ALFRED, FIND THE ATTACKER

A primer on AI & ML applications in the IT Security Domain

WHOAMI

Matthias Meidinger

Network Engineer

M.Sc. Student

🐦 @mat_zilla

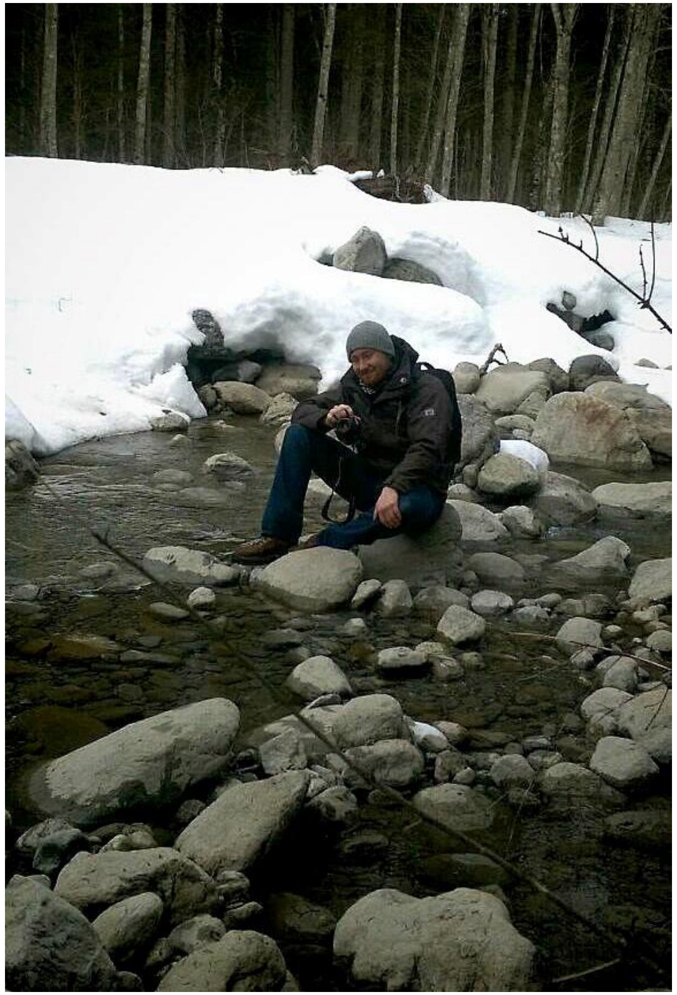

DISCLAIMER

„What is it“, not „How does it work“

Simplified Concepts

Details? Ask me!

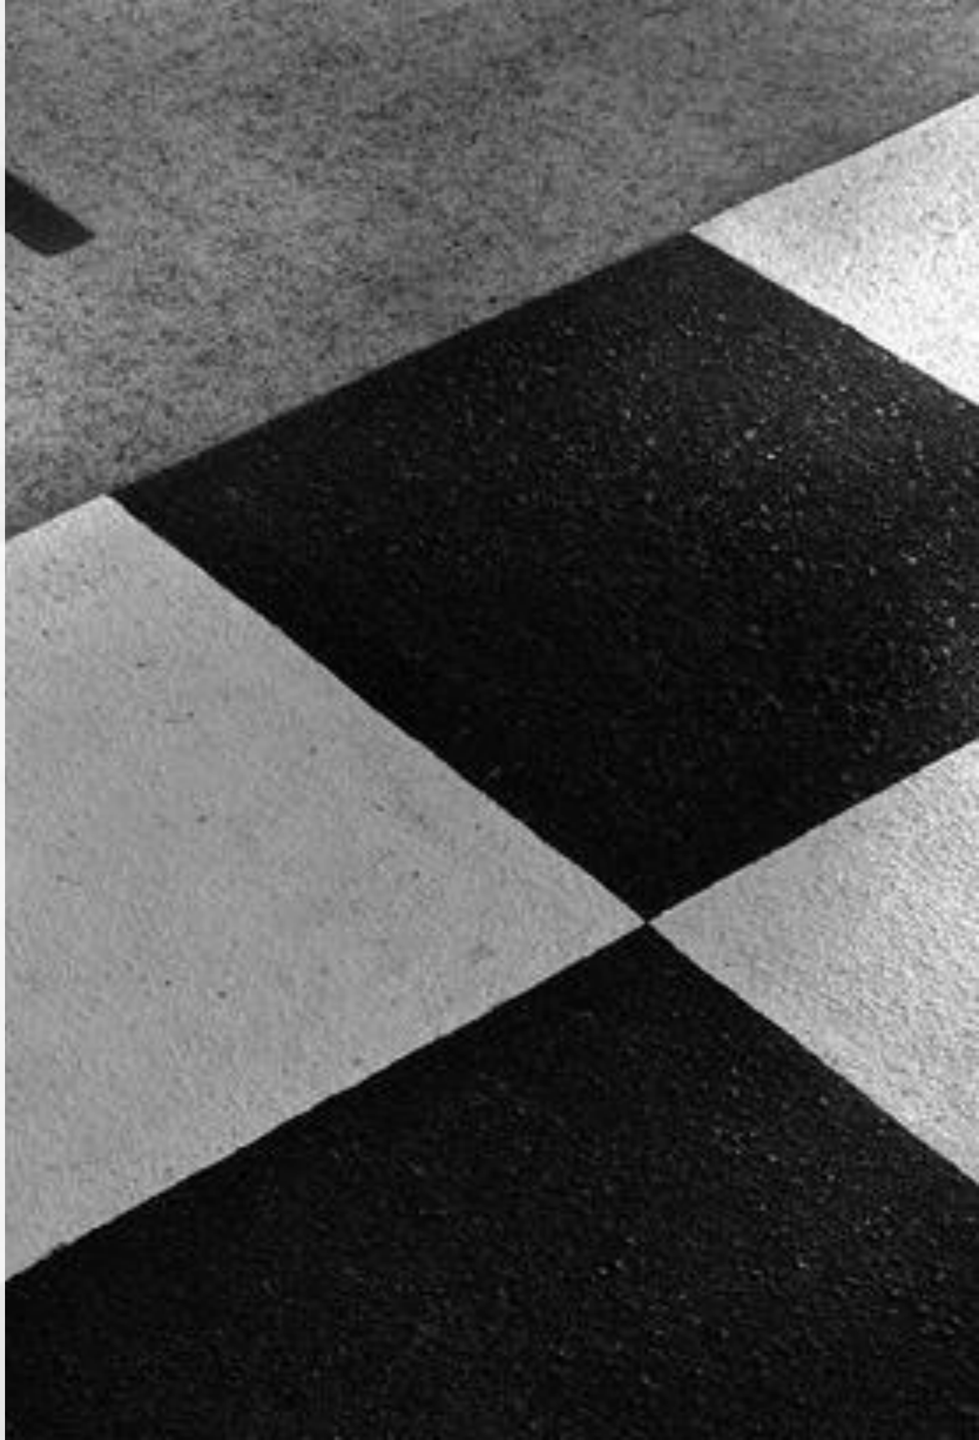

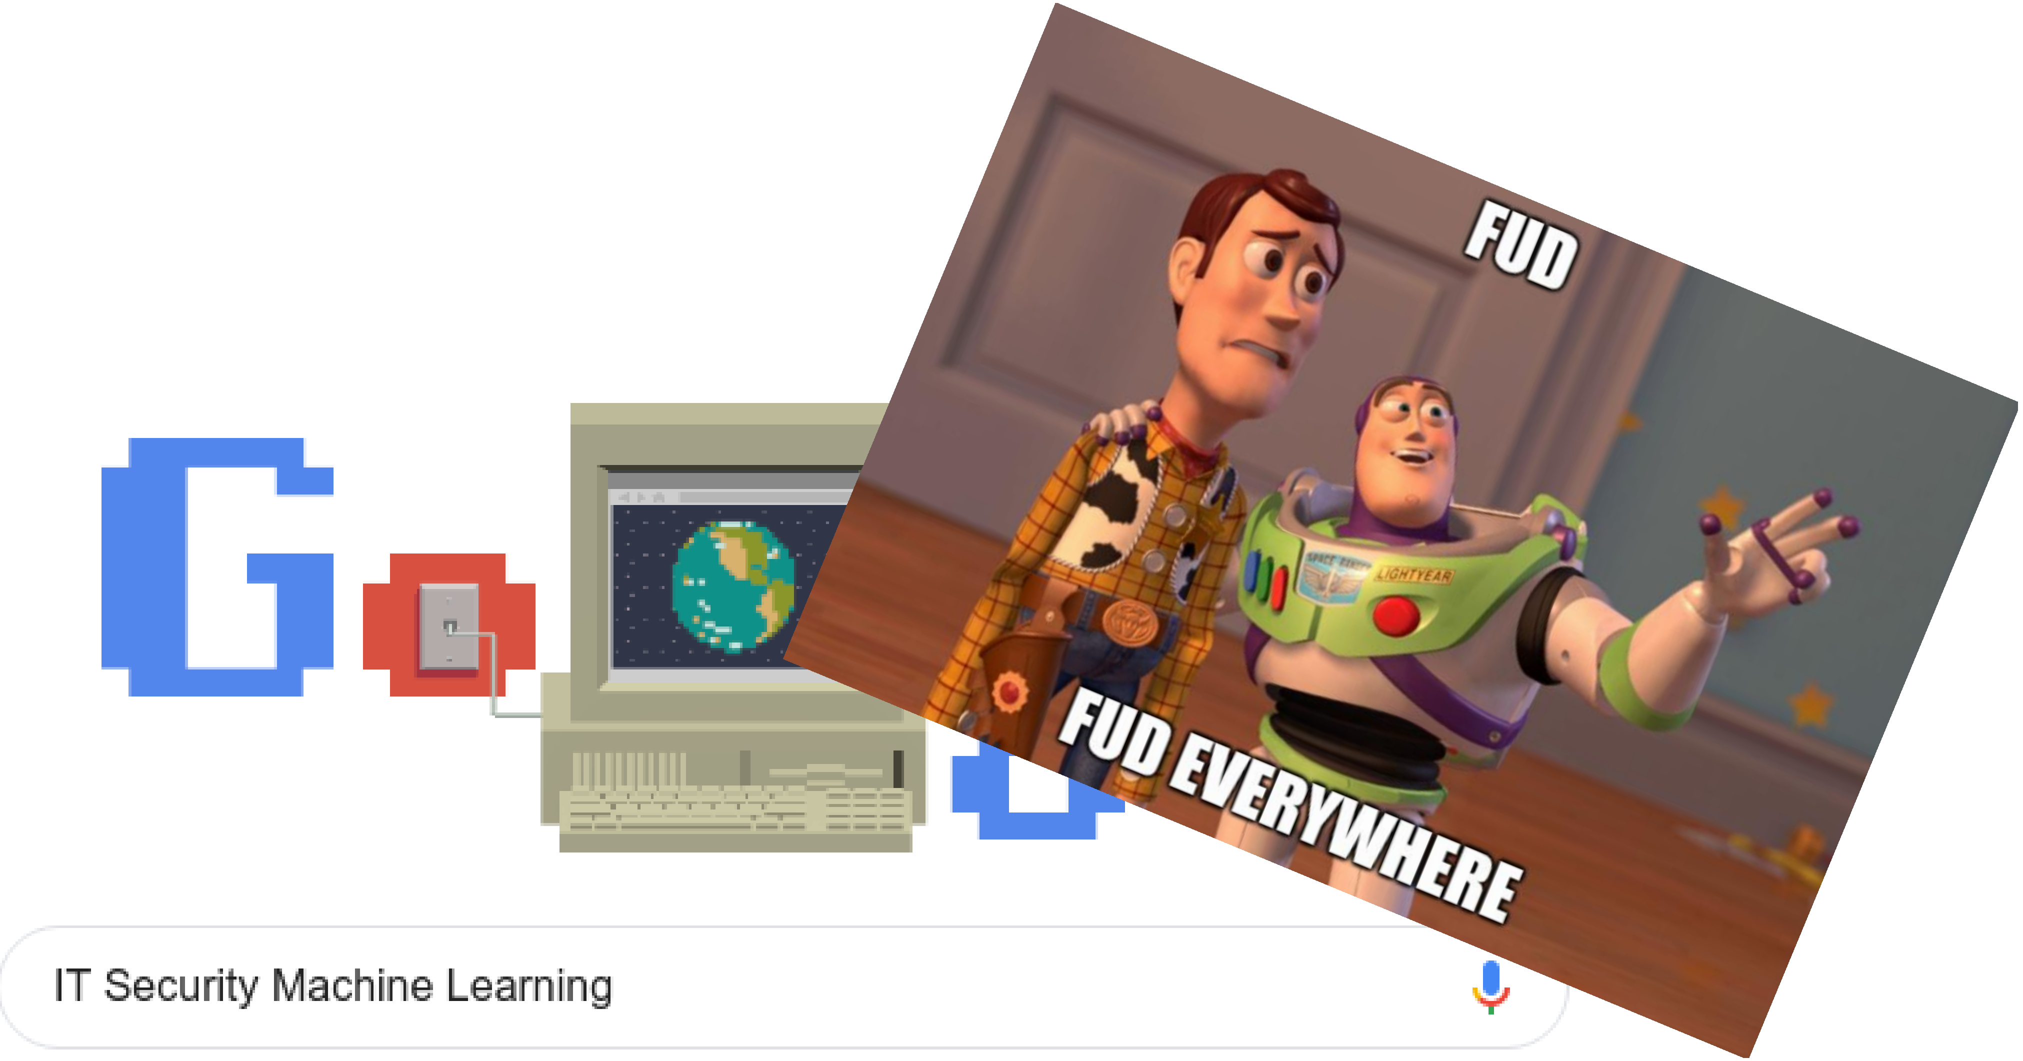

IT Security Machine Learning

“Cyber AI across the cloud, enterprise, and industrial”

“(…) built not just using AI, but entirely from AI.”

“AI-powered technology that answers the question ‘Am I under attack?’ proactively”

“(…) delivering superior AI-native fraud detection and claims handling approaches via our SaaS+ model.”

“AI driven technology prevents attacks before they can damage your devices(…)

BASICS

DEEP… WHAT?

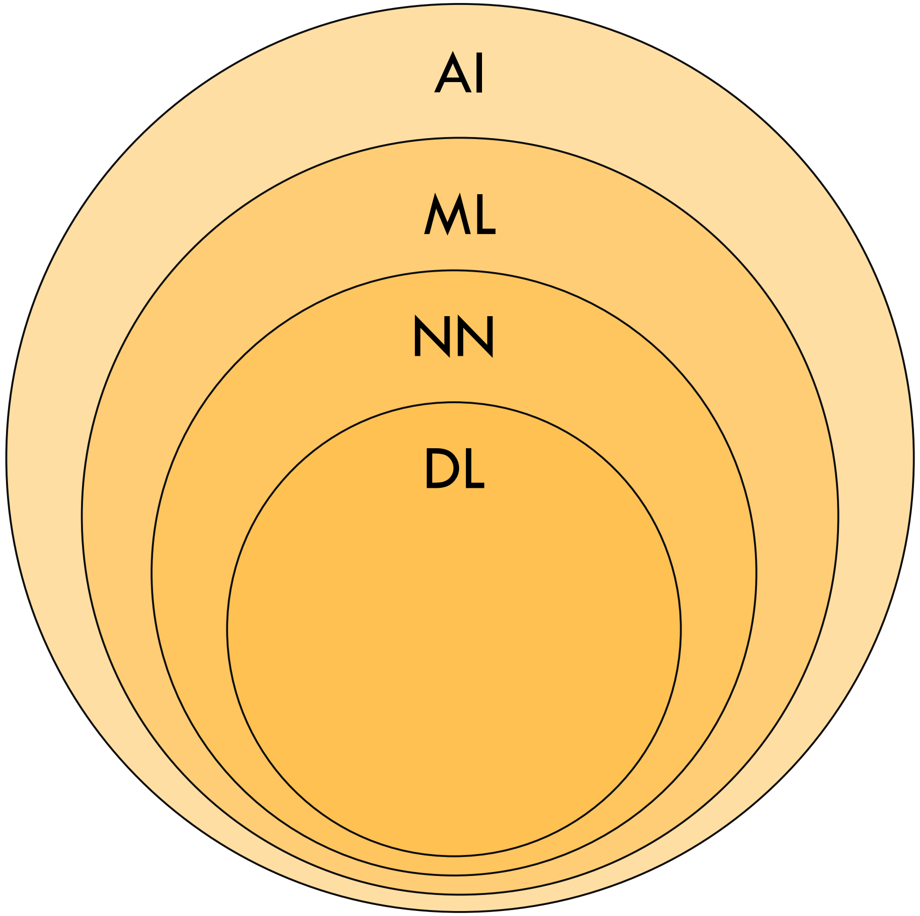

- Artificial Intelligence
  - Interpret & Learn from data
- Machine Learning
  - Self-driven Pattern recognition
- Neural Networks
  - Inspired by Neurons, layered
- Deep Learning
  - NNs with many Layers

TIMELINE

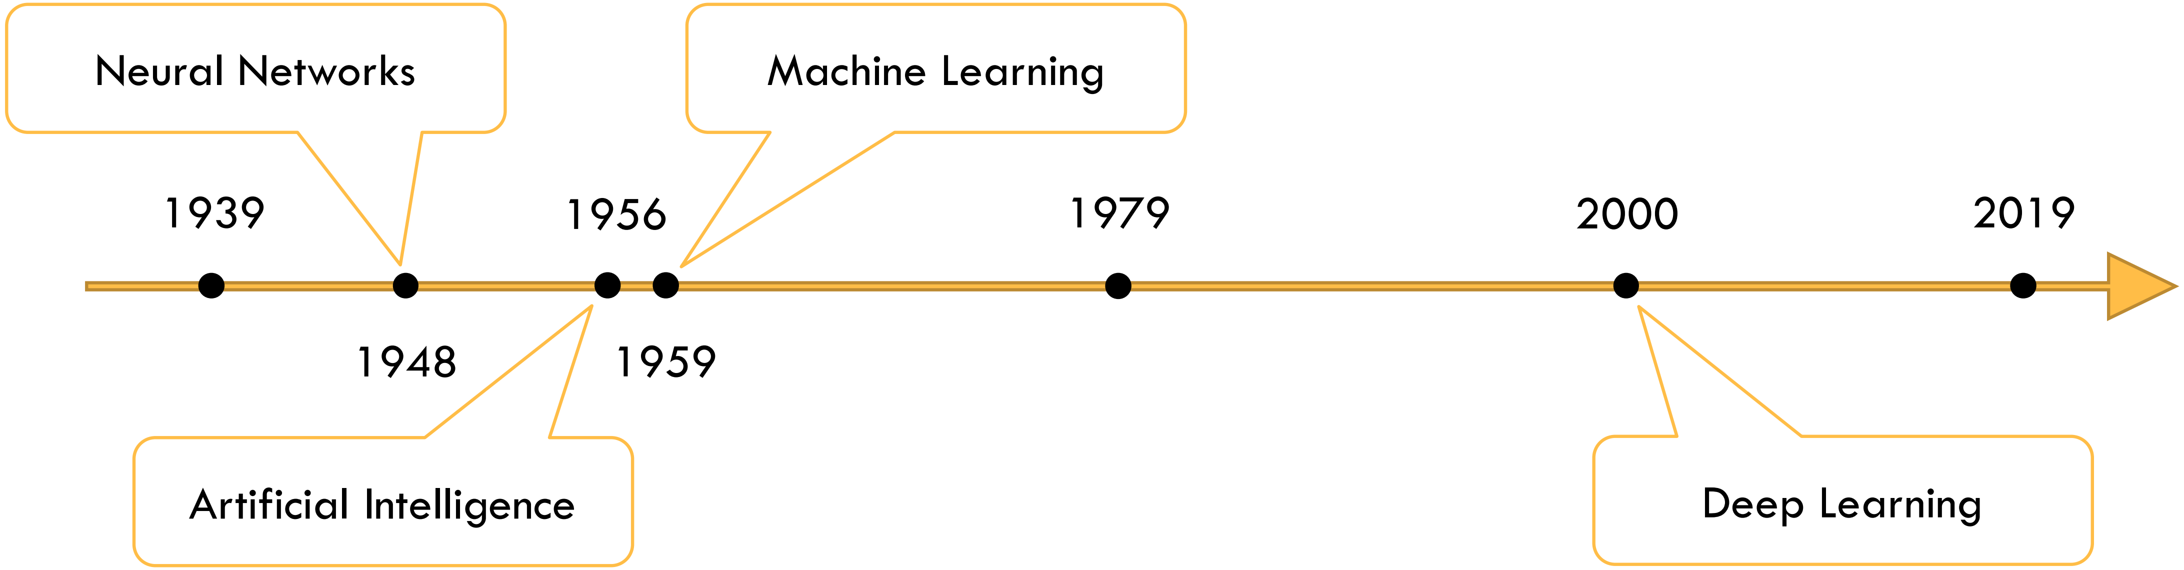

Neural Networks
Machine Learning
1939
1956
1979
2000
2019
1948
1959
Artificial Intelligence
Deep Learning

METRICS

True Positive Rate: How many correct items were detected? (Detection Rate)

False Positive Rate: How many false alerts?

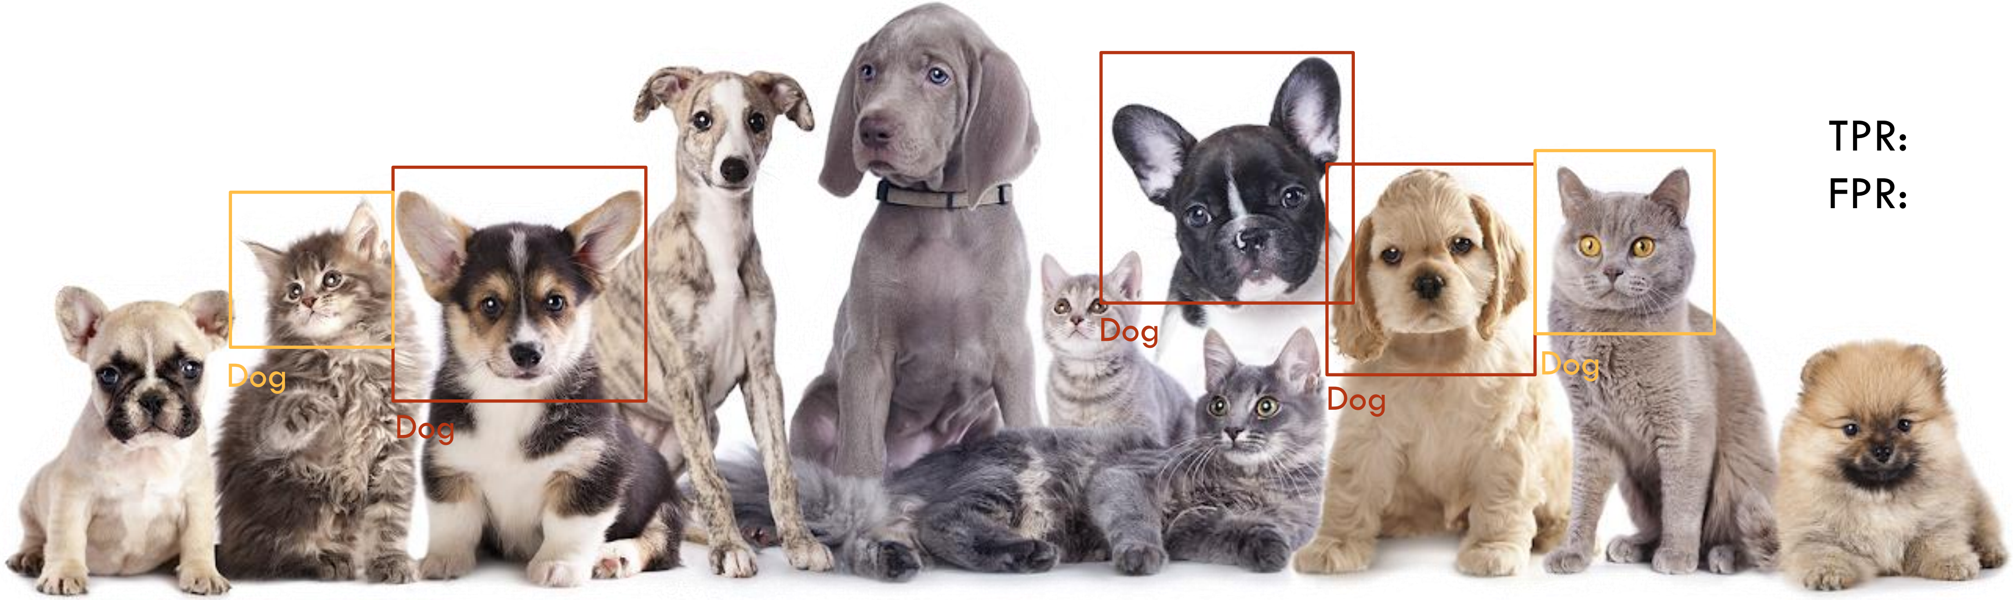

TPR: 3/7 (43%)

FPR: 2/4 (50%)

7 Dogs & 4 Cats

CLASSIFICATION & REGRESSION

Predict a label

Spam vs. Ham

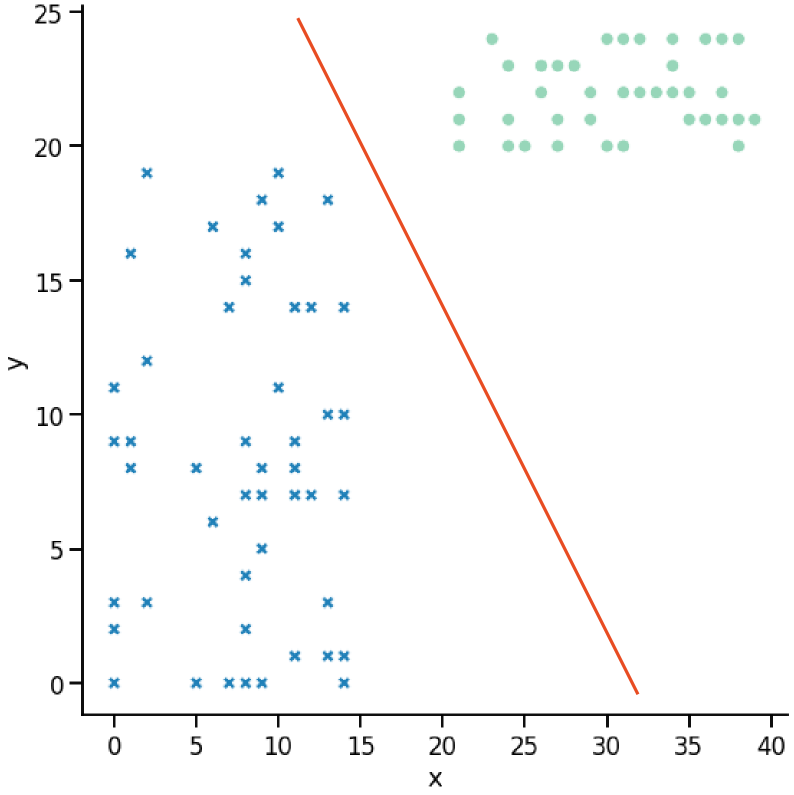

CLASSIFICATION & REGRESSION

Predict a label

Spam vs. Ham

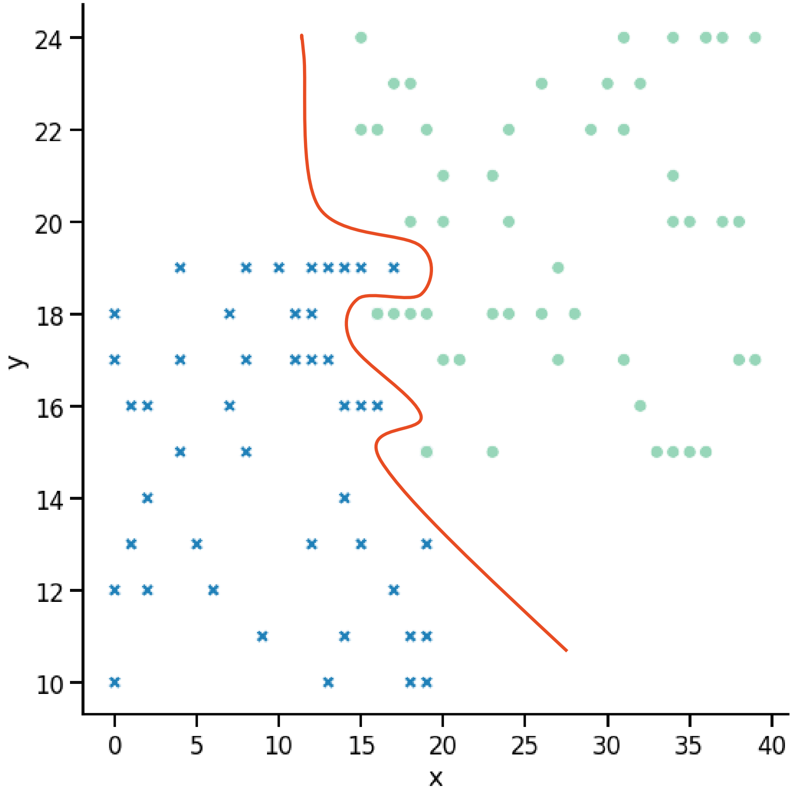

CLASSIFICATION & REGRESSION

Predict a label

Spam vs. Ham

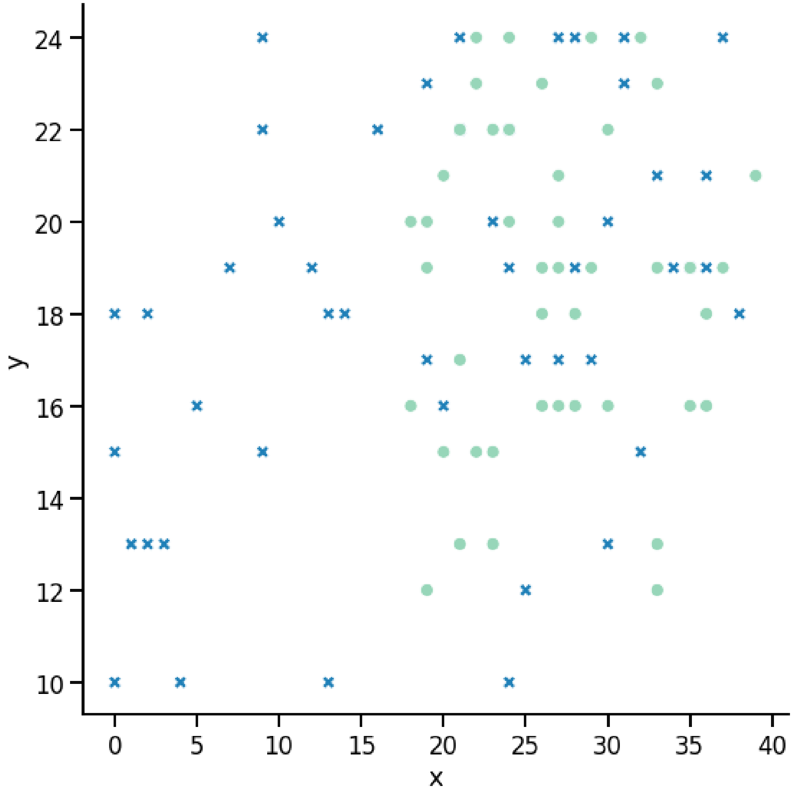

CLASSIFICATION & REGRESSION

Should I pack my jacket?

Walk the dog at 6 am

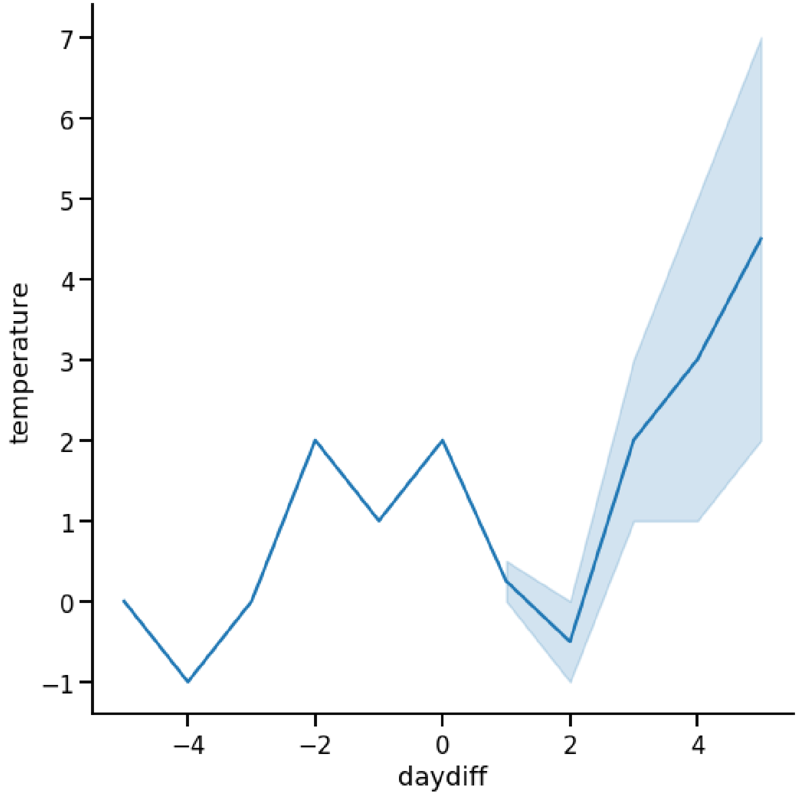

OUTLIER DETECTION

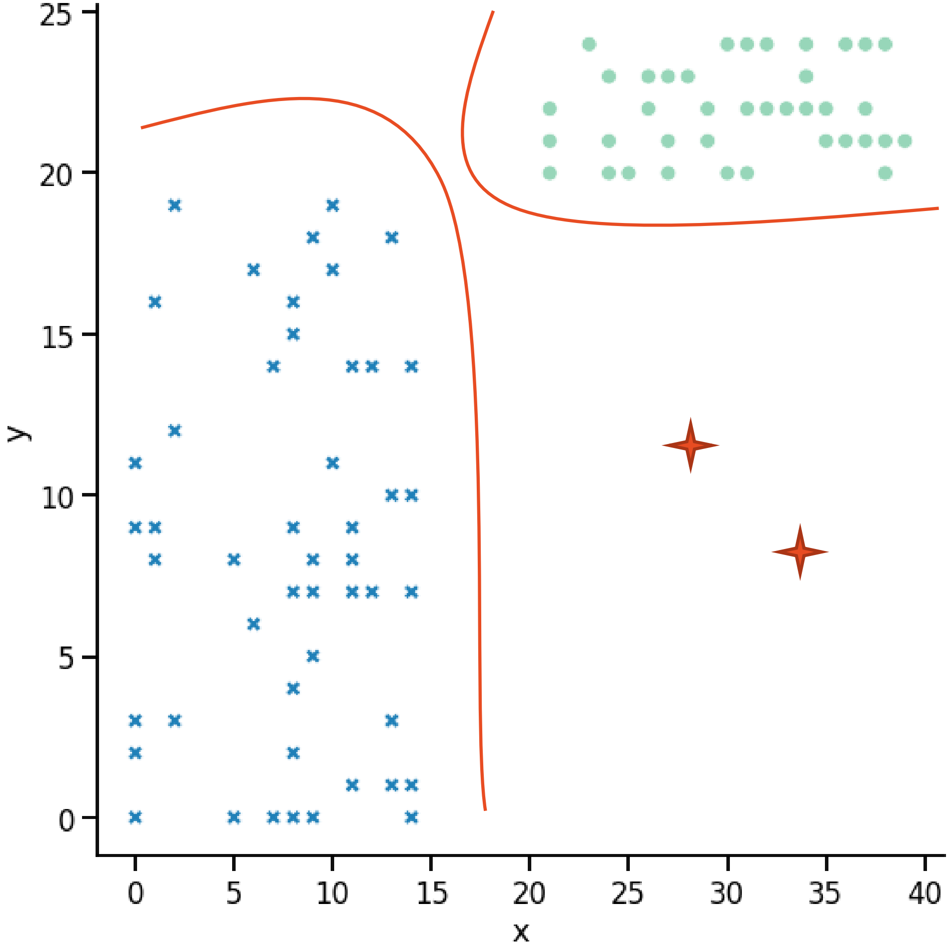

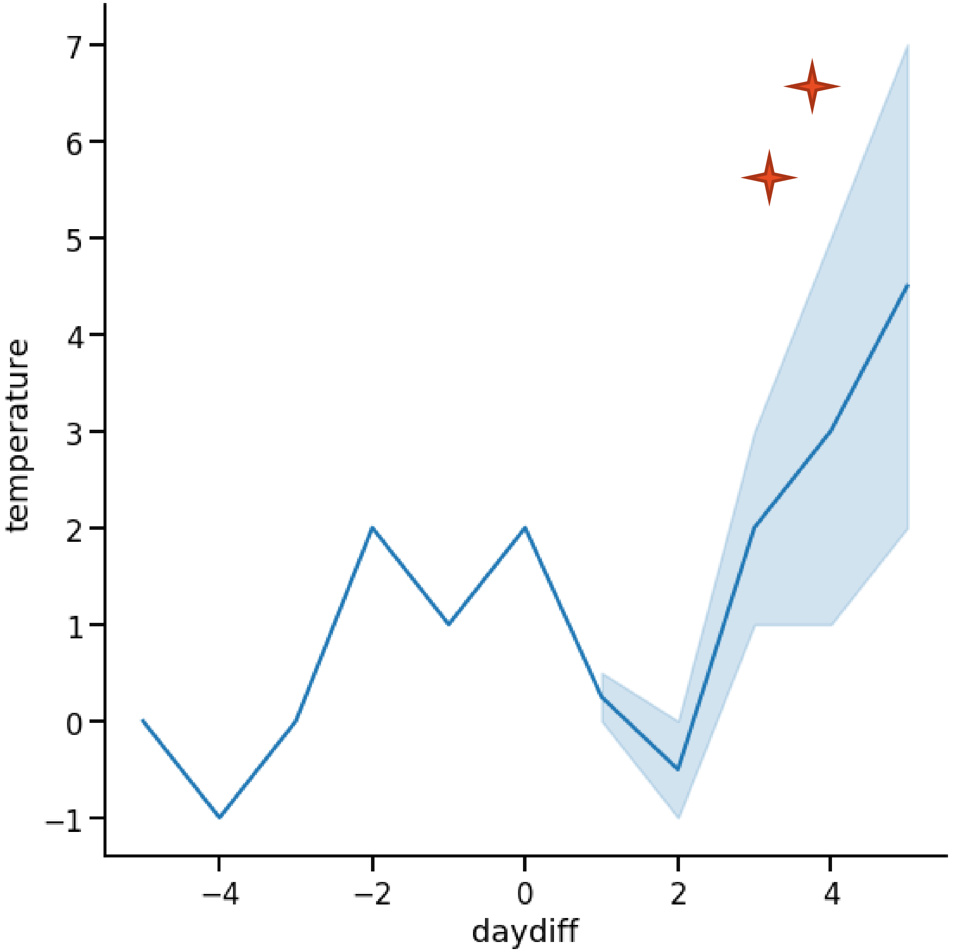

SUPERVISED & UNSUPERVISED

- “This is spam, this isn‘t”
- Labeled training data
  - Expensive to create
  - Hard to collect

SUPERVISED & UNSUPERVISED

- „Find groups & patterns in this!“
- Unlabeled training data
  - Easy to collect
- Hard to train, more False Positives

MACHINE LEARNING & DEEP LEARNING

- Everything besides „Deep Learning“
- Linear Models, Clustering, Ensembles

- Quick training, less HW needed
- Feature Engineering!

MACHINE LEARNING & DEEP LEARNING

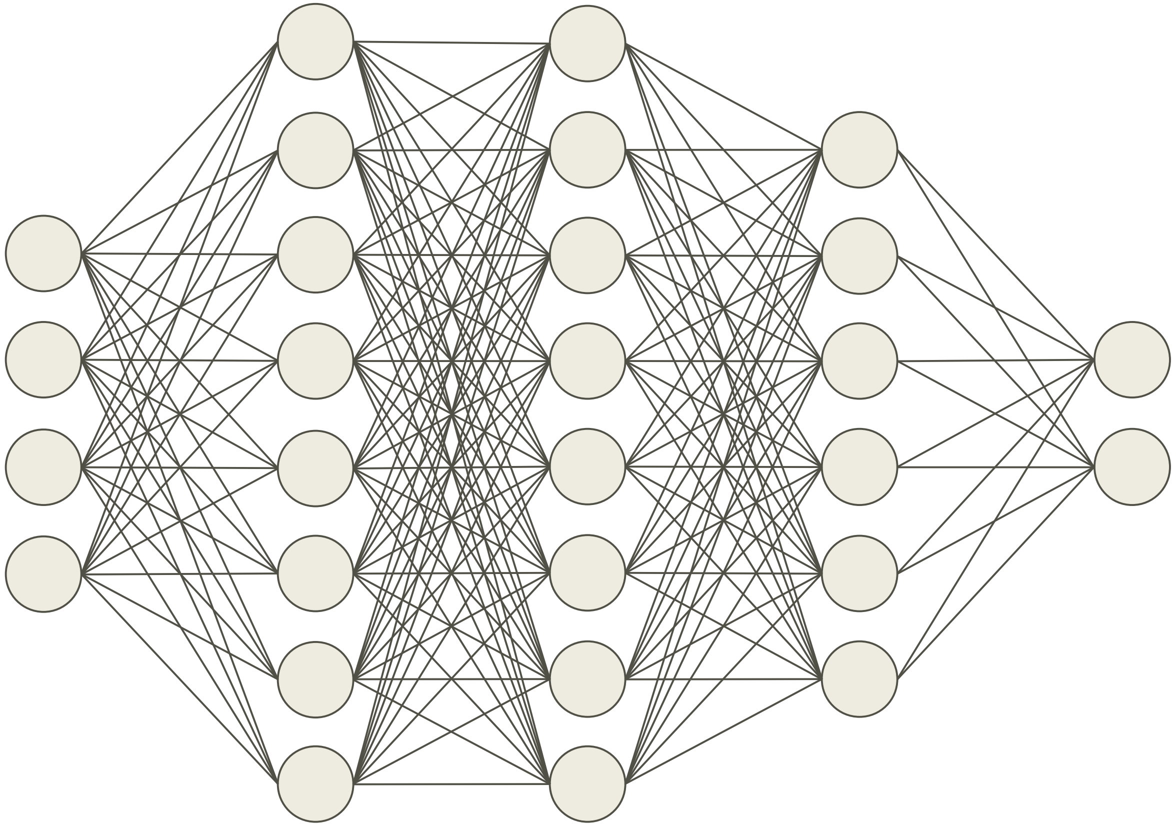

- Representation Learning
- Multi-Layer Neural Networks

- Expensive & HW-Intensive
- Long training
- No Feature Engineering

PRACTICAL APPLICATIONS

AV-SYSTEMS / ENDPOINT PROTECTION

- Regression:
  - Line count in Mail attachments

- Classification:
  - File Access
  - Syscall types

FRAUD PROTECTION

- Regression:
  - Buy history
  - Country of Origin

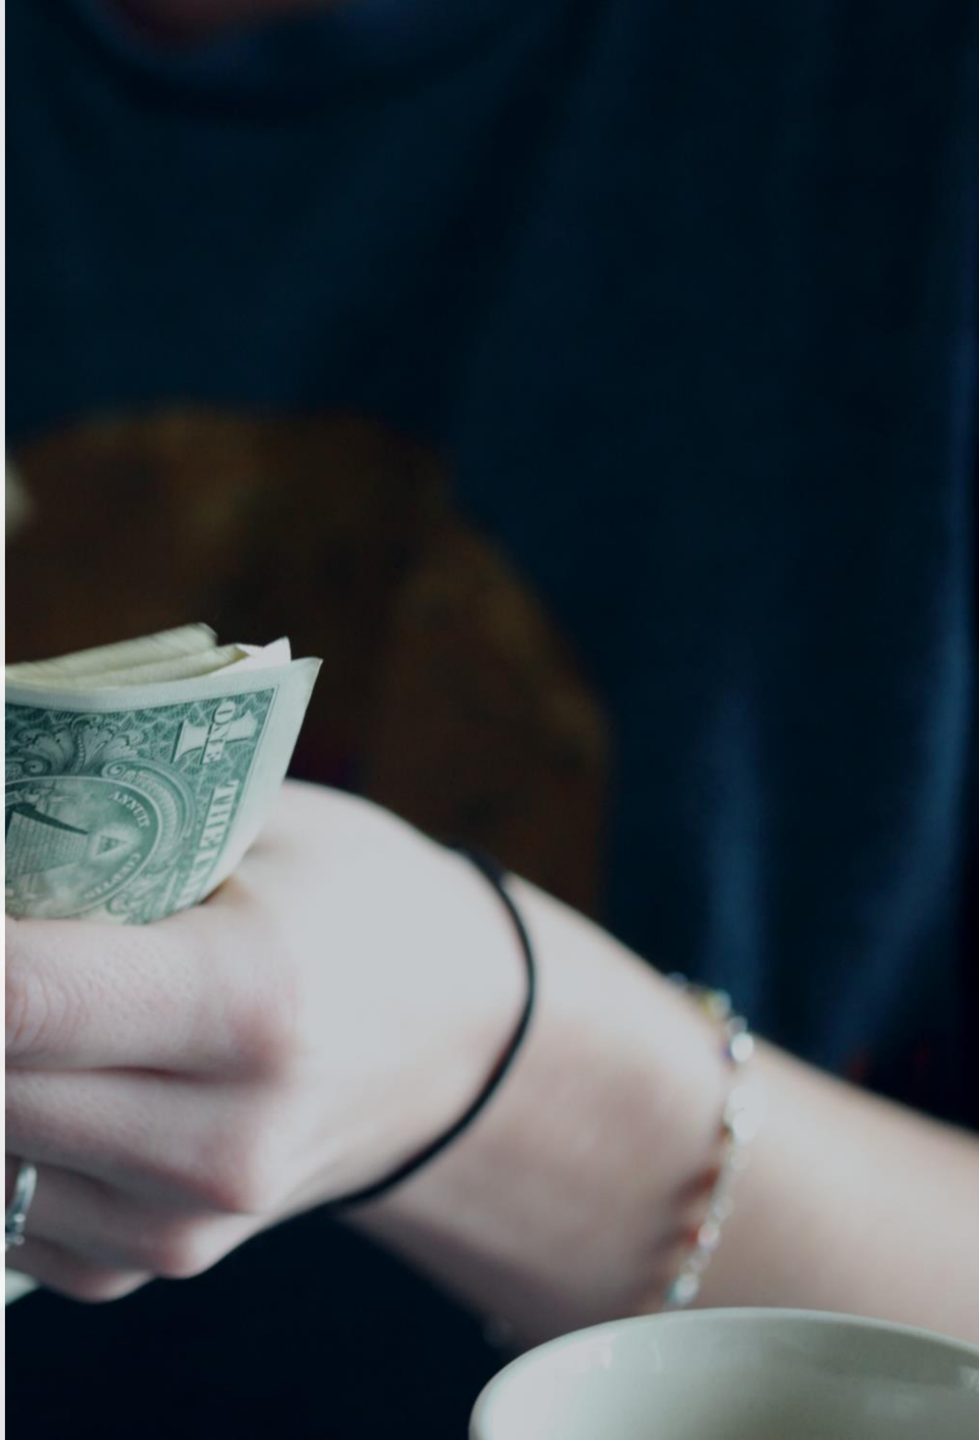

USER BEHAVIOR ANALYSIS

- Regression:
  - Login Times

- Classification
  - Common Shares
  - System Access (CRM vs Ticket)

APPLICATION

- Regression:
  - HTTP Communication
  - Call Pattern

- Classification:
  - Android Malware
  - SQLi, XSS

NETWORK

- Regression:
  - TCP Sessions
  - Session Flags

- Classification:
  - Scanning
  - DoS
  - Christmas Tree Packets

GETTING YOUR FEET WET

MACHINE LEARNING TESTBENCH

- Python 3
- Multi ML frameworks
- Multiple Datasets
- Splunk integration
- Supervised & unsupervised algorithms

MACHINE LEARNING TESTBENCH

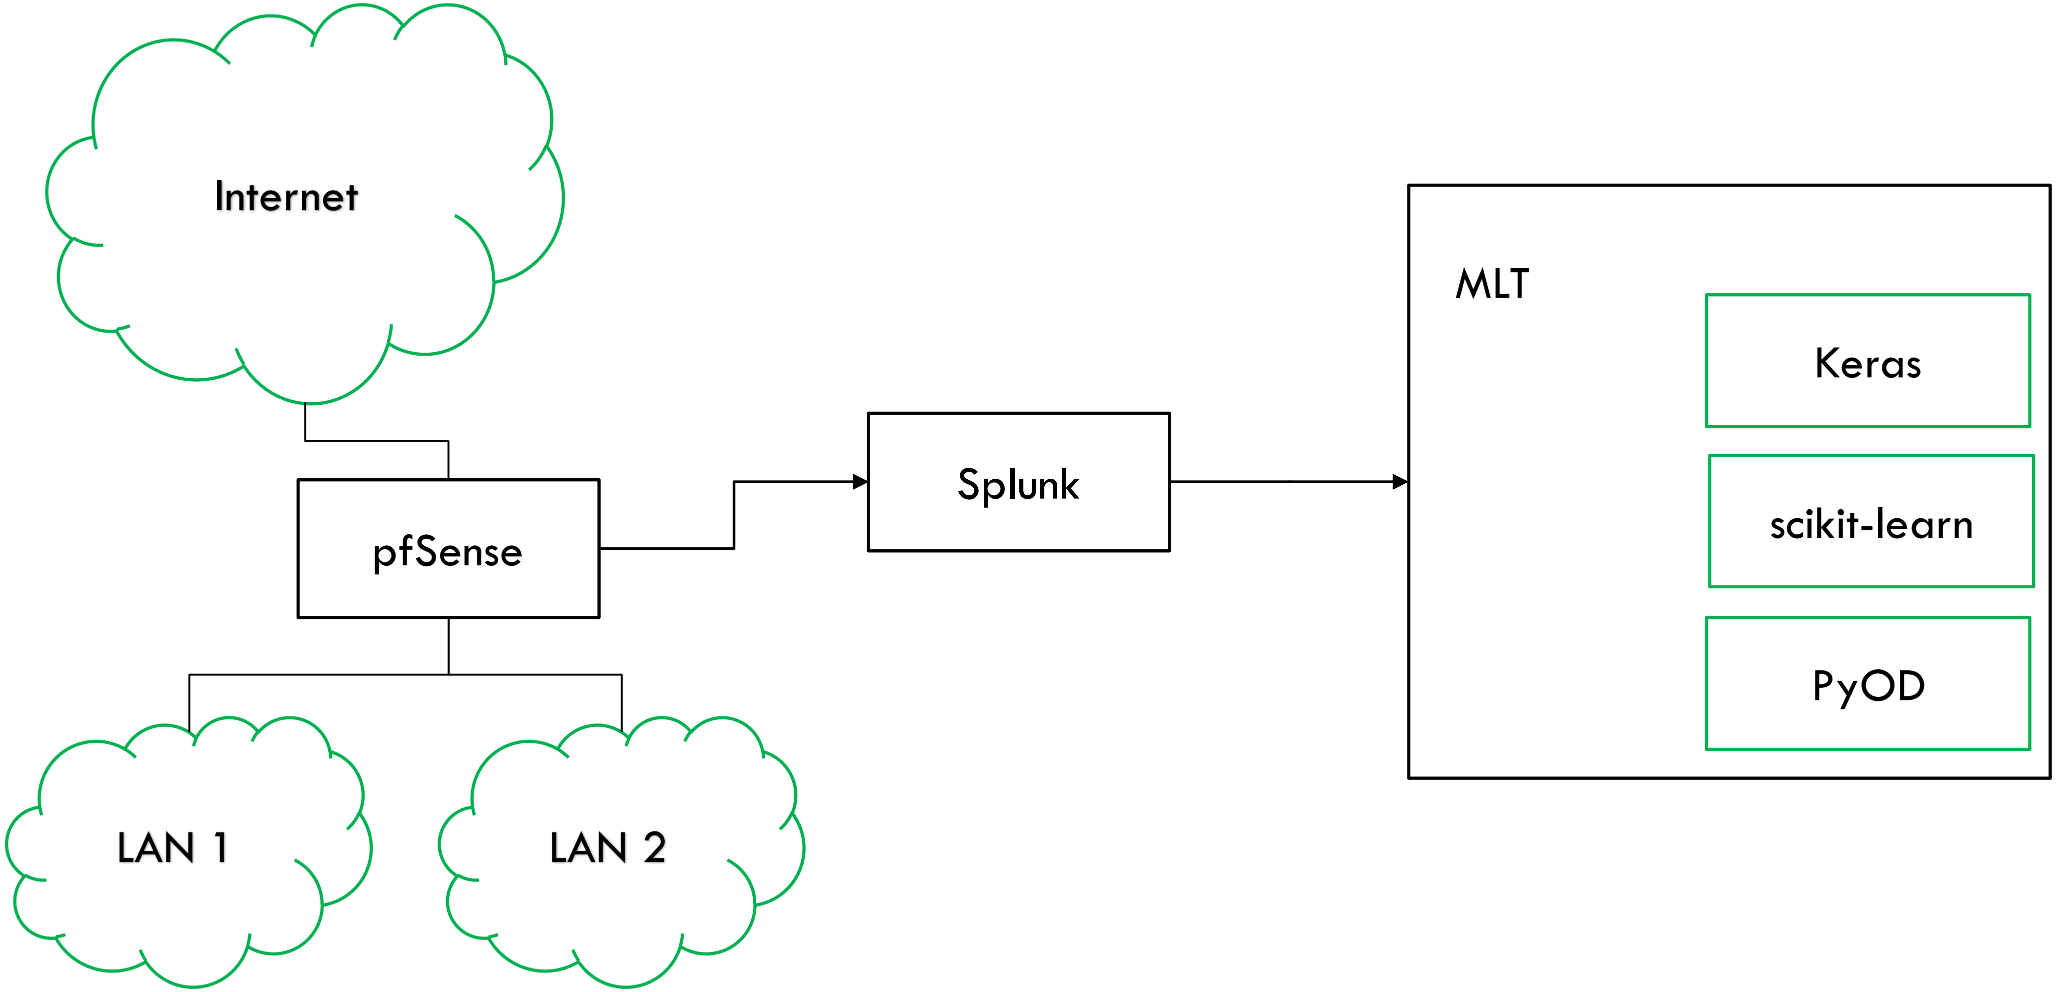

DEMO

THE FINE PRINT

- Things left out:
  - Explorative Analysis
  - Dataset Preparation
  - Feature Selection
  - Normalization
  - …

Ask me if you‘re interested

HOW DO I…?

- What are you trying to do?
- What data and features do you have?
- „Do you even realtime?“
- High detection rates or low false positive rates?

KEY TAKEAWAYS

- ML & DL: Handle with care
- Frameworks do heavy lifting
- Existing data, new insights!

FIN

@mat_zilla

https://github.com/Maddosaurus/Alfred

https://github.com/Maddosaurus/MLT