Detecting and preventing domain name abuse in the .eu TLD

Lieven Desmet – SecAppDev 2019 – Leuven, BE

Lieven.Desmet@cs.kuleuven.be – @lieven_desmet

DistriNet

EURid

About me

- Senior Research Manager at KU Leuven
  - (Web) Application & DNS Security
  - Security Analytics

- Board member of intigriti
- Board member of OWASP BE chapter

- *Joint research with EURid, registry of .eu*

Malicious use of domain names

- Domain names are often abused by cyber criminals
  - Spam, botnet C&C infrastructure, phishing, malware, …
- To avoid blacklisting, malicious actors often deploy a hit-and-run strategy
  - 60% are only active for 1 day after registration [Hao et al]

[Hao et al] “Understanding the Domain Registration Behavior of Spammers” IMC 2013

Research hypothesis:

“Malicious actors register domains in bulk, and do so for longer periods of time.”

Goal of this research

- “Can we identify such bulk behavior based on commonalities between individual registrations or registration patterns?”

- Understand the malicious domain registration ecosystem
- To detect and prevent malicious registrations

Outline of the talk

- Longitudinal campaign analysis
- Registration-time prediction of malicious intent
- Detecting and preventing abuse in .eu

Longitudinal campaign analysis

Domain name registrations in the .eu TLD

- **.eu** – 8th largest ccTLD (European Economic Area)
  - 3.8 million domain names
- Dataset used in this research:
  - 824,121 new registrations over 14 months (Apr 2015 – May 2016)
  - 20,870 registrations end up on blacklists (2.5%)

Available registration data

- Basic registration information
  - domain name, registrar, registration date and time, name servers
- Contact information of the registrant
  - company name, name, language, email address, phone, fax, as well as postal address

Dataset enrichments

- Maliciousness of a domain name
  - Spamhaus DBL
  - SURBL multi list
  - Google Safe Browsing
- Geolocation information of name servers
  - MaxMind GeoLite2 Free database

Example campaign (c_11)

- Multiple fake registrant details
  - Combinations of

    2 email accounts,

    3 phone numbers,

    4 street addresses

- **8 months active (Jun 3, 2015 – Feb 3, 2016)**
- **1,275 blacklisted registrations** ***(= 53.96%)***

Registration details used by c_11

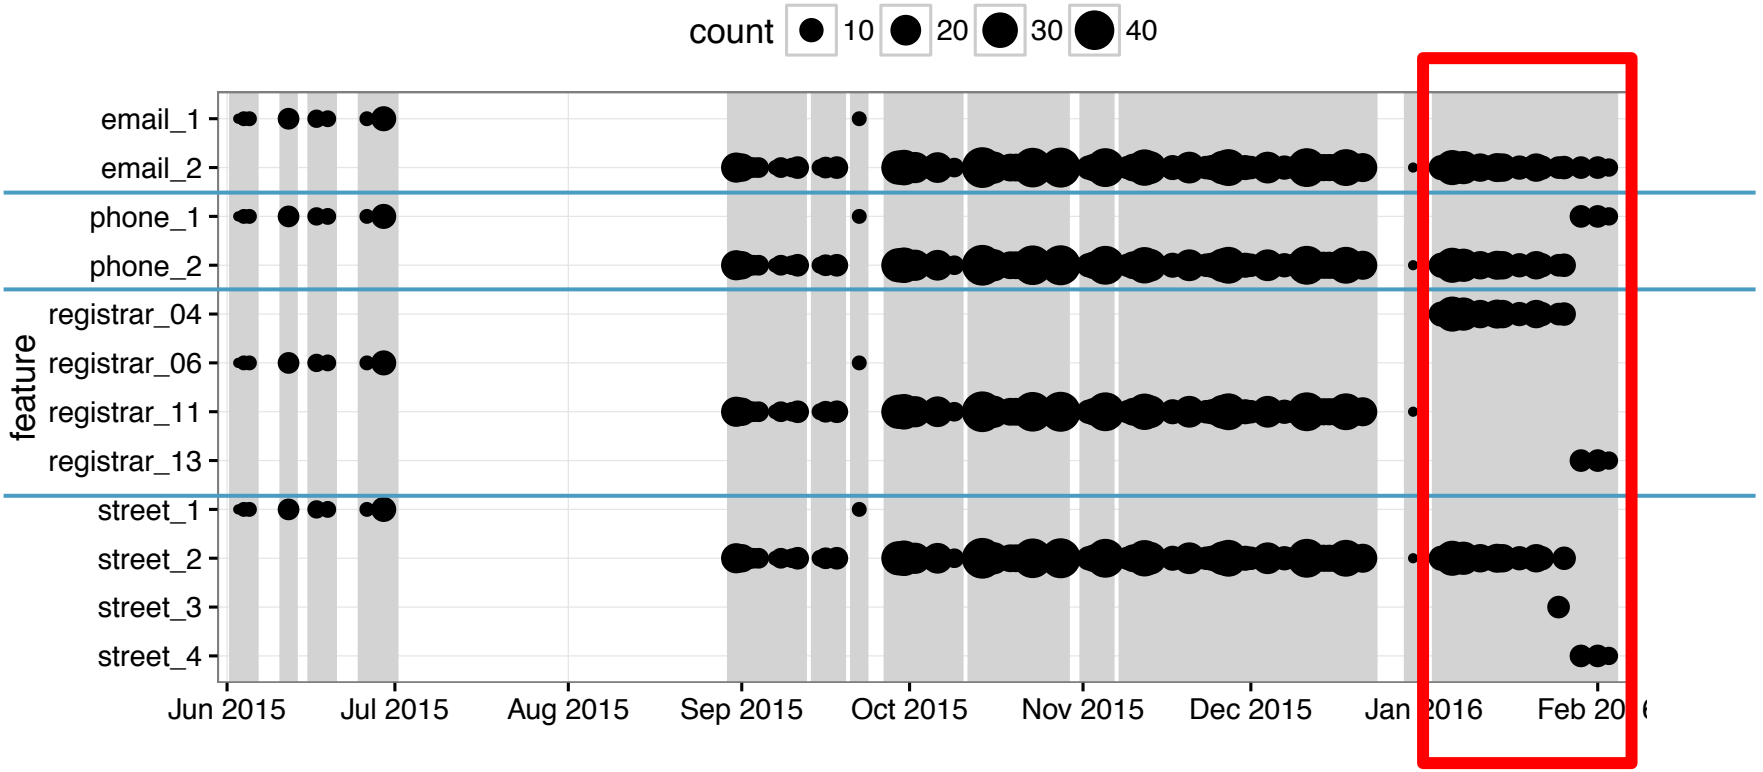

DistriNet

Activity of identified campaigns

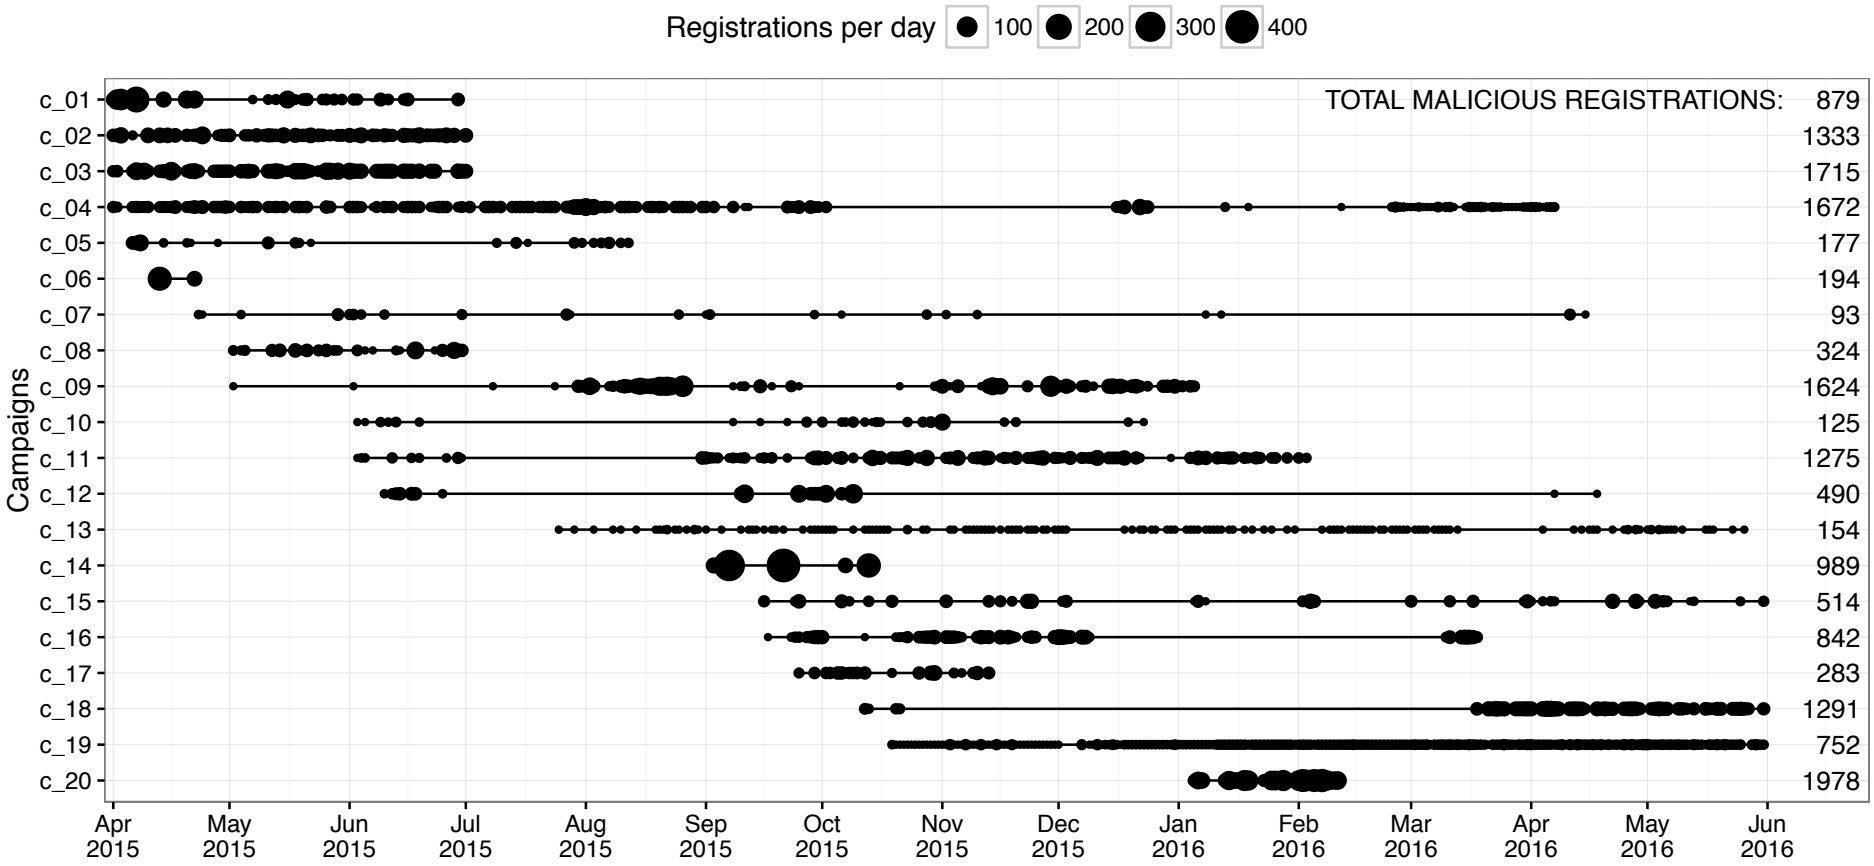

DistriNet

Campaign selection criteria

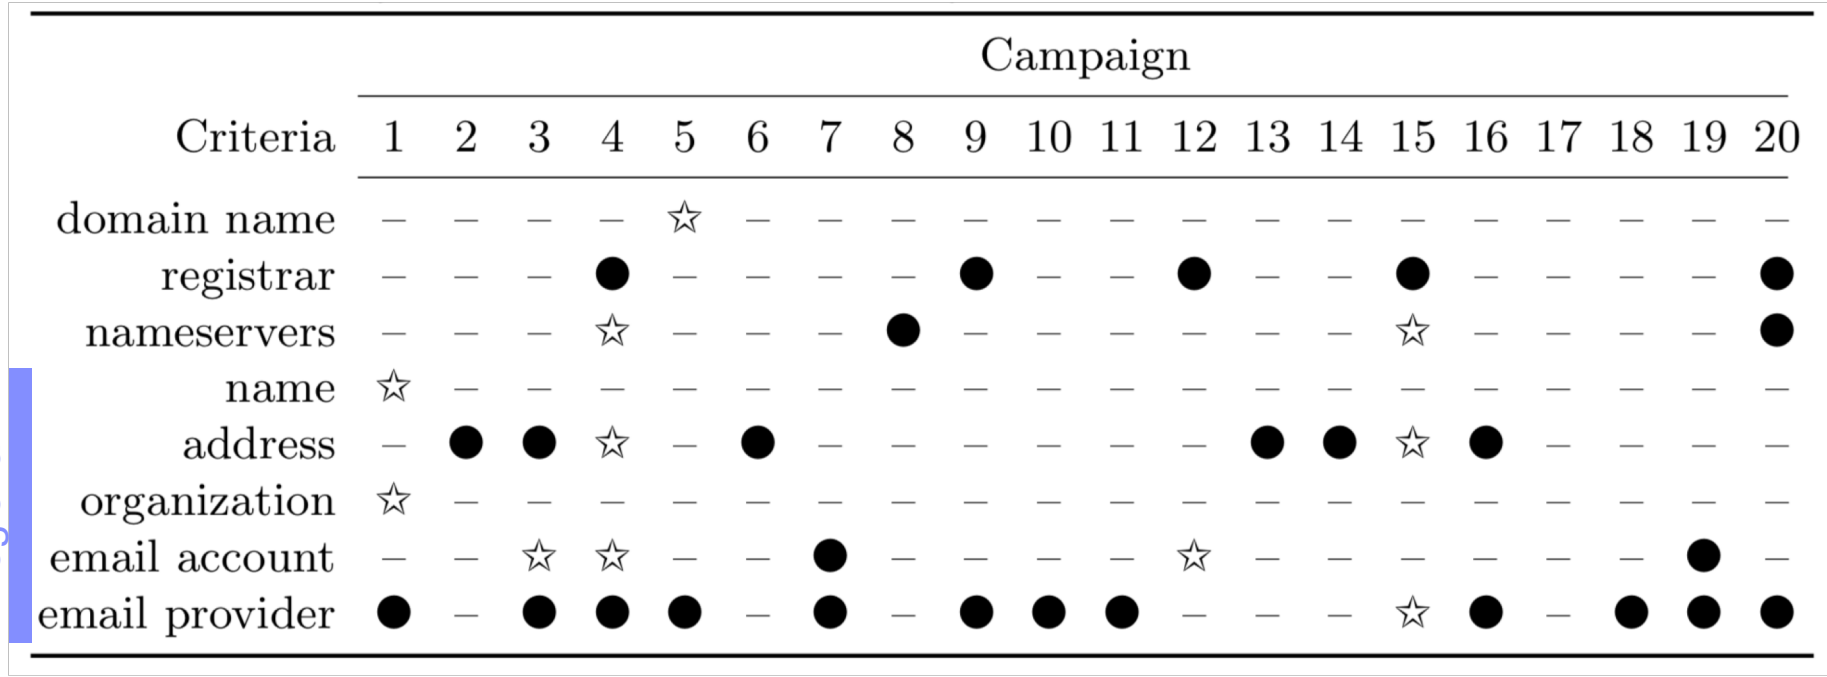

	Criteria	Campaign 1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Registrant	domain name	–	–	–	–	☆	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	registrar	–	–	–	●	–	–	–	–	●	–	–	●	–	–	●	–	–	–	–	●
	nameservers	–	–	–	☆	–	–	–	●	–	–	–	–	–	–	☆	–	–	–	–	●
	name	☆	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	address	–	●	●	☆	–	●	–	–	–	–	–	–	●	●	☆	●	–	–	–	–
	organization	☆	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	email account	–	–	☆	☆	–	–	●	–	–	–	–	☆	–	–	–	–	–	–	●	–
	email provider	●	–	●	●	●	–	●	–	●	●	●	–	–	–	☆	●	–	●	●	●

● represents a string match, and ☆ a regular expression pattern

Insights in malicious domain registrations

DistriNet

Insight 1: Hit-and-run strategies

- Small window of opportunity:
  - Domain rendered useless once blacklisted
  - 73% is blacklisted 5 days after registration

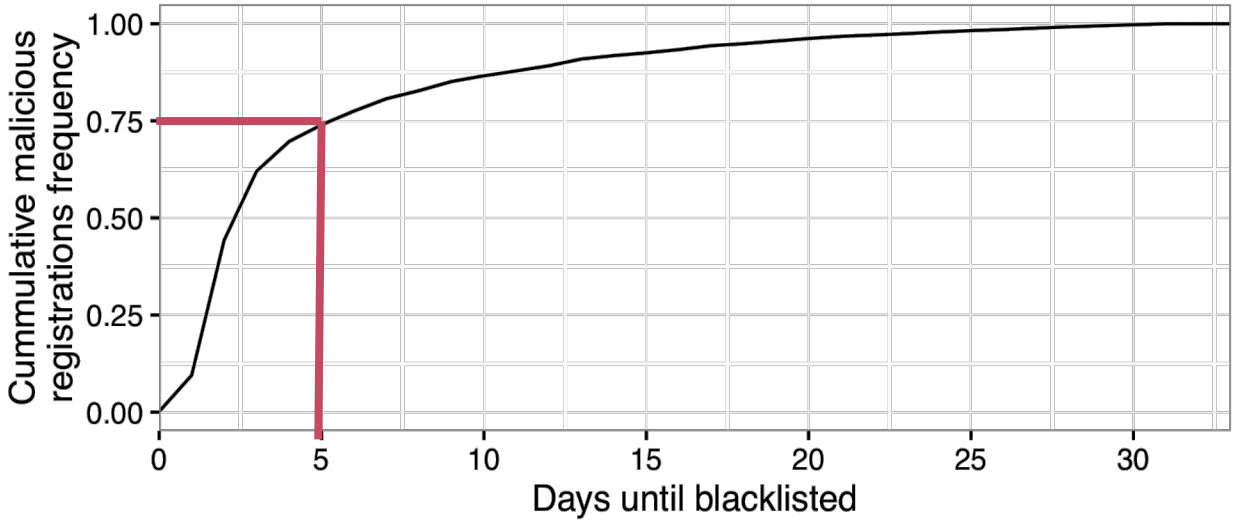

DistriNet

Insight 2: Campaigns are primarily linked to spam

	Abuse types					Blacklist sources		
Campaign	Spam	Botnet	Malware	Phishing	Unwanted	pamhaus	SURBL	Google SB
c_01	100.00%						100.00%	
c_02	100.00%					100.00%	27.53%	
c_03	100.00%					99.48%	86.82%	
c_04	99.88%		0.12%	1.38%		99.64%	76.26%	
c_05	83.05%					12.99%	77.97%	
c_06	100.00%					87.63%	12.37%	
c_07	91.40%					91.40%	1.08%	
c_08	100.00%					100.00%	3.70%	
c_09	99.63%		0.12%	1.97%		99.26%	28.45%	
c_10	99.20%			1.60%		78.40%	90.40%	
c_11	85.18%		0.08%			16.00%	77.02%	
c_12	99.59%			0.20%		99.39%	74.29%	
c_13	96.75%					81.82%	19.48%	
c_14	100.00%					84.43%	86.05%	
c_15	97.28%					73.35%	33.46%	
c_16	100.00%			0.12%		100.00%	43.71%	
c_17	100.00%					100.00%	8.83%	
c_18	99.85%			0.15%		99.77%	28.04%	
c_19	72.07%	27.93%				100.00%		
c_20	99.29%		0.96%			99.14%	7.58%	
All malicious	93.68%	1.27%	0.85%	3.22%	0.57%	81.07%	50.04%	1.81%

DistriNet

Insight 3: Variety in intensity and duration

DistriNet

Insight 4: Some campaigns are fully automated

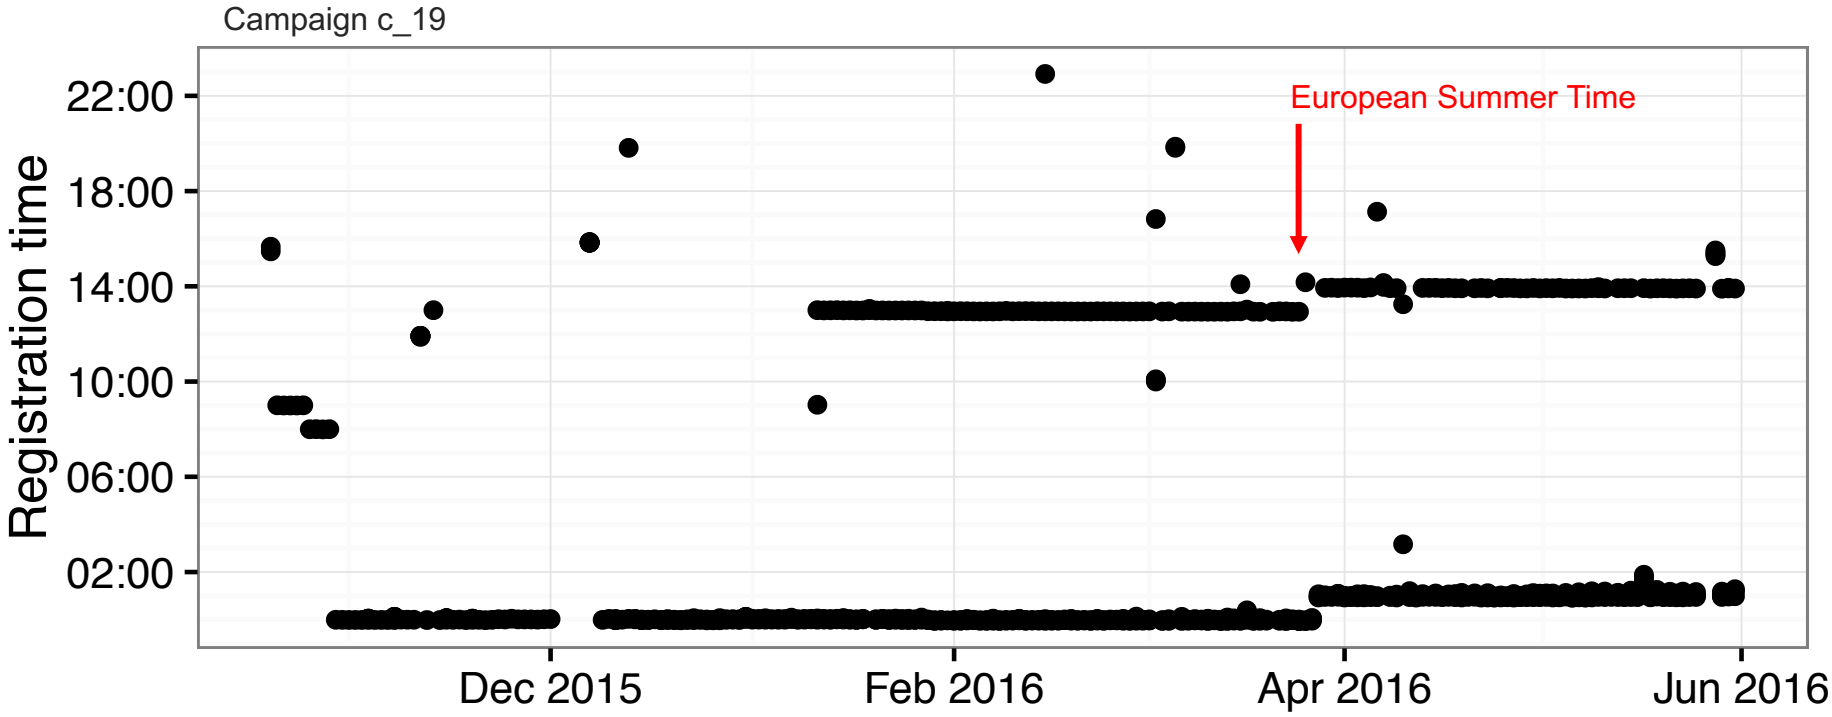

DistriNet

Insight 5: Some campaigns align with regular business activity patterns (1)

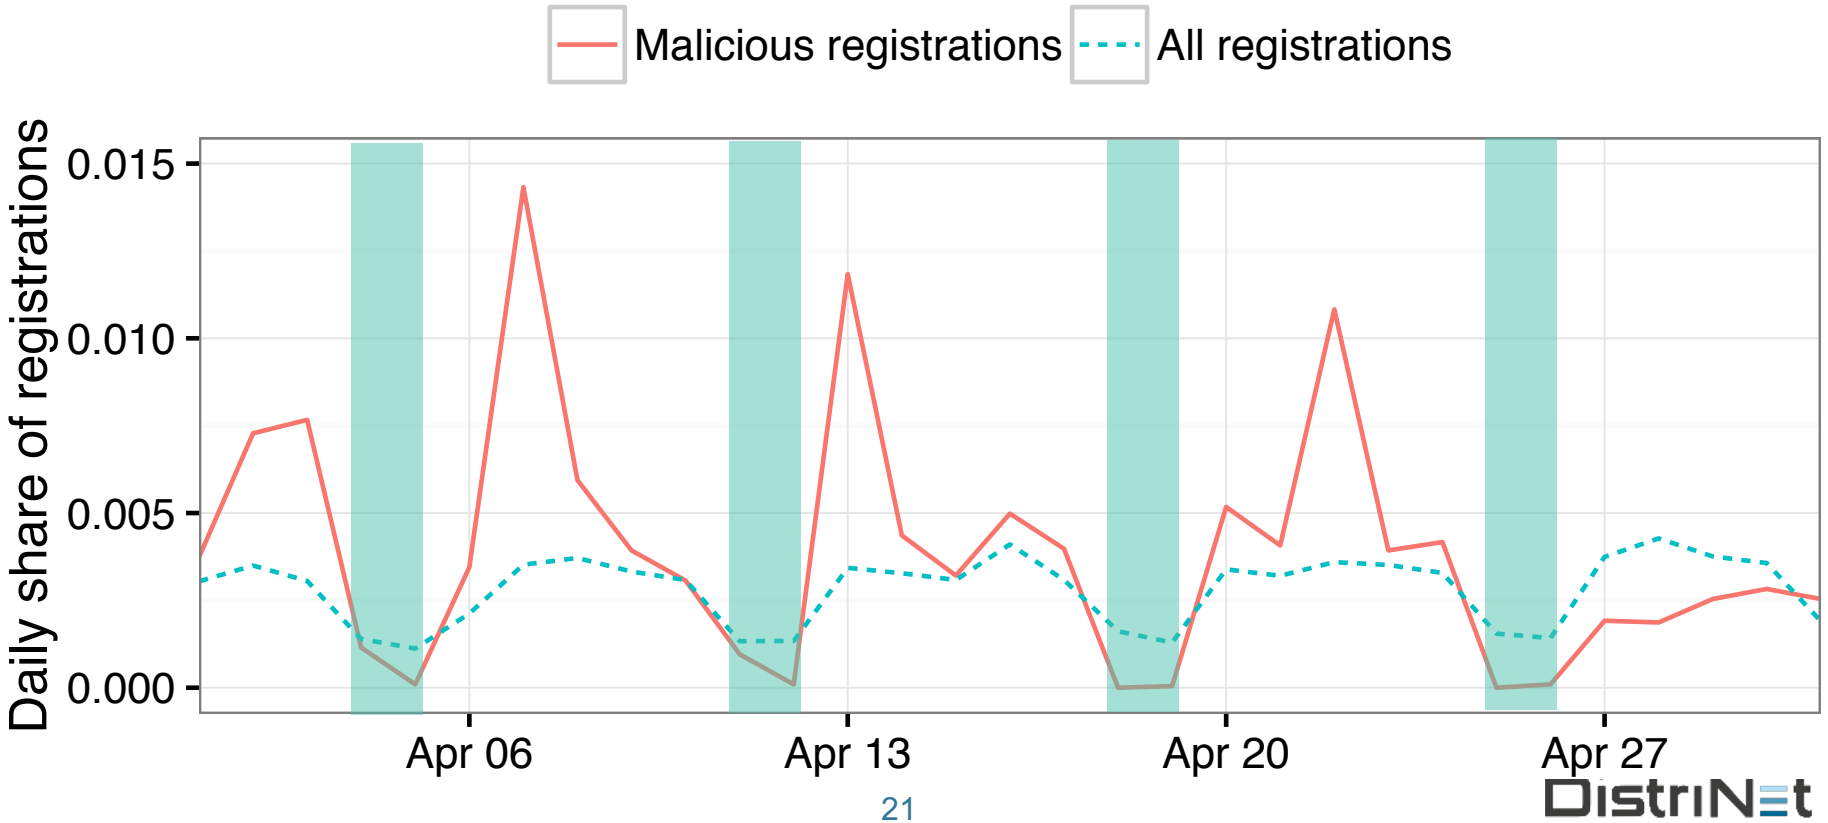

DistriNet

Insight 5: Some campaigns align with regular business activity patterns (2)

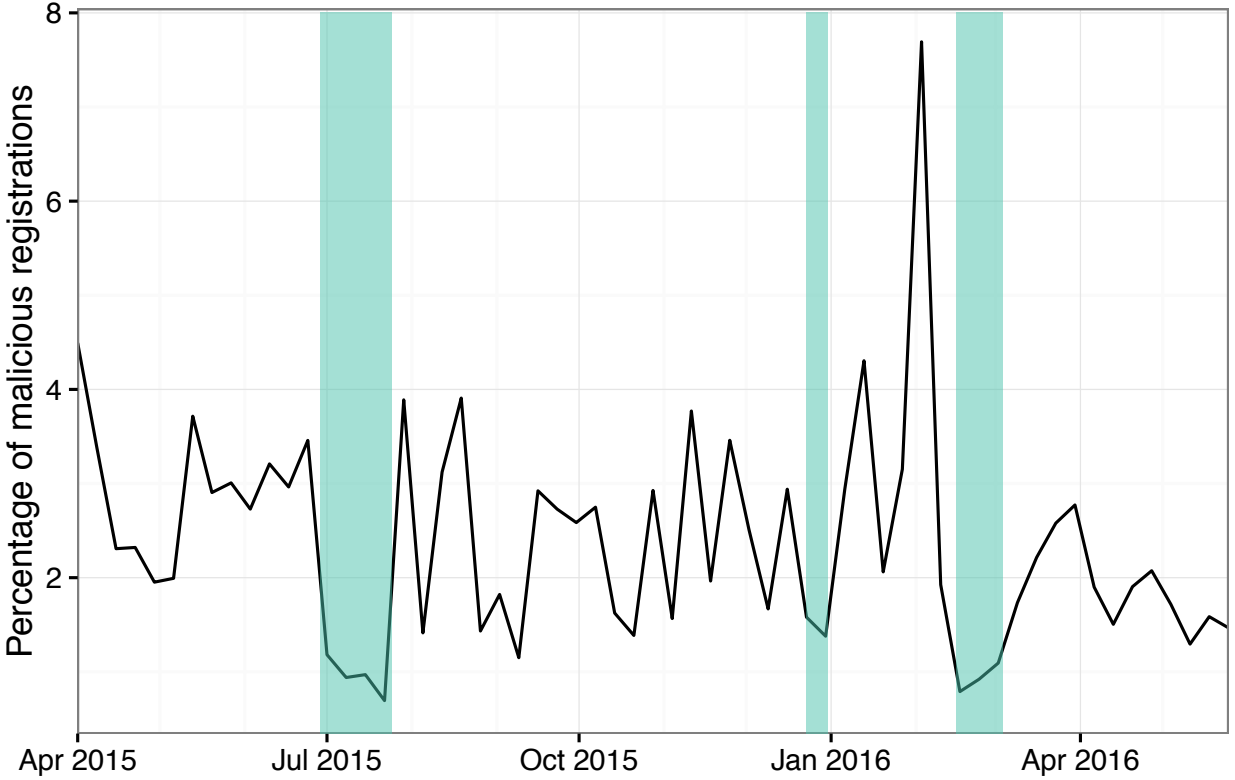

DistriNet

Insight 5: Some campaigns align with regular business activity patterns (3)

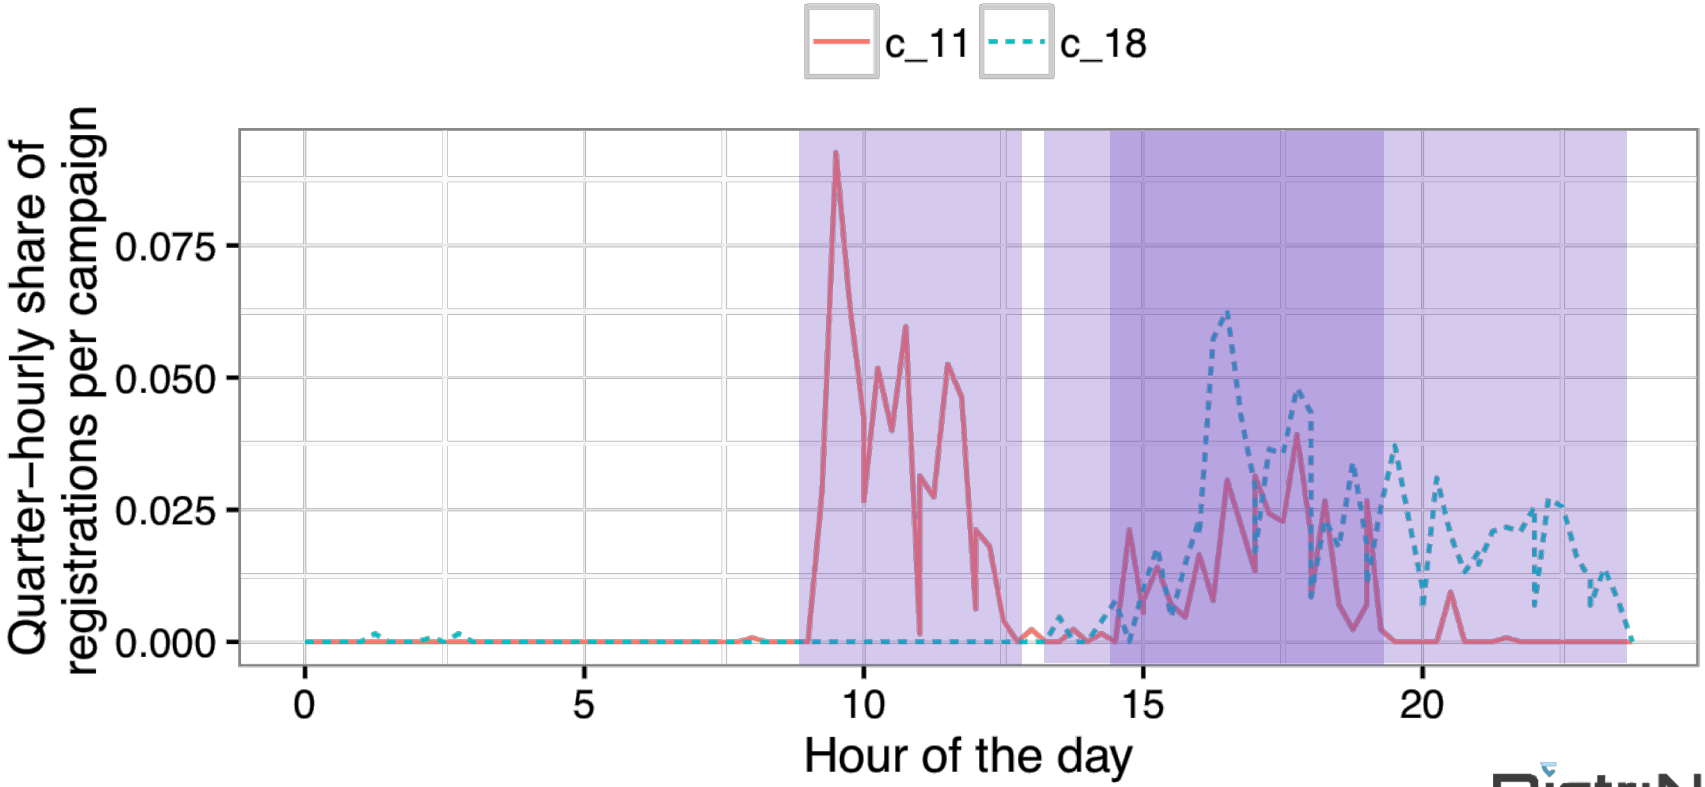

DistriNet

Insight 6: Top facilitators for malicious registrations

	Nb of malicious	Contribution Malicious	Contribution Benign	Contribution Toxicity
1. registrar_5	10,353	49.61%	2.27%	36.25%
2. registrar_3	3,004	14.39%	2.64%	12.41%
3. registrar_7	2,327	11.15%	0.46%	38.67%
1. gmail.com	4,221	20.23%	24.79%	2.08%
2. yahoo.com	3,348	16.04%	1.49%	21.85%
3. aol.com	2,134	10.23%	0.31%	46.28%

Insight 7: Adaptive campaign strategies

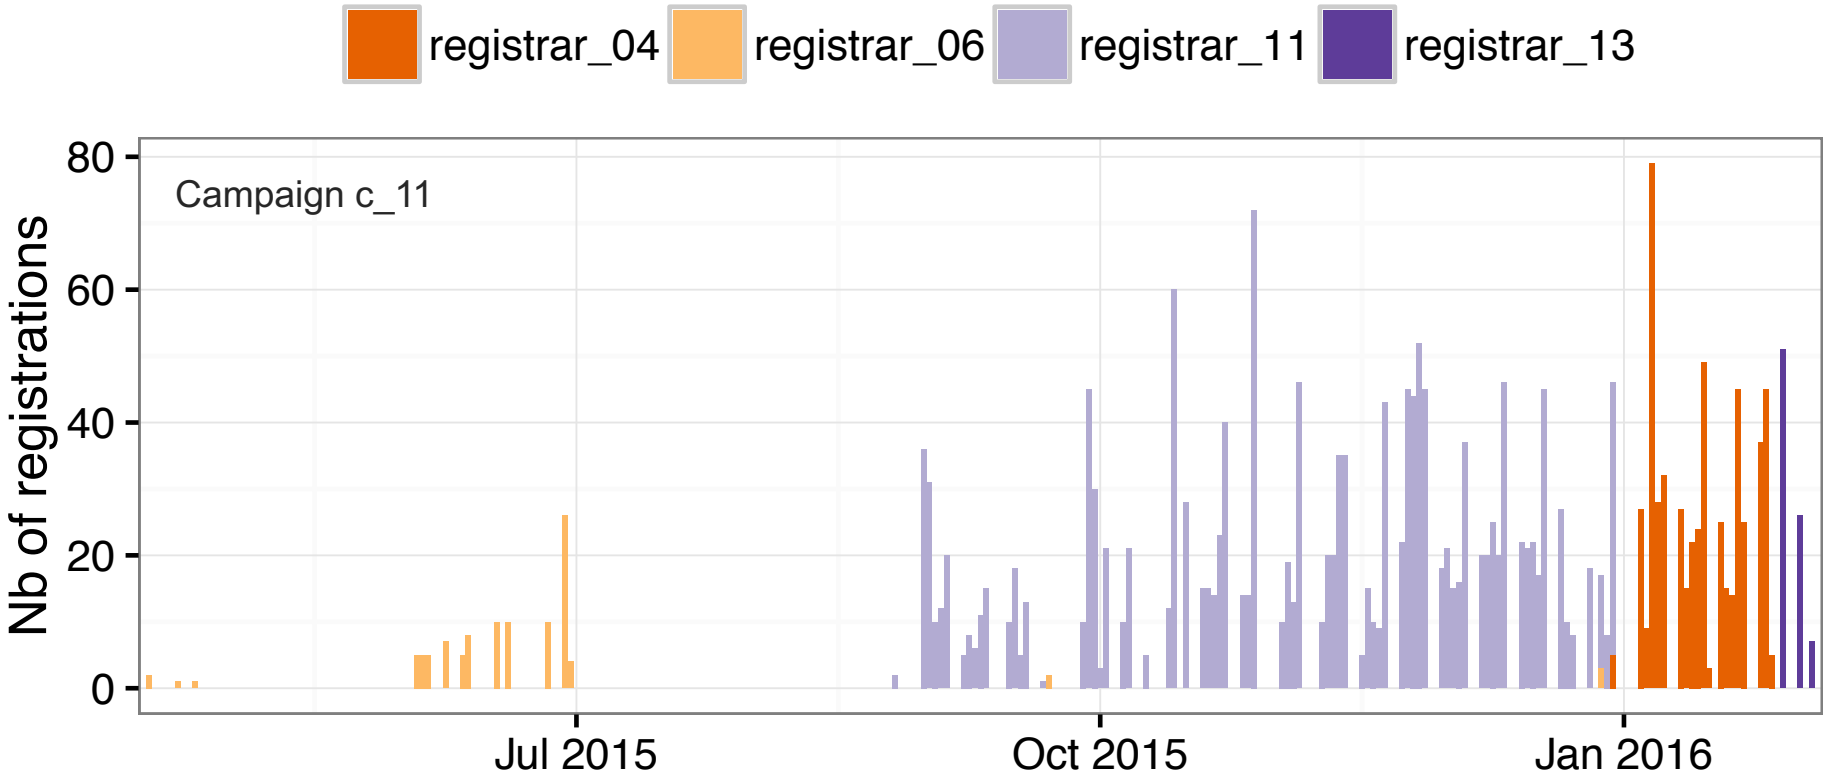

DistriNet

Insight 8: Campaigns vs blacklists

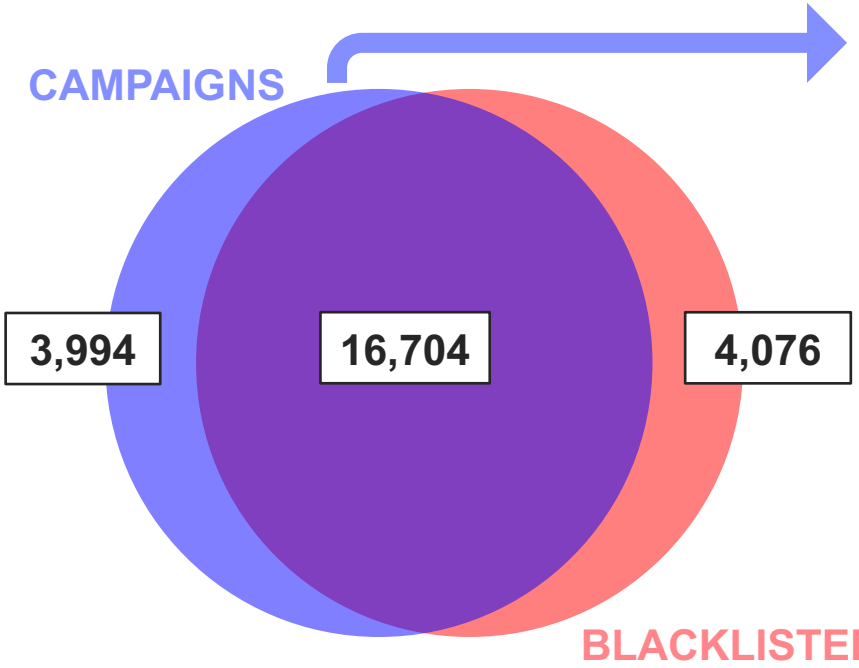

- Manual analysis of non-blacklisted domains
- Result: < 1% false positives
- About 20% extra on top of existing blacklists

DistriNet

Automated validation of campaign selection

Manual campaign identification process

- Start from maliciously flagged registrations
- Group registrations based on similarities between registration details
- Start heuristics:
  - Peaks in malicious registrations
  - Strong discrepancies between malicious and benign domains

a) Days with high number of malicious registrations

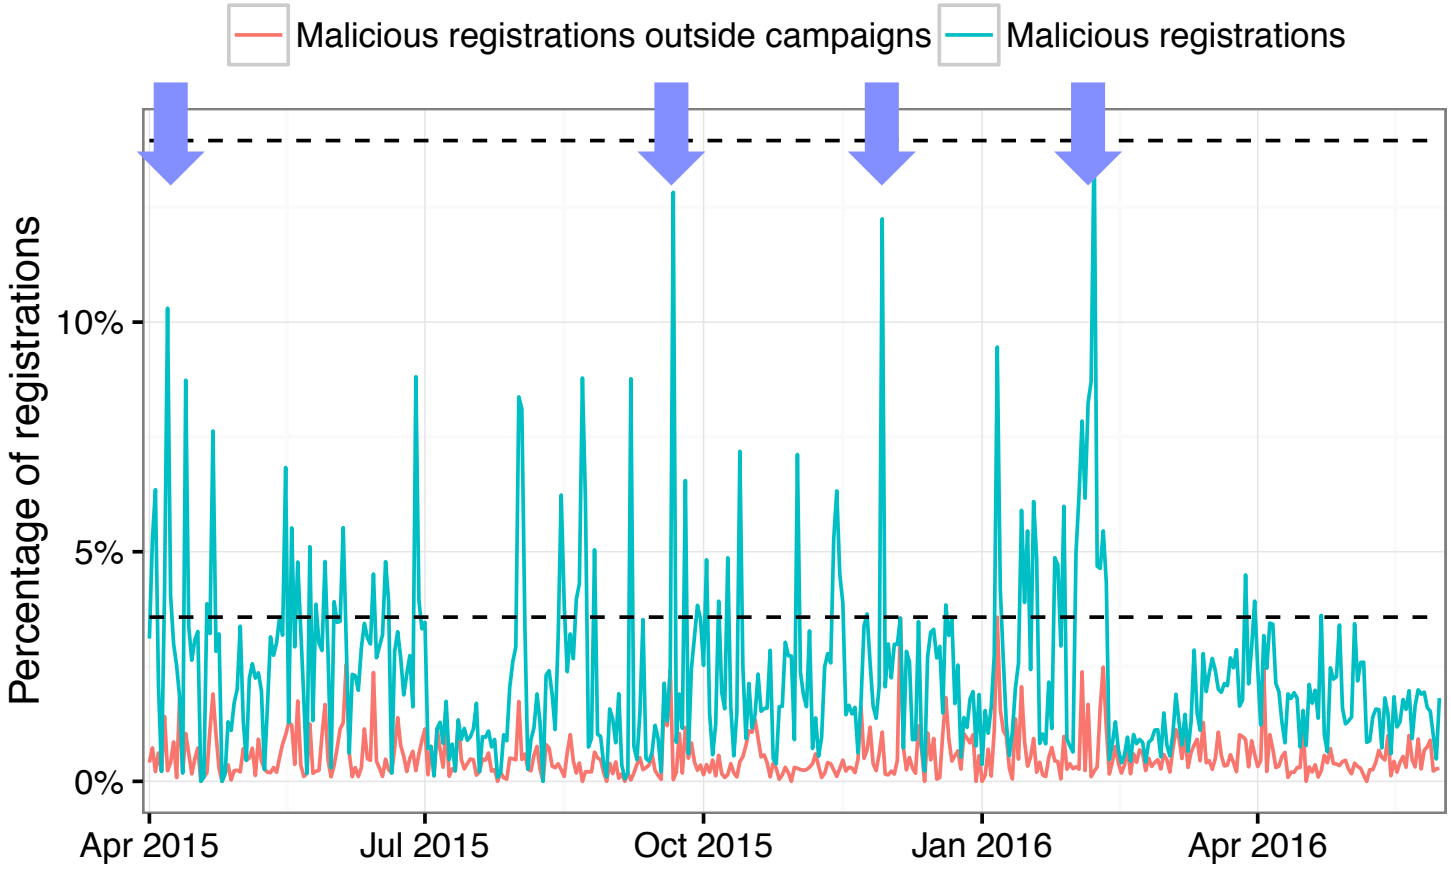

DistriNet

b) Frequent combinations of registration details

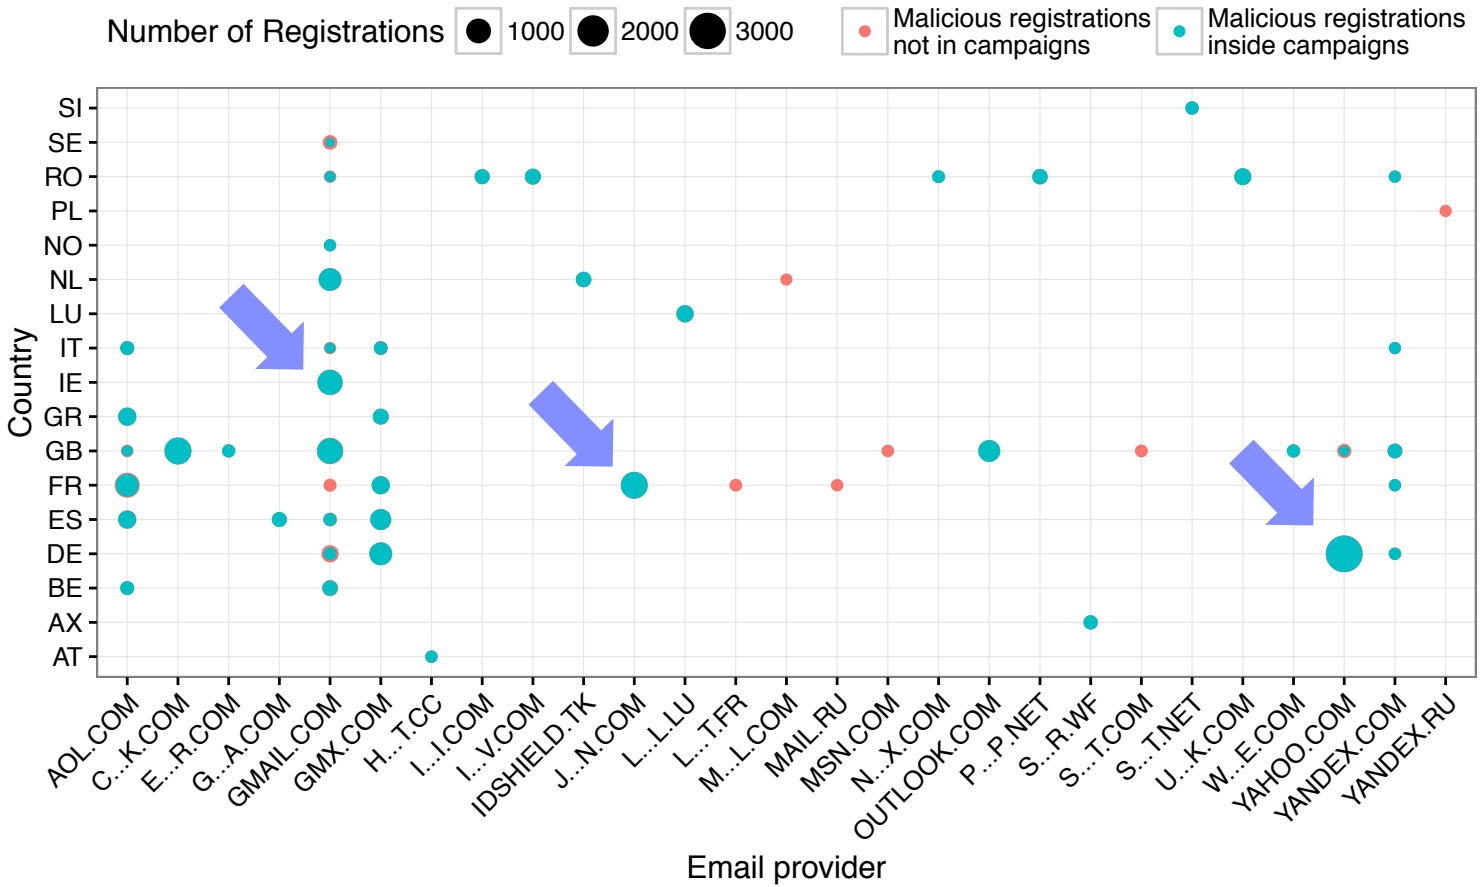

DistriNet

Campaign validation: clustering algorithm

- Machine learning technique to group registrations based on similarities between registration details
  - Agglomerative clustering of blacklisted registrations
  - Iteratively merge two closest clusters
- 30 largest (of 432) clusters represent 92% of campaign registrations

Finding 1: Some campaigns are linked to each other

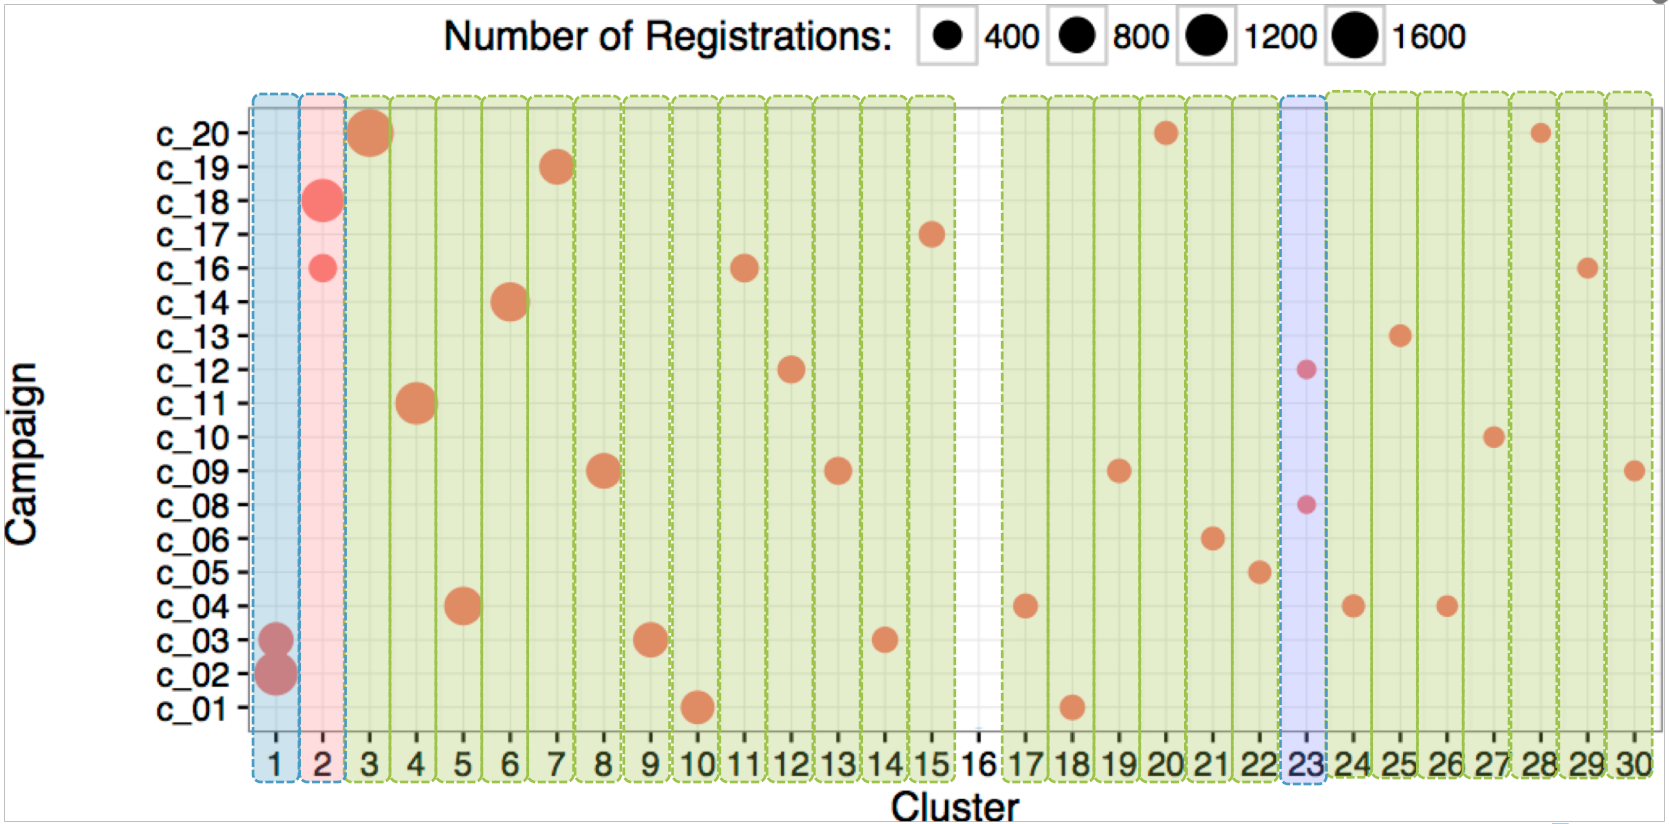

DistriNet

Finding 1: Some campaigns are linked to each other

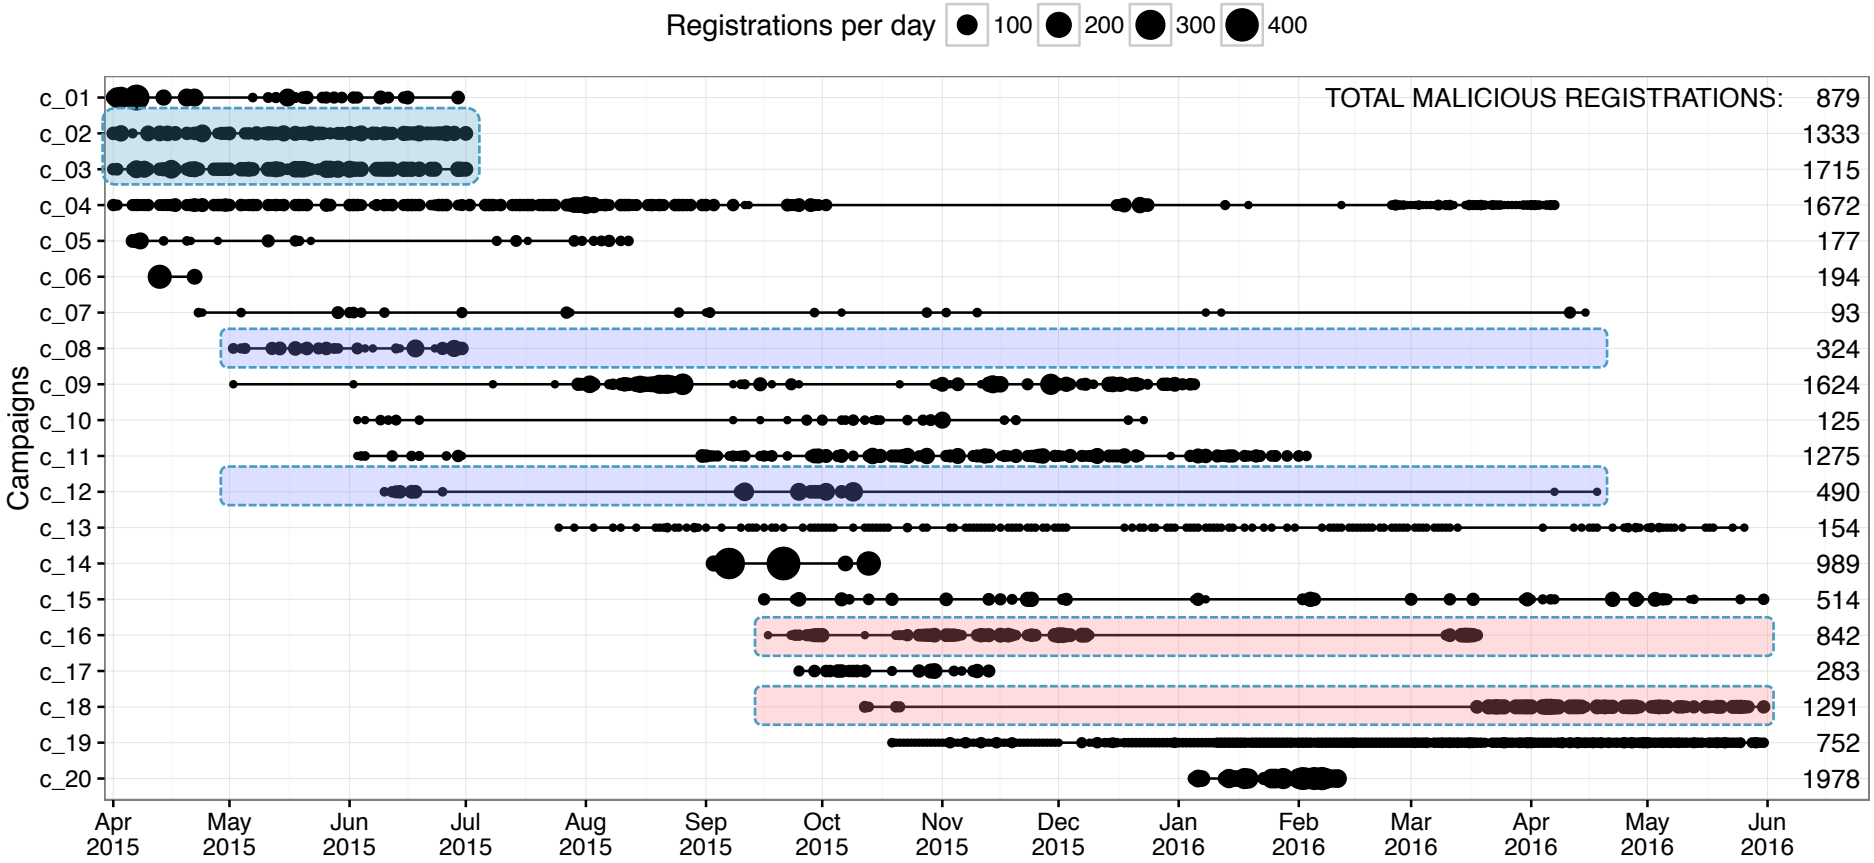

DistriNet

Finding 2: Some registrations were missed during campaign analysis

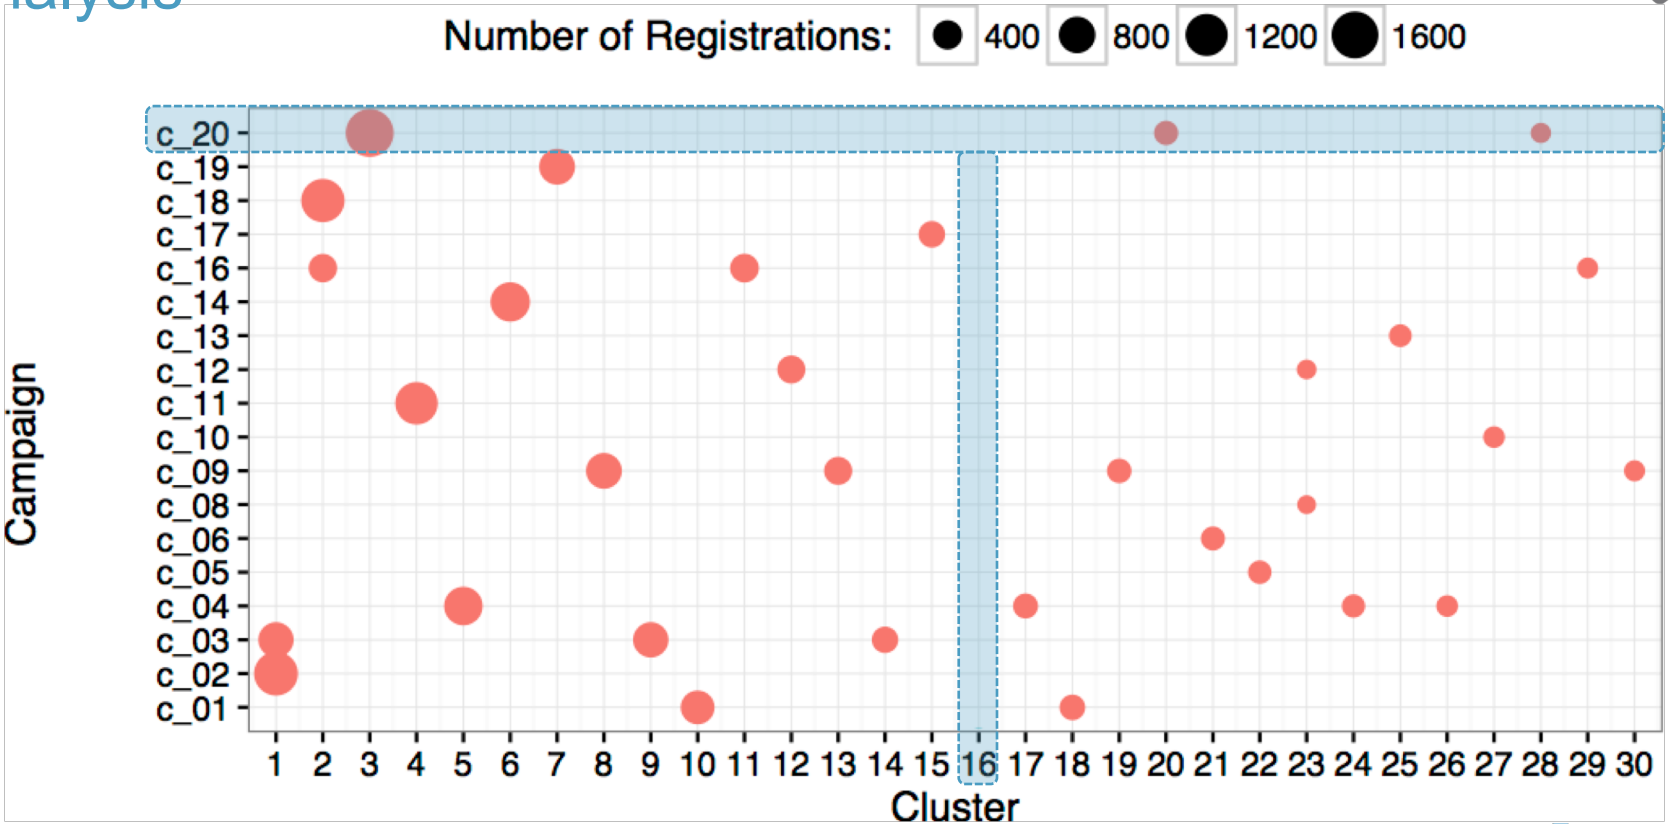

DistriNet

Finding 2: Some registrations were missed during campaign analysis

› Timing patterns further confirms that cluster 16 is part of campaign c_20

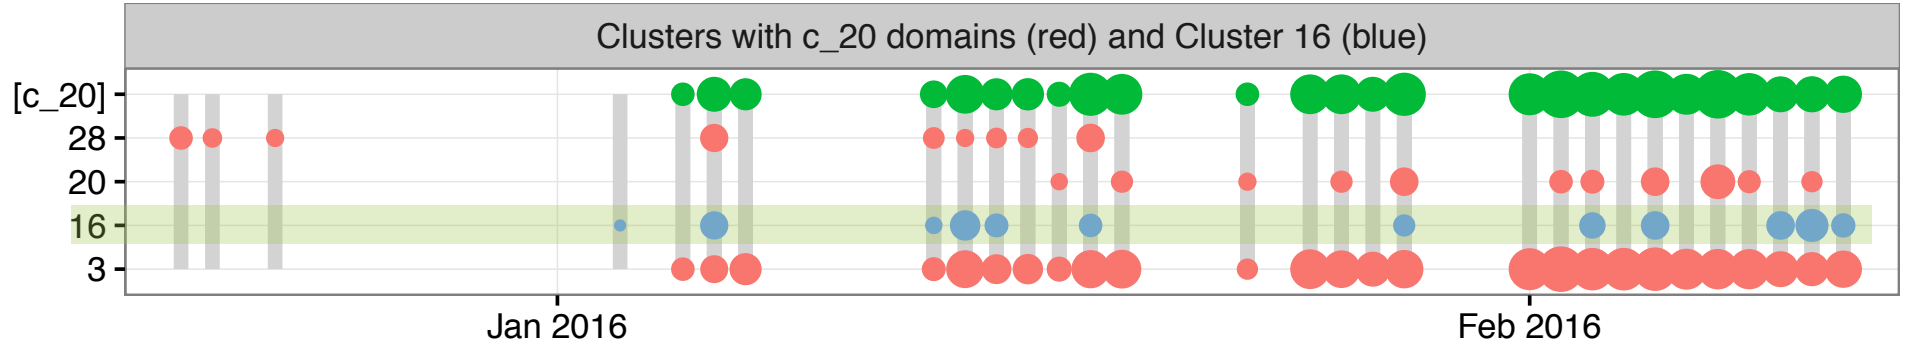

DistriNet

Finding 3: Advanced campaigns are not part of large clusters

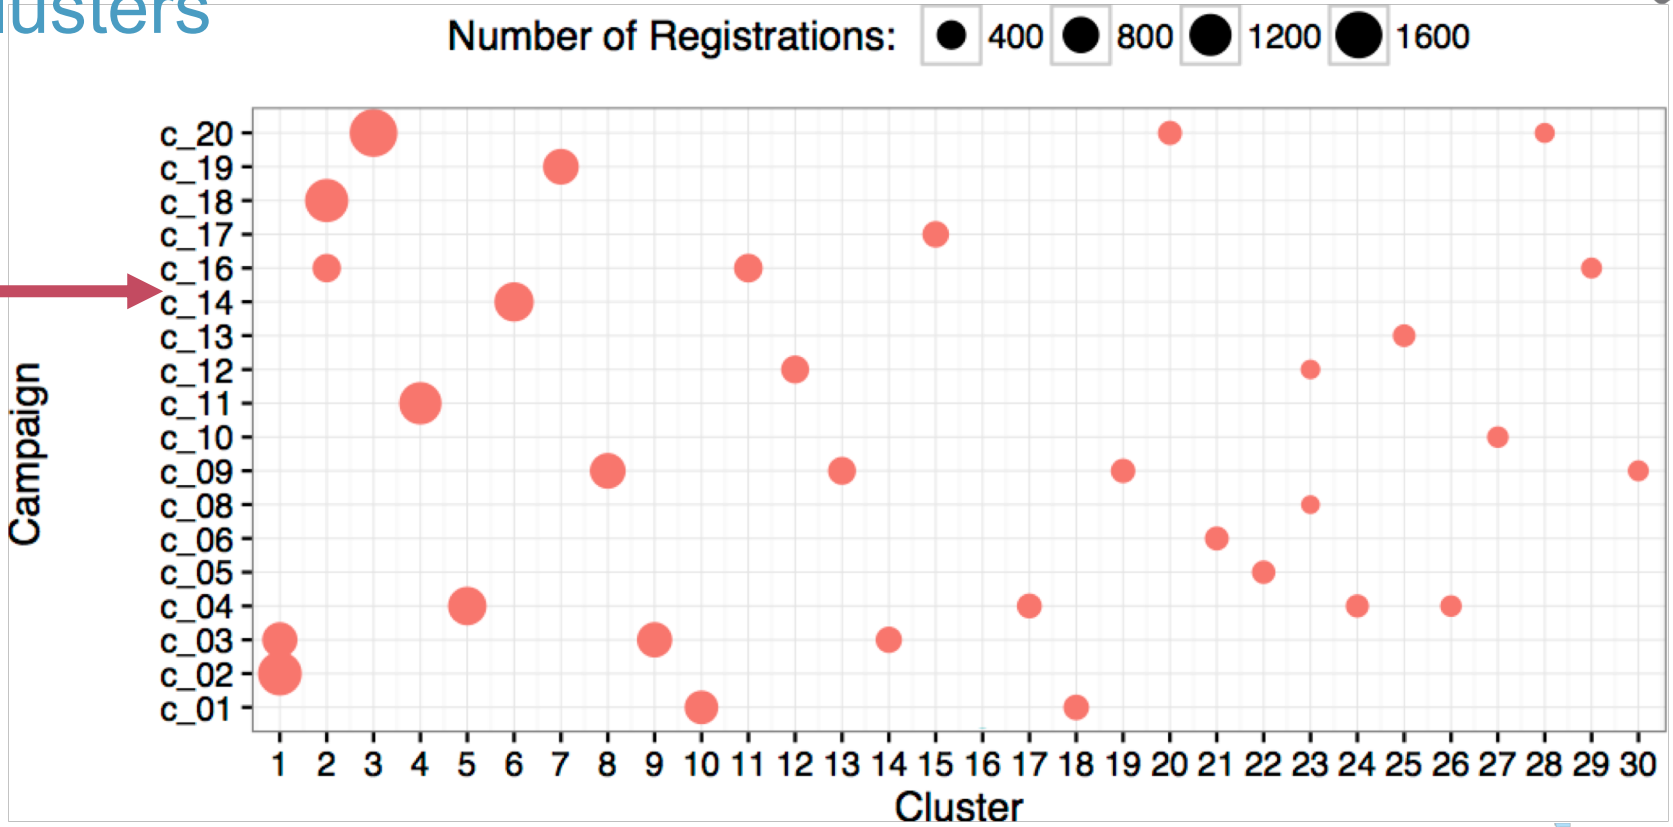

15

Example of an advanced campaign (c_15)

- Registrant details:
  - 98 fake registrants
  - Generated by Laravel Faker tool
- Domain names:
  - Consist out of 2-3 Dutch words
  - Dutch words are reused across registrants
- Batches of 8, 16, 24 or 32 registrations

- **8+ months active (Sep 16, 2015 – May 31, 2016)**
- **514 blacklisted registrations** ***(= 26.95%)***

Registration-time prediction of malicious intent

“Could newly registered domains with malicious intent be detected or prevented at registration time?”

Pro-active detection and prevention*

- Based on previous domain registrations, prediction models are trained
- For each new registration, the system predicts if the domain will be used for malicious activity
- Domains with malicious intent can be
  - Early detected
  - Delayed
  - Prevented from being registered

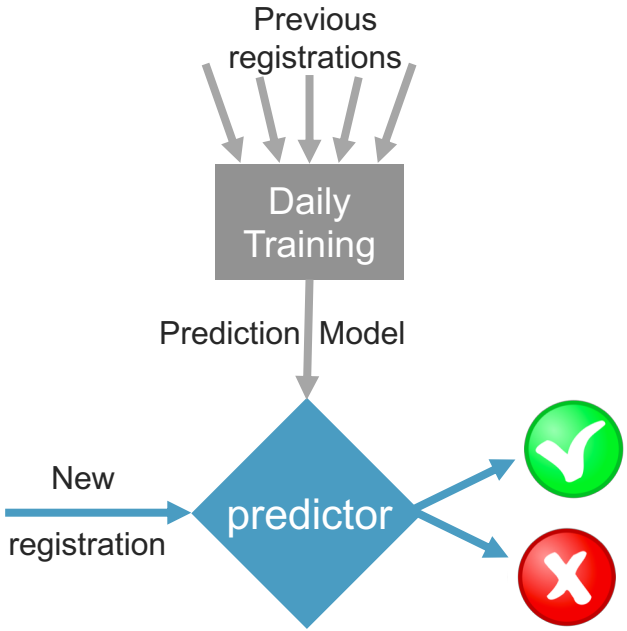

* Patent pending

DistriNet

Underlying assumptions/rationales for our predictors

- Similarity-based agglomerative clustering
  - Domains belonging to the same campaign have very similar registration details
- Reputation-based classification
  - Domains using registration facilitators with a bad reputation (e.g. email providers or registrars), are likely to be malicious as well

Dataset: validation and testing period

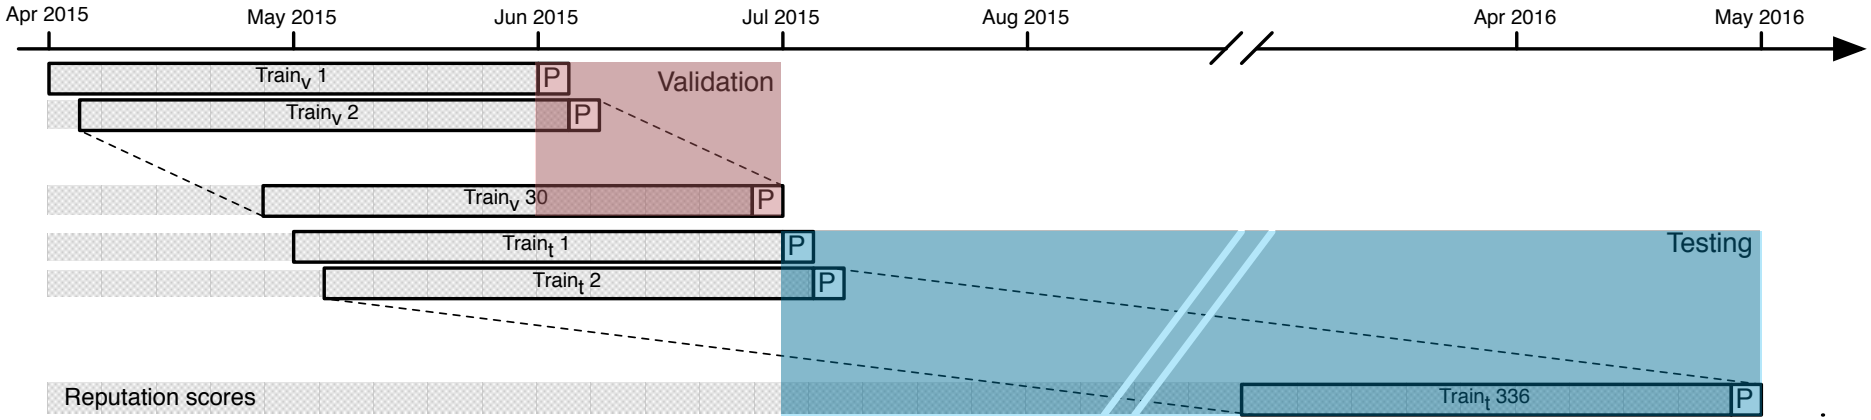

- Validation period: 1 month (June 2015)
- Testing period: 11 month (July 2015 – May 2016)

Evaluation criteria

- $precision = \frac{TP}{TP+FP}$
- $recall = TPR = \frac{TP}{TP+FN}$

Primary criteria for imbalanced problems

- $fall\ out = FPR = \frac{FP}{FP+TN}$

Reputation-based classifier

Reputation features of “facilitators”

- Facilitators:
  - Technical facilitators: registrar, name servers
  - Communication means: email provider and phone number
- Reputation score:
  - Percentage of malicious registrations of a facilitator in the past period
- Reputation window:
  - 15, 30, 60 days and all-time

Inherent dataset challenges

- Class imbalance: *only* 2.5% is malicious
  - Skew towards benign registrations

- Ground truth labeling [Vissers et al] :
  - Blacklist delays: up to 5 days till 73% is flagged
  - Incompleteness: 19% remains undetected

[Vissers et al] “Exploring the ecosystem of malicious domain registrations in the .eu TLD” RAID 2017

Addressing dataset challenges

- Class imbalance:
  - Standard machine learning tactic: class subsampling
- Blacklist delays
  - Prune *benign* registrations of last 5 days
- Incompleteness
  - Prune *benign outliers* for registrants which have a reputation score above a chosen threshold

Parameter tuning

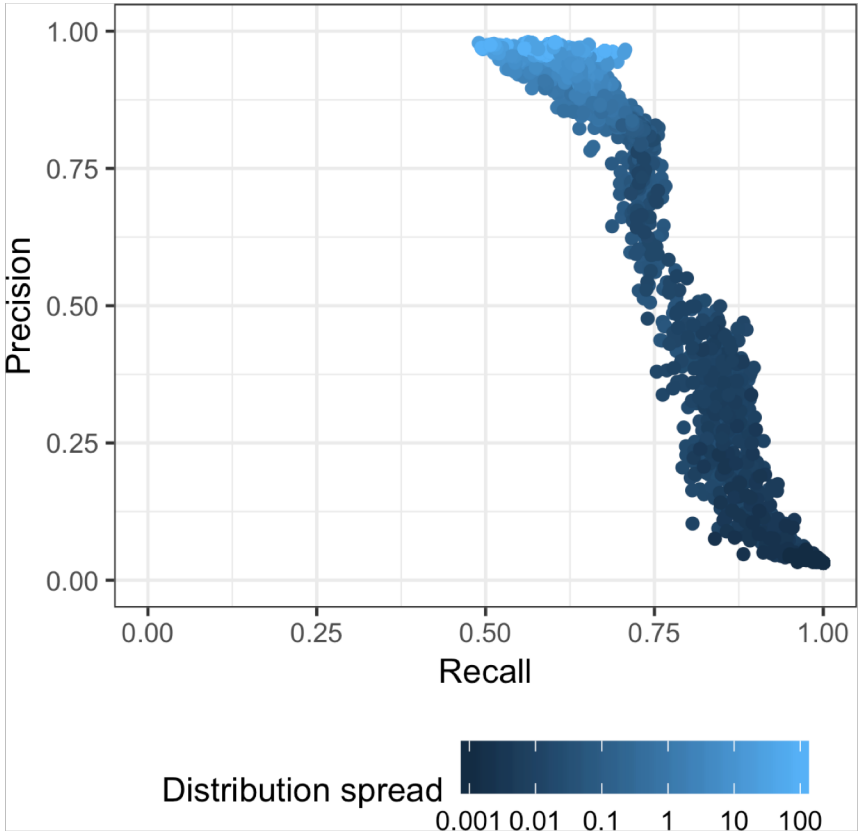

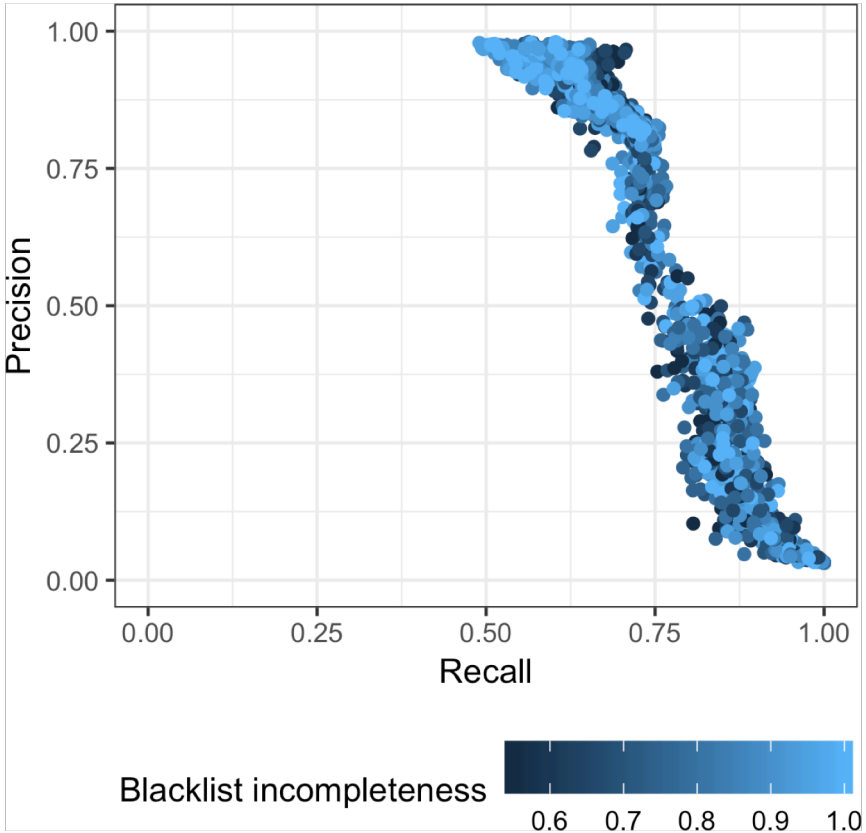

DistriNet

Similarity-based clustering

Agglomerative clustering, based on similarity of registration data

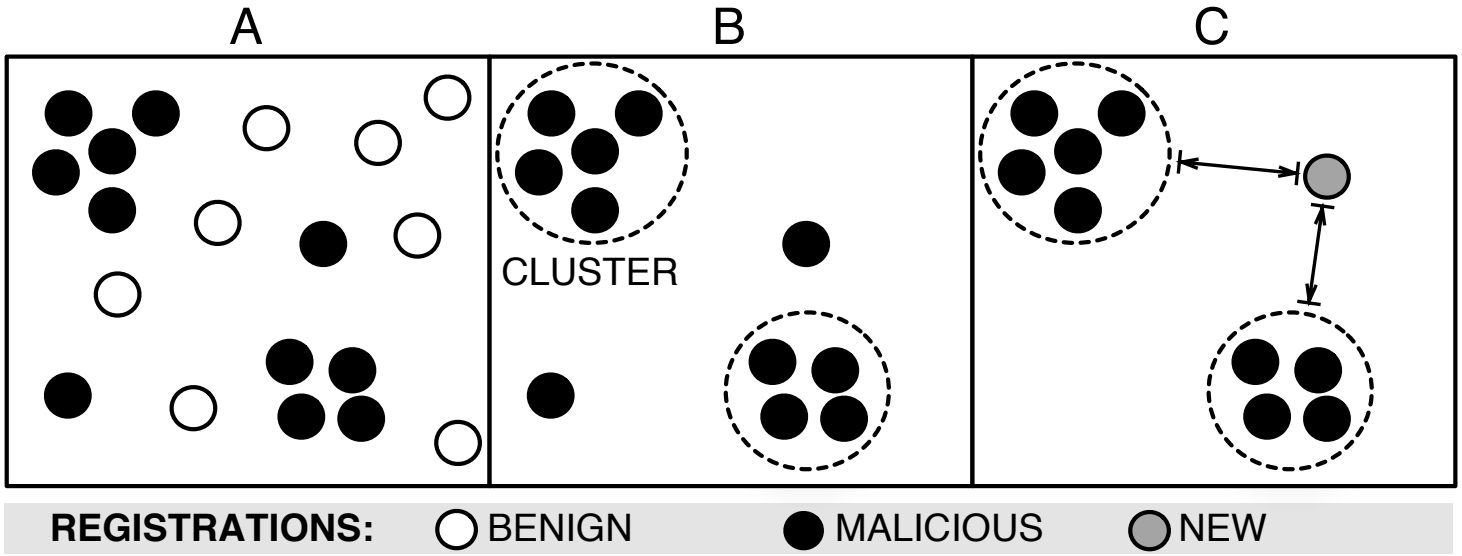

Inter-domain distance ?

- Pairwise distance:
  - Per feature: feature distance between registrations
  - Compose to a pairwise distance
- Maximum intra-cluster distance?

Pairwise feature distance

- String features:
  - normalized Levenshtein distance
- Numerical features:
  - Euclidian distance
- Categorical features:
  - equality (or fraction for multi-valued)

$$f_{ij} = dist(f_i, f_j)$$

Information-gain based feature weights

- Pair-wise distance is weighted feature distance

$$d_{ij} = d_{ij} = \sum_{f} w_f \times f_{ij}$$

- Information gain: the ability of a feature to partition the registrations in malicious and benign

$$w_f = IG\left(\begin{array}{c} \text{Malicious reg.} \mid \text{Benign reg.} \\ \left[\begin{array}{ccc|ccc} f_{11} & \cdots & f_{1n} & \cdots & f_{1m+n} \\ f_{21} & \cdots & f_{2n} & \cdots & f_{2m+n} \\ \vdots & \ddots & \vdots & \cdots & \vdots \\ f_{n1} & \cdots & f_{nn} & \cdots & f_{nm+n} \end{array}\right] \end{array} \right)$$

DistriNet

Feature weights over time

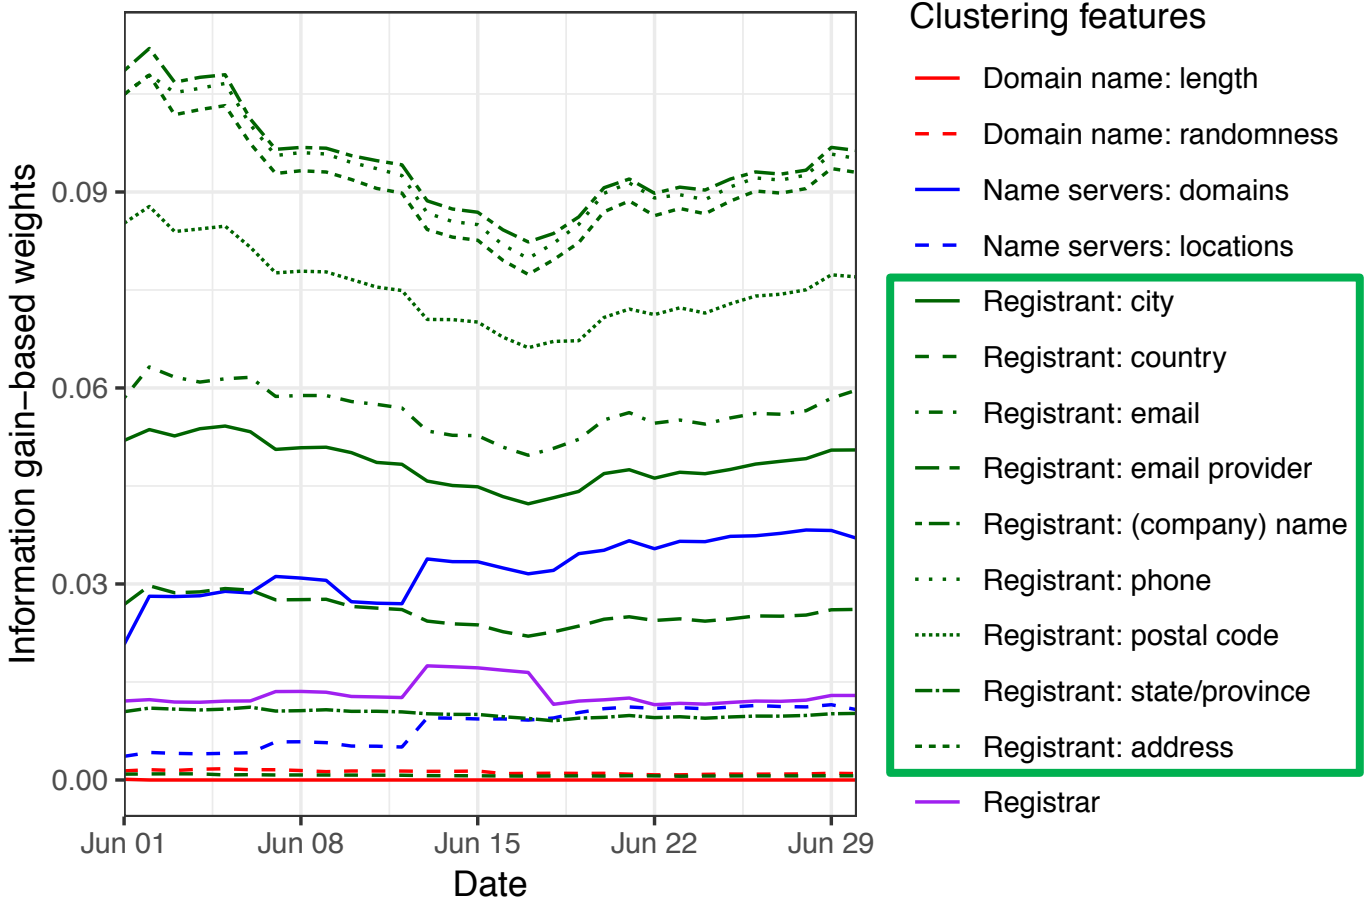

DistriNet

Maximum intra-cluster distance

› Closest distance of a registration to malicious domain

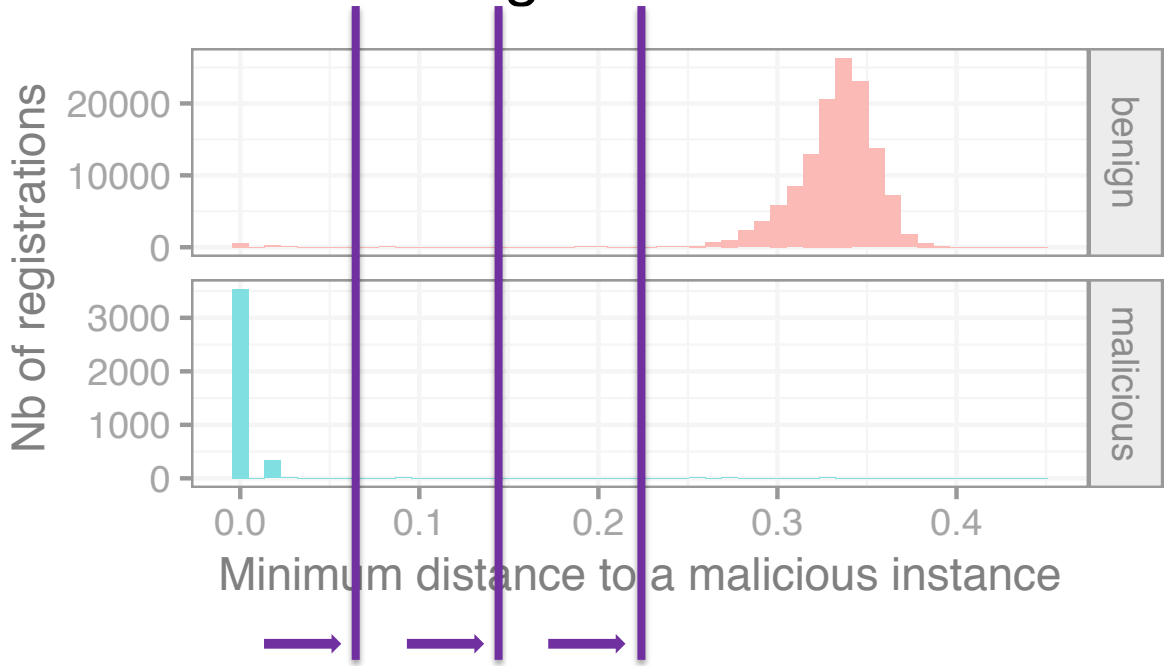

DistriNet

Parameter tuning

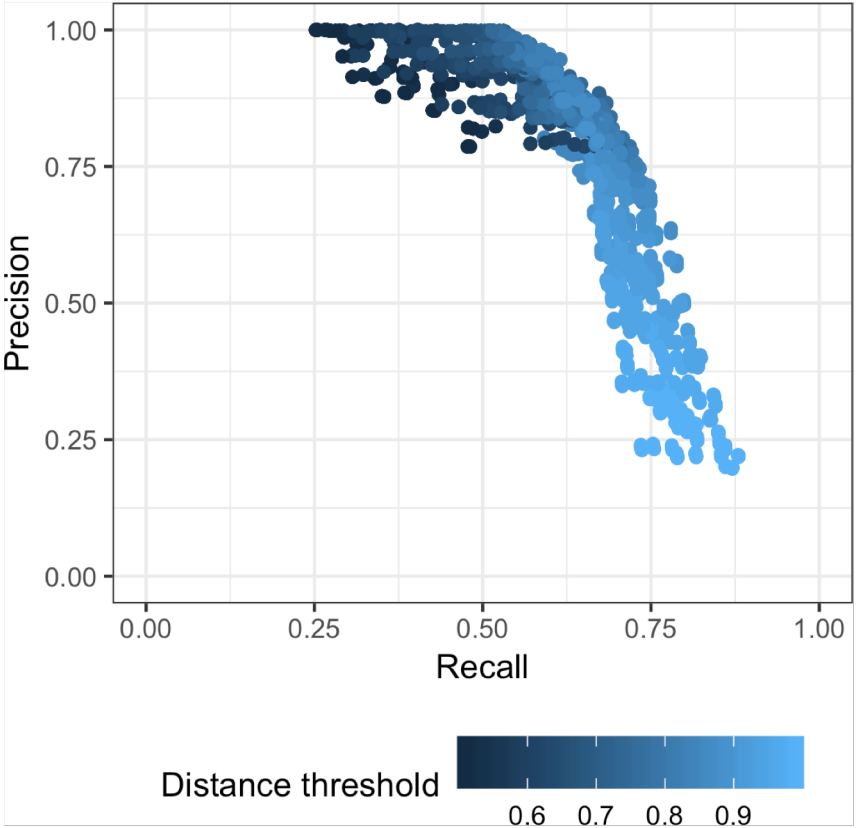

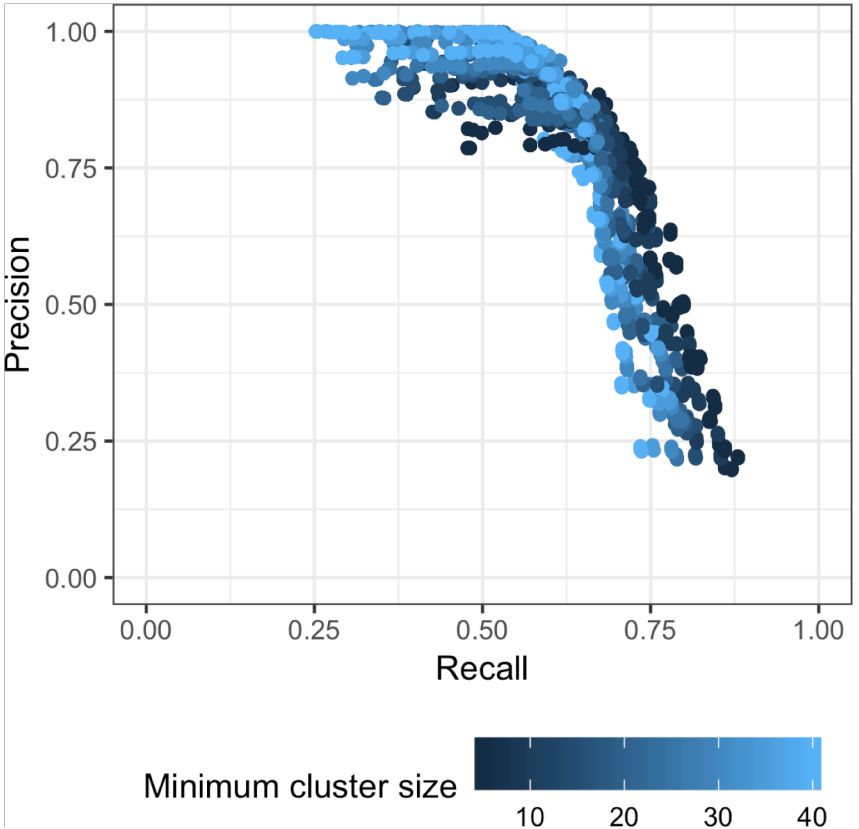

DistriNet

Final prediction model

Majority voting ensemble

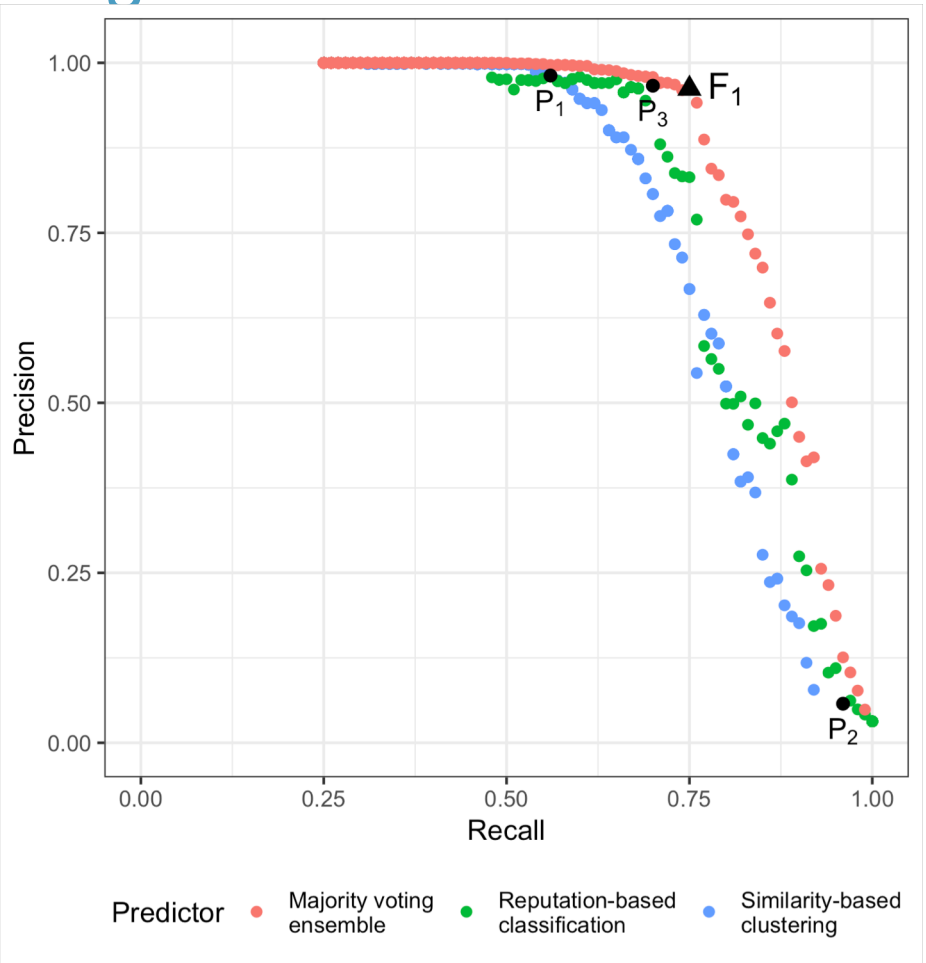

Evaluation on the testing set

- Ground truth-based evaluation
  - Recall: 66.23%
  - Precision: 84.57
  - False positive rate: 0.30%
- Campaign-based evaluation
  - 17 out of the 20 campaigns are well predicted,
  - Campaign recall: 76.68%
  - Campaign precision: 87.20%

Campaign-based evaluation

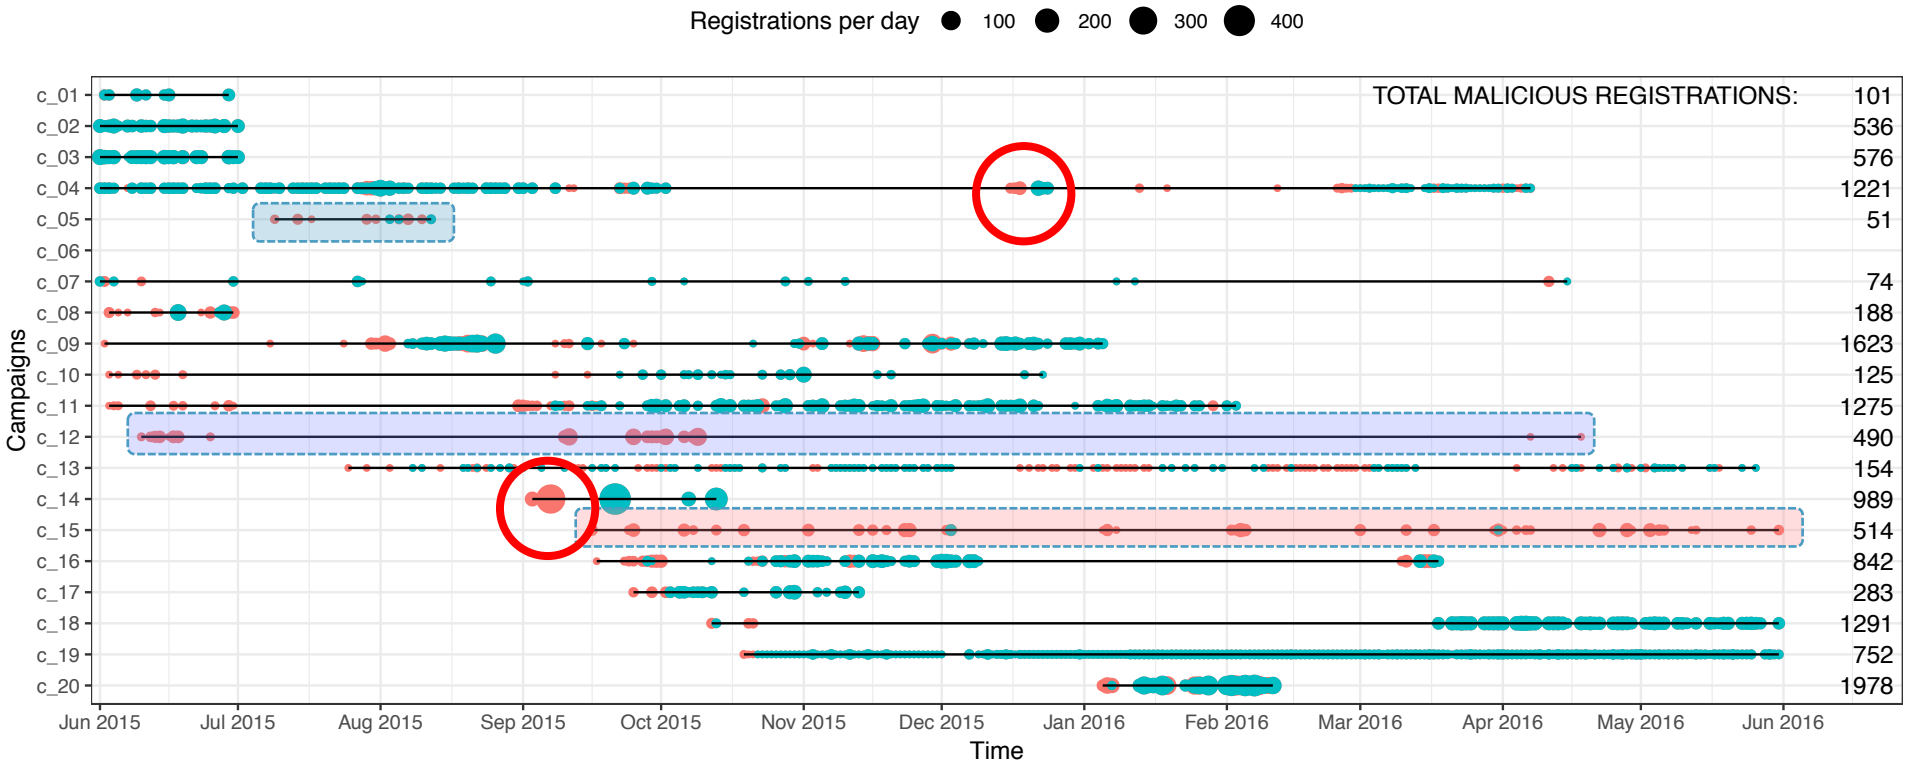

DistriNet

Detecting and preventing abuse in .eu

Operational requirements

- Adaptive and autonomous
- Stability and trust in the predictions

- Real-time predictions

- Feedback to the business/security/legal teams
- Focus on low number of false positives

Performance tactics

Speeding up similarity-based clustering

- Calculating the distance matrix is computational expensive:
  - About *1 billion* (!!!) Levenshtein distances per daily model

$$w_f = IG\left(\begin{array}{c} \text{Malicious reg.} \mid \text{Benign reg.} \\ \left[\begin{array}{ccc|ccc} f_{11} & \cdots & f_{1n} & \cdots & f_{1m+n} \\ f_{21} & \cdots & f_{2n} & \cdots & f_{2m+n} \\ \vdots & \ddots & \vdots & \cdots & \vdots \\ f_{n1} & \cdots & f_{nn} & \cdots & f_{nm+n} \end{array}\right] \end{array} \right)$$

Speeding up the training phase (2)

- Cached Levenshtein distances
  - Speed up of factor 20
- Static weights instead of IG-based
  - Speed up of factor 40
- Complete linkage clustering
  - Prediction time in linear instead of quadratic

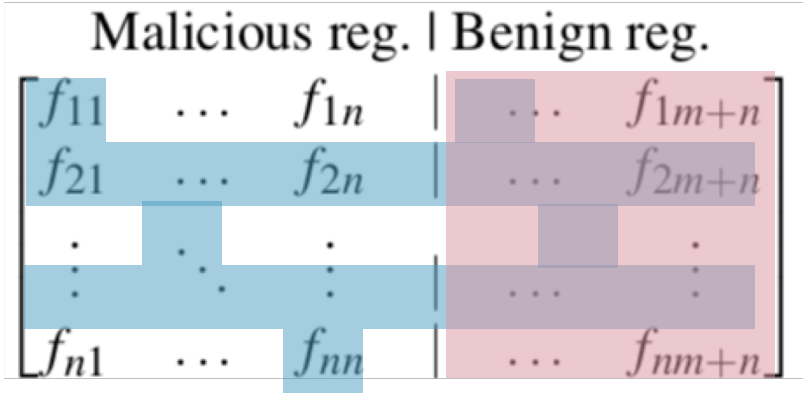

DistriNet

Feedback to the teams

Need for understandable predictions

- Legal teams need insights why a domain registration has been prevented
- Business and security teams need to be able to observe changing adversarial ecosystem

- Explicit choice for simple machine learning techniques, that provide sufficient insights in their models

Operational results

“1 picture …”

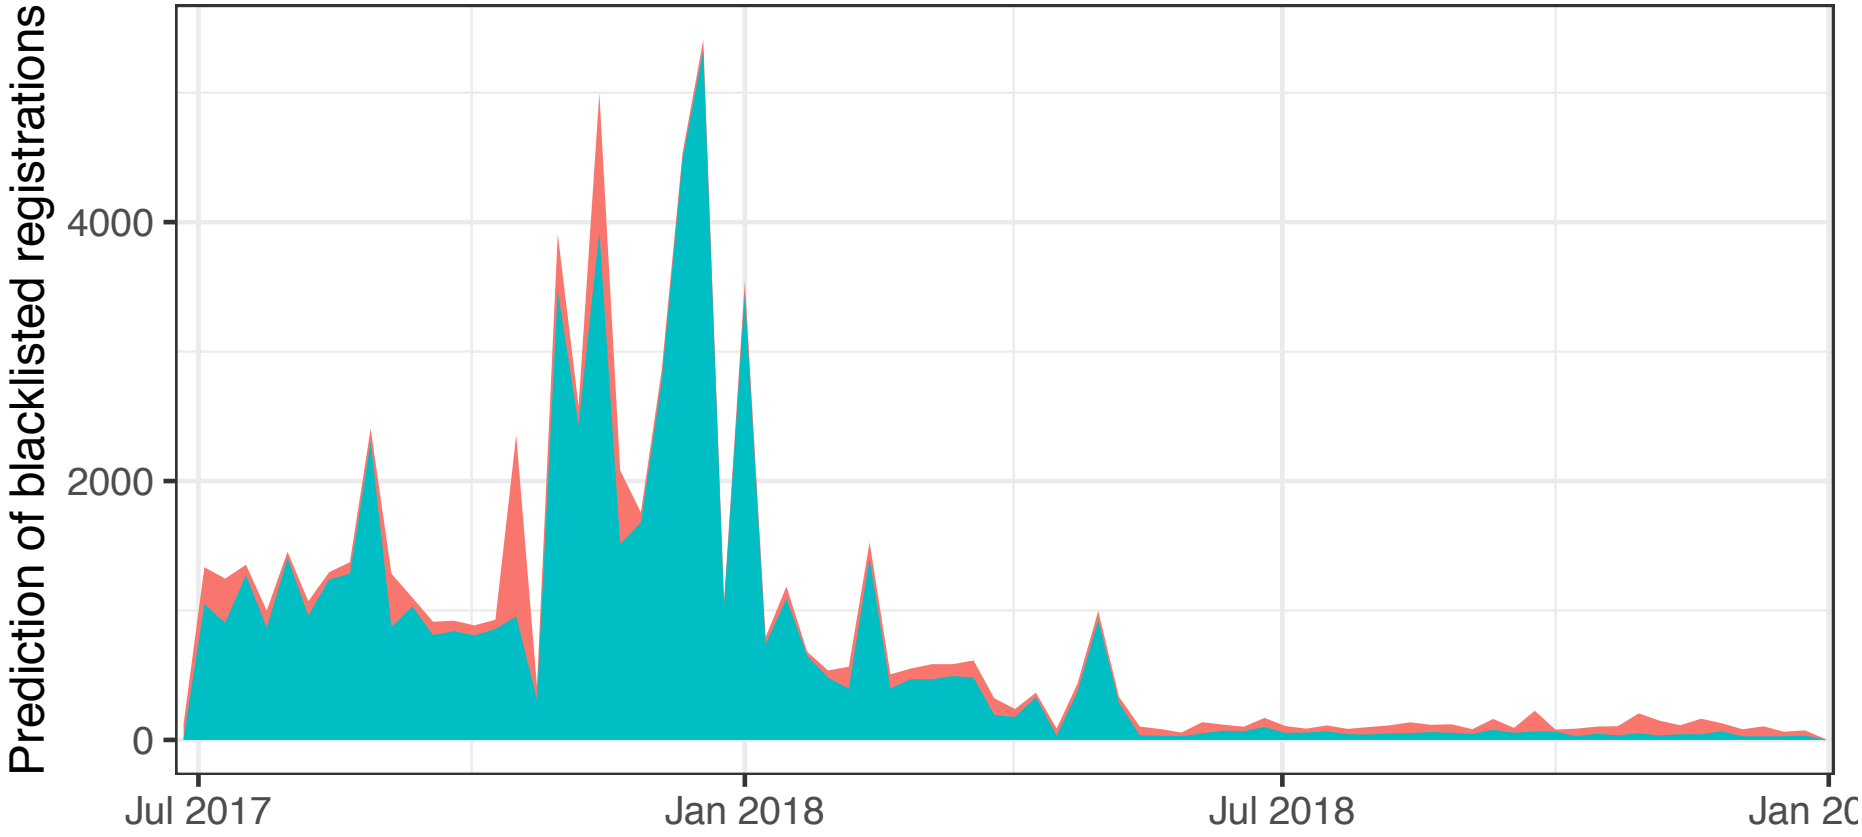
Prediction of blacklisted registrations
4000
2000
0
Jul 2017
Jan 2018
Jul 2018
Jan 20

.eu Home News Register a domain name My .eu Become a registrar WHOIS About us Contact Us

Over 25 000 domain names suspended with ties to identity fraud

Tweet

« Back to the news page

On 29 January 2018, EURid susp

With actions as such, our focus i
enforcement, both on a nationa
towards building the most trustv
illegal activity online. "With our th
names for potential abuse, leadi
EURid Legal Manager.

In 2017, we suspended 20 126 d
enforcement.

Predictive Algorithms

Through the use of historical data and self-learning algorithms, we are working to predict at the time of registration whether or not a domain name might be used in an abusive way in an effort to prevent such malicious domain names from becoming active in the first place.

.eu Home News Register a domain name My .eu Bec

Over 11 000 abusive domain names suspended

Tweet

« Back to the news page

On 21 June 2018, EURid suspended 11 760 domain names that were registered with non-eligible registration data, of which some have been reported for abuse.

With actions as such, our focus is on the safety of online consumers. Via close collaborative efforts with law enforcement and our partners, both on a national and European level, as well as with our registrar channel, we continue to work towards building the most trustworthy online domain name space, taking a stand against

onitor our domain names for potential
pared to 2017, where we suspended 20 126
spensions thus far in 2018." said Geo Van

Learn more about the ways we're building a trustworthy .eu and .ею domain name space at trust.eurid.eu.

As part of the EURid's Trust & Security program, 58,966 domains were suspended in 2018.

Operational results

- Period: July 2017 – December 2018 (18 months)
  - Recall: 85.51%
  - Precision: 72.04%
  - False positive rate: 2.86%
- Very big campaigns (October 2017 - March 2018)
- Incomplete ground truth

Most prominent malicious actors

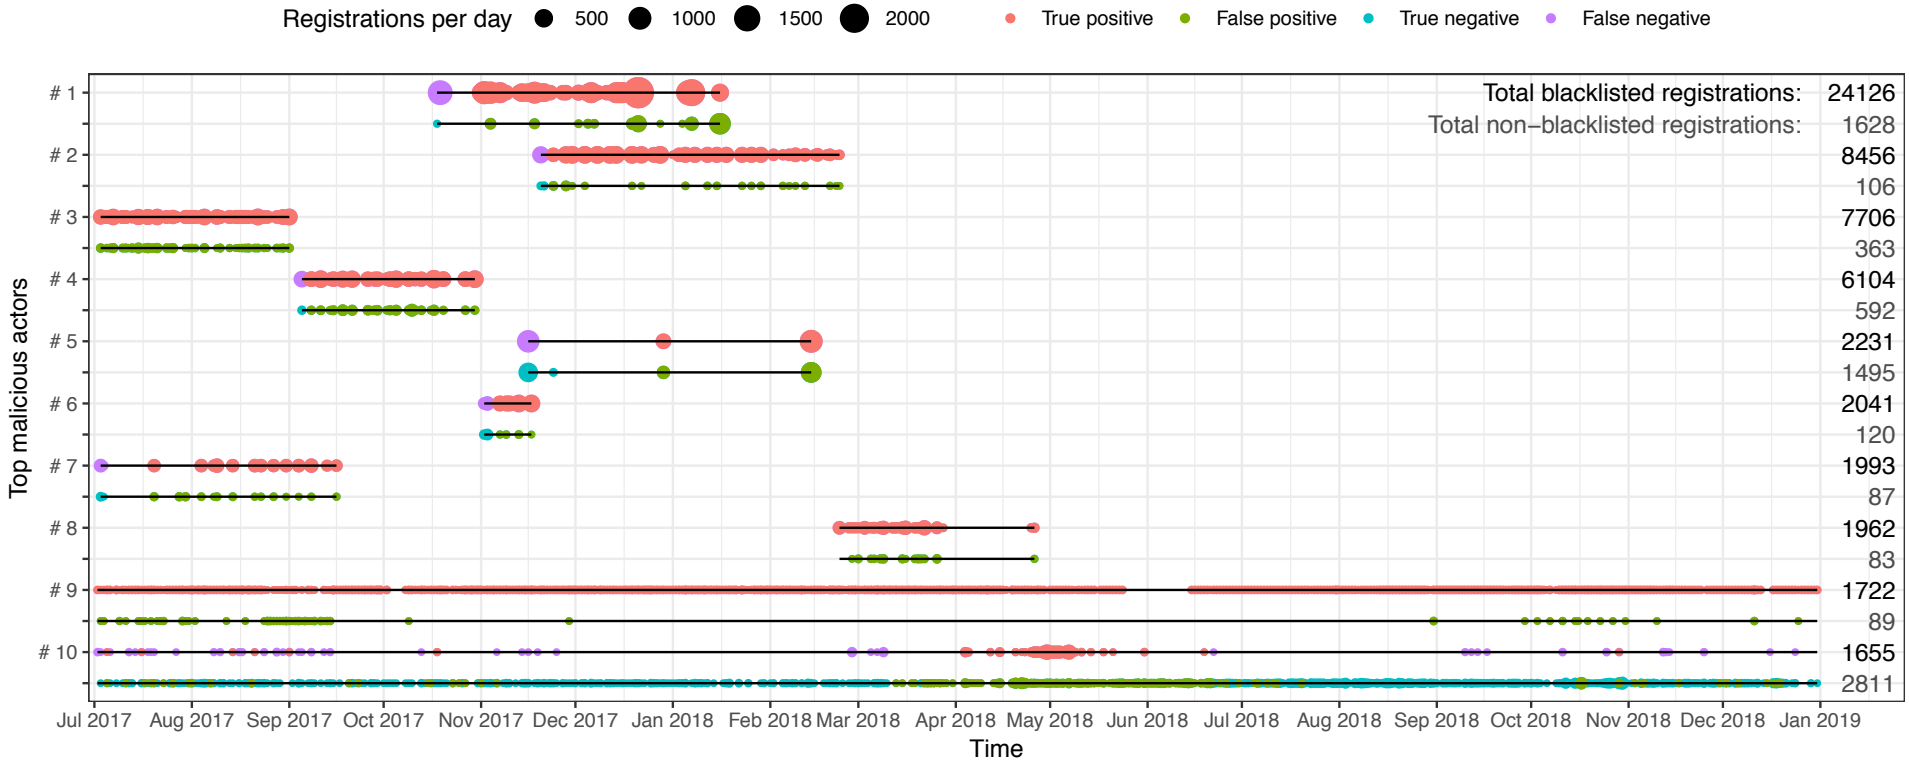

DistriNet

Key takeaways

Rather small set of bad actors

- Up to 20 campaigns are responsible for 80% of malicious registrations

- Top facilitators:
  - About half of the malicious registrations via 1 registrar
  - 1 public email provider are malicious with a high toxicity

Cyber criminals are “human”….

- Lazy
  - Reuse same fake registrants
  - Use generators for registrant details
- Work force
  - Work 9 to 5
  - Take week-ends, holidays
  - Make mistakes (e.g. typos)
- Adapt over time

Registration-time detection and prevention

- Two prediction models:
  - Reputation-based classification
  - Similarity-based clustering

- Captures the majority of malicious domain registrations
- Operates at a low false-positive rate

Machine learning for security

- Data cleansing
  - Tactics to work with incomplete ground truth
- Feature engineering
  - Understand the problem domain well!
- Correct interpretation of the predictions is challenging
  - But often needed in the operational context
  - No shame to start with simple models and gradually enhance

Detection and prevention at .eu

- From research to operations
  - Trade-offs between performance and security
  - Drive for low-false positive rate and understandibility
- Ground truth incompleteness is challenging
- Interesting to see how this will impact the security landscape

Detecting and preventing domain name abuse in the .eu TLD

Lieven Desmet – SecAppDev 2019 – Leuven, BE

Lieven.Desmet@cs.kuleuven.be – @lieven_desmet

DistriNet

EURid