation)
- Explore to include other abuse lists
- Start delaying

More information

Exploring the ecosystem of malicious domain registrations in the .eu TLD

Thomas Vissers[1], Jan Spooren[1], Pieter Agten[1], Dirk Jumpertz[2], Peter Janssen[2], Marc Van Wesemael[2], Frank Piessens[1], Wouter Joosen[1], and Lieven Desmet[1]

[1] imec-DistriNet, KU Leuven, Belgium
{firstname.lastname}@cs.kuleuven.be,
[2] EURid VZW, Belgium
{firstname.lastname}@eurid.eu

Abstract. This study extensively scrutinizes 14 months of registration data to identify large-scale malicious campaigns present in the .eu TLD. We explore the ecosystem and modus operandi of elaborate cybercriminal entities that recurrently register large amounts of domains for one-shot, malicious use. Although these malicious domains are short-lived, by incorporating registrant information, we establish that at least 80.04% of them can be framed in to 20 larger campaigns with varying duration and intensity. We further report on insights in the operational aspects of this business and observe, amongst other findings, that their processes are only partially automated. Finally, we apply a post-factum clustering process to validate the campaign identification process and to automate the ecosystem analysis of malicious registrations in a TLD zone.

Keywords: malicious domain names, campaigns, DNS security

1 Introduction

The Domain Name System (DNS) is one of the key technologies that has allowed the web to expand to its current dimensions. Virtually all communication on the web requires the resolution of domain names to IP addresses. Malicious activities are no exception, and attackers constantly depend upon functioning domain names to execute their abusive operations. For instance, phishing attacks, distributing spam emails, botnet command and control (C&C) connections and malware distribution: these activities all require domain names to operate.

Widely-used domain blacklists are curated and used to stop malicious domain names[3] shortly after abusive activities have been observed and reported. As a consequence, attackers changed to a hit-and-run strategy, in which malicious domain names are operational for only a very small time window after the initial registration, just for a single day in 60% of the cases [11]. Once domain names

[3] We use the term *malicious domain name* whenever we refer to a domain name that is registered to be bound to a malicious service or activity.

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-66332-6_21

https://link.eurid.eu/prediction1

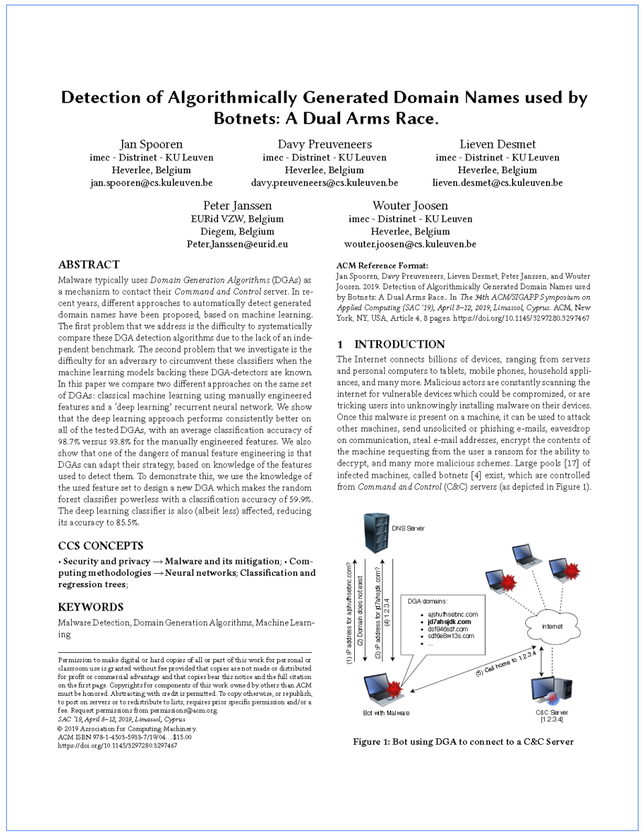

Detection of Algorithmically Generated Domain Names used by Botnets: A Dual Arms Race.

Jan Spooren
imec - Distrinet - KU Leuven
Heverlee, Belgium
jan.spooren@cs.kuleuven.be

Davy Preuveneers
imec - Distrinet - KU Leuven
Heverlee, Belgium
davy.preuveneers@cs.kuleuven.be

Lieven Desmet
imec - Distrinet - KU Leuven
Heverlee, Belgium
lieven.desmet@cs.kuleuven.be

Peter Janssen
EURid VZW, Belgium
Diegem, Belgium
Peter.Janssen@eurid.eu

Wouter Joosen
imec - Distrinet - KU Leuven
Heverlee, Belgium
wouter.joosen@cs.kuleuven.be

ABSTRACT

Malware typically uses *Domain Generation Algorithms* (DGAs) as a mechanism to contact their *Command and Control* server. In recent years, different approaches to automatically detect generated domain names have been proposed, based on machine learning. The first problem that we address is the difficulty to systematically compare these DGA detection algorithms due to the lack of an independent benchmark. The second problem that we investigate is the difficulty for an adversary to circumvent these classifiers when the machine learning models backing these DGA-detectors are known. In this paper we compare two different approaches on the same set of DGAs: classical machine learning using manually engineered features and a 'deep learning' recurrent neural network. We show that the deep learning approach performs consistently better on all of the tested DGAs, with an average classification accuracy of 98.7% versus 93.8% for the manually engineered features. We also show that one of the dangers of manual feature engineering is that DGAs can adapt their strategy, based on knowledge of the features used to detect them. To demonstrate this, we use the knowledge of the used feature set to design a new DGA which makes the random forest classifier powerless with a classification accuracy of 59.9%. The deep learning classifier is also (albeit less) affected, reducing its accuracy to 85.5%.

CCS CONCEPTS

• **Security and privacy** → **Malware and its mitigation**; • **Computing methodologies** → **Neural networks**; **Classification and regression trees**;

KEYWORDS

Malware Detection, Domain Generation Algorithms, Machine Learning

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SAC '19, April 8–12, 2019, Limassol, Cyprus
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-5933-7/19/04…$15.00
https://doi.org/10.1145/3297280.3297467

ACM Reference Format:
Jan Spooren, Davy Preuveneers, Lieven Desmet, Peter Janssen, and Wouter Joosen. 2019. Detection of Algorithmically Generated Domain Names used by Botnets: A Dual Arms Race. In *The 34th ACM/SIGAPP Symposium on Applied Computing (SAC '19), April 8–12, 2019, Limassol, Cyprus.* ACM, New York, NY, USA, Article 4, 8 pages. https://doi.org/10.1145/3297280.3297467

1 INTRODUCTION

The Internet connects billions of devices, ranging from servers and personal computers to tablets, mobile phones, household appliances, and many more. Malicious actors are constantly scanning the internet for vulnerable devices which could be compromised, or are tricking users into unknowingly installing malware on their devices. Once this malware is present on a machine, it can be used to attack other machines, send unsolicited or phishing e-mails, eavesdrop on communication, steal e-mail addresses, encrypt the contents of the machine requesting from the user a ransom for the ability to decrypt, and many more malicious schemes. Large pools [17] of infected machines, called botnets [4] exist, which are controlled from *Command and Control* (C&C) servers (as depicted in Figure 1).

Figure 1: Bot using DGA to connect to a C&C Server

https://link.eurid.eu/prediction2

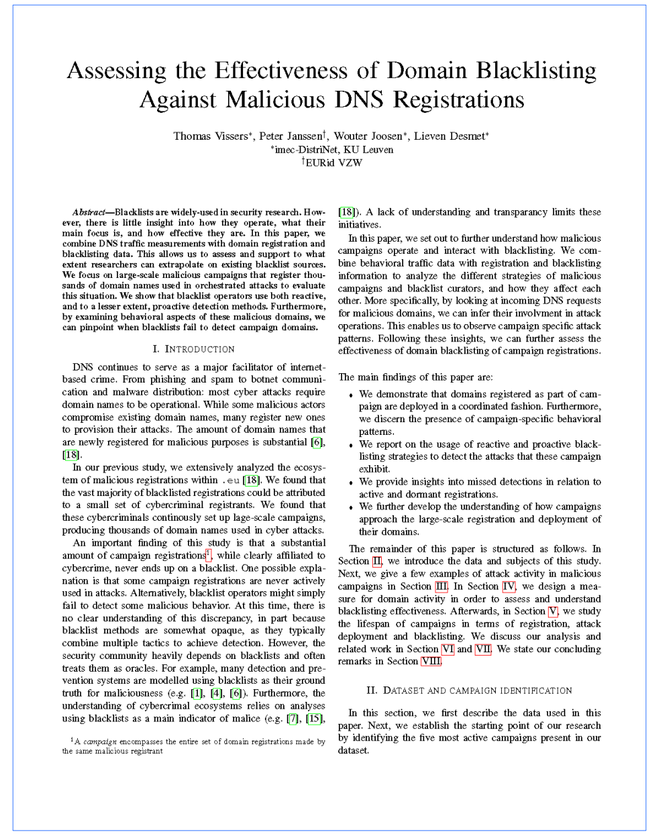

Assessing the Effectiveness of Domain Blacklisting Against Malicious DNS Registrations

Thomas Vissers*, Peter Janssen†, Wouter Joosen*, Lieven Desmet*
*imec-DistriNet, KU Leuven
†EURid VZW

***Abstract*—Blacklists are widely-used in security research. However, there is little insight into how they operate, what their main focus is, and how effective they are. In this paper, we combine DNS traffic measurements with domain registration and blacklisting data. This allows us to assess and support to what extent researchers can extrapolate on existing blacklist sources. We focus on large-scale malicious campaigns that register thousands of domain names used in orchestrated attacks to evaluate this situation. We show that blacklist operators use both reactive, and to a lesser extent, proactive detection methods. Furthermore, by examining behavioral aspects of these malicious domains, we can pinpoint when blacklists fail to detect campaign domains.**

I. Introduction

DNS continues to serve as a major facilitator of internet-based crime. From phishing and spam to botnet communication and malware distribution: most cyber attacks require domain names to be operational. While some malicious actors compromise existing domain names, many register new ones to provision their attacks. The amount of domain names that are newly registered for malicious purposes is substantial [6], [18].

In our previous study, we extensively analyzed the ecosystem of malicious registrations within .eu [18]. We found that the vast majority of blacklisted registrations could be attributed to a small set of cybercriminal registrants. We found that these cybercriminals continously set up lage-scale campaigns, producing thousands of domain names used in cyber attacks.

An important finding of this study is that a substantial amount of campaign registrations[1], while clearly affiliated to cybercrime, never ends up on a blacklist. One possible explanation is that some campaign registrations are never actively used in attacks. Alternatively, blacklist operators might simply fail to detect some malicious behavior. At this time, there is no clear understanding of this discrepancy, in part because blacklist methods are somewhat opaque, as they typically combine multiple tactics to achieve detection. However, the security community heavily depends on blacklists and often treats them as oracles. For example, many detection and prevention systems are modelled using blacklists as their ground truth for maliciousness (e.g. [1], [4], [6]). Furthermore, the understanding of cybercrimal ecosystems relies on analyses using blacklists as a main indicator of malice (e.g. [7], [15], [18]). A lack of understanding and transparancy limits these initiatives.

In this paper, we set out to further understand how malicious campaigns operate and interact with blacklisting. We combine behavioral traffic data with registration and blacklisting information to analyze the different strategies of malicious campaigns and blacklist curators, and how they affect each other. More specifically, by locking at incoming DNS requests for malicious domains, we can infer their involvment in attack operations. This enables us to observe campaign specific attack patterns. Following these insights, we can further assess the effectiveness of domain blacklisting of campaign registrations.

The main findings of this paper are:

- We demonstrate that domains registered as part of campaign are deployed in a coordinated fashion. Furthermore, we discern the presence of campaign-specific behavioral patterns.
- We report on the usage of reactive and proactive blacklisting strategies to detect the attacks that these campaign exhibit.
- We provide insights into missed detections in relation to active and dormant registrations.
- We further develop the understanding of how campaigns approach the large-scale registration and deployment of their domains.

The remainder of this paper is structured as follows. In Section II, we introduce the data and subjects of this study. Next, we give a few examples of attack activity in malicious campaigns in Section III. In Section IV, we design a measure for domain activity in order to assess and understand blacklisting effectiveness. Afterwards, in Section V, we study the lifespan of campaigns in terms of registration, attack deployment and blacklisting. We discuss our analysis and related work in Section VI and VII. We state our concluding remarks in Section VIII.

II. Dataset and campaign identification

In this section, we first describe the data used in this paper. Next, we establish the starting point of our research by identifying the five most active campaigns present in our dataset.

[1] A *campaign* encompasses the entire set of domain registrations made by the same malicious registrant

https://link.eurid.eu/prediction3

Coming soon :

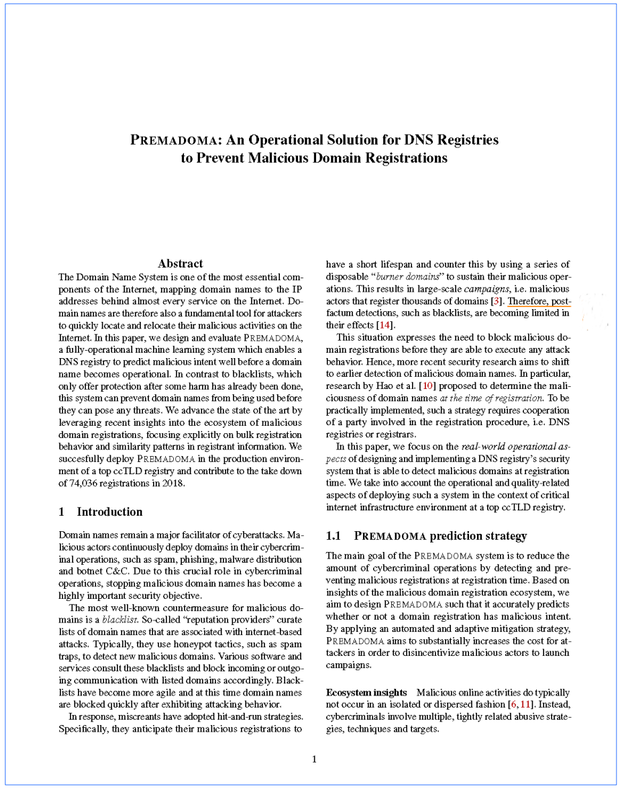

PREMADOMA: An Operational Solution for DNS Registries to Prevent Malicious Domain Registrations

Abstract

The Domain Name System is one of the most essential components of the Internet, mapping domain names to the IP addresses behind almost every service on the Internet. Domain names are therefore also a fundamental tool for attackers to quickly locate and relocate their malicious activities on the Internet. In this paper, we design and evaluate PREMADOMA, a fully-operational machine learning system which enables a DNS registry to predict malicious intent well before a domain name becomes operational. In contrast to blacklists, which only offer protection after some harm has already been done, this system can prevent domain names from being used before they can pose any threats. We advance the state of the art by leveraging recent insights into the ecosystem of malicious domain registrations, focusing explicitly on bulk registration behavior and similarity patterns in registrant information. We succesfully deploy PREMADOMA in the production environment of a top ccTLD registry and contribute to the take down of 74,036 registrations in 2018.

1 Introduction

Domain names remain a major facilitator of cyberattacks. Malicious actors continuously deploy domains in their cybercriminal operations, such as spam, phishing, malware distribution and botnet C&C. Due to this crucial role in cybercriminal operations, stopping malicious domain names has become a highly important security objective.

The most well-known countermeasure for malicious domains is a *blacklist*. So-called "reputation providers" curate lists of domain names that are associated with internet-based attacks. Typically, they use honeypot tactics, such as spam traps, to detect new malicious domains. Various software and services consult these blacklists and block incoming or outgoing communication with listed domains accordingly. Blacklists have become more agile and at this time domain names are blocked quickly after exhibiting attacking behavior.

In response, miscreants have adopted hit-and-run strategies. Specifically, they anticipate their malicious registrations to have a short lifespan and counter this by using a series of disposable "*burner domains*" to sustain their malicious operations. This results in large-scale *campaigns*, i.e. malicious actors that register thousands of domains [3]. Therefore, post-factum detections, such as blacklists, are becoming limited in their effects [14].

This situation expresses the need to block malicious domain registrations before they are able to execute any attack behavior. Hence, more recent security research aims to shift to earlier detection of malicious domain names. In particular, research by Hao et al. [10] proposed to determine the maliciousness of domain names *at the time of registration*. To be practically implemented, such a strategy requires cooperation of a party involved in the registration procedure, i.e. DNS registries or registrars.

In this paper, we focus on the *real-world operational aspects* of designing and implementing a DNS registry's security system that is able to detect malicious domains at registration time. We take into account the operational and quality-related aspects of deploying such a system in the context of critical internet infrastructure environment at a top ccTLD registry.

1.1 PREMADOMA prediction strategy

The main goal of the PREMADOMA system is to reduce the amount of cybercriminal operations by detecting and preventing malicious registrations at registration time. Based on insights of the malicious domain registration ecosystem, we aim to design PREMADOMA such that it accurately predicts whether or not a domain registration has malicious intent. By applying an automated and adaptive mitigation strategy, PREMADOMA aims to substantially increases the cost for attackers in order to disincentivize malicious actors to launch campaigns.

Ecosystem insights Malicious online activities do typically not occur in an isolated or dispersed fashion [6, 11]. Instead, cybercriminals involve multiple, tightly related abusive strategies, techniques and targets.

1

DAG VAN DE DOMAINNAAM – 24 September 2019

Powered by EURid

Thanks