Machine Learning in Cyber Security

May 7, 2021

Benjamin Süß

Warning

Warning

- Machine Learning won't solve all problems of Cyber Security.
- Machine Learning can't replace all tasks of Cyber Security.
- Machine Learning is just another tool for Cyber Security.
- This presentation only covers a tiny part of Machine Learning.

Table of contents

Introduction

Introductory book "Machine Learning"

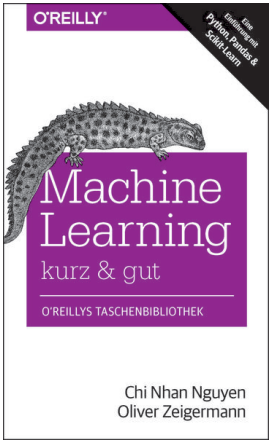

- ISBN: 978-3-96009-052-6
- 180 pages (German)
- €15

"Applying Machine Learning to Cyber Security"

IT Security Conference

The Honeynet Project Workshop 2019

Innsbruck, Austria July 1st–3rd, 2019

The Honeynet Project Workshop is a technical international security conference focused on deception and cyber intelligence. The aim of the workshop is to bring the security commun to learn the tools, tactics and motives involved in computer and network attacks, and share learned.

Every year, we select a country from one of our 40 chapters to meet.
This year, we're coming to Innsbruck. Please join us for:

- **One day of briefings** on Monday.
- **Two days of high-end trainings** and hands-on experience on Tuesday and Wednes
- **A lot of fun!**

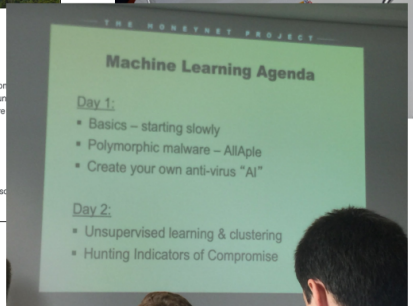

Creating a model

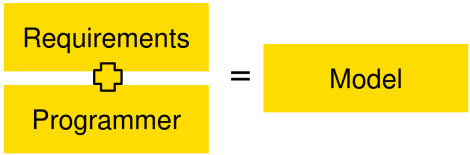

The model is some code.

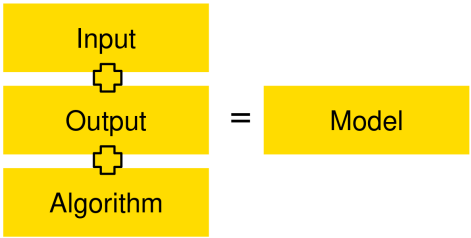

The model is ???

Setup for exercises

- 2019 workshop: Jupyter Notebook + Python 2, including `pandas`, `scikit-learn` and `matplotlib`
- Nov 2020: JupyterLab + Python 3, including `pandas`, `scikit-learn` and `matplotlib`

Terms in Machine Learning

Features and labels

Definition (Feature)

- Attributes of an object
- Like a "fingerprint"
- Part of the input data

Definition (Label)

- Class of an object
- Part of the output data

High bias and high variance

Definition (High bias)

- Model is too unspecific
- aka "underfitting"

Definition (High variance)

- Model is too specific
- aka "overfitting"

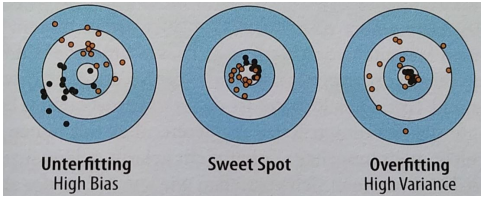

Supervised vs. unsupervised learning

Definition (Supervised learning)

- Input data is known
- Output data is **known**
- The magic AI creates a model **based on training**
- We compare data processed by the model with pre-classified data (**accuracy**)
- We tune parameters to improve the model

Definition (Unsupervised learning)

- Input data is known
- Output data is **unknown**
- The magic AI creates a model **based on magic**
- We look at the output data, created by the model
- We tune parameters to (hopefully) see new things

Supervised learning

Some types of supervised learning

We know the input and output data!

- **Classification**: A simple binary decision (This is a dog vs. This isn't a dog).
- **Regression**: Learning a function by evaluating points of it (What is the weight of the dog in the picture?)

How to choose correctly?

- Output is discrete data (e.g., 0 or 1) $\rightarrow$ (Likely) classification
- Output is continuous data (e.g., any real number) $\rightarrow$ (Likely) regression

Supervised learning in Cyber Security

- Classification: Malware detection, Spam filtering
- Regression: Risk assessment

Classification exercises

Use case: **Malware detection**

- Exercises 1.x: Introduction to SVC
- Exercises 2.x: SVC and real data
- Exercises 3.x: Malware classification from A to Z

What to do?

1. Get or prepare a clean data set.
2. Split the data set into training data and testing data.
3. Use cross validation to validate your results.
4. Select relevant features.
5. Use all samples you can get.
6. Fine tune your results (probabilities, thresholds, ROC curve).

Conclusions

- Provide clean input data and select relevant features.
- Choose classifiers that fit your use case.
- Validate your results.
- Fine tune!

Unsupervised learning

Some types of unsupervised learning

We know the input but not the output data!

- **Clustering**: Sorting data based on similar properties (e.g., big dogs, small dogs)
- **Principal component analysis (PCA)**: Grouping properties to reduce complexity without losing information (e.g., length + width + height $\rightarrow$ size)
- **Anomaly detection**:
  - Outlier detection: Detecting something that doesn't belong to the others (e.g., cat picture in a set of dog pictures)
  - Novelty detection: Detecting something new when there are no outliers (e.g., we look at dog pictures to find new information)

Unsupervised learning in Cyber Security

- Clustering: Malware labelling, Log analysis, Attribution on forensics
- PCA: Object recognition in forensics
- Anomaly detection: Intrusion detection, (obviously) anomaly detection

Unsupervised learning exercises

Use cases: **Cluster unknown data, find outliers, process text**

- Exercises 4.1 – 4.3: Clustering with K-means
- Exercises 4.4 – 4.6: Outlier detection with LocalOutlierFactor and more
- Exercise 5.0: Introduction to CountVectorizer
- Exercises 5.1 – 5.3: Tokenizing and clustering + fine tuning

Conclusions

- Machine learning does something in unsupervised learning. The result may be good or wrong.
- The quality of this approach is subjective.
- Validation is necessary but not always possible or easy.
- Applying unsupervised learning on unknown data may be helpful however, automatizing this might be impossible.

Things to remember

7 things to remember

1. Use clean input data
2. Select relevant features
3. Look for over- and underfitting
4. Use more input data to improve the model
5. Use an algorithm that fits your use case
6. Fine tune your approach
7. Validate your results