# **Cyber Security** Risk and ML applications

Chris G. Willcocks Durham University

# **Motivation**

Why this lecture?

- Risk calculations can be very difficult
  - Equip with basic tools to do it
- Many modern security applications use machine learning
- Many final year security projects are based on machine learning

Why machine learning in security?

- Machine learning is function fitting
  - Fast (packet inspection)
  - Probabilistic
  - Cheap
  - Generalises well to new scenarios/threats



# Lecture content





### Covered today

#### Risk

- Qualitative risk
- Quantitative risk
  - SLE, ARO, ALE, Bayesian Risk

## ML applications

- Security datasets
- Discriminative security models
- Threat detection
- Conditional generative models and metalearning security tasks
- PassGAN & briefly ethical research



#### **Definition:** Risk

The definition of risk varies based on application, but it is generally defined:

risk = asset value  $\cdot p$ (threat occurance)  $\cdot$  severity

Two ways to compute risk:

- Quantitative risk
- Qualitative risk

What about time?

Source: The Security Risk Assessment Handbook

#### **Qualitative Risk**

Traffic light grid gives immediate impression of where effort should be focused.

#### Advantages:

- Simple
- Not much effort
- Easy to understand

#### Disadvantages:

- Subjective results
- Subjective asset value
- Subjective recommendations
- Difficult to track improvements

	G	Catastrophic	5	5	10	15		
	ev	Significant	4	4	8	12		
	e r	Moderate	3	3	6	9	12	15
	i t	Low	2	2	4	6	8	10
	У	Negligible	1	1	2	3	4	5
Catastrophic		STOP		1	2	3	4	5
Jnacceptable		URGENT ACTION		Improbable	Remote	Occasional	Probable	Frequent
Undesirable		ACTION						
Acceptable		MONITOR		Likelihood				
Desirable	NO ACTION			Likeiillood				

Una

#### **Definition:** SLE

Single Loss Expectancy (SLE)

This is the amount that would be lost in a single occurrence of an incident:

 $SLE = asset value \cdot exposure factor$ 



#### **Definition:** ARO and ALE

Annual Rate of Occurrence (ARO) Annual Loss Expectancy (ALE)

Consider the annual rate of events.

annual loss expectancy =  $SLE \cdot ARO$ 



#### **Small example** (in practice this would be much larger)

Asset	Security Goal	Vulnerability	<b>SLE</b> (£/incident)	<b>ARO</b> (incidents/yr)	<b>ALE</b> (£/year)
Confidential emails	Confidentiality	Hacker MITM	£100,000	0.5	£50,000
Non-confidenti al emails (business details)	Integrity Reputation	Employee breach	£10,000	3	£30,000
	Availability	DDoS	£20,000	5	£100,000
Database	Integrity	Hardware failure	£10,000	0.5	£5,000
	Confidentiality	Hacker breach	£50,000	0.2	£10,000

#### **Definition:** safeguard value

Quantifying the value of safeguarding the risk (the value of the countermeasure):

Safeguard value = (ALE before - ALE after) - annual cost of countermeasure

Vulnerability	Counter- measure	<b>ALE Before</b> (£/year)	<b>ALE After</b> (£/year)	<b>Countermeasure</b> (£/year)	<b>Safeguard</b> value (£/year)
Phishing	Security training	£70,000	£5,000	£5,000	£60,000
DDoS	24/7 Network monitoring	£100,000	£10,000	£70,000	£20,000
Physical break in 24/7 CCTV + physical security		£10,000	£1,000	£80,000	- £71,000

#### Advantages

- Objective
- Expressed as a real number
- Help make sensible decisions
- Easy to understand
- Decisions are traceable
- Credible
- Basis for cost-benefit analysis

#### Disadvantages

- Complex
- Confusing to non-technical readers, sometimes even resulting in a lack of trust
- False sense of accuracy

Source: The Security Risk Assessment Handbook

#### Expectations & Monte Carlo sampling

Quantitative risk assessment requires expected value of the annual rate of occurrence. We can gather this empirically but it's very sensitive to sampling process (e.g. location, time, threat conditions).

We can improve our data with better priors, for example a Bayesian risk assessment with conditional probability:





$$p(\text{covid} \mid \text{symptoms}) = \frac{p(\text{symptoms} \mid \text{covid}) \cdot p(\text{covid})}{p(\text{symptoms})}$$

#### Security Datasets

There's a huge amount of public security datasets available:

- 1) <u>https://github.com/shramos/Awe</u> <u>some-Cybersecurity-Datasets</u>
- 2) <u>https://github.com/jivoi/awesome</u> <u>-ml-for-cybersecurity</u>

But what can we do with this data?

#### Three types of ML model

Discriminative models:  $p(Y \mid X)$ Conditional generative models:  $p(X \mid Y)$ Generative models: p(X,Y)

#### **Discriminative model**



#### while(training):

```
X,T = random.sample(dataset)
Y = DNN(X)
loss = ((Y-T)**2).mean() # error
loss.backward() # calculate grads
DNN.params -= 0.01*DNN.grad # optimise
```

#### **Example:** sentiment analysis

https://huggingface.co/distilbert-baseuncased-finetuned-sst-2-english

DNN("I love kittens")  $\rightarrow$  positive DNN("I hate people!")  $\rightarrow$  negative



# **Threat object detection**

#### **Object detection**

X = input images
Y = region proposals (boxes)



- Human operators get distracted with cluttered X-ray (miss threats)
- Every commercial flight has <u>certified explosive detection</u> <u>systems (EDS)</u>

**Example:** baggage security

Data (124.78 GB)

Link to Kaggle competition

State-of-the-art detection models:

GitHub link to YOLOv4 CSP

Note: the above repo links to sub repositories



Towards Automatic Threat Detection: A Survey of Advances of Deep Learning within X-ray Security Imaging (Akcay and Breckon)

#### Extracting network features



#### **Counterfeit classification**

#### Check whether products are genuine



#### SLOW SLOW SLOW SLOW SLOW SLOW....



#### ML models



#### Inferring new unseen risks & tasks

Consider a massive model trained on the internet that tries to predict the next token given the previous words.



Estimate the next most likely thing, "he was motivated by..."

#### Large-scale language model

#### Try zero-shot examples with bart:

- 1) Come buy our product! sales, phishing, spam
- Your credit card is about to expire!
   Visit this link to renew it: http://dodgeylink.com phishing, sales, spam
- John Doe robbed Jess in Newcastle in 2024 and then stole jewellery from Goldsmiths.

criminal, innocent, neutral





#### Adversarial models

Adversarial models can be used to generate more criminal data.

#### **Example:**

PassGAN (considered unethical)



Also make virus code, harmful traffic...

#### Remember which side you're on



Just because you can technically do it, doesn't mean it's ethical research.

Just because you can build a dangerous open source weapon, doesn't mean you should.

#### Key points

- The threat landscape is always evolving
- Remember how easy most tools/ threats are (only a little effort)
- Security covers all levels and infrastructure of a system
  - The weakest link
- Hierarchically assess the risk
  - Understand the **enemy**
  - Understand the **platform**
  - Understand the **people**
- Network with the broader security community <u>and practice</u>

#### Key points

- Don't be careless or manage in a way that promotes carelessness
- KISS!

