Anomaly-based intrusion detection: challenges and possible strategies from unknowns to APT detection

Andrea Ceccarelli

With the contribution of:

A. Bondavalli, T. Puccetti, T. Zoppi, and BsC and MsC students from the University of Florence.

RCL
RESILIENT COMPUTING LAB

UNIVERSITÀ DEGLI STUDI FIRENZE
DIMAI
DIPARTIMENTO DI MATEMATICA E INFORMATICA "ULISSE DINI"

Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

UNIVERSITÀ DEGLI STUDI FIRENZE
DIMAI
DIPARTIMENTO DI MATEMATICA E INFORMATICA "ULISSE DINI"

Florence, Italy

710,000 The current population of the Metropolitan Area of Florence

5km2 The size of the concentrated area where 95% of Florence's tourism flows through

10-16 M The average yearly tourists in Florence






UNIVERSITÀ DEGLI STUDI FIRENZE
DIMAI
DIPARTIMENTO DI MATEMATICA E INFORMATICA "ULISSE DINI"

Università degli Studi di Firenze

60.000 students, 2500 foreigners
12 faculties, more than 150 degree courses
2.300 professors and researchers

750 research fellows
100 temporary researchers
1.400 PhD students
1.700 technicians and administrative people

The RCL Group is part of the
Dipartimento di Matematica ed Informatica (DiMaI)
Viale Morgagni, 65
50134 – Firenze , Italy
http://www.dimai.unifi.it/

UNIVERSITÀ DEGLI STUDI FIRENZE
DIMAI DIPARTIMENTO DI MATEMATICA E INFORMATICA "ULISSE DINI"

Meet RCL in Florence!








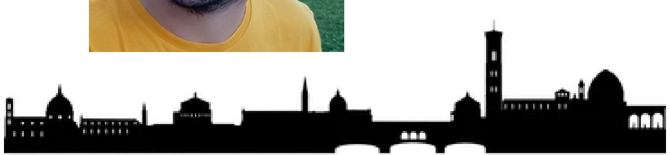


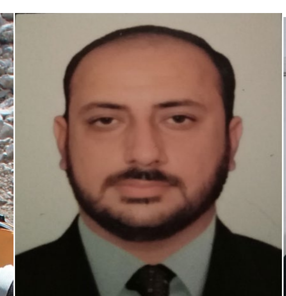
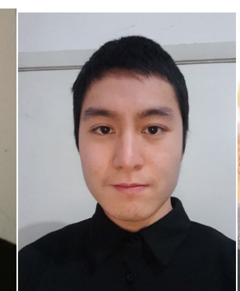


Main Expertises of the Group

Design of Critical Systems and Infrastructures

- Dependable and Secure Architectures
- Intrusion, Error, Anomaly Detection
- Monitoring, Analysis, Diagnosis

V&V and Assessment

- Threat/Hazard Analysis, Risk Assessment
- Modelling and Simulation
- Fault Injection, Robustness Testing
- Quantifying Safety of AI Systems

UNIVERSITÀ DEGLI STUDI FIRENZE
DIMAI
DIPARTIMENTO DI MATEMATICA E INFORMATICA "ULISSE DINI"

Research Projects since 2022 – Funders and Timeline

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

eurostars™

Chips JU

Regione Toscana

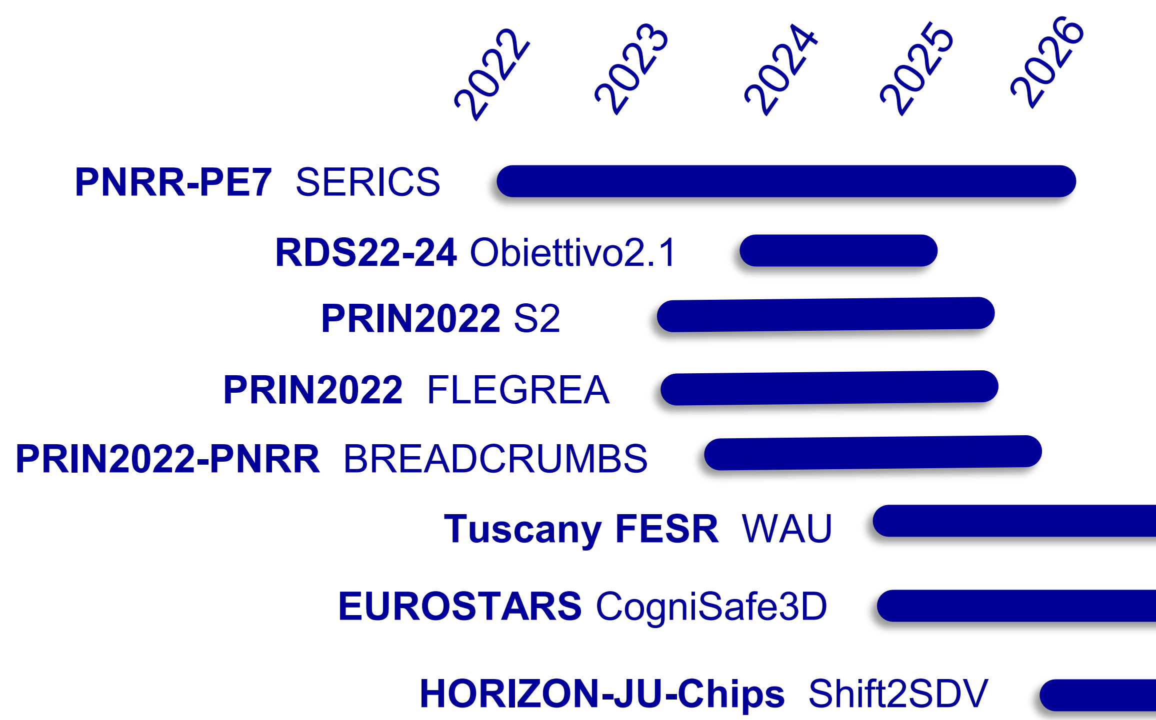

Main relations with Industries

Rete Ferroviaria Italiana (**RFI**): 2018--2024

- Support to the design, implementation and V&V of embedded railway systems, HMIs and communication protocols, with **full compliance to EN 50126/28/29/59 SIL 4**

Resitech SRL, an SME focused on safety-critical embedded systems, mainly automotive and railway

- Was our Academic Spinoff
- Regular interactions and collaborations on research subjects

Aruba S.p.A.

- Support to security assessment

Many training courses on

- Safety Critical Systems
- Fault-Tolerant Architectures
- Risk Assessment, safety standards

Presentation Outline

Some Basics on Threats and Anomalies

Building an Anomaly-Based Intrusion Detection

Detecting unknowns

What’s next: towards detection of APT

Wrap-Up and Concluding Remarks

Presentation Outline

Some Basics on Threats and Anomalies

Building an Anomaly-Based Intrusion Detection

Detecting unknowns

What’s next: towards detection of APT

Wrap-Up and Concluding Remarks

Threats to Security

Security builds around three properties

- **Availability**: readiness for correct service
- **Confidentiality**: the absence of unauthorized disclosure of information
- **Integrity**: absence of improper system alterations

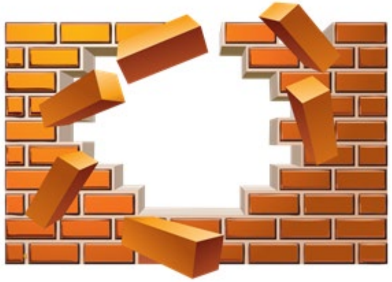

Attacks aim at damaging at least one of the three attributes

Definition from: Avizienis, A., Laprie, J. C., Randell, B., & Landwehr, C. (2004). Basic concepts and taxonomy of dependable and secure computing. IEEE transactions on dependable and secure computing, 1(1), 11-33.

ENISA's Threat Landscape - analysed incidents by threat type

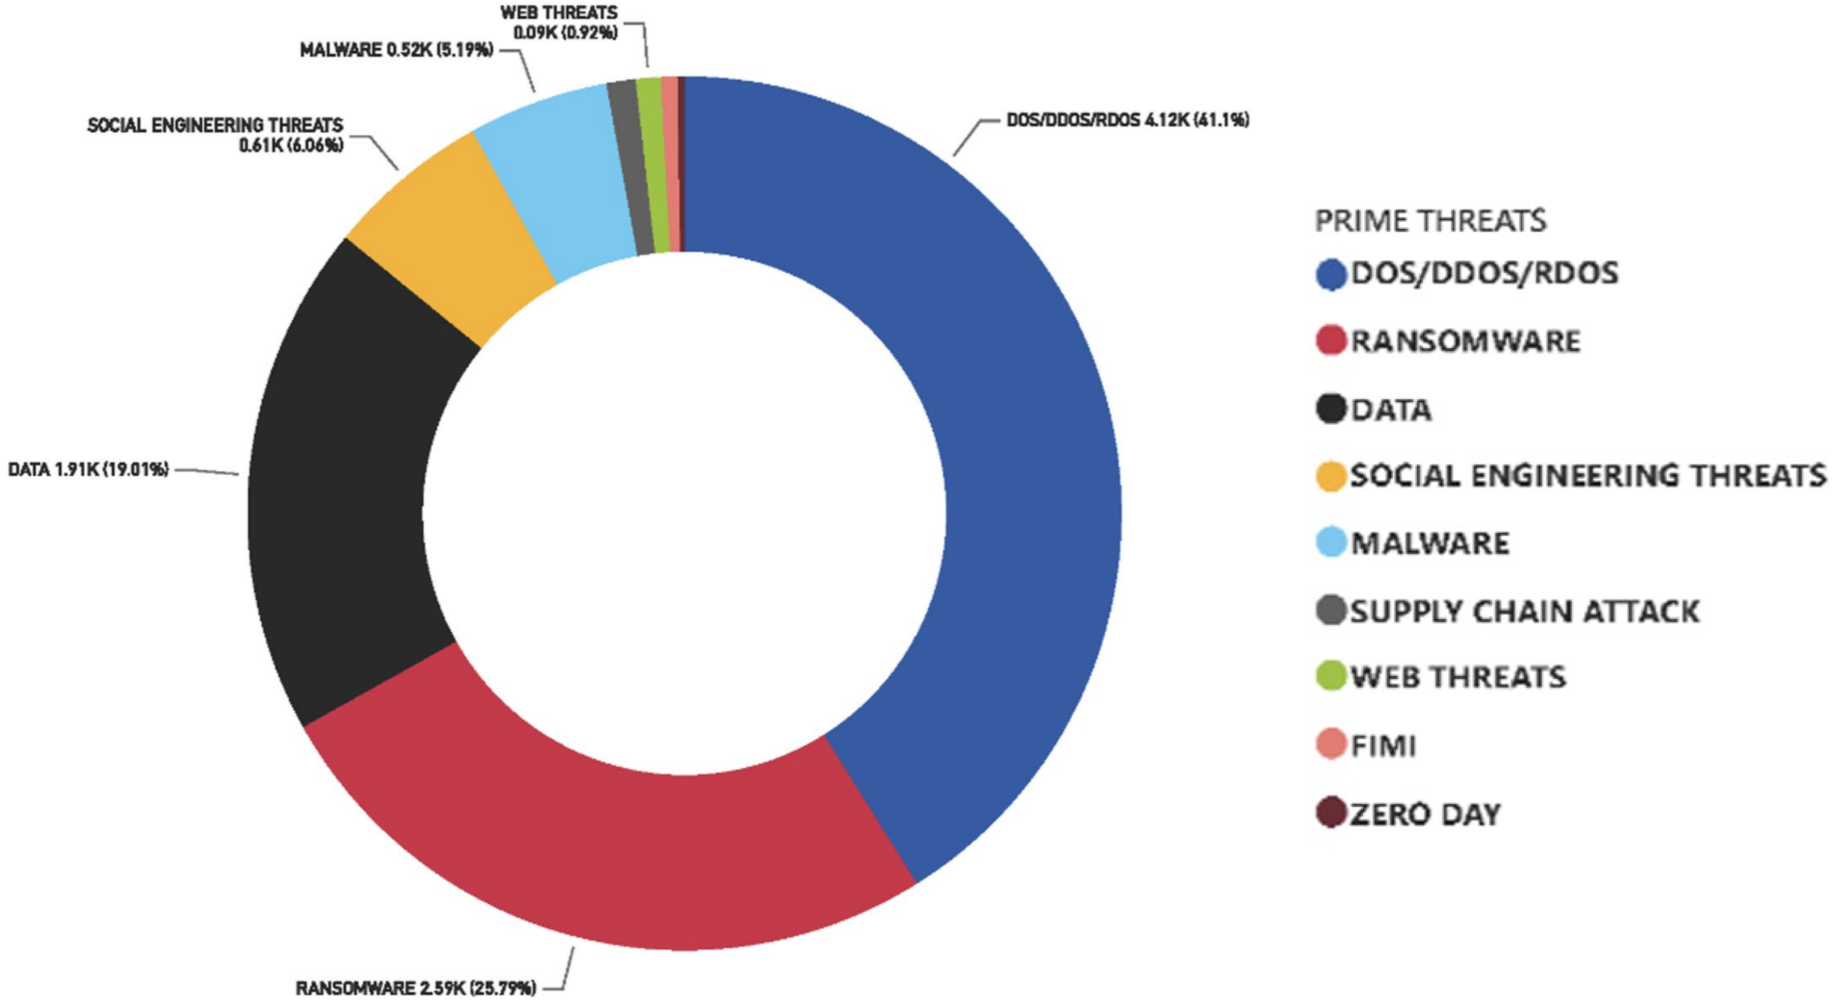

https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024

How to defend

Means to realize intrusion detections:

Rule-based, Invariant-Based, Signature-based

our focus!

Anomaly-based (under the underlying assumption that attacks have a visible effect on monitored system indicators)

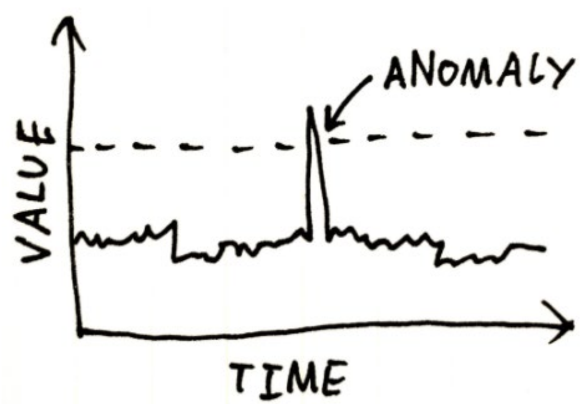

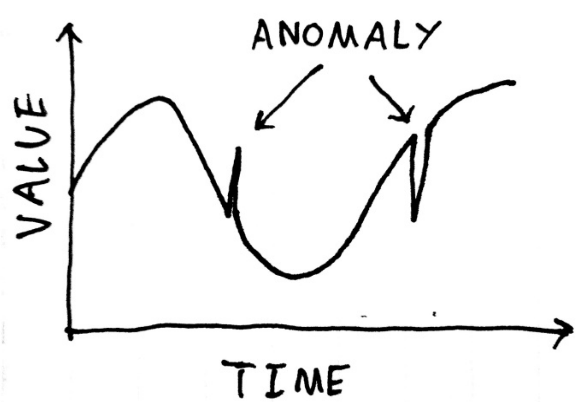

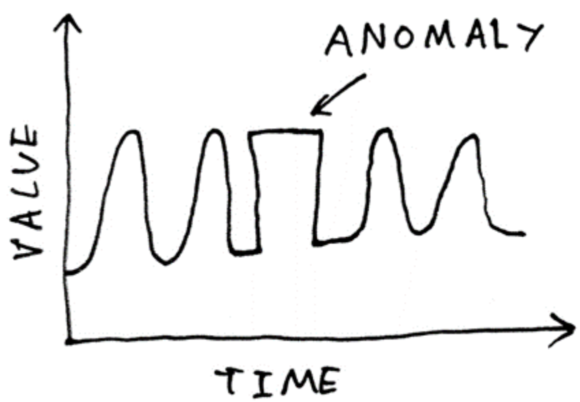

First things first: what is anomaly detection?

Anomaly detection refers to the problem of finding patterns in data that do not conform to an expected behaviour

Chandola, Varun, Arindam Banerjee, and Vipin Kumar. "Anomaly detection: A survey." ACM computing surveys (CSUR) 41.3 (2009): 15.

Searching for anomalies

Anomalies in data can be symptoms of attacks or errors

- **Dependability**: software errors, misconfigurations
- **Security**: malware, attacks (e.g., DDoS/Ping Flood)

our focus:

Finding anomalies requires an **anomaly-based intrusion detection system**

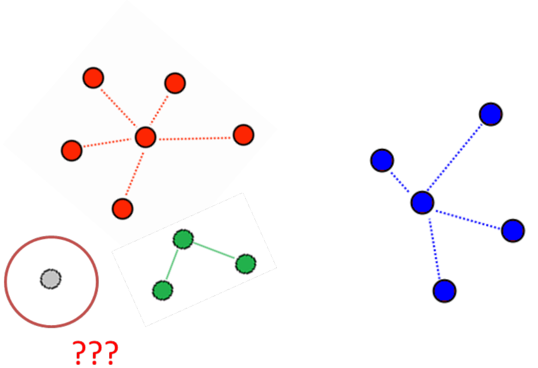

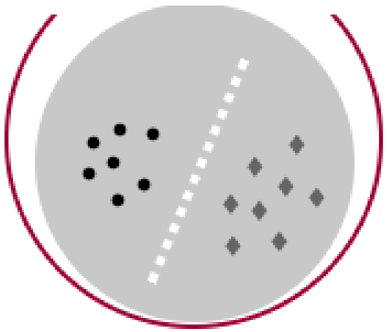

Presentation Outline

Some Basics on Threats and Anomalies

Building an Anomaly-Based Intrusion Detection

Detecting unknowns

What's next: towards detection of APT

Wrap-Up and Concluding Remarks

Paradigm Shift: from rules identification…

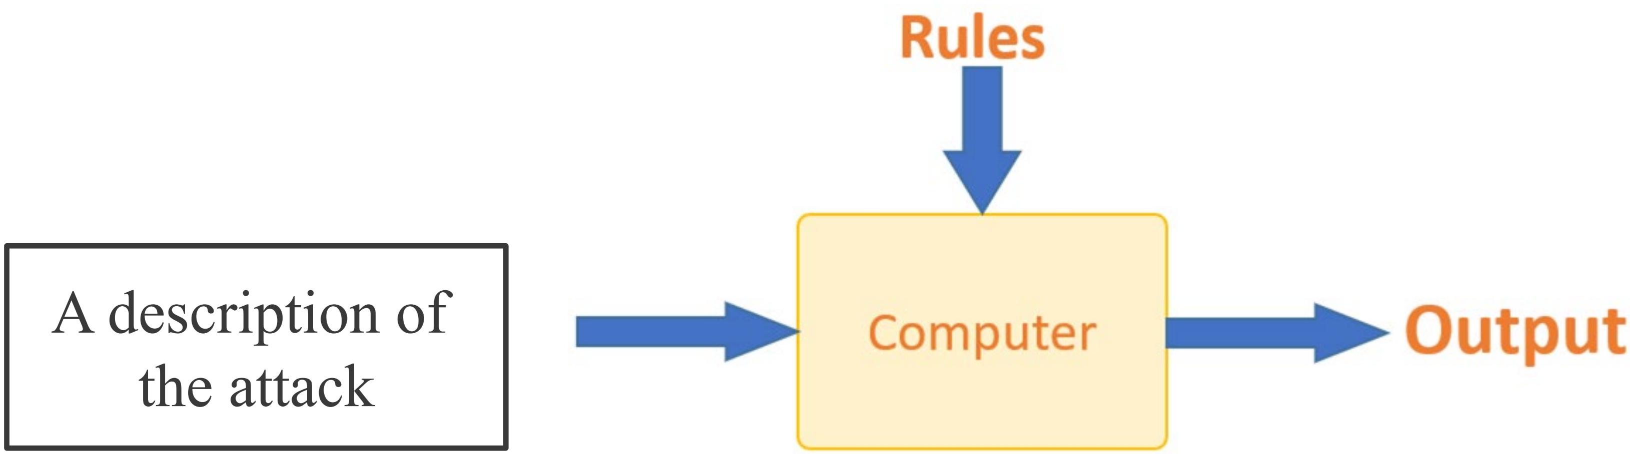

Traditional Programming

… to training and testing!

Next, short review of:
1- datasets
2- classifiers
3- evaluation

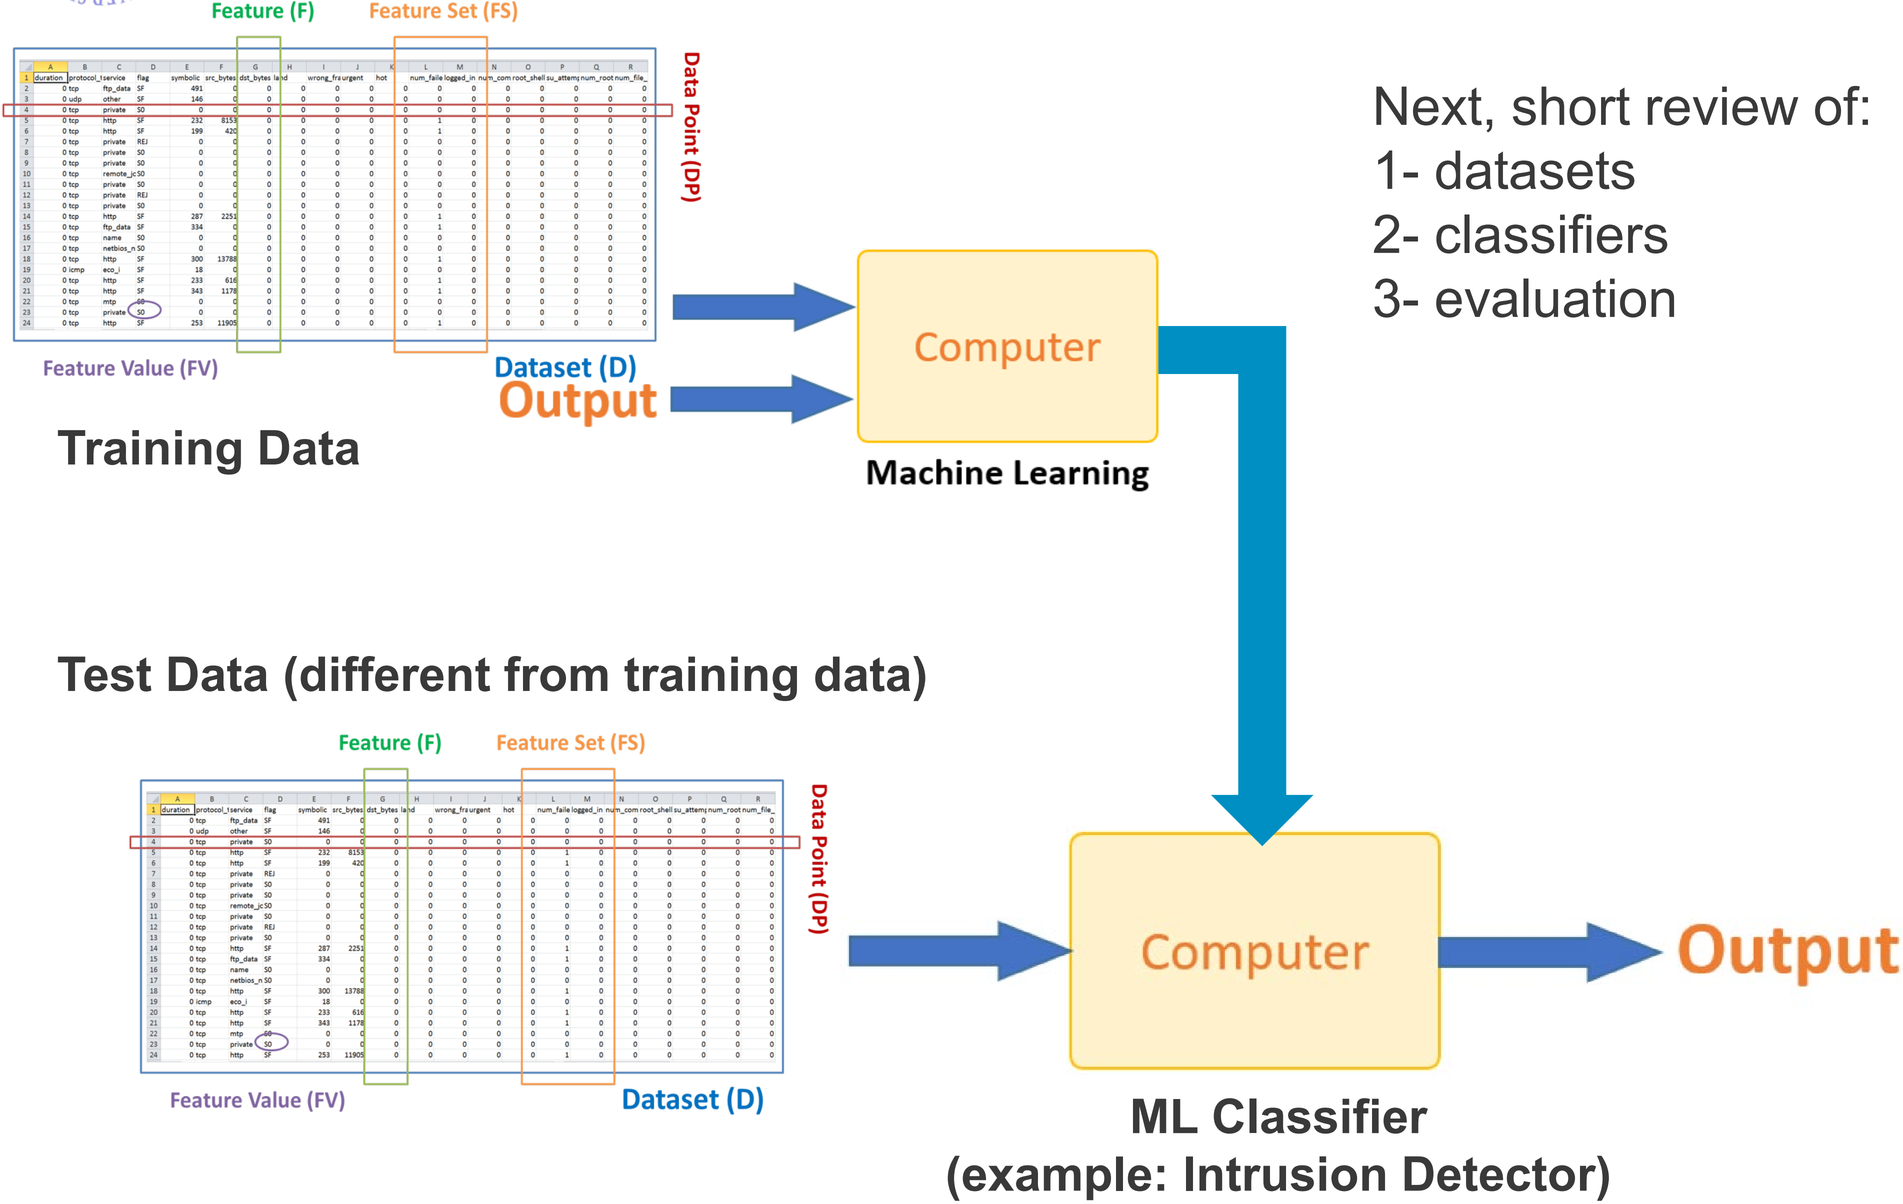

General Structure of a Dataset

and a label!

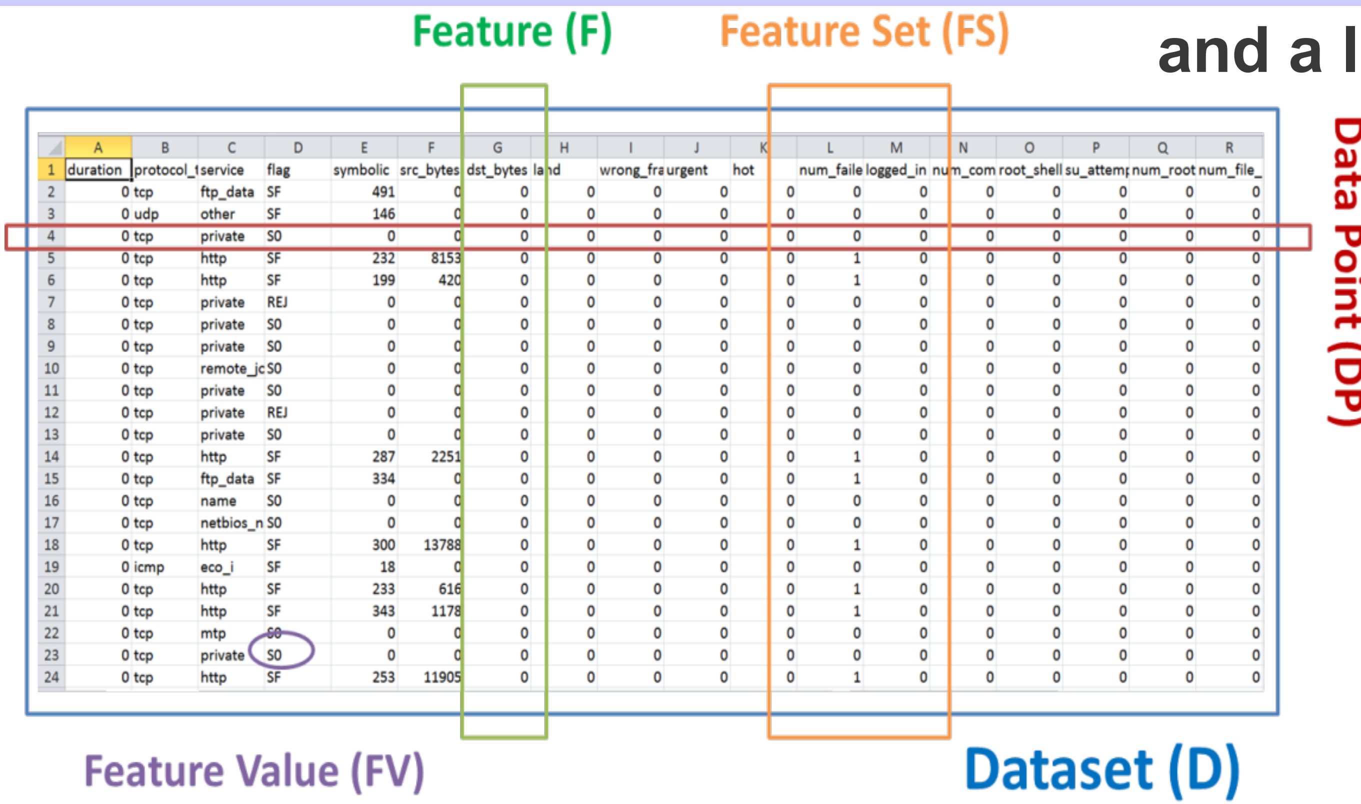

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
1	duration	protocol_t	service	flag	symbolic	src_bytes	dst_bytes	land	wrong_fra	urgent	hot	num_faile	logged_in	num_com	root_shell	su_attemp	num_root	num_file_
2	0	tcp	ftp_data	SF	491	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	udp	other	SF	146	0	0	0	0	0	0	0	0	0	0	0	0	0
4	0	tcp	private	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	0	tcp	http	SF	232	8153	0	0	0	0	0	1	0	0	0	0	0	0
6	0	tcp	http	SF	199	420	0	0	0	0	0	1	0	0	0	0	0	0
7	0	tcp	private	REJ	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	0	tcp	private	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9	0	tcp	private	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10	0	tcp	remote_jc	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0	tcp	private	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0	tcp	private	REJ	0	0	0	0	0	0	0	0	0	0	0	0	0	0
13	0	tcp	private	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14	0	tcp	http	SF	287	2251	0	0	0	0	0	1	0	0	0	0	0	0
15	0	tcp	ftp_data	SF	334	0	0	0	0	0	0	1	0	0	0	0	0	0
16	0	tcp	name	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
17	0	tcp	netbios_n	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
18	0	tcp	http	SF	300	13788	0	0	0	0	0	1	0	0	0	0	0	0
19	0	icmp	eco_i	SF	18	0	0	0	0	0	0	0	0	0	0	0	0	0
20	0	tcp	http	SF	233	616	0	0	0	0	0	1	0	0	0	0	0	0
21	0	tcp	http	SF	343	1178	0	0	0	0	0	1	0	0	0	0	0	0
22	0	tcp	mtp	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
23	0	tcp	private	S0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
24	0	tcp	http	SF	253	11905	0	0	0	0	0	1	0	0	0	0	0	0

pcap

session summaries

syscall traces

system indicators

network indicators

(2009) NSL-KDD

(2011) CTU-13

(2012) ISCX12

(2015) UNSW-NB15

(2017) AndMal17

(2017) Netflow-IDS

(2018) CICIDS18

(2020) SDN20

General Structure of a Dataset

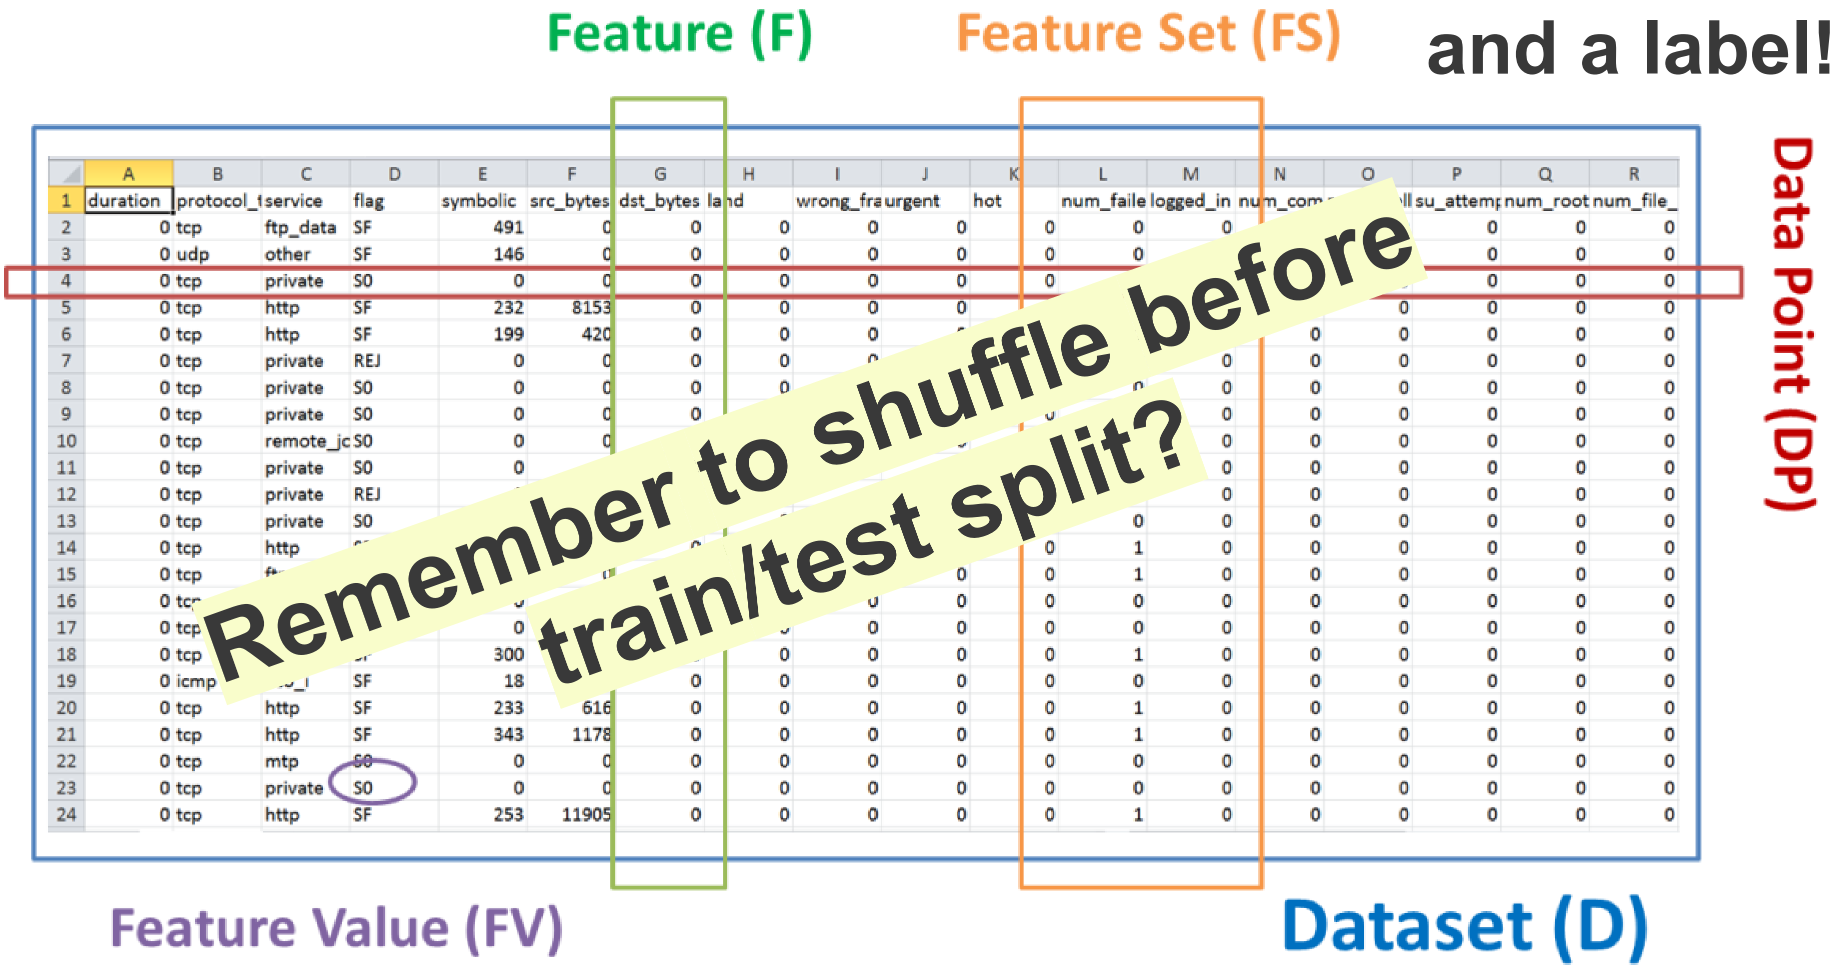

(2009) NSL-KDD
(2011) CTU-13
(2012) ISCX12
(2015) UNSW-NB15
(2017) Netflow-IDS
(2017) AndMal17
(2018) CICIDS18
(2020) SDN20

Mapping of Attacks and Datasets (2020)

Attack Category ENISA Rank	Malware 1	Web Attack 2	Web Application 4	Spam / Phishing 3, 5	(D)Dos 6	BotNet 7	Data Breaches 8
NSL-KDD	u2r		r2l		DoS		Probe
CTU-13						BotNet	
ISCX12		BruteForce			DoS, DDoS		Infiltration
UNSW-NB15	Worms	Fuzzers	Backdoor, Exploits, Shellcode		DoS		Analysis, Reconnaissance
UGR16				Blacklist, Spam	DoS	BotNet	Scan
NGIDS-DS	Malware, Worms		Backdoor, Exploits, Shellcode		DoS		Reconnaissance
Netflow-IDS				Mailbomb	Neptune, Portsweep		
AndMal17	Ransomware, Scareware			SMS, Adware			
CIDDS-001		BruteForce			DoS		PortScan, PingScan
CICIDS17		BruteForce			DoS (Slowloris, Goldeneye)		PortScan
CICIDS18		BruteForce (FTP, SSH)			DoS, DDoS	Bot	Infiltration
SDN20		BruteForce	Exploits		DoS, DDoS		Probe

different features different systems Same attack, different visible effects

T. Zoppi, et al. "Towards a general model for intrusion detection: An exploratory study." *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Cham: Springer Nature Switzerland, 2022.

Classifiers: supervised vs unsupervised

Supervised: labels attack/normal are available in the training set (and are used)

Unsupervised: no labels are used during training

	Known attacks	Unknown attacks
Supervised	Very Good!	Potentially Bad
Unsupervised	Average	

Supervised Algorithms: Examples

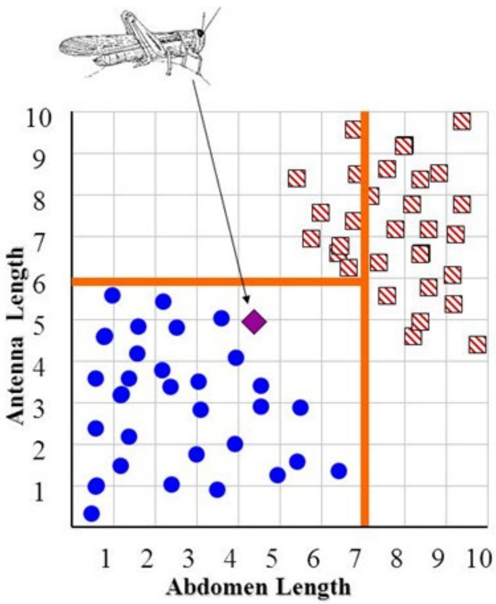

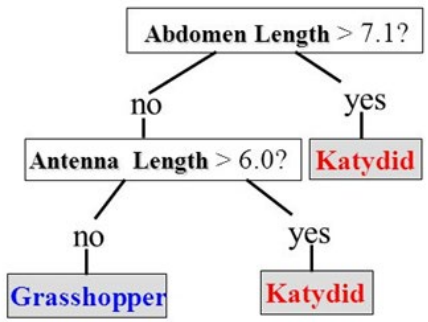

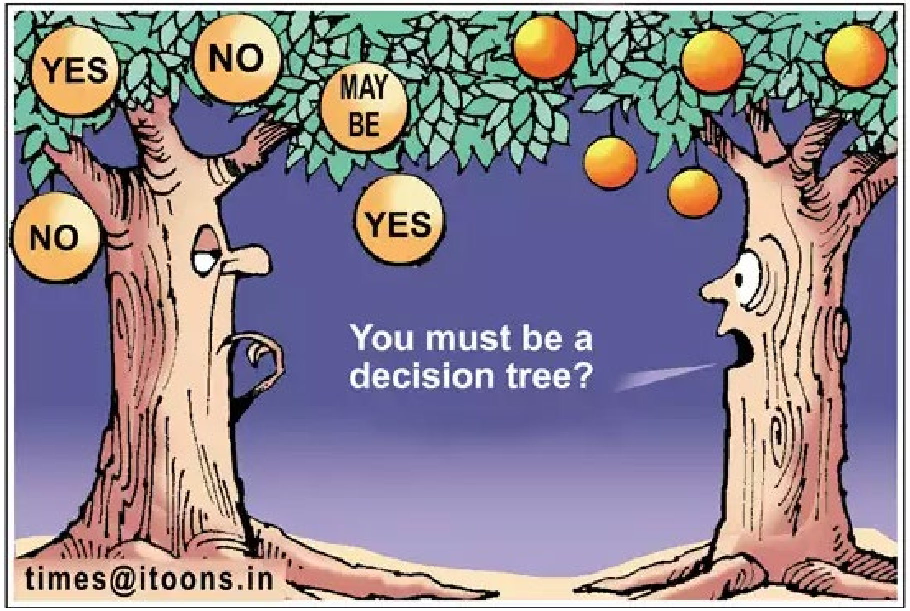

Supervised Algorithms: Examples

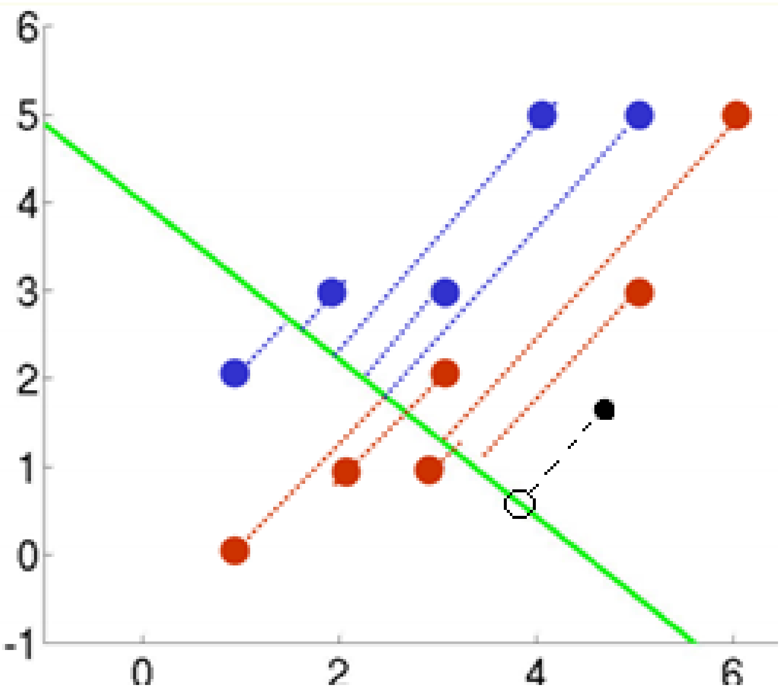

Linear Discriminant Analysis
(dimensionality reduction)

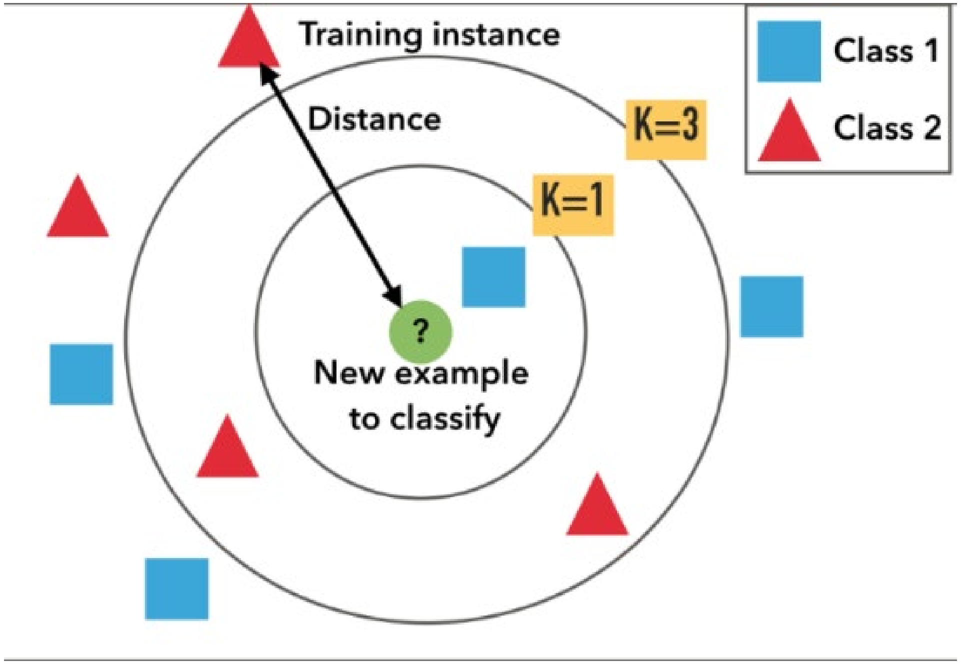

kNN

Unsupervised Algorithms: Examples

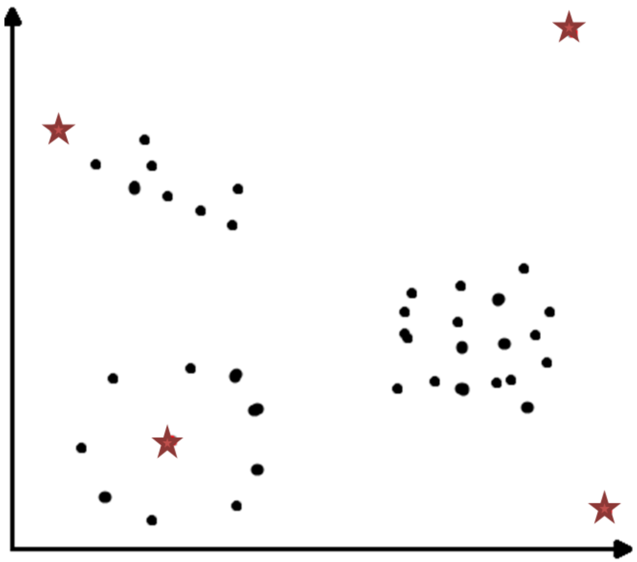

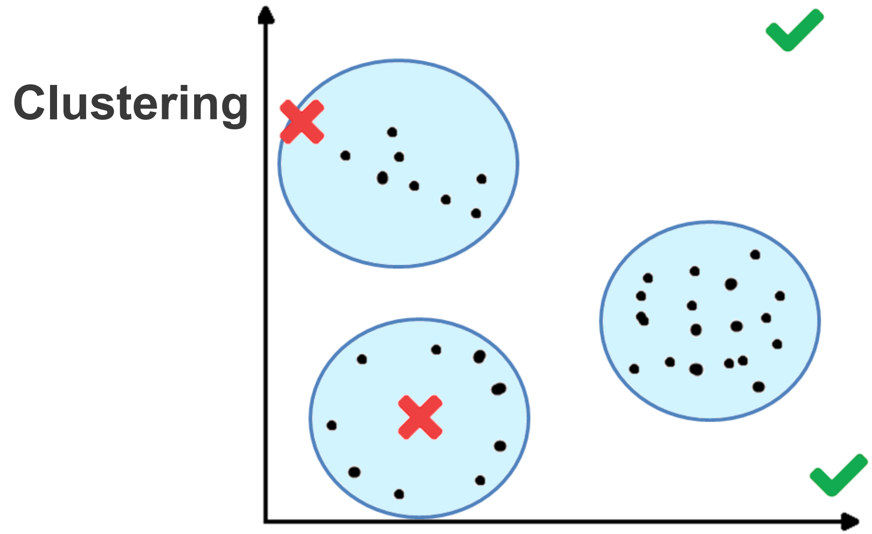

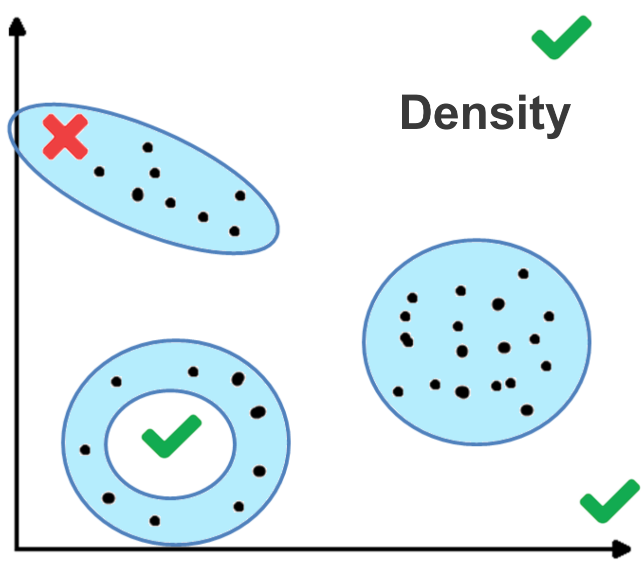

Unsupervised Algorithms: Examples

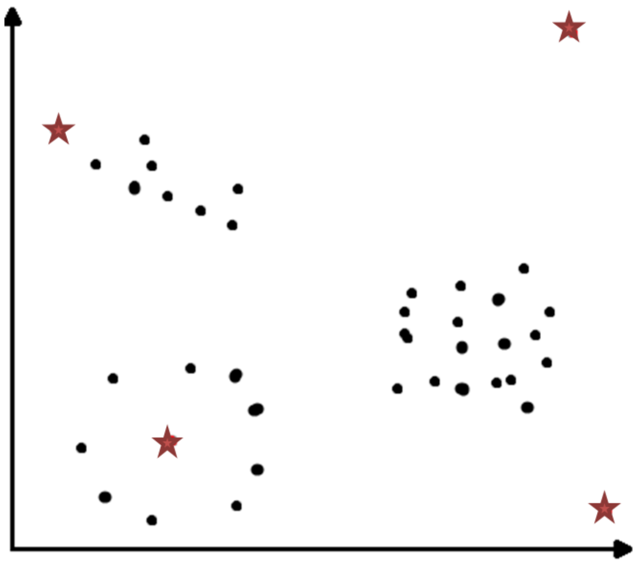

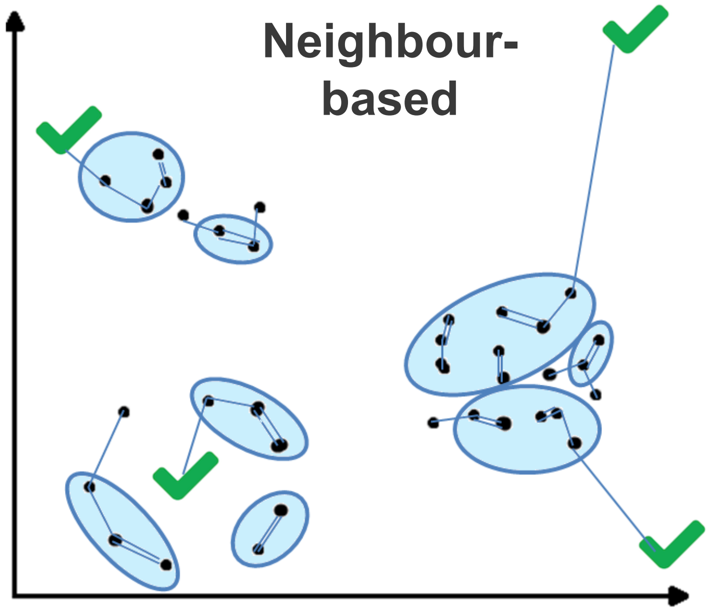

Angle-
based

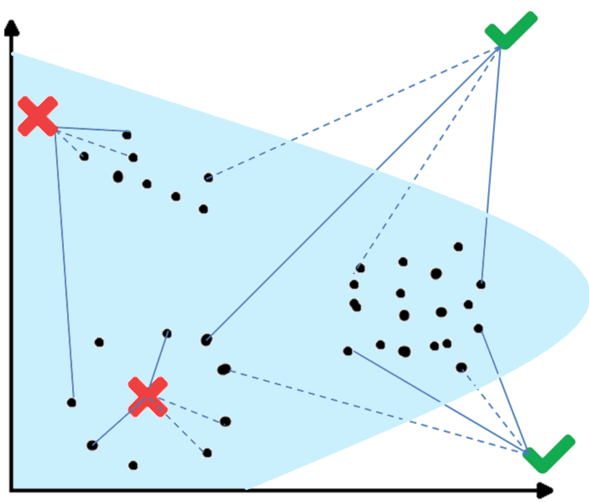

Deep Neural Networks?

Nowadays DNNs are very popular as they work well in many applications

However, they struggle when classifying tabular data and especially IDS datasets

T. Zoppi, et al. "Anomaly-based error and intrusion detection in tabular data: No DNN outperforms tree-based classifiers." Future Generation Computer Systems 160 (2024): 951-965.

Therefore, in this talk we will skip DNNs and focus on non-DNN algorithms

Evaluation of an IDS

The trained model is used for testing.

- The model outputs a **numeric score** that allows to decide on the «class» of the data point
- To decide attack/normal (binary classification), numeric score is converted into a boolean score

If **Ground Truth** (label) is available, it is possible to calculate Metric Scores

How to evaluate an anomaly detector

The suitability and the effectiveness of anomaly detectors are usually evaluated and compared depending on specific metrics

- True Positives (TP)
- True Negatives (TN)
- False Positives (FP)
- False Negatives (FN)

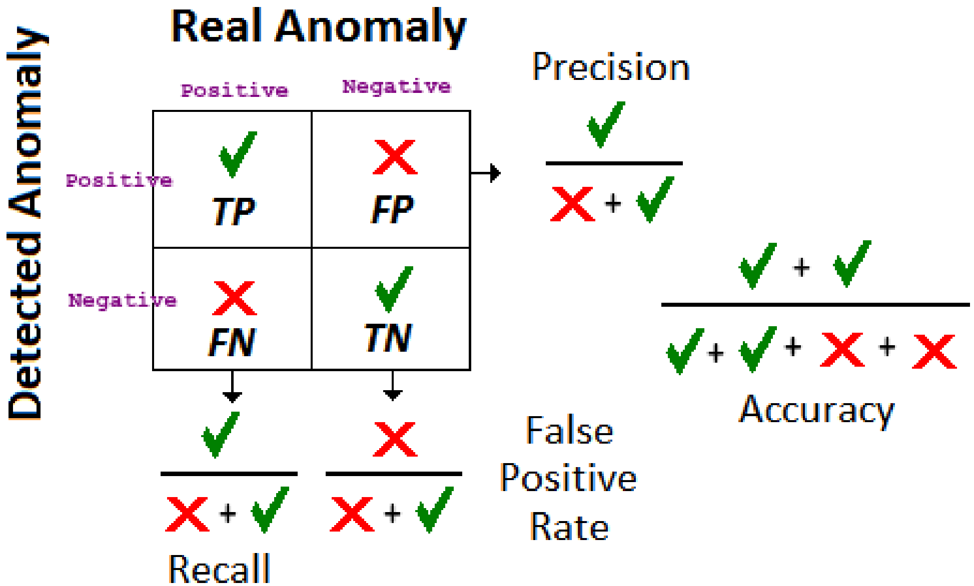

Scoring Metrics: problems?

However…. Most likely, you will have unbalanced test sets: metrics need to be used with caution!

Example

A test set with 1% of normal and 99% of attacks

A useless IDS that always answers "attack", gets accuracy 99%, precision 99%, recall 100%!

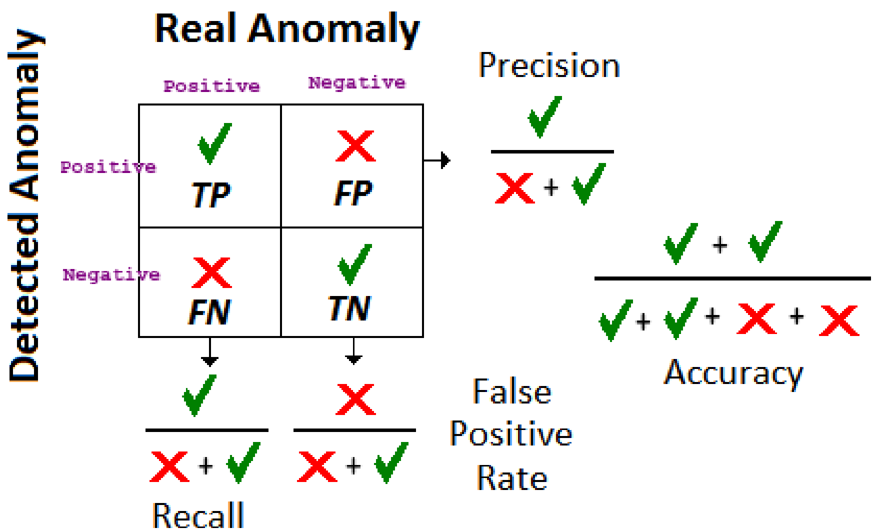

Alternatives

Matthews Correlation Coefficient (MCC)

$$\mathrm{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

Ranges from -1 to 1: 1 “perfect”, -1 “perfectly wrong”, 0 random guessing

Or: clearly declare the class balance, specify the normal/anomaly ratio, specify FPR, …

One more perspective: attack latency

We want to **promptly** detects attacks

but what does it mean «promptly»?
just a matter of response time?

In practise, we may want to understand relations between latency and detection capability, for example:

attackers should be detected within X seconds from their first action!

timestamp	features	label
Tue, 24 Sep 2024 10:59:18	...	normal
Tue, 24 Sep 2024 10:59:53	...	normal
Tue, 24 Sep 2024 11:00:00	...	normal
Tue, 24 Sep 2024 11:00:10	...	normal
...	...	...
Tue, 25 Sep 2024 00:00:10	...	attack
Tue, 25 Sep 2024 00:30:00	...	attack
Tue, 25 Sep 2024 00:31:00	...	attack
Tue, 25 Sep 2024 00:31:30	...	attack
...	...	...

SotA Datasets
Days of normal data points, followed by many attacks executed in sequence.
Not good to answer the question above!

Introducing attack latency

Many attack are not "send 1 packet, immediate effect". We measure latency as a time interval, or as the number of data points between two data points x_i "attack started" and x_d "attack detected".

- **Average Latency = $\boldsymbol{\Delta L}$** $= \frac{\sum_{i=0}^{N} \Delta l_i}{N}$
- **Sequence Detection Rate SDR** (as there is the case in which x_d never occur)

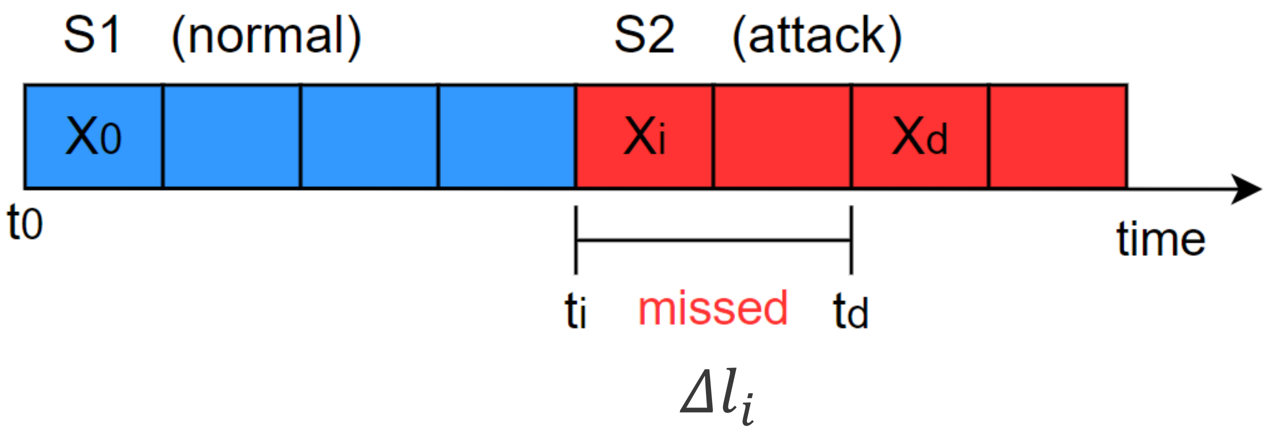

A bit more on the SDR

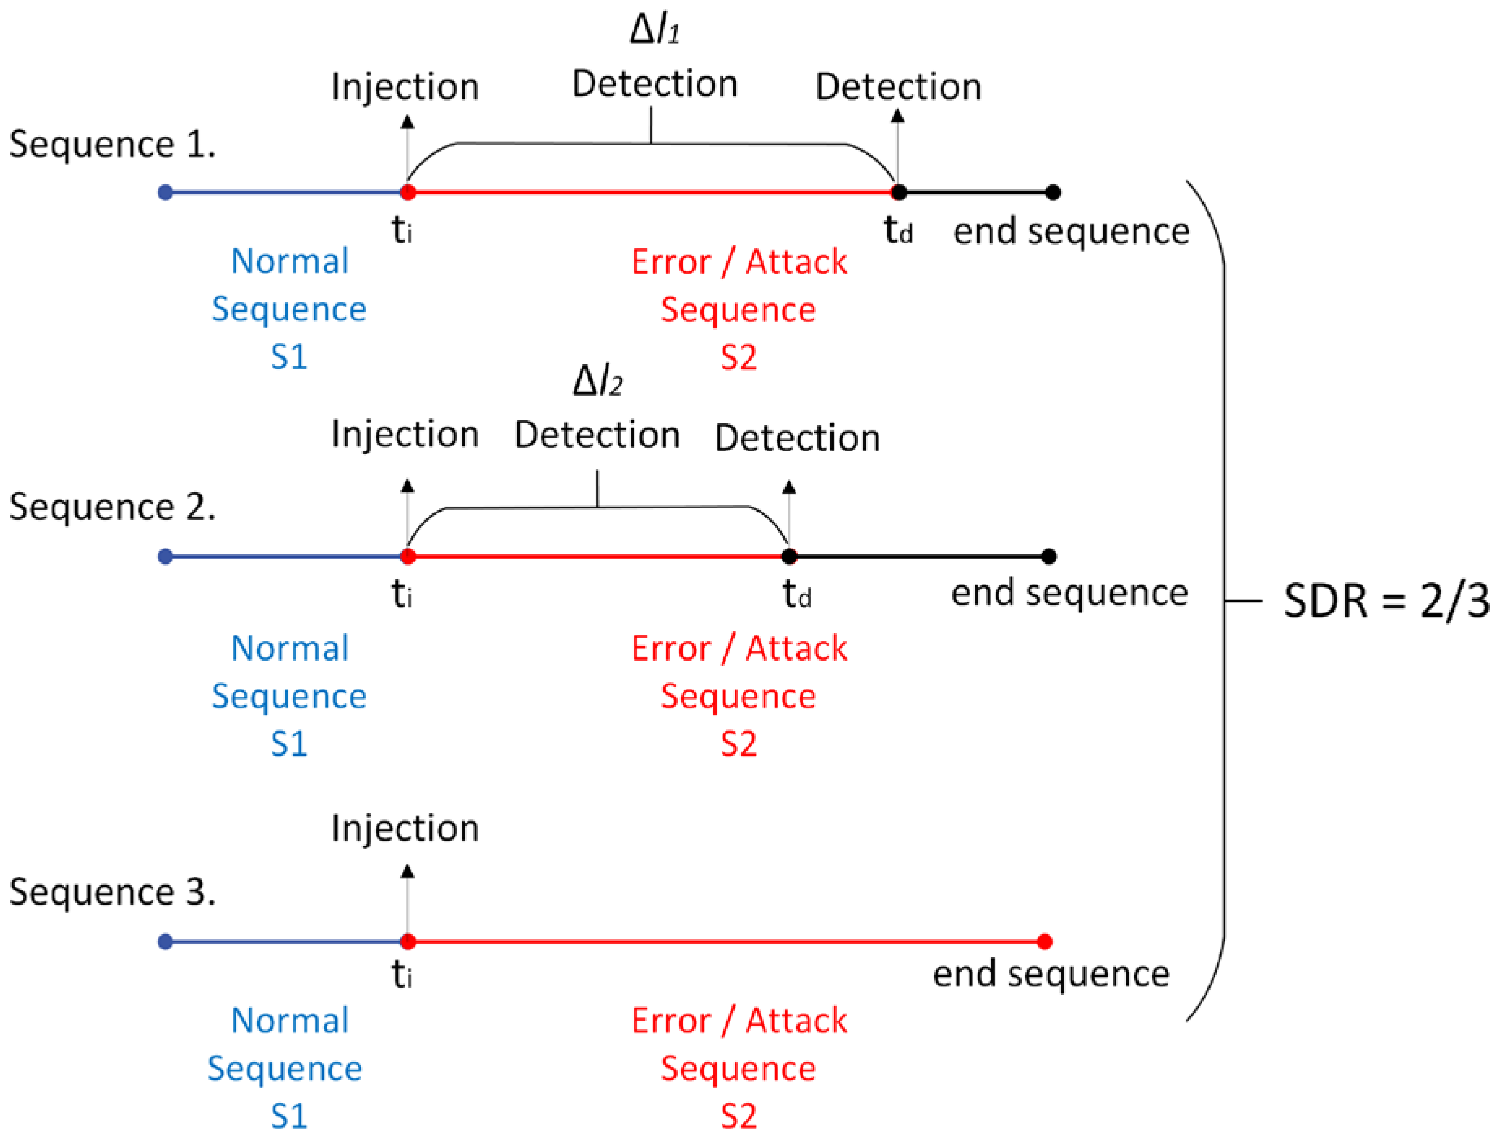

Putting everything in use: create a suitable dataset

SPaCe prototype

Regione Toscana

- Onboard system for metro carriage surveillance

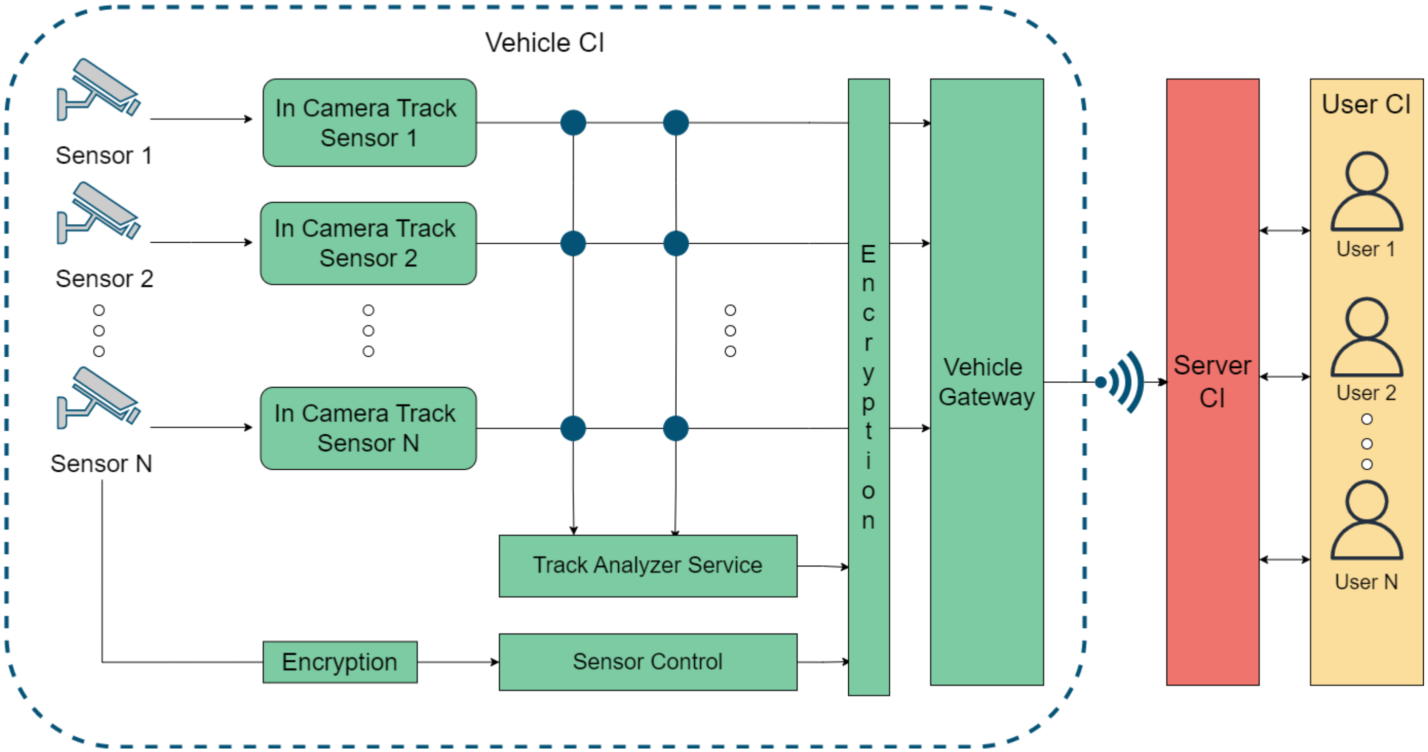

ROSPaCe data collection procedure

Attacker Tool:

- Metasploit
- Nmap
- Scapy Custom Scripts

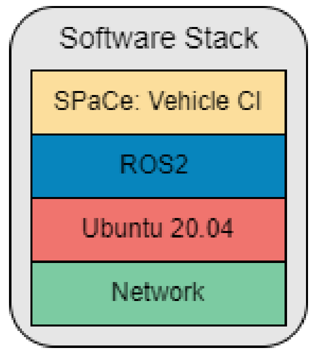

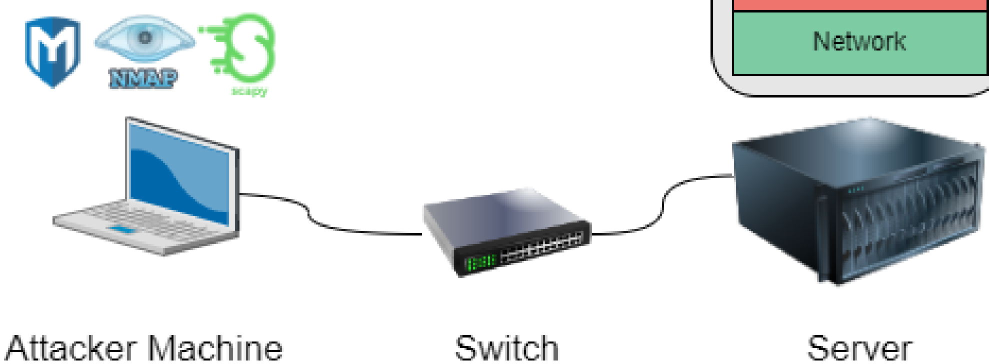

ROSPaCe: Intrusion Detection Dataset for a ROS2-Based Cyber-Physical System and IoT Network

Tommaso Puccetti[1], Simone Nardi[2], Cosimo Cinquilli[1], Tommaso Zoppi[3] & Andrea Ceccarelli[1]

6 different attacks:

- 2 discovery attacks
- 4 DoS attacks

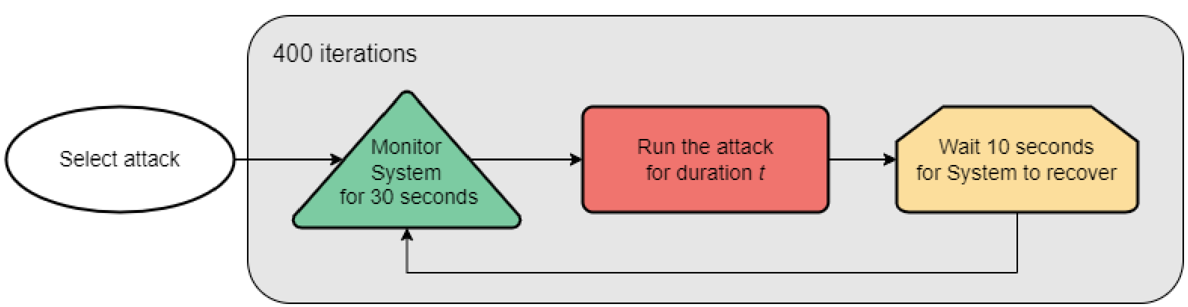

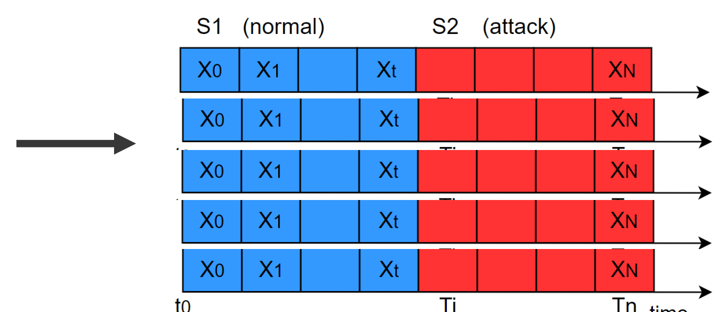

Some results: with «traditional» metrics

XGBOOST			LSTM CD		
Accuracy	Recall	F1	Accuracy	Recall	F1
0.927	0.991	0.952	0.879	0.911	0.953

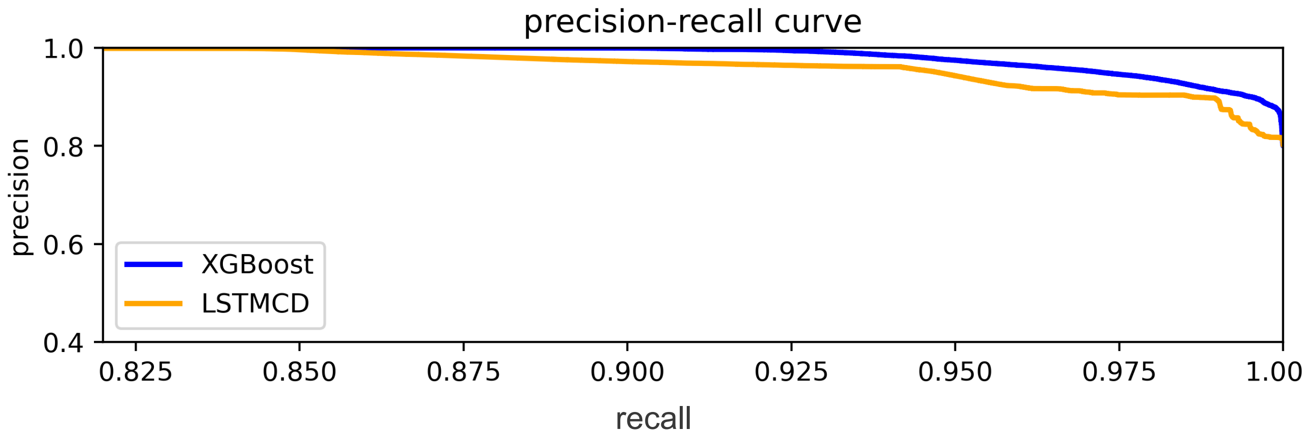

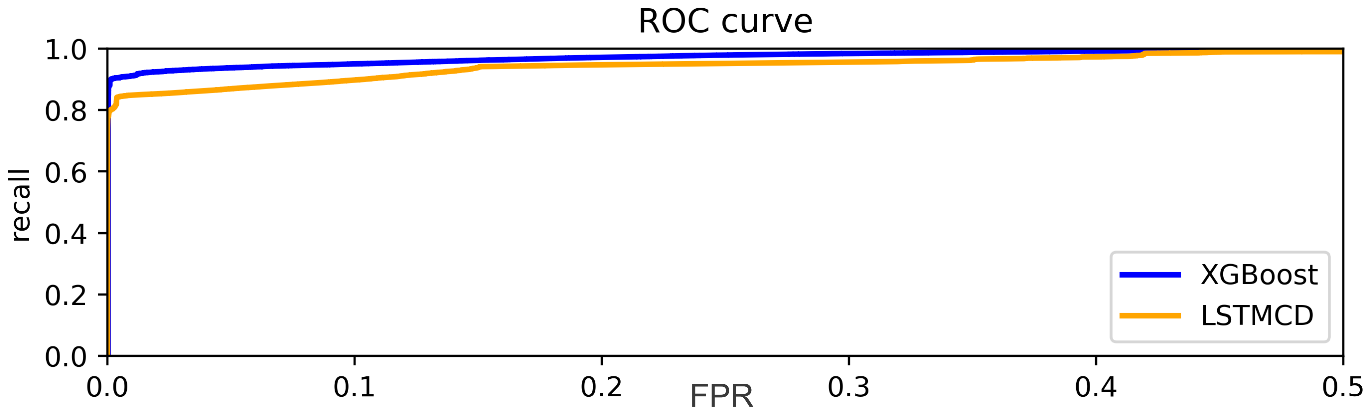

Average latency (versus FPR)

Not such a nice curve, because of undetected sequences

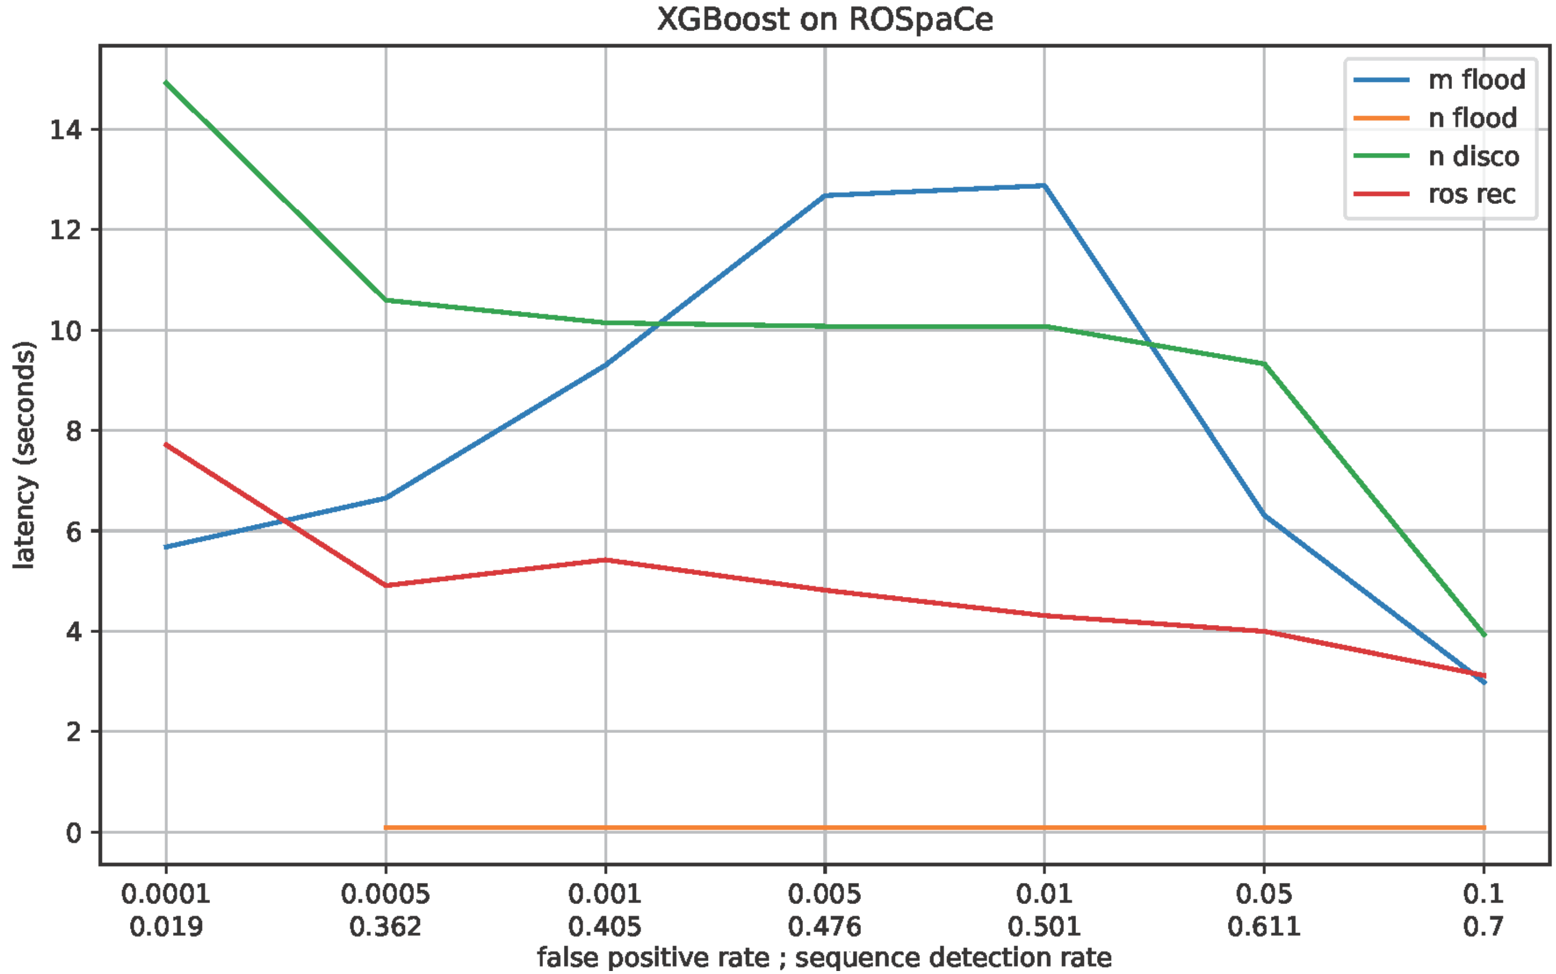

Presentation Outline

Recap on Anomalies and Intrusions

Building an Anomaly-Based Intrusion Detection

Detecting unknowns

What’s next: towards detection of APT

Wrap-Up and Concluding Remarks

AND… What if something unknown pops up?

Research and Practice found ways to defend against specific attacks

- Mostly rule, signature-based or supervised learning

But what about unknowns attacks (zero days), attack variants, … ?

- No rule / signature available
- Anomaly detectors much less efficient

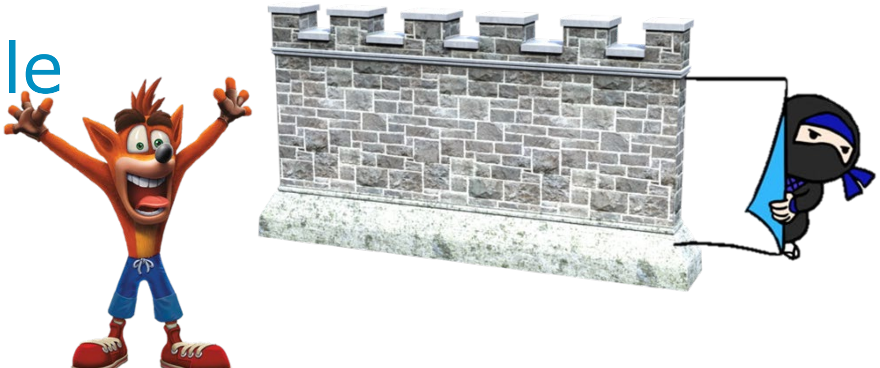

Back to Supervised and Unsupervised strategies

Supervised algorithms are very good in detecting known issues, but have essentially no means to detect unknowns

Detection capability of unsupervised does not change “much” when detecting both known and unknown events

zero-days!

	Known Attacks	Unknown Attacks
Supervised	Very Good!	Potentially Bad
Unsupervised	Average	

How to test against unknowns?

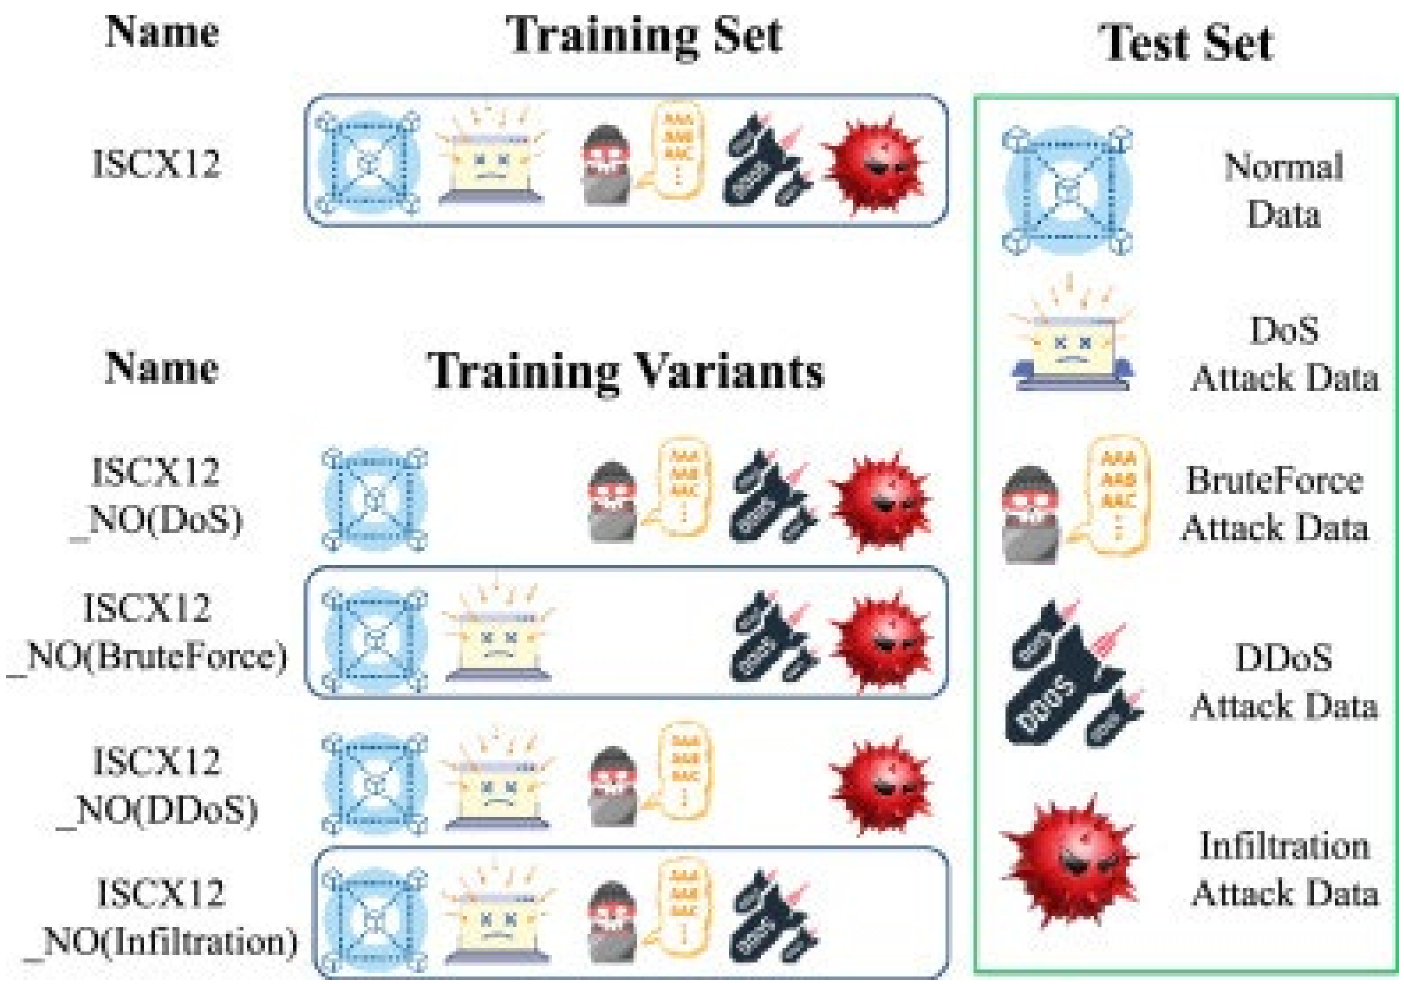

Zoppi, Tommaso, et al. "Which algorithm can detect unknown attacks? Comparison of supervised, unsupervised and meta-learning algorithms for intrusion detection." *Computers & Security* 127 (2023): 103107.

Variants of attack datasets…

Name	Year	# Data Points	Features		Attacks		# Variants
			Ord.	Cat.	#	%	
ADFANet	2015	132 002	5	6(0)	3	11.3	3
AndMal17	2017	100 000	77	5(0)	4	15.5	4
CICIDS17	2017	500 000	77	5(1)	5	79.7	5
CICIDS18	2018	200 000	77	5(1)	8	26.2	8
CIDDS	2015	400 000	5	7(2)	4	14.4	4
IoT-IDS	2019	210 425	8	1(1)	8	42.3	8
ISCX12	2012	600 000	4	10(3)	4	43.5	4
NSLKDD	2009	148 516	37	5(3)	4	40.7	4
SDN20	2020	205 167	63	5(1)	5	66.6	5
UGR16	2016	207 256	4	6(2)	5	3.3	5
UNSW-NB15	2015	165 461	38	6(5)	8	6.5	8

Here you see details of some of the datasets we used

the more attacks a dataset contains, the more variants

… and the results!

Unsupervised algorithm getting better than supervised, when unknowns increase

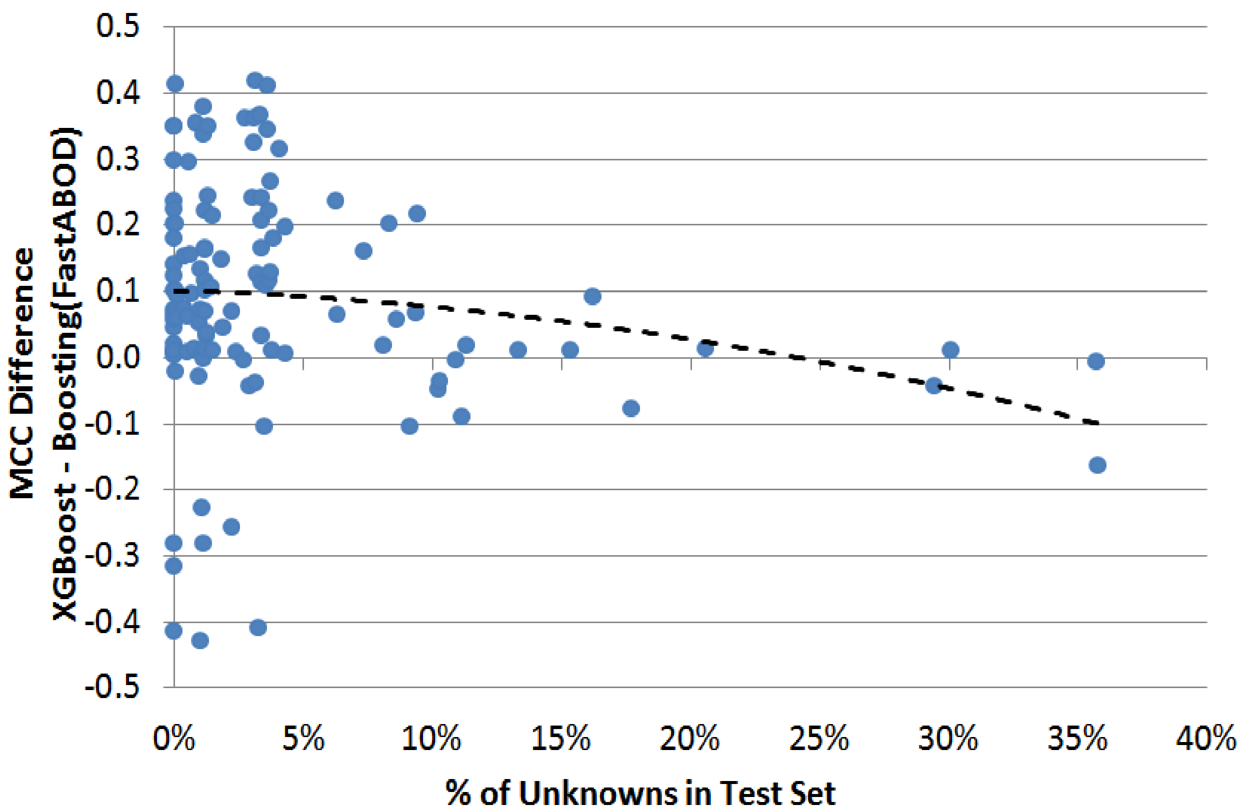

Ensemble: take the best from supervised and unsupervised

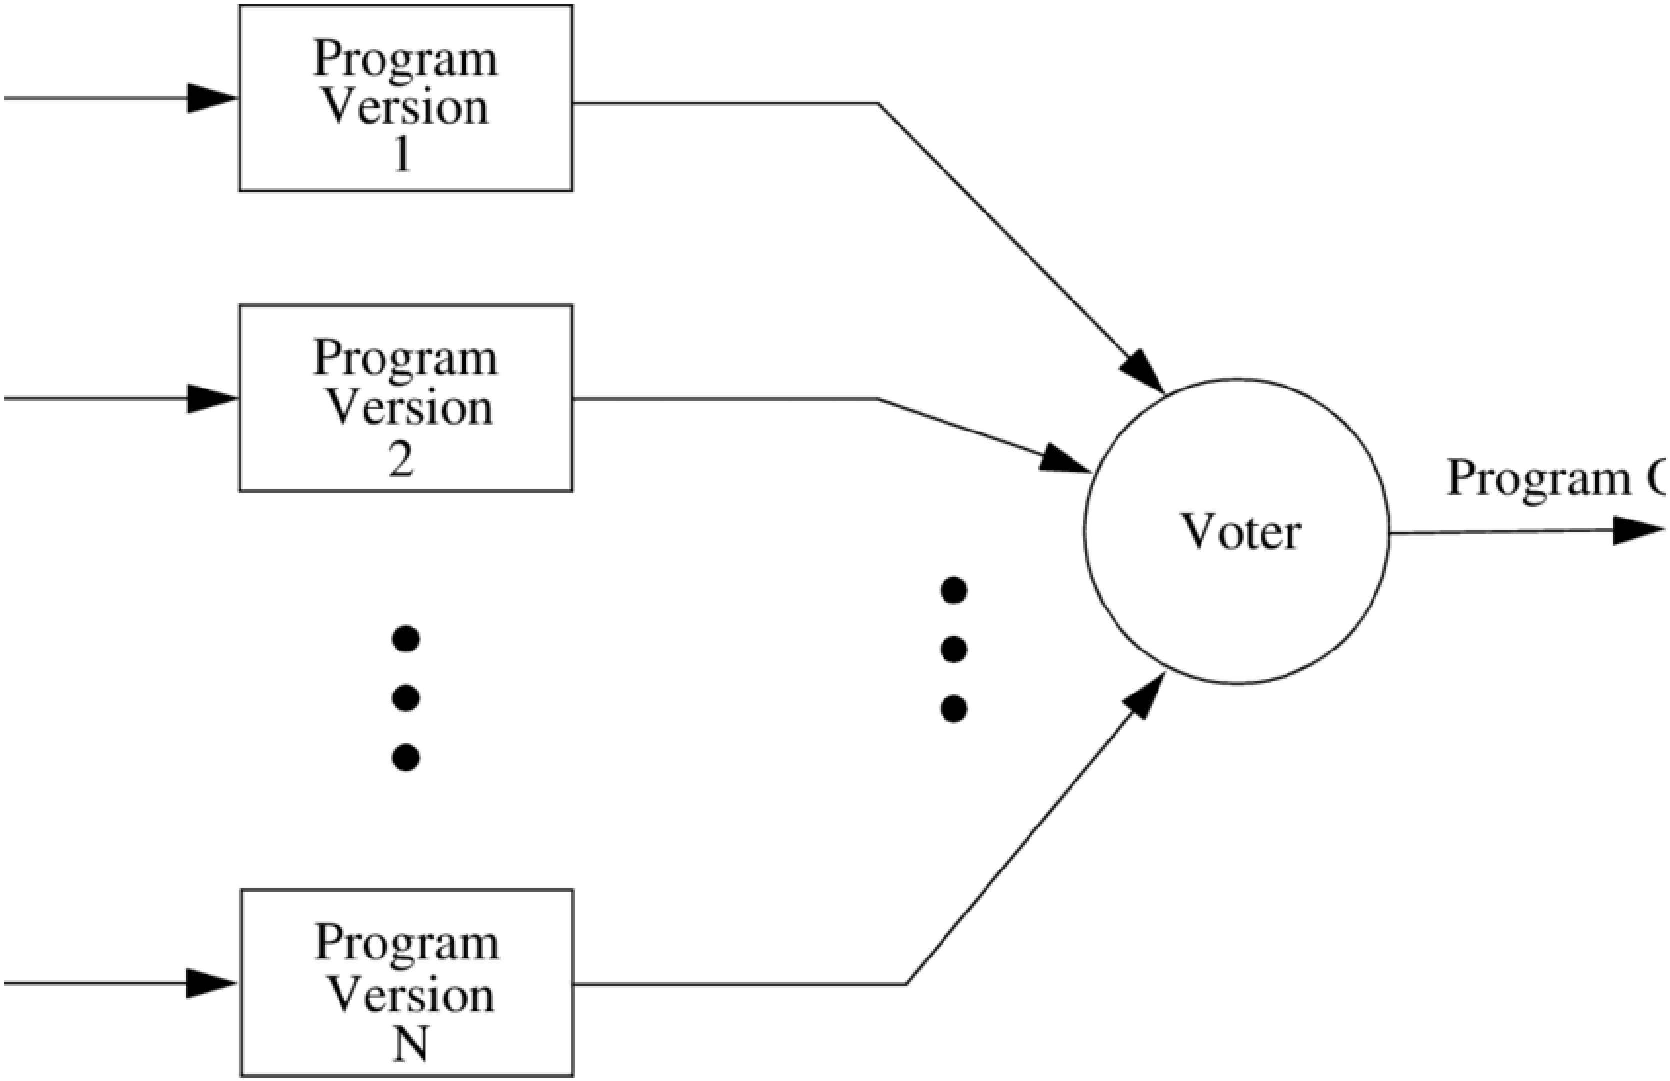

Meta-Learning (I)

Base-learning process: train more learning algorithms, to be used for classification at a first stage

Results of base learners are provided alongside with other features to the ***meta-layer***

Brazdil P, Giraud-Carrier C, Soares C, Vilalta R (2009) Metalearning: applications to data mining. Springer, Berlin.

Bagging

Bagging combines **base-learners of the same type** by submitting bootstrap replicas of the training set

- Individual learners execute the **same algorithm**, but are fed with **different training subsets** created through random sampling with replacement i.e., *Bootstrap AGGregatING*
- The unified result of the ensemble is derived by **majority voting** the individual results of base-learners

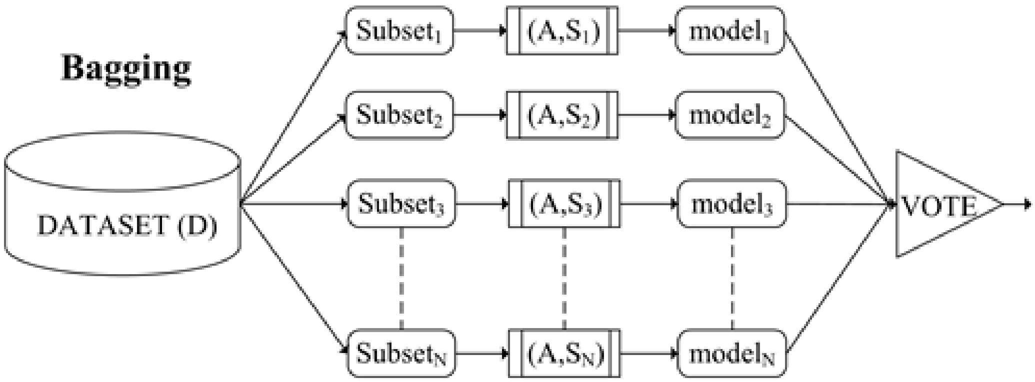

Boosting

Relies on the concept of “Weak Learner” (WL)

- A WL is good in classifying some items, wrong on others
- Subsequent WLs are trained with hard-to classify regions of training set

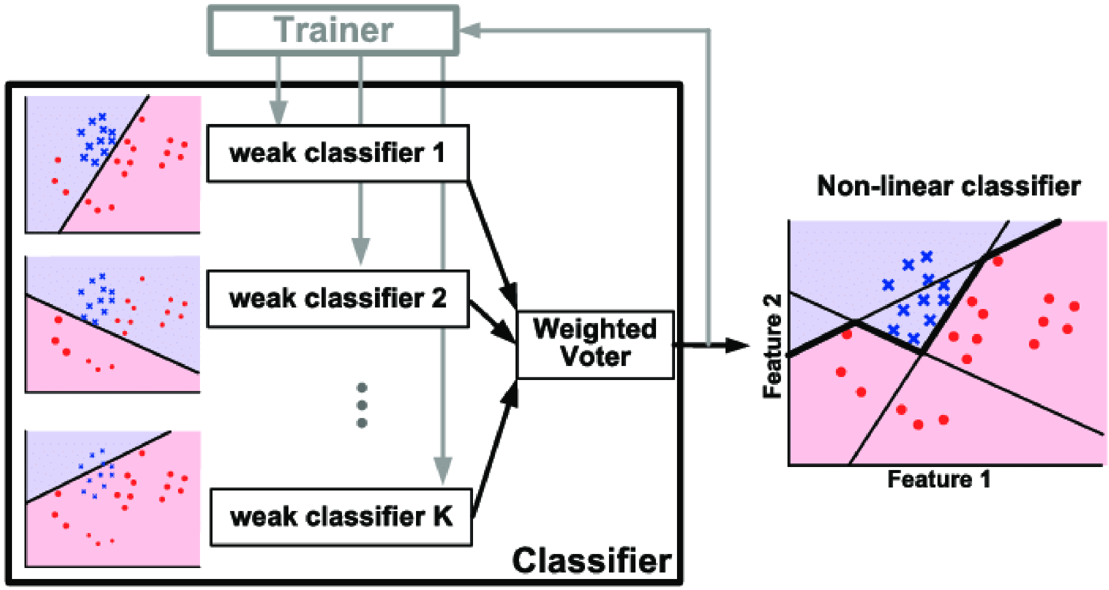

Nowadays, XGBoost (eXtreme Gradient Boosting) is the go-to algorithm for classifying tabular data

Wang, Zhuo, Jintao Zhang, and Naveen Verma. "Realizing low-energy classification systems by implementing matrix multiplication directly within an ADC." *IEEE transactions on biomedical circuits and systems* 9.6 (2015): 825-837.

Stacking different models

Stacking uses yet another machine learner to "vote"

This builds a two-layer structure with

- A base-layer (with diverse base-learners A_1 - A_N), and
- A meta-layer, with a single classifier A_{meta} that delivers a unique result

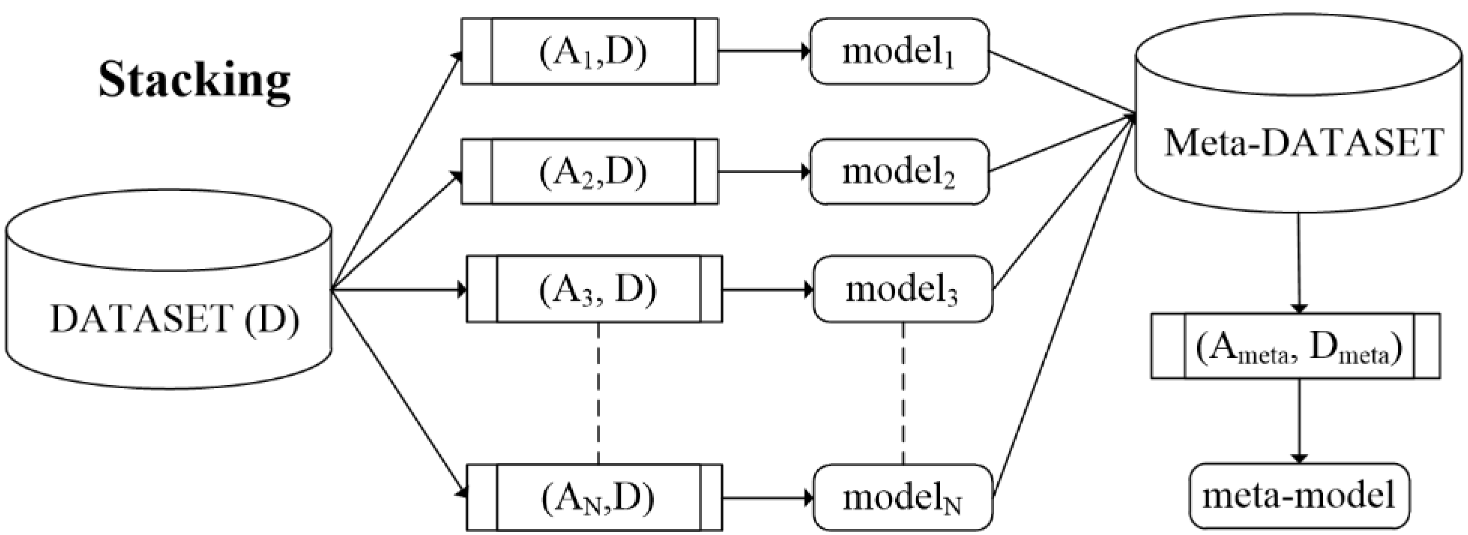

An IDS stacker

A Stacker with

- Unsupervised base-level learners (3, 4, 5)
- A Supervised Meta-level learner (6)

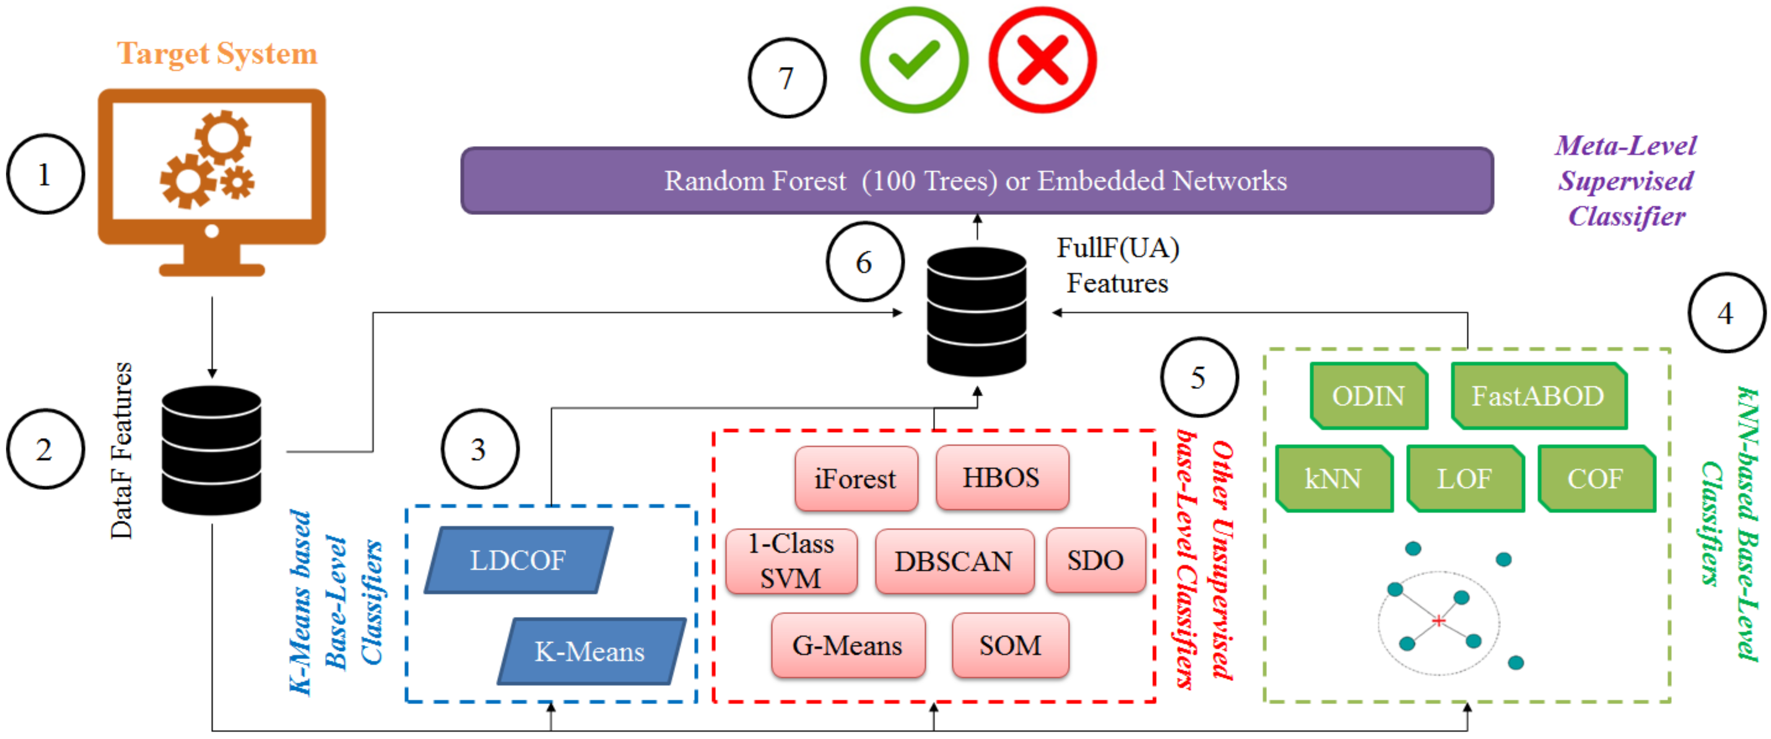

Zoppi, T., Ceccarelli, A. (2021) "Prepare for trouble and make it double! Supervised–Unsupervised stacking for anomaly-based intrusion detection." *Journal of Network and Computer Applications* 189: 103106.

Evaluation of the Stacker

Comparison between MCC Stacker vs **supervised**

Each dataset, we take the best supervised algorithm

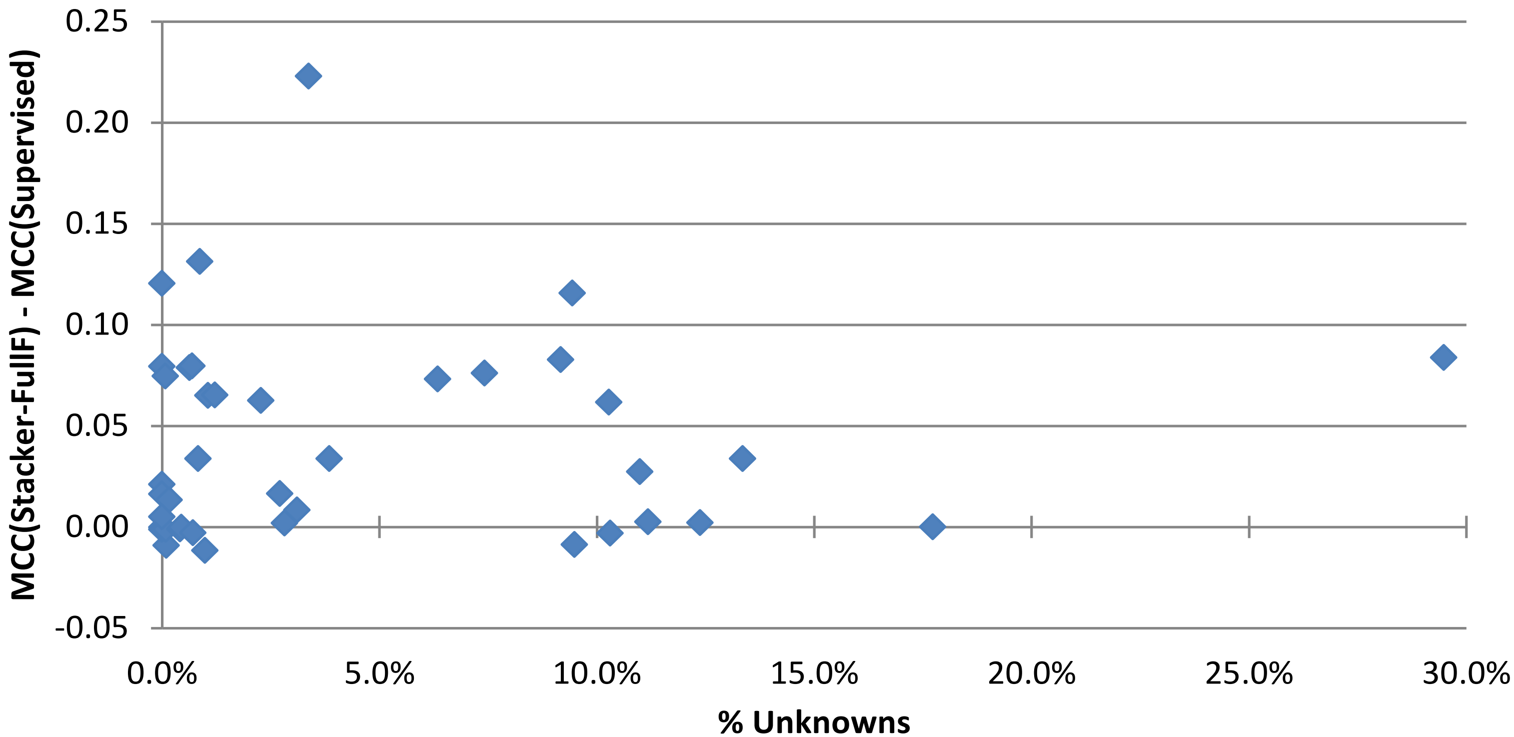

Presentation Outline

Recap on Anomalies and Intrusions

Building an Anomaly-Based Intrusion Detection

Detecting unknowns

What’s next: towards detection of APT

Wrap-Up and Concluding Remarks

Advanced Persistent Threats

Advanced, well-financed attack campaign with a full spectrum of intelligence-gathering techniques.

Persistent, from highly determined and persistent attackers. One of the attackers' goals is maintaining long-term access to the target.

Threats executed by coordinated human actions rather than mindless automated code.

Reconaissance, Scanning ,

Exploitation, Maintaing access

Anomaly detectors for APTs

A shift of perspective:

- not just «detect an attack»,

  but

- interrupt the attack path before the goal is reached

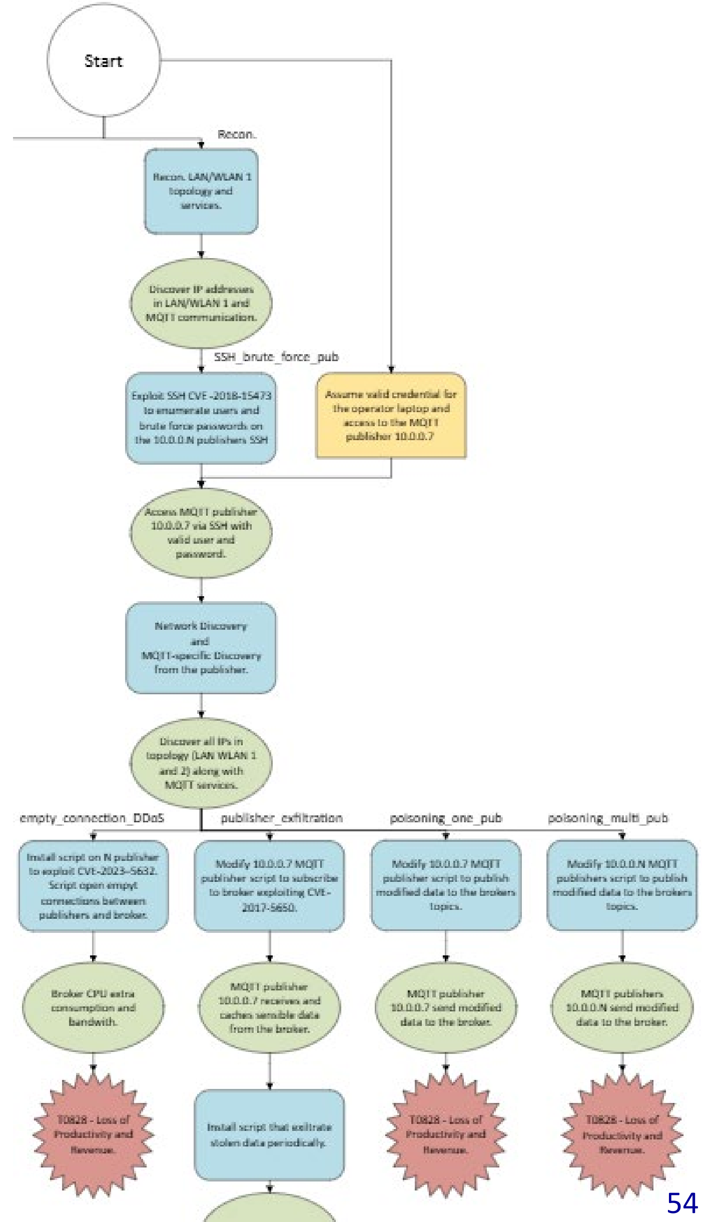

What is missing with respect to everything we have seen:

- Above all, datasets!
- Then, algorithms for time series exists (even if *maybe* not so much applied to IDS *yet*)

(Again another) datasets review

dataset	year	dom	apt	type	data	lat
KDD/NSL-KDD	1999	ent	no	real	net	No
ADFA LD/WD	2014	ent	no	real	log	No
ISCX	2012	ent	no	real	log	No
CICIDS17	2017	Ent	no	real	net	No
CICIDS18	2018	Ent	no	semi	net	No
InSDN	2020	Ent	no	semi	net	No
IoT-IDS	2019	Iot	no	real	net	No
LANL Dataset	2019	Iot	no	real	net	No
ROSPaCe	2024	cps	no	real	net, log	yes
Modbus	2016	cps	no	real	net	No
SWaT	2020	cps	no	semi	net, log	No
BATADAL	2018	cps	no	synth	log	Yes
VASTs	2018	ent	no	semi	net, log	No
DAPT2020	2020	ent	yes	semi	net, log	No
Unraveled	2023	ent	yes	Semi	Net	No
Linux-APT	2024	ent	yes	semi	net, log	No
Next slides	**2025**	**cps**	**yes**	**semi**	**net**	**Yes**

Let’s try to build a dataset

Industrial network traffic dataset DoS/DDoS-MQTT-IoT (publish/subscribe)

Simulate Network environment using DDoShield-IoT

Can replay dataset .pcap file and simulate network normal behavior ← **and we can craft attack!**

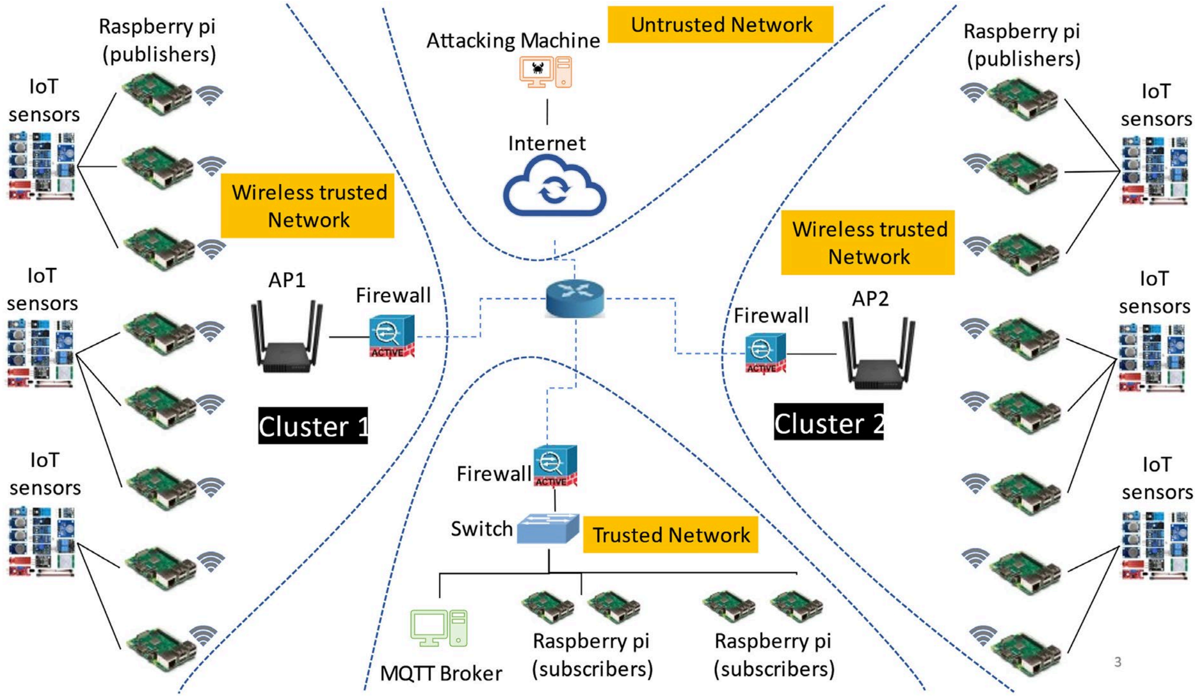

Alatram, Alaa, et al. "DoS/DDoS-MQTT-IoT: A dataset for evaluating intrusions in IoT networks using the MQTT protocol." *Computer Networks* 231 (2023): 109809.

De Vivo, Simona, et al. "DDoShield-IoT: A Testbed for Simulating and Lightweight Detection of IoT Botnet DDoS Attacks." *2024 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. IEEE, 2024.

Design and implement the attack paths

MITRE | ATT&CK

CVE

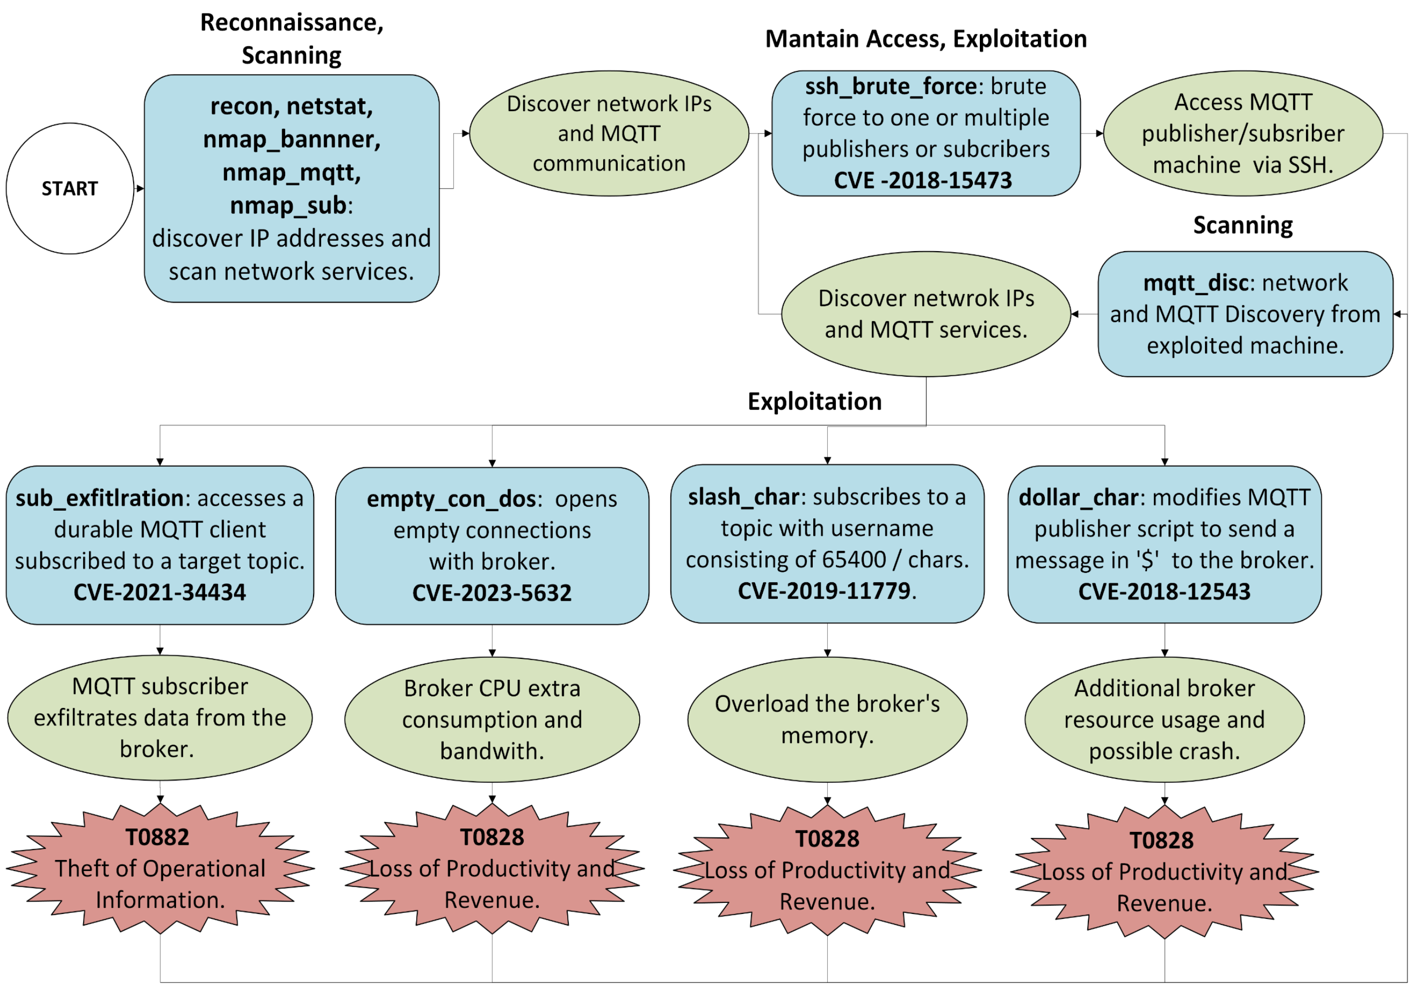

MUR FLEGREA - *Federated Learning for Generative Emulation of Advanced Persistent Threats*

Dataset composition

Using the replay functionality of DDoShield-IoT, we re-create normal traffic; plus, we inject attacks in the simulated system, and we log attack data

We merge normal+attack data, to create attack paths

Dataset composed of 11.904.459 packets (88% normal data)

attack	#	average duration	minimum duration	maximum duration	average length	minimum length	maximum length
empty_con_dos	83	149.68	84.71	524.74	675.05	30	2120.5
dollar_char	71	601.5	86	1260.97	5468.16	637	11448
nmap_10	15	1040.36	1034.69	1045.12	44417.37	43927	44720
ssh_brute	24	140.41	118.72	194.36	2219.91	110	2706
nmap_banner	24	253.37	244.99	258.76	1181.47	612	2266
nmap_mqtt	24	258.95	250.69	271.18	1001.65	87	2378
slash_char	60	21.05	14.6	26.15	537.3	395	721
nmap_sub	10	61.66	55.65	67.11	627.5	237	770
netstat	28	56.84	38.70	67.88	233.27	22	370
sub_exfiltration	10	18.23	13.24	20.71	2355.6	2224	2425

Train-test; analyze results

XGBoost: not too good but just our first try

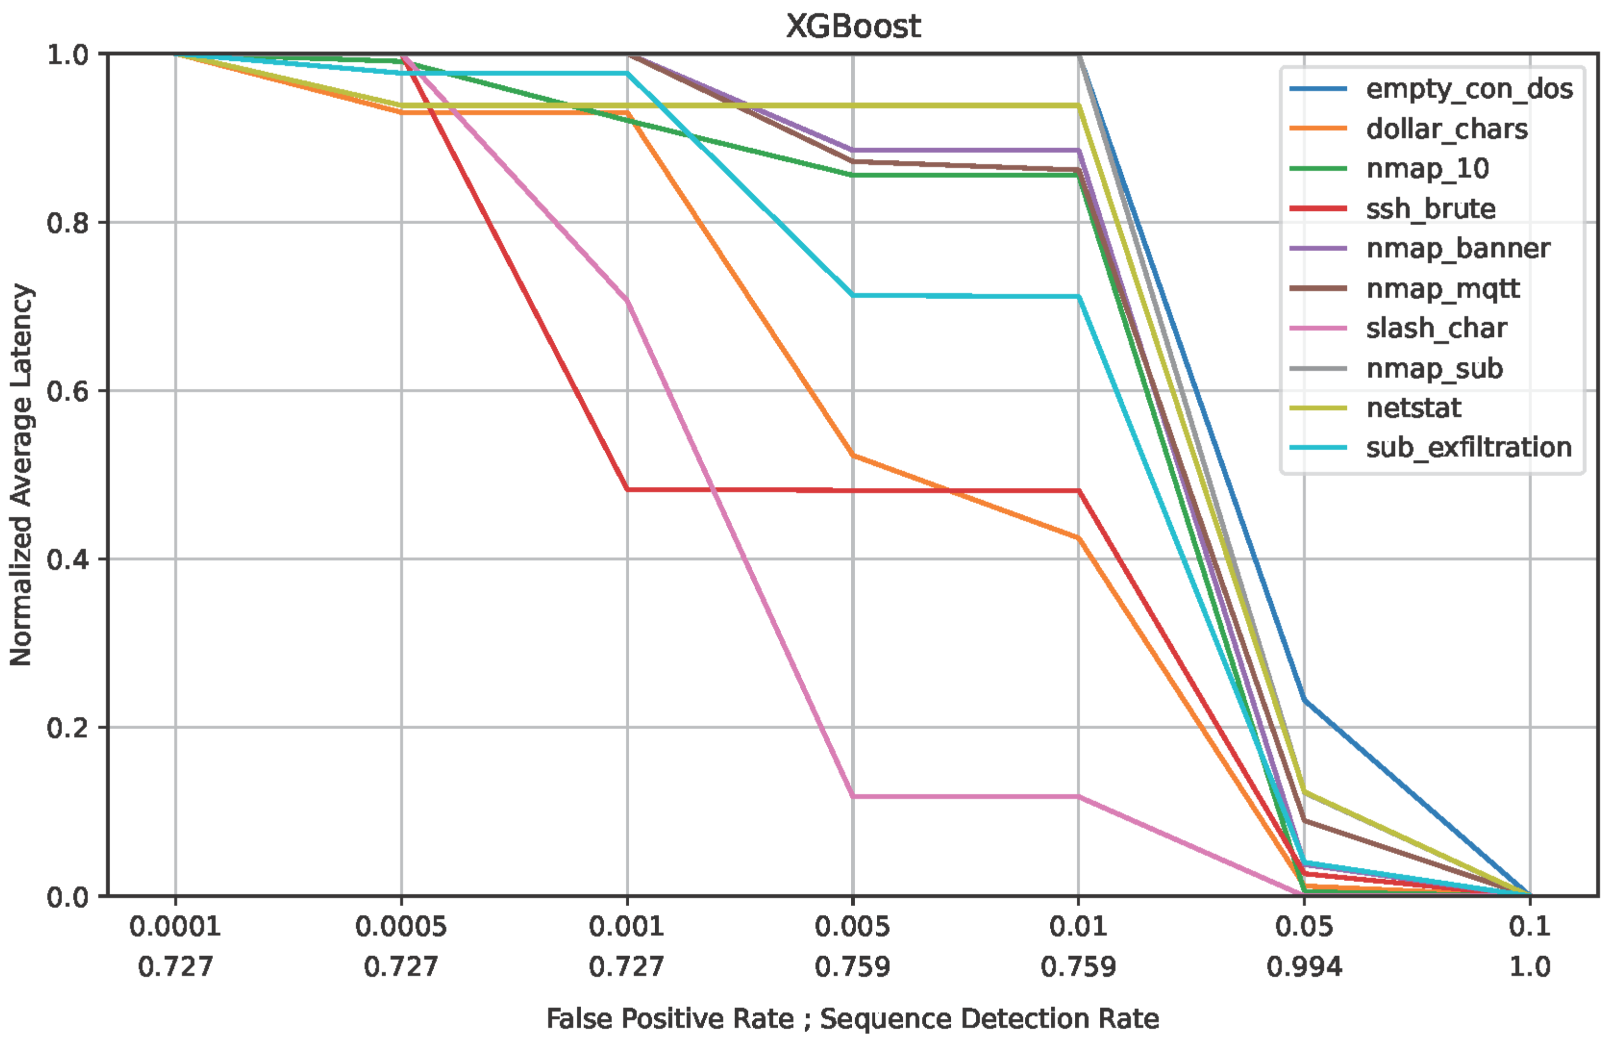

Presentation Outline

Recap on Anomalies and Intrusions

Building an Anomaly-Based Intrusion Detection

Detecting unknowns

What's next: towards detection of APT

Wrap-Up and Concluding Remarks

(Finally!) Wrapping Up…

This talk went through different ways to build anomaly-based IDS

- Using ensembles of algorithms
- Accounting for zero-day attacks
- Showing new frontiers for IDSs

Future Works

We are always open to ideas and collaborations

- And criticisms as well!

Overall, we feel that unknown and complexity will become more and more relevant in the near future

- Systems are more and more complex, thus a complete characterization of errors / attacks and related paths becomes impossible!

So… be prepared to fight complex attacks!

- Maybe using ensembles?

Selection of our recent works (mentioned through the talk)

– Puccetti, T., et al. (2024) "ROSPaCe: Intrusion Detection Dataset for a ROS2-Based Cyber-Physical System and IoT Networks." *Scientific Data* 11.1 (2024): 481.
– Zoppi, T., et al. (2024) "Anomaly-based error and intrusion detection in tabular data: No DNN outperforms tree-based classifiers." *Future Generation Computer Systems* 160: 951-965.
– Zoppi, T., et al. (2023) “Which algorithm can detect unknown attacks? Comparison of supervised, unsupervised and meta-learning algorithms for intrusion detection”, *Computers & Security*, 127, 103107.
– Zoppi, T., Ceccarelli, A. (2021) "Prepare for trouble and make it double! Supervised–Unsupervised stacking for anomaly-based intrusion detection." *Journal of Network and Computer Applications* 189: 103106.
– Zoppi, T., et al. (2021) “Unsupervised Algorithms to Detect Zero-Day Attacks: Strategy and Application” *IEEE Access*, 9, 90603-90615
– Zoppi, T., et al. (2021) "Unsupervised anomaly detectors to detect intrusions in the current threat landscape" *ACM/IMS Transactions on Data Science* 2.2: 1-26.

Get in touch!

https://rcl.unifi.it
resilientcomputinglab@gmail.com

Andrea Ceccarelli

andrea.ceccarelli@unifi.it

International cooperations

MsC studies

PhD Grants

Postdocs