Abuse Prevention and Early Warning System (APEWS)

Predictive Model

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

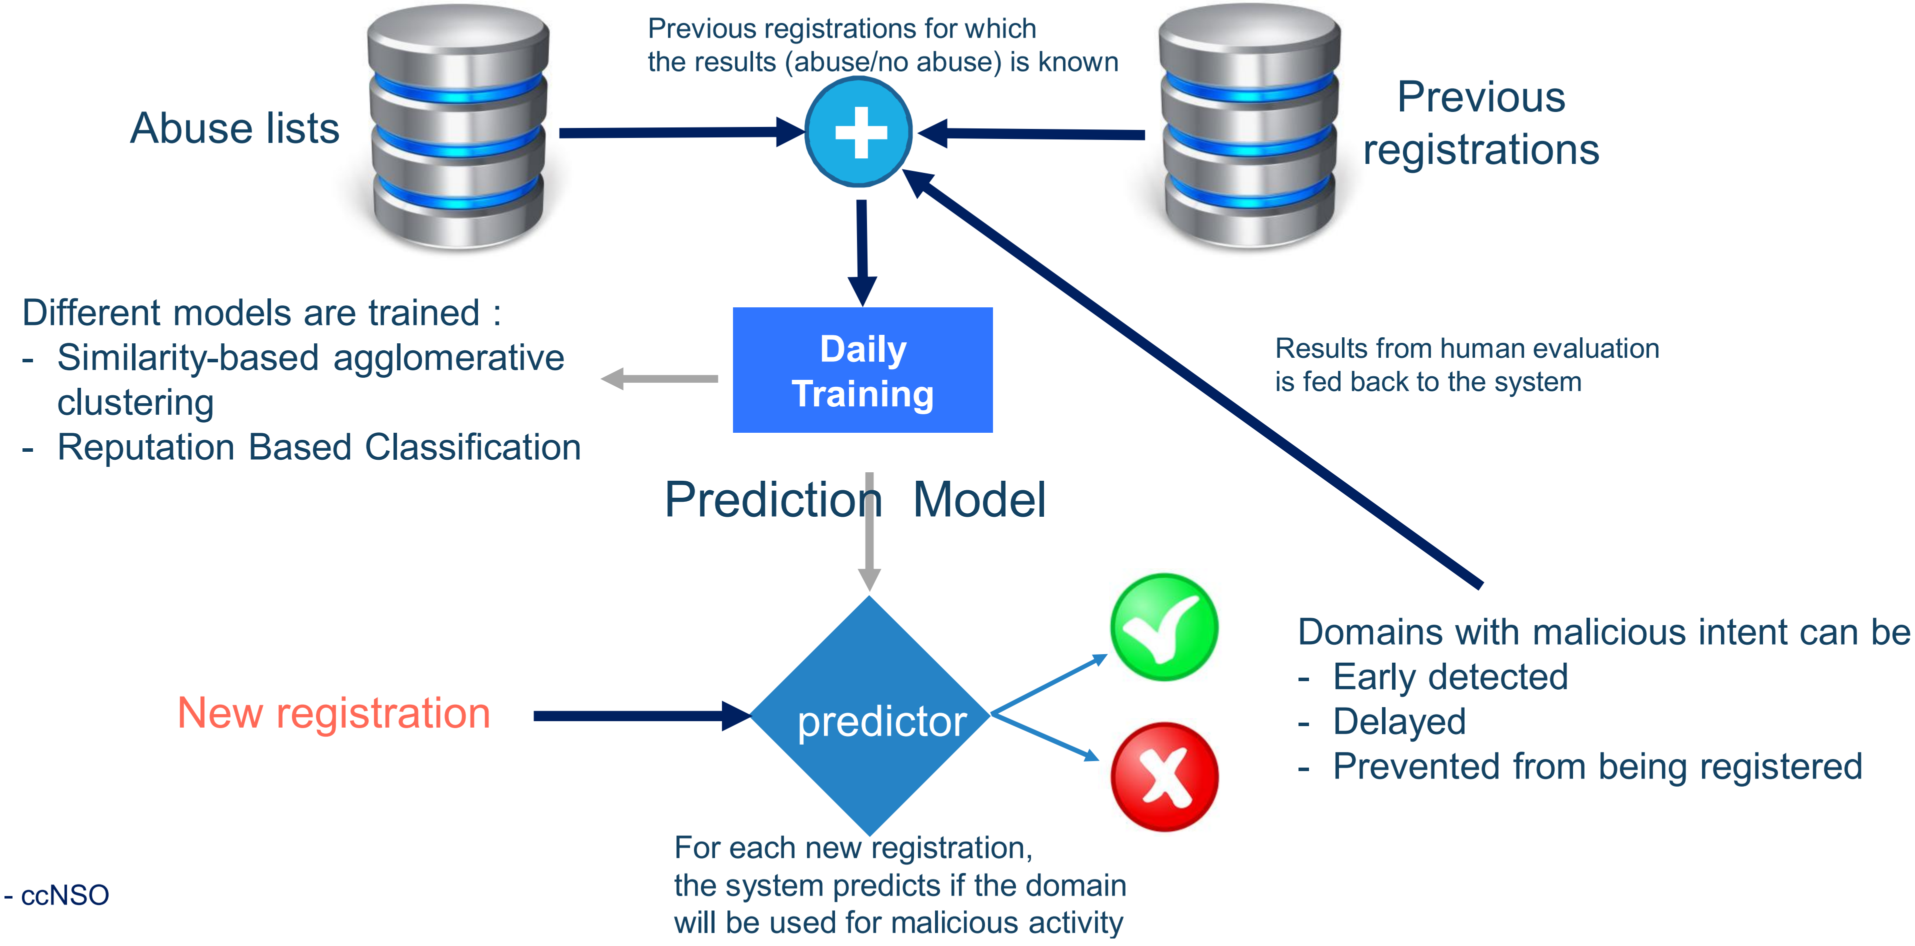

Similarity Based Clustering

- 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

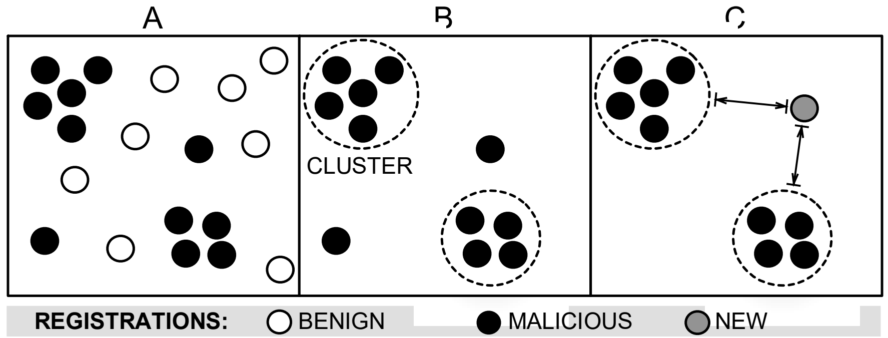

Results test phase

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)

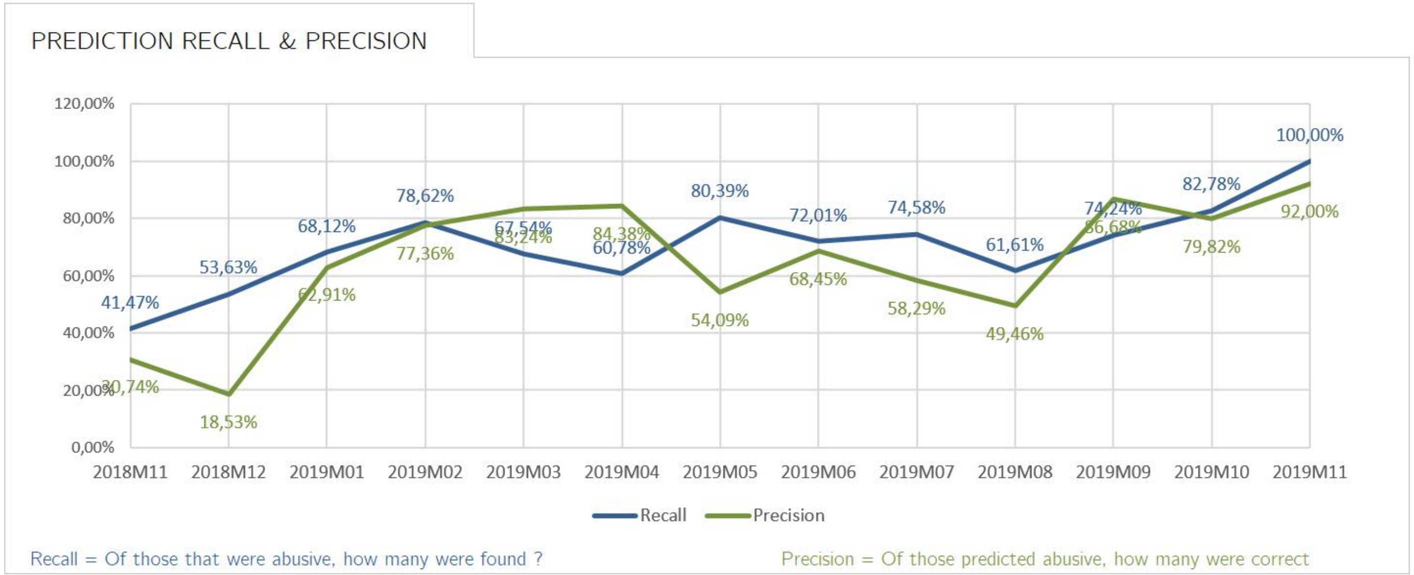

RESULTS OKT 2019

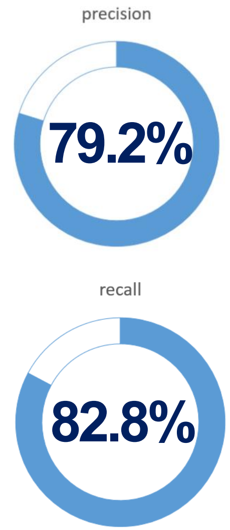

$$Recall = \frac{TP}{TP + FN}$$

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)

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 ?

PREDICTION CORRECTNESS (2)

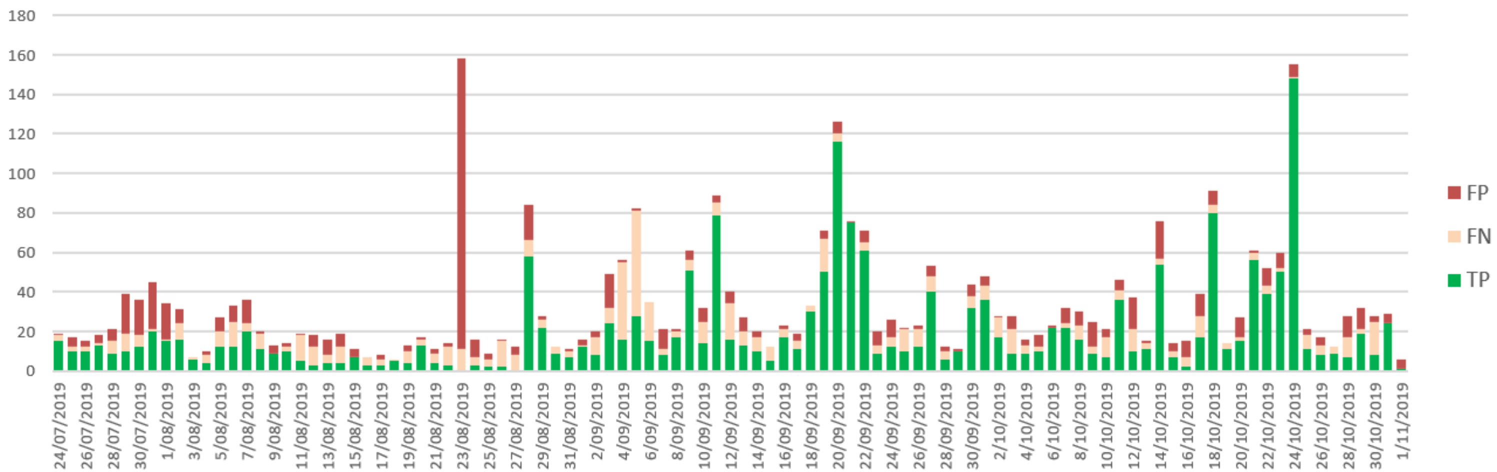

TP : Nbr of DNs that were correctly predicted as abusive in the last 100 days

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.

The Accuracy trap

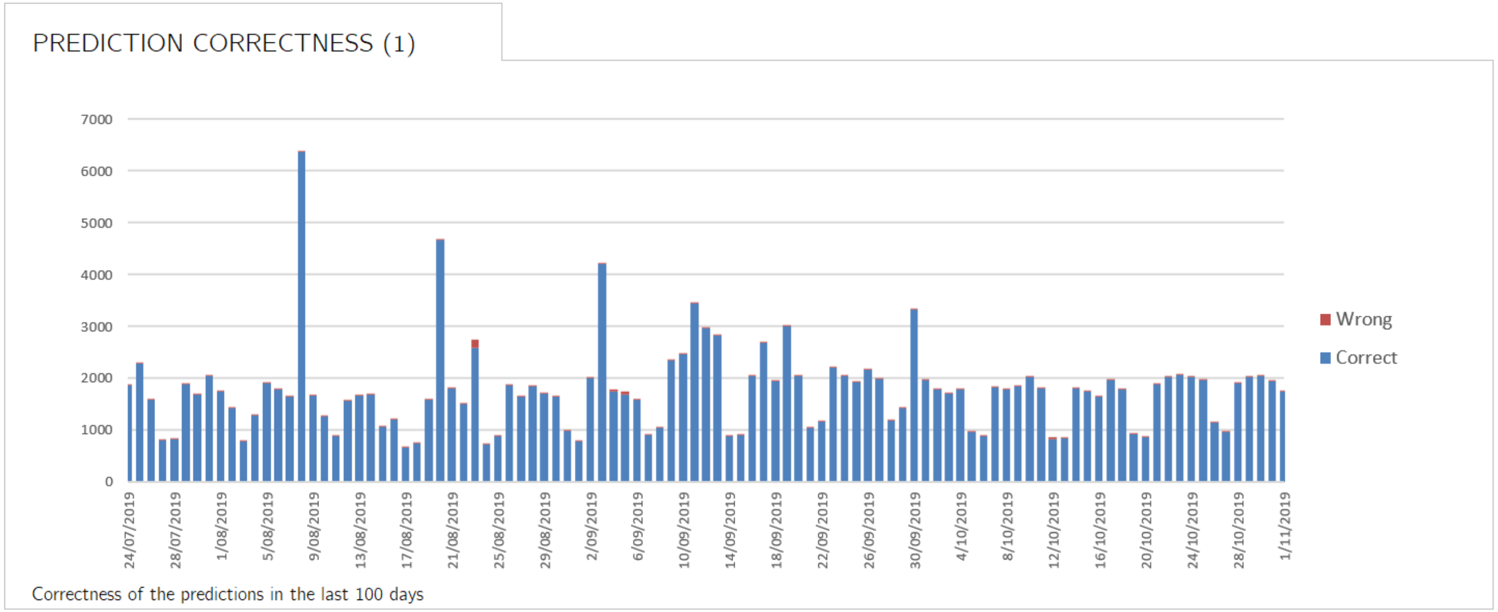

RESULTS OKT 2019

Precision

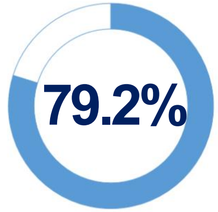

Recall

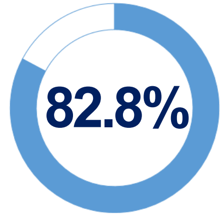

Accuracy

Pct of the prediction that was correct : 99.33%

But ... if we would always predict *no abuse*, accuracy would be 98.53% !
Typical for unbalanced data.

Effectiveness

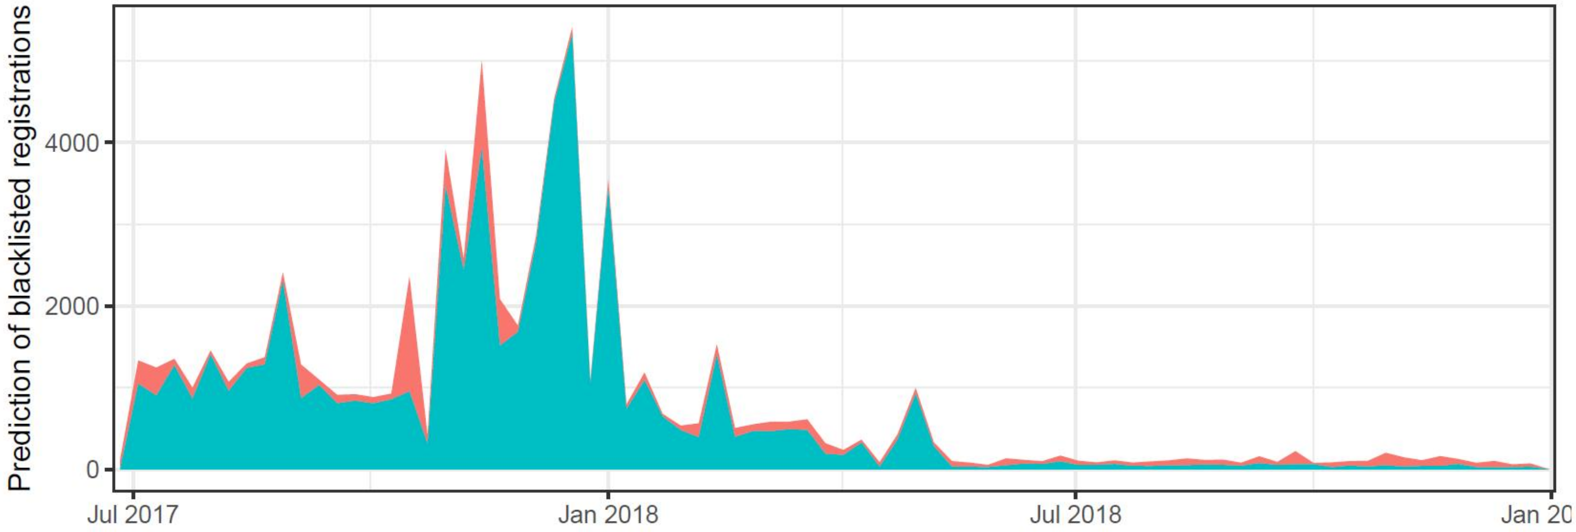

Figure 8: The weekly prediction of blacklisted registrations for the selected ensemble predictor during operations. The red area plots the total number of blacklisted registrations on that week, whereas the green area represents the predictions.

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


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

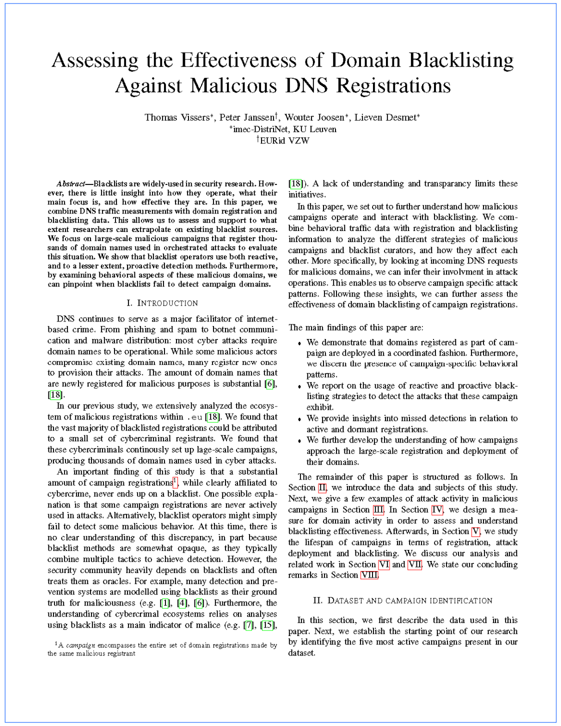
Assessing the Effectiveness of Domain Blacklisting Against Malicious DNS Registrations


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

https://link.eurid.eu/prediction1 https://link.eurid.eu/prediction2 https://link.eurid.eu/prediction3 https://link.eurid.eu/prediction4

Thanks

marc.vanwesemael@eurid.eu