riscure

Lowering the Bar: Deep Learning for Side Channel Analysis

Guilherme Perin, Baris Ege,
Jasper van Woudenberg @jzvw
Blackhat, August 9, 2018

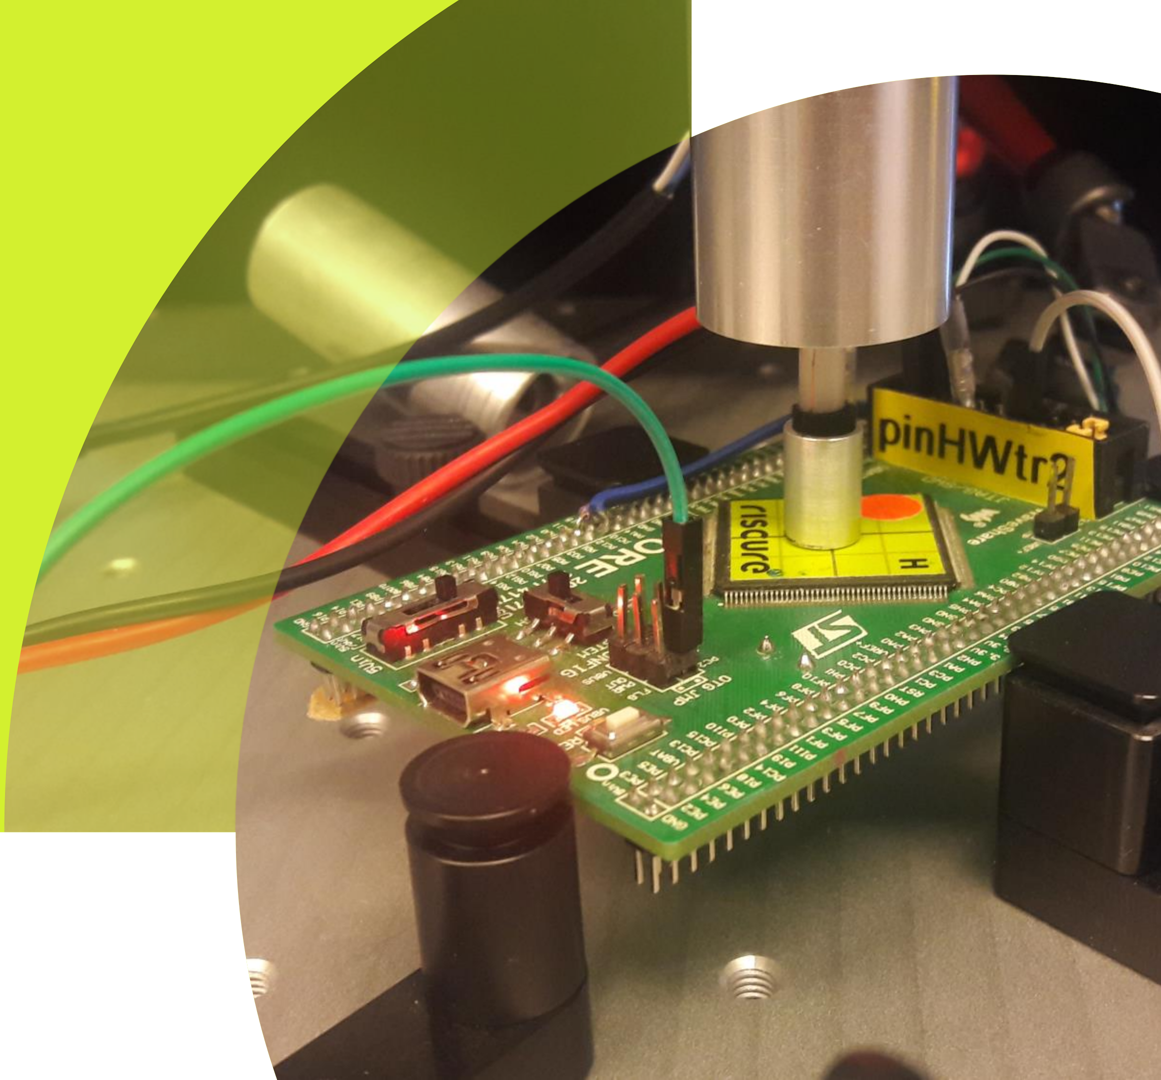

black hat USA 2018

Before

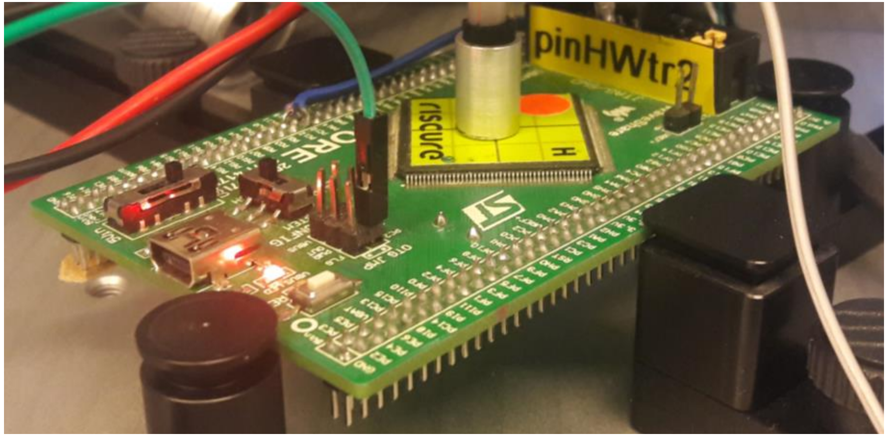

Signal processing

Leakage modeling

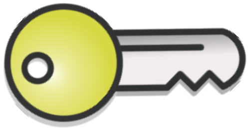

After

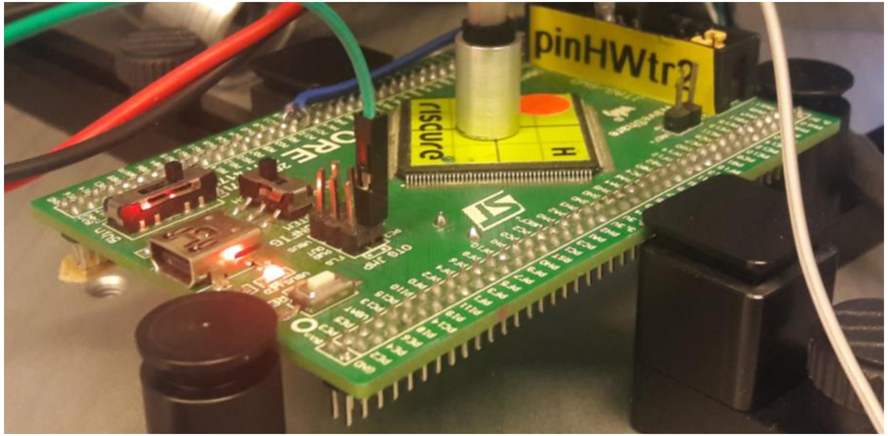

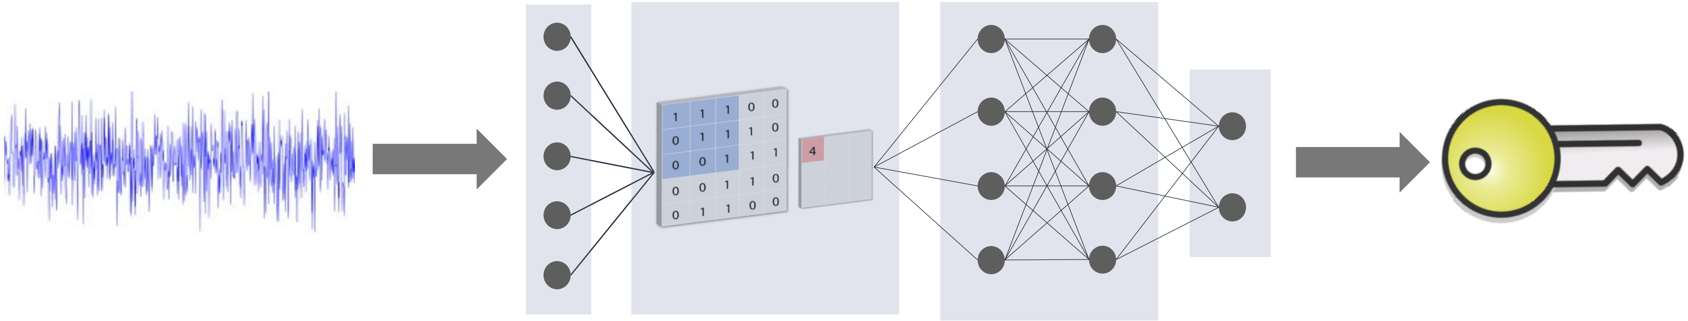

Power / EM side channel analysis

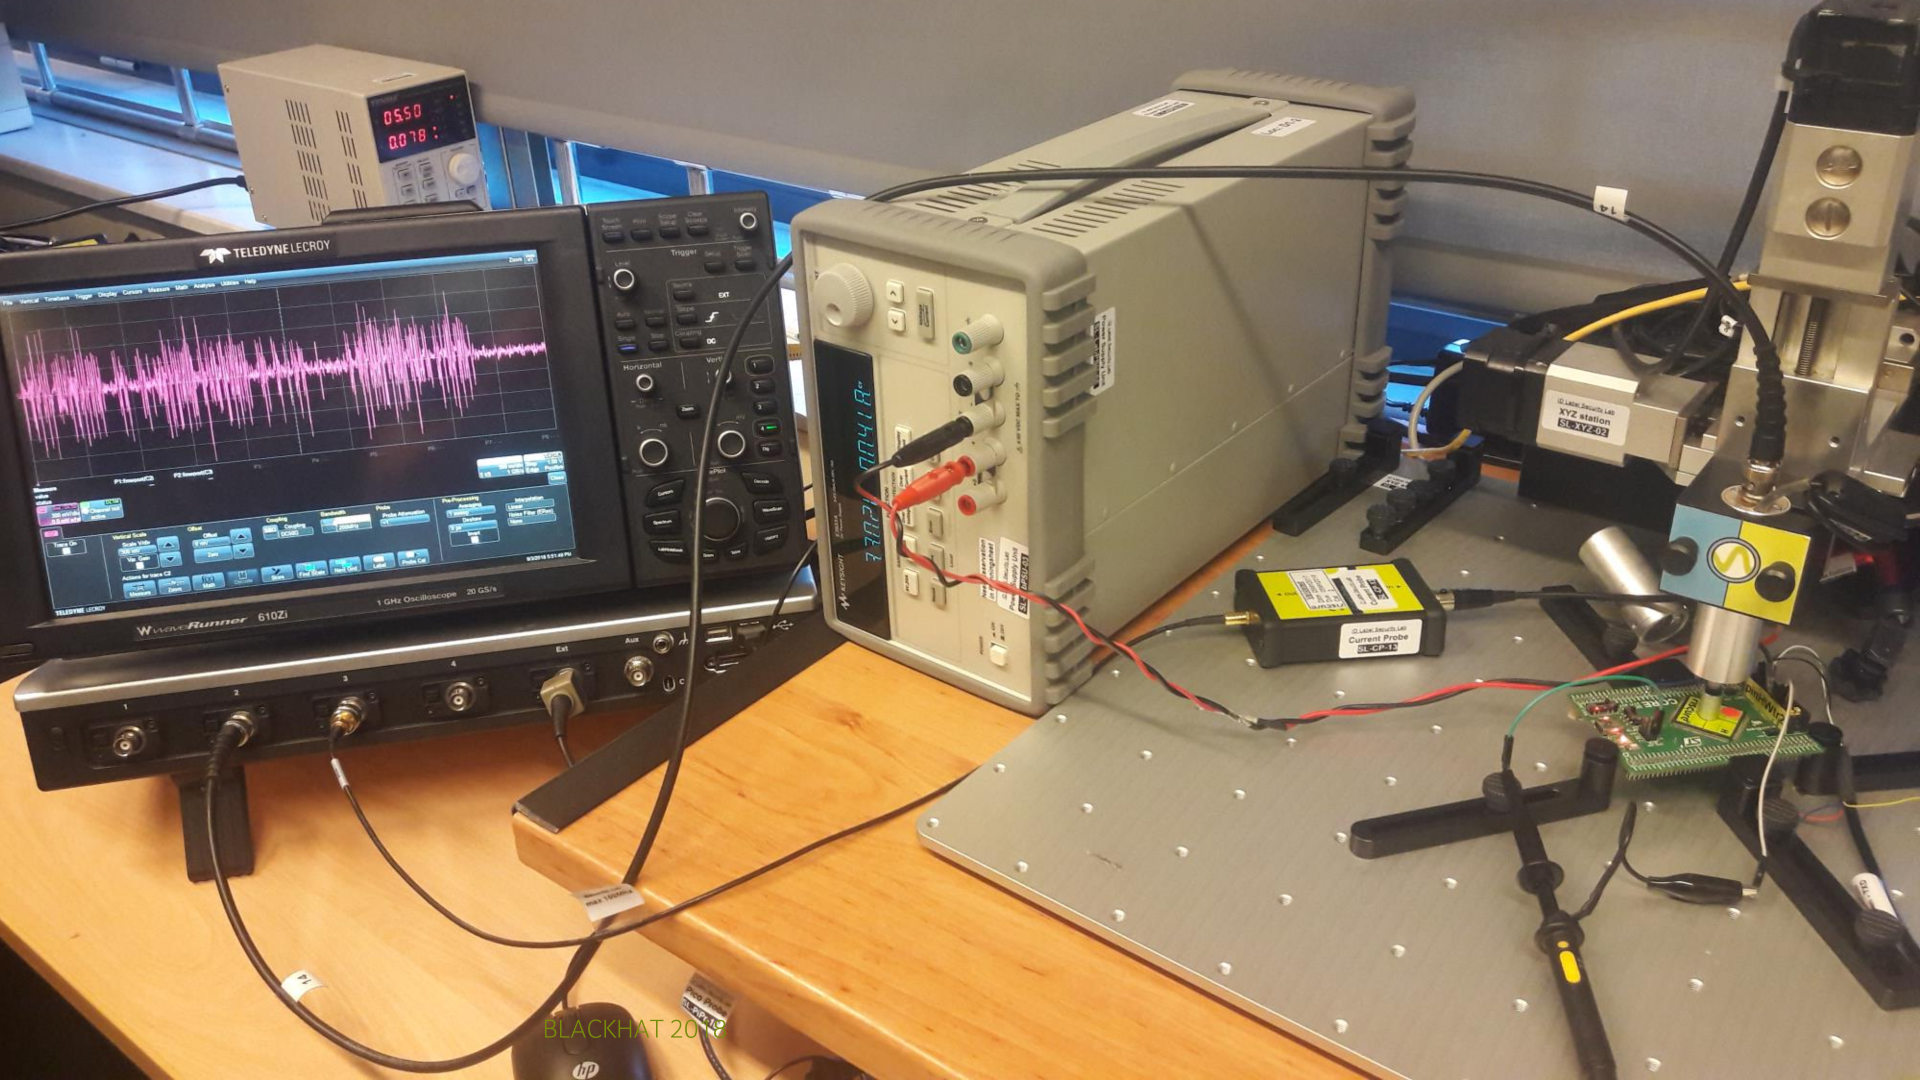

BLACKHAT 2018

Power analysis

Some crypto algorithm

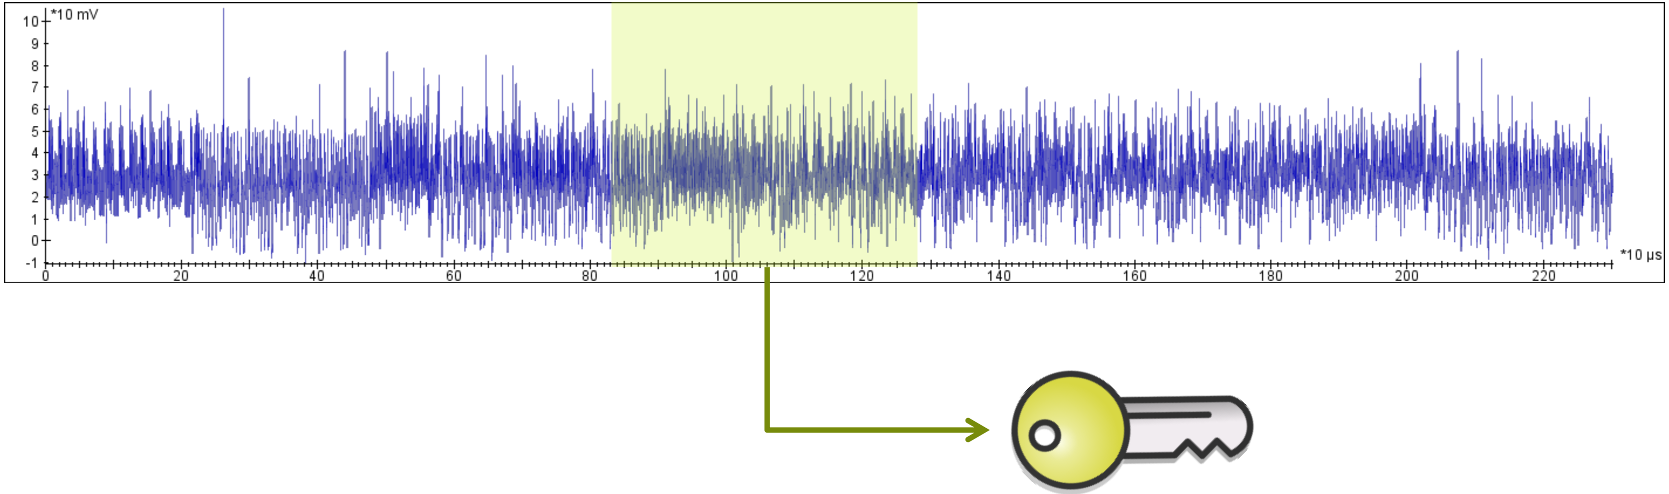

Example (huge) leakage

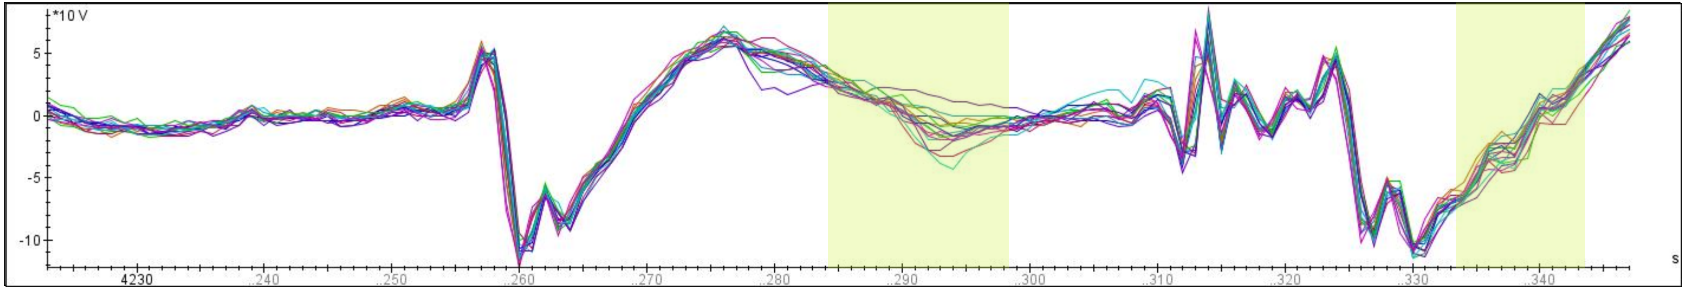

Data leakage

Noise

Signal processing (demo)

Raw trace

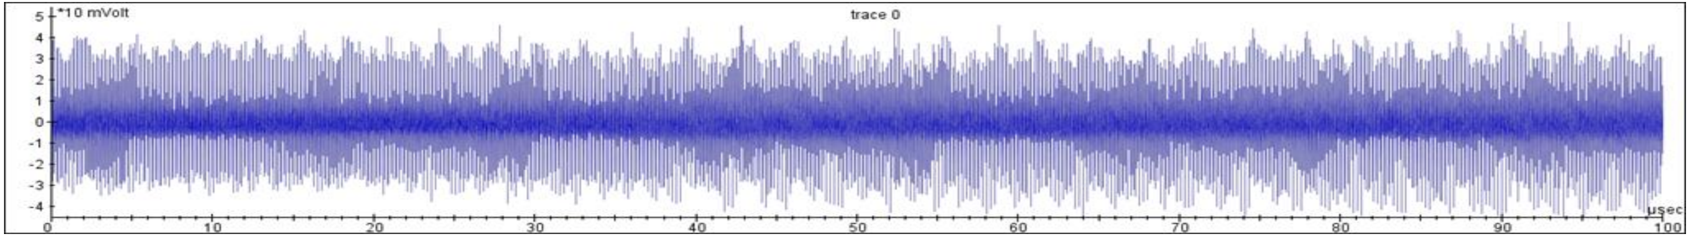

Processed trace

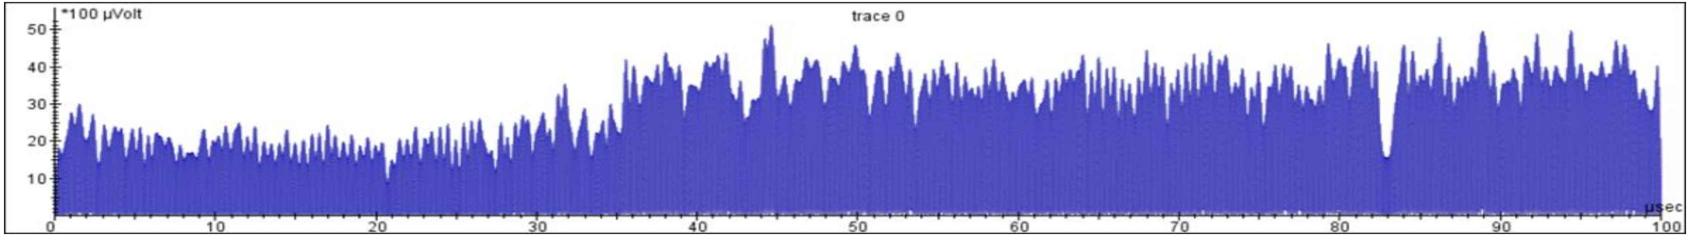

Misalignment (demo)

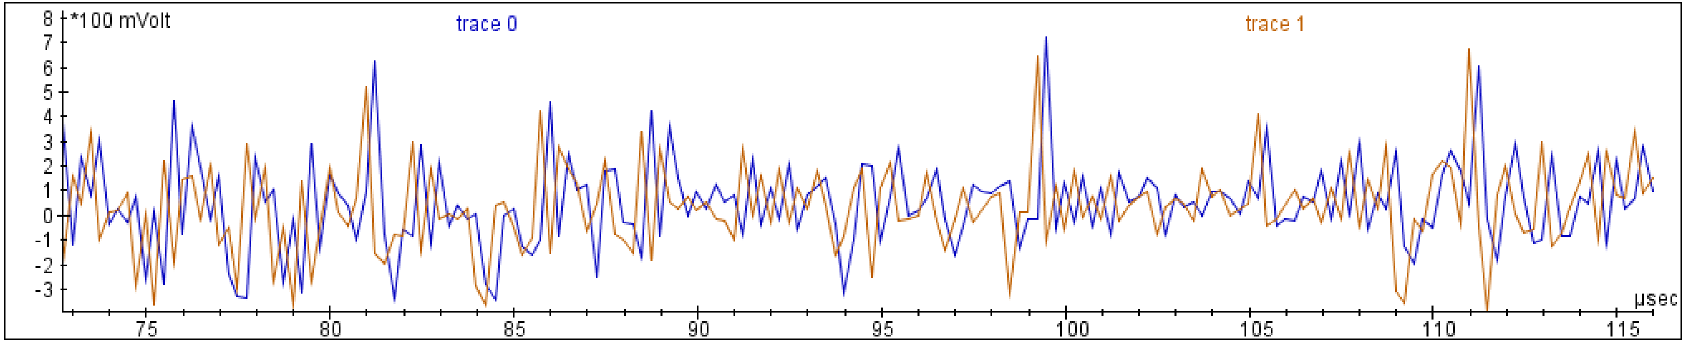

AES-128 first round attack

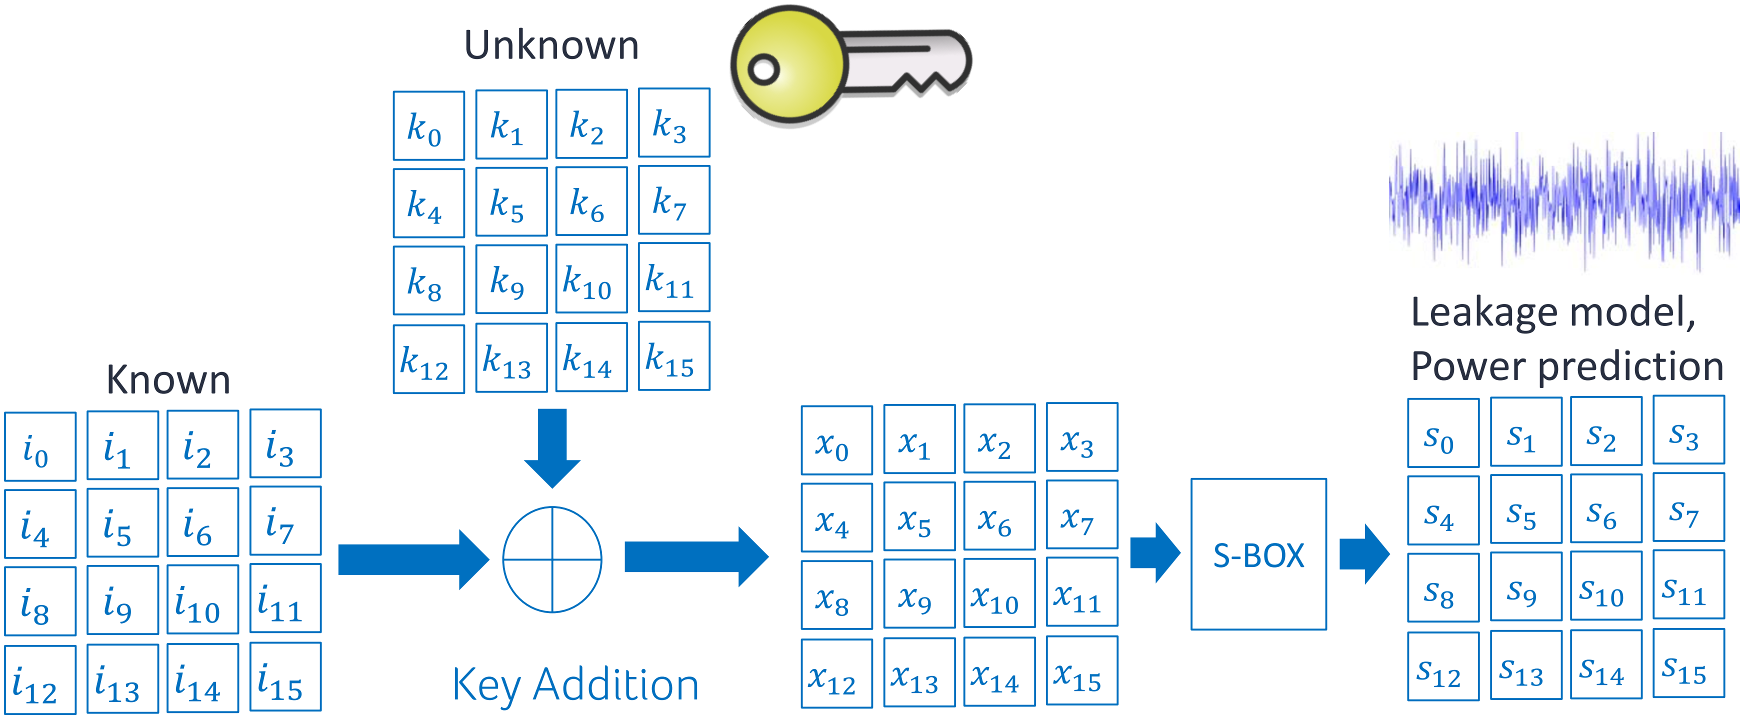

Points of interest selection

Correlation, T-test, Difference of Means

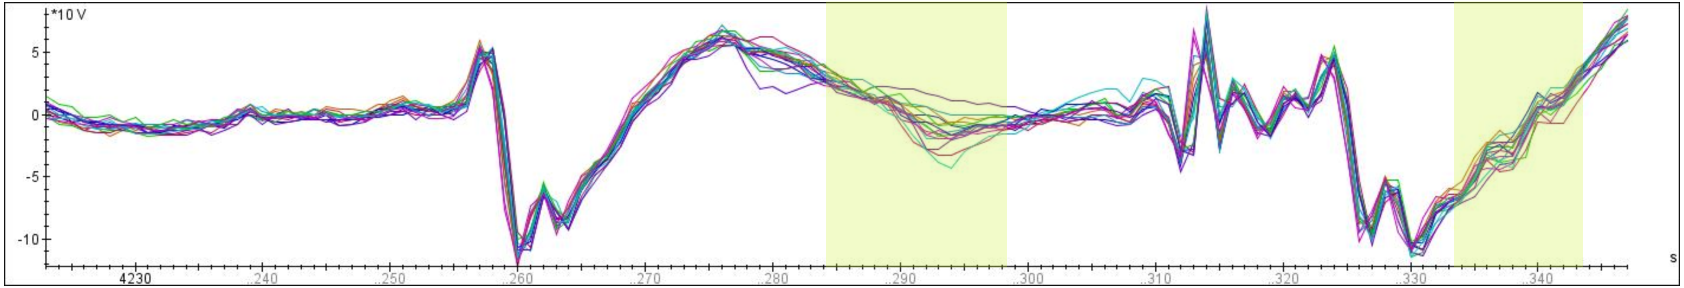

Data leakage Noise

Samples showing **statistical dependency** between intermediate (key-related) data and power consumption.

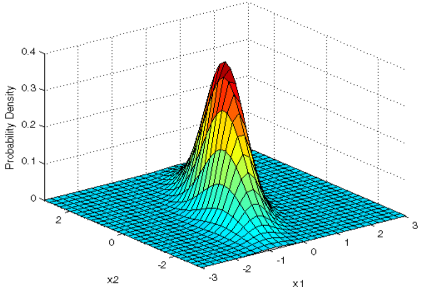

Concept of Template Analysis

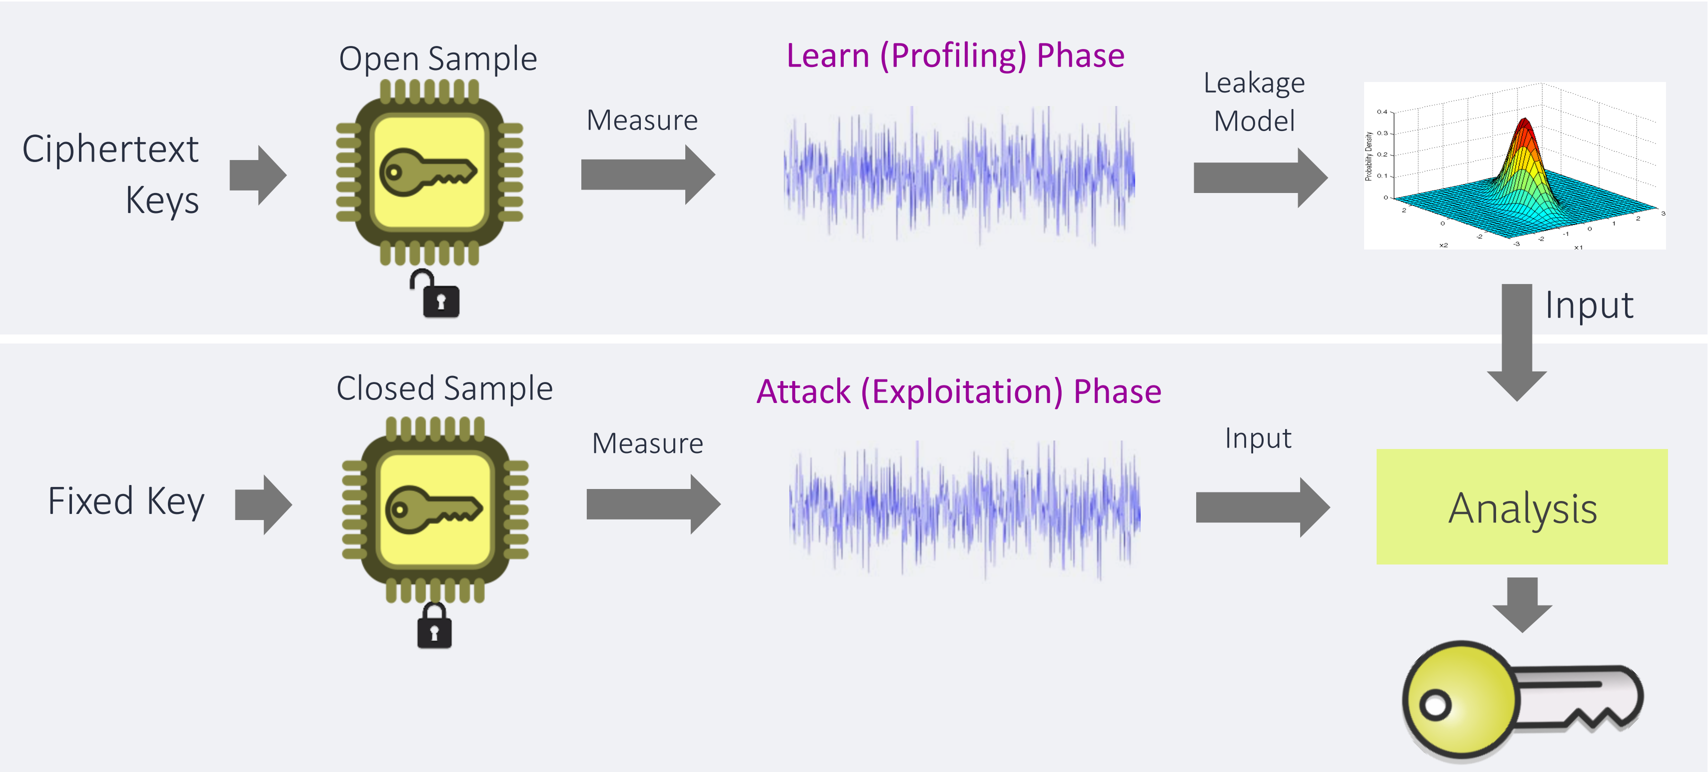

Key recovery

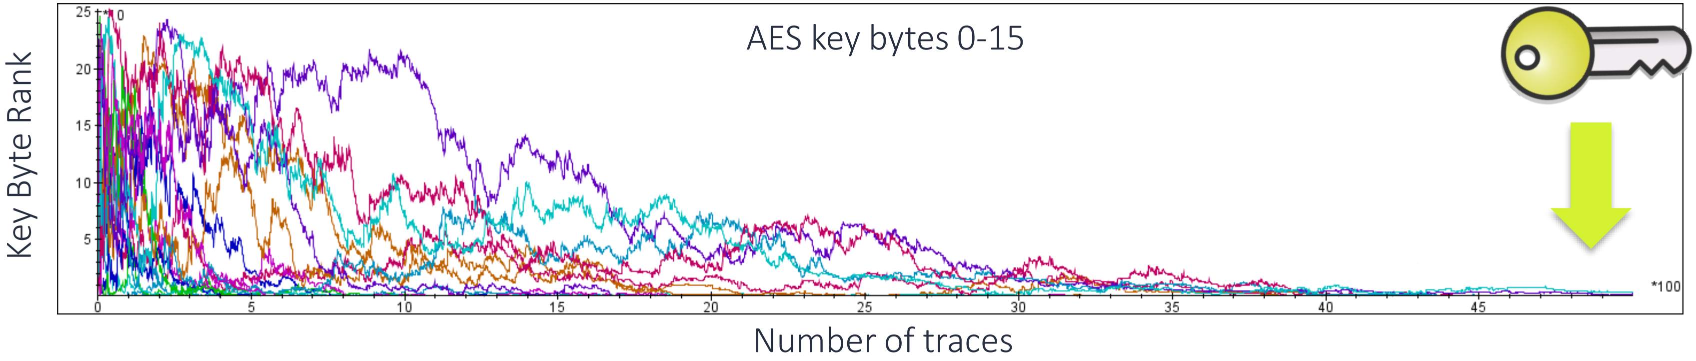

The actual process

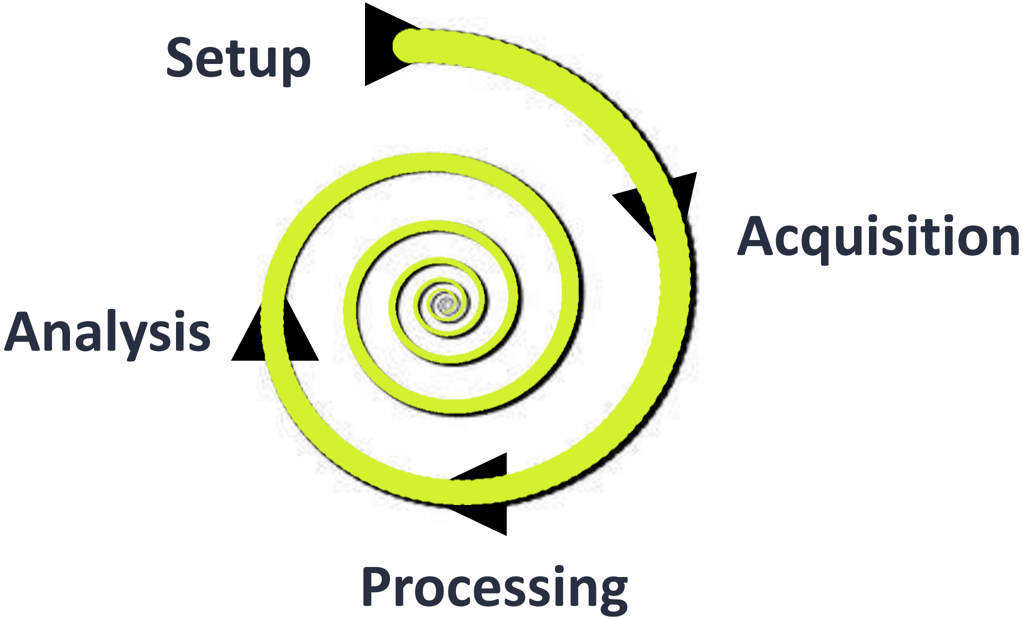

Deep learning background

Deep Learning

Data
with labels

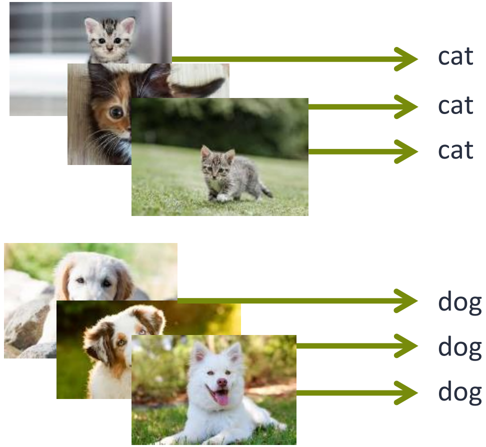

Deep Learning

Data
with labels

Train a machine to classify these data

machine

Cat (%)
Dog (%)

Error function

BACK-PROPAGATION ALGORITHM

Deep Learning

Data
with labels

Train a machine to
classify these data

Test the machine
on new data

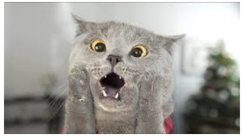

Trained
machine

Cat (%)

Dog (%)

Deep Learning

Data
with labels

Train a machine to classify these data

Change parameters

No

Test the machine on new data

Is classification accuracy good enough?

Yes

We are done!

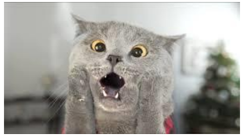

Trained machine

Cat

Machine = Deep Neural Network

Convolutional Neural Networks (CNNs)

Input Layer
(the size is equivalent to the number of samples)

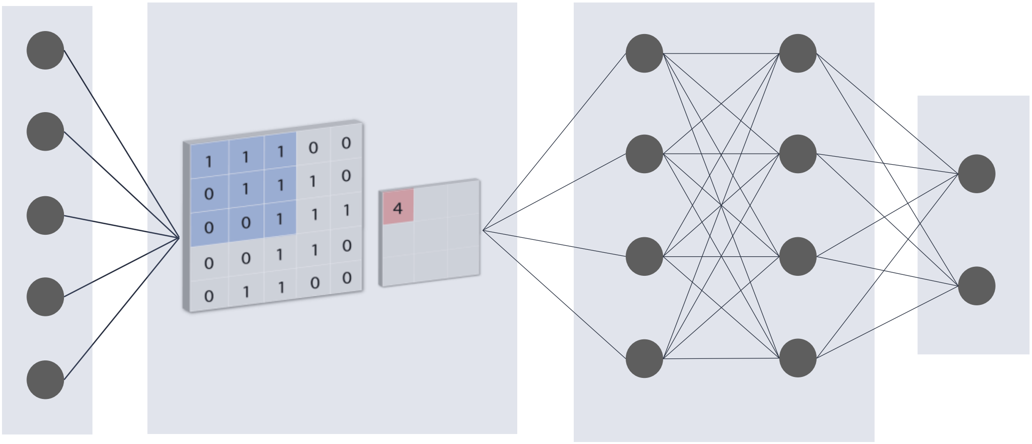

Output Layer
(the size is equivalent to the number of classes)

Conv. Layers
(feature extractor + encoding)

Dense Layers
(classifiers)

The **convolutional layers** are able to detect the features independently of their positions

Creating training/test/validation data sets

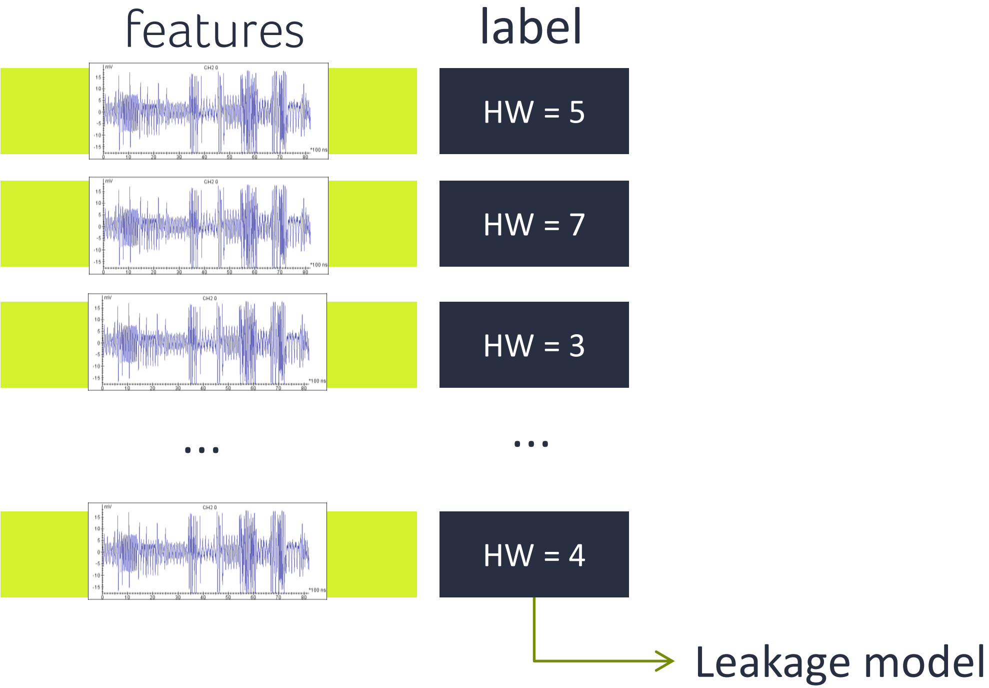

Classification

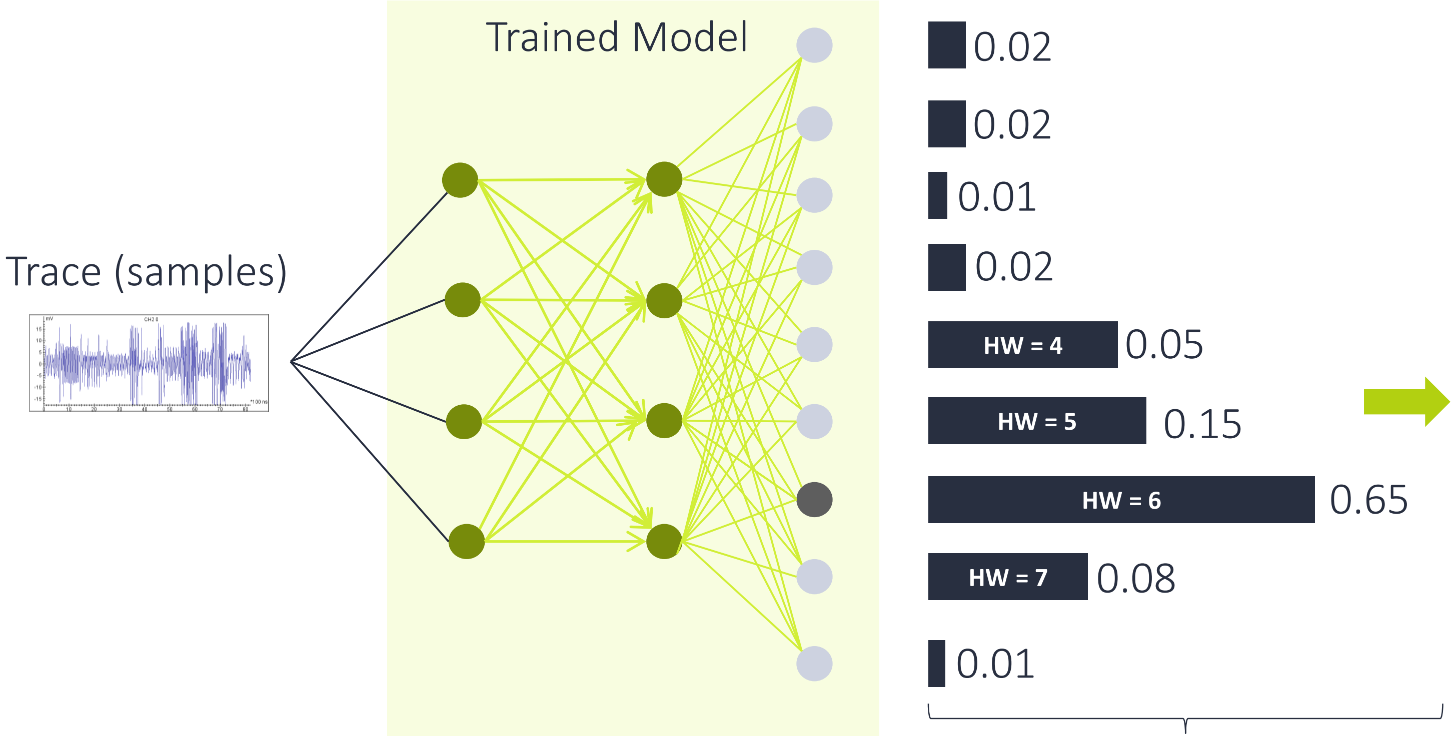

Softmax ($\sum p_i = 1$)

Key enumeration using output probabilities (Bayes)

Deep learning on side channels in practice

BLACKHAT 2018

Step 1: Define initial hyper-parameters (demo)

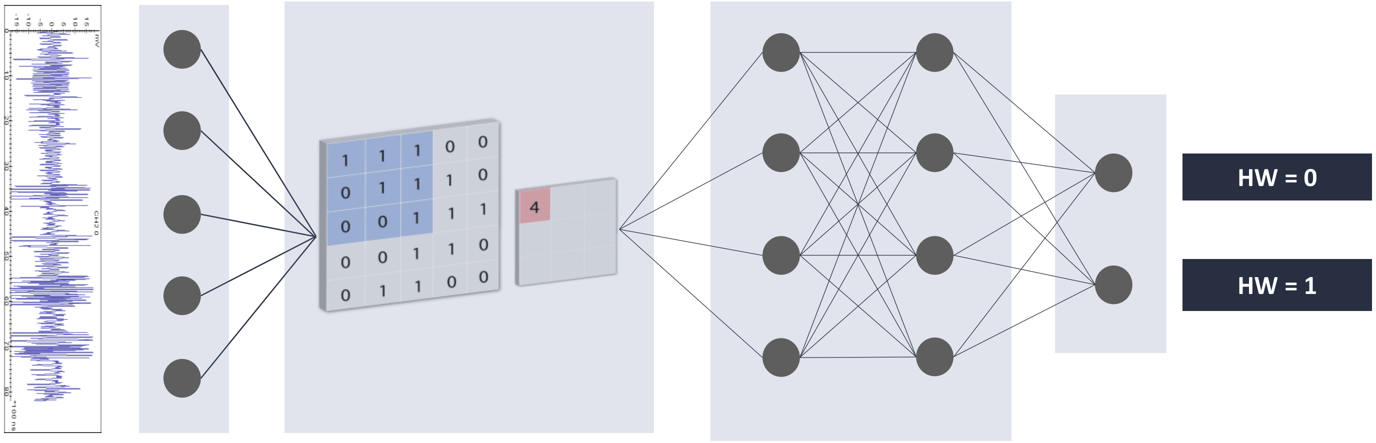

Step 2: Make sure it's capable of learning

- Increase the number of training traces and observe the training and validation accuracy
- Overfitting too fast?
  - Training accuracy: 100% | Validation accuracy: low
  - Neural network is too big for the number of traces and samples

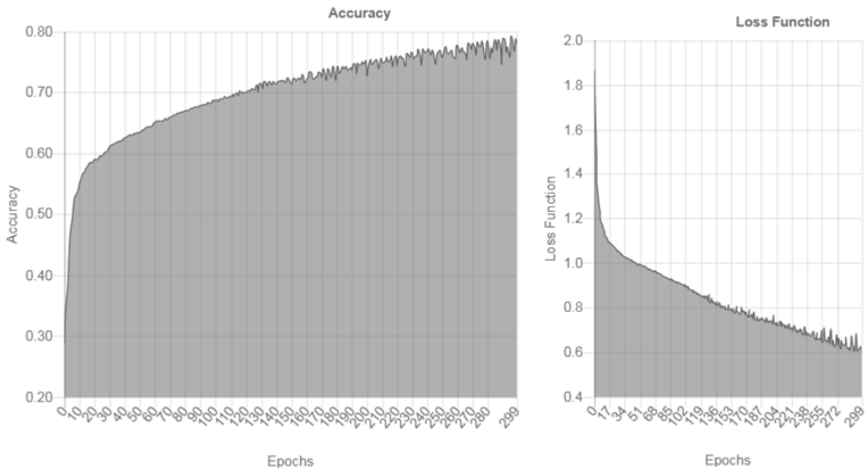

Step 3: Make it generalize

Make sure the training accuracy/recall is increasing

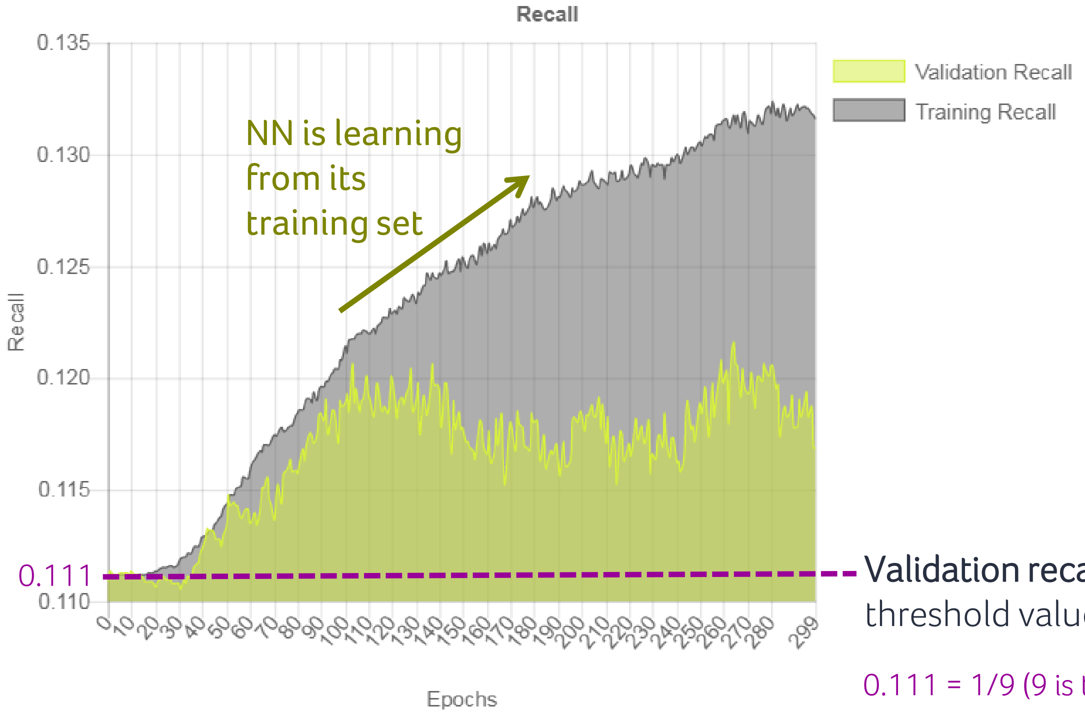

Validation recall *stays* above the minimum threshold value = **model is generalizing**

0.111 = 1/9 (9 is the number of classes – HW of a byte)

Step 3: Make it generalize

Regularization techniques:

- L1, L2 (penalty applied to the weights)
- Dropout
- **Data Augmentation (+traces)**
- Early Stopping

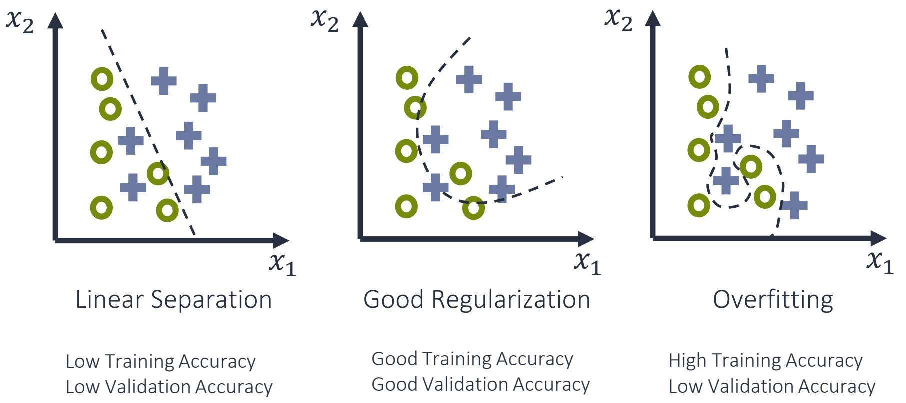

Step 4: Key Recovery

In this analysis, we only need slightly-above coin flip accuracy!

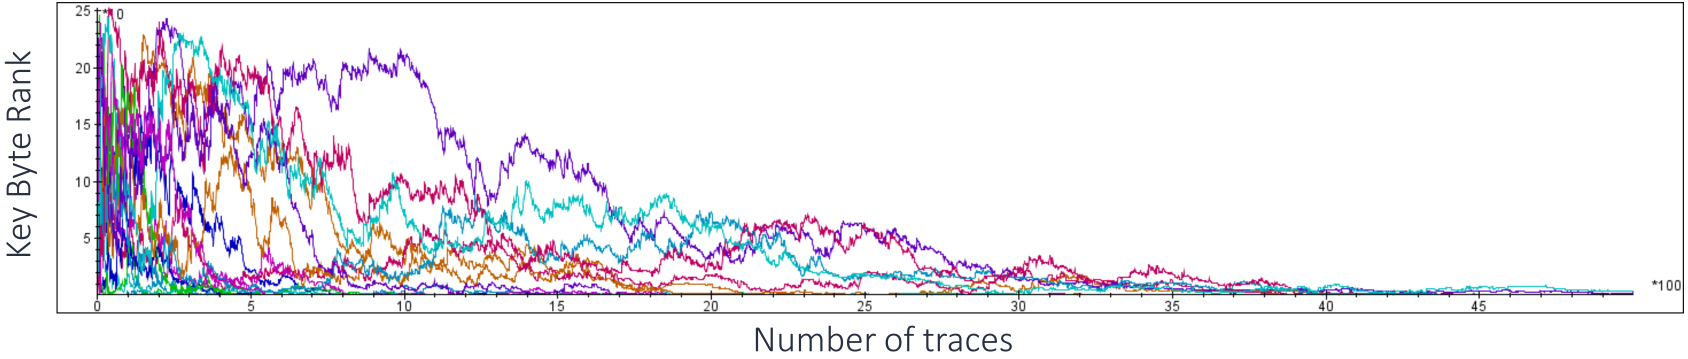

Getting keys from the thingz!

Piñata AES-128 with misalignment (demo)

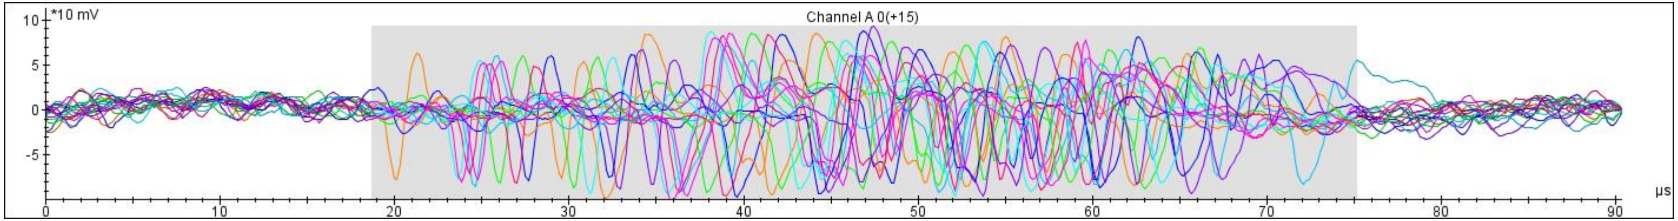

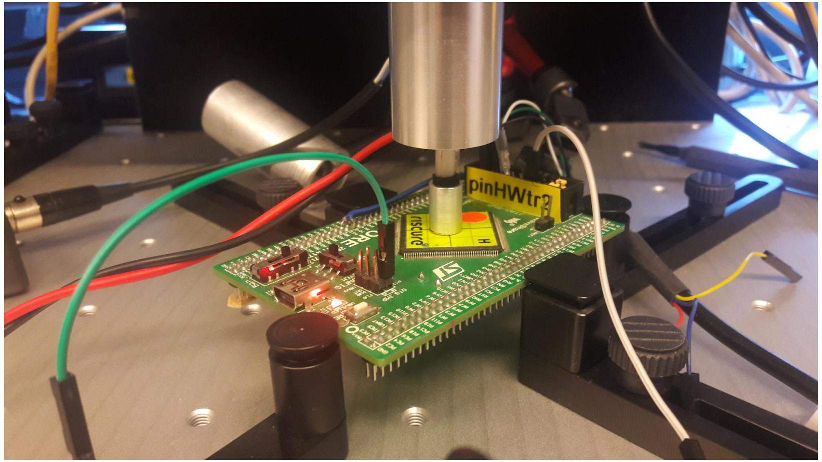

Bypassing Misalignment with CNNs

Neural Network: Input Layer > ConvLayer > 36 > 36 > 36 > Output Layer

Training/validation/test sets: 90000/5000/5000 traces of 500 samples

Leakage Model: HW of S-Box Out (Round 1) → 9 classes

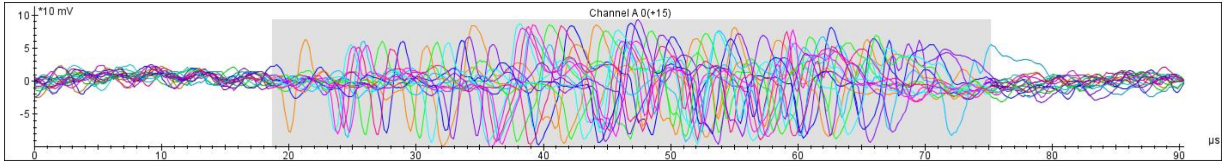

Use Data Augmentation as regularization technique to improve generalization

Results for key byte 0:

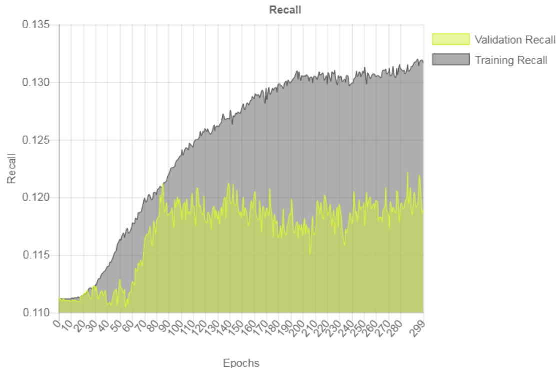

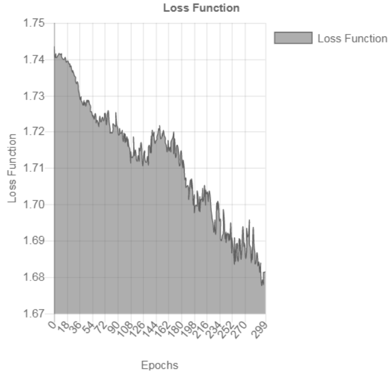

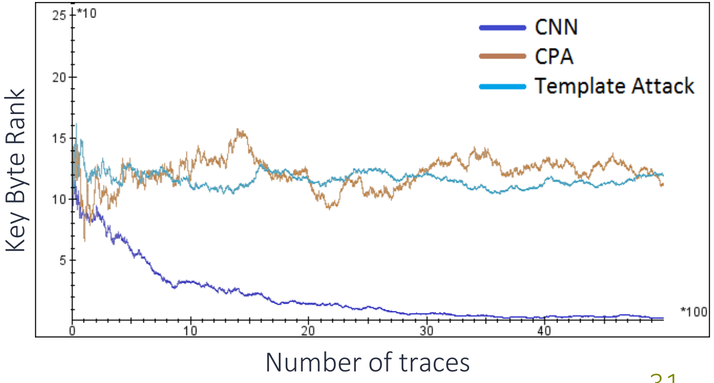

Breaking protected ECC on Piñata

Supervised deep learning attack:

- Curve25519, Montgomery ladder, scalar blinding
- Messy signal
- Brute-force methods for ECC are needed if test accuracy < 100%
- Need to get (almost) all bits from one trace!

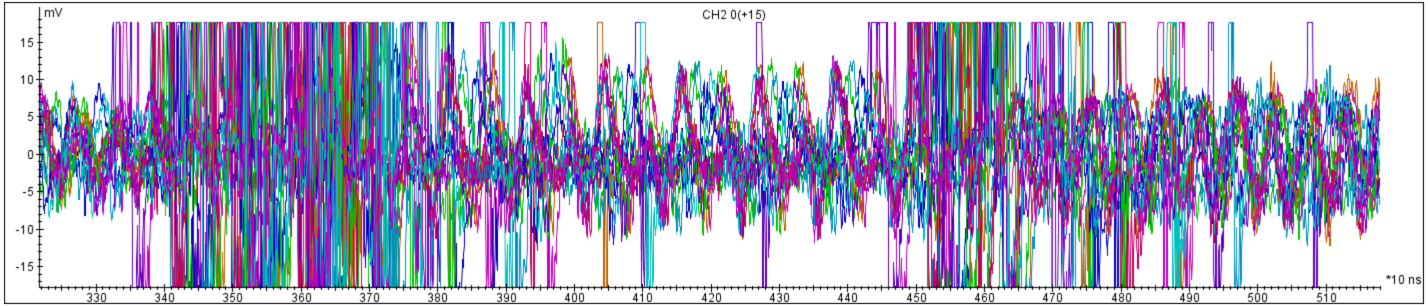

Breaking protected ECC

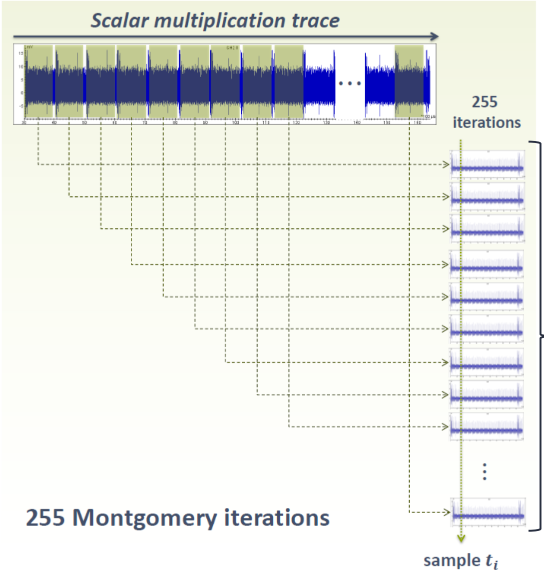

Misaligned traces

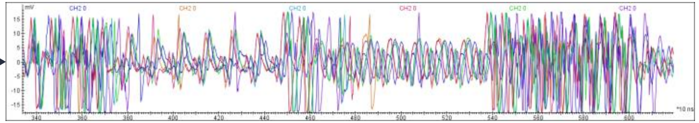

Unsupervised/Supervised Horizontal Attack: 60% success rate
Deep learning: 90% success rate
Deep learning(**+ data augmentation**): 99.4% success rate
Data augmentation: 25k → 200k traces.

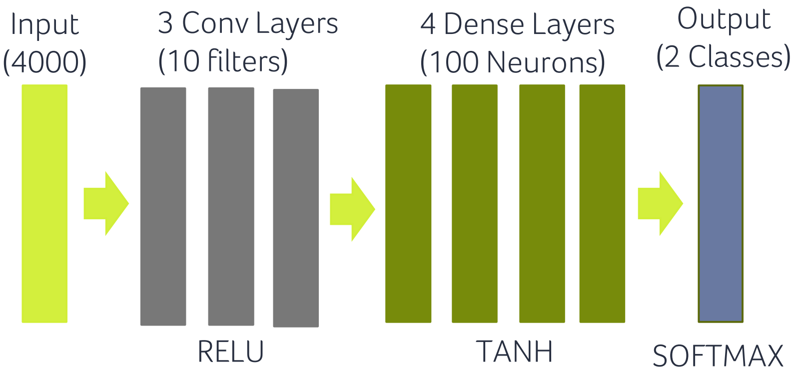

Breaking AES with First-Order Masking (demo)

- Target published in 2013 (http://www.dpacontest.org/v4/)
- 40k traces available
- AES-256 (Atmel ATMega-163 smart card)
- Countermeasure: Rotating S-box Masking (RSM)

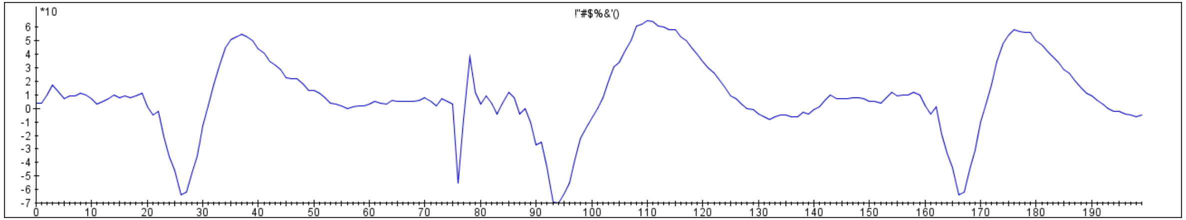

Breaking AES with First-Order Masking

$x \oplus m$
$y \oplus m$

Remove the relationship between
power consumption (EM) and predictable data

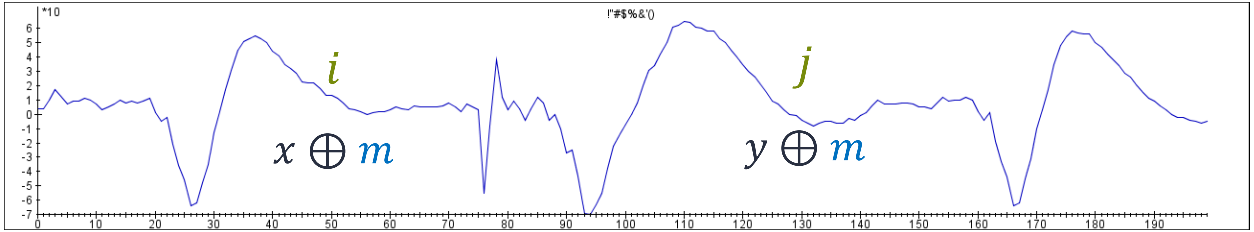

Combine data: $(x \oplus m) \oplus (y \oplus m) = x \oplus y$

Combine samples: $t[i] \times t[j]$

Brute-force $i, j \rightarrow$ Quadratic complexity

Breaking AES with First-Order Masking

Challenge: Training key == validation key

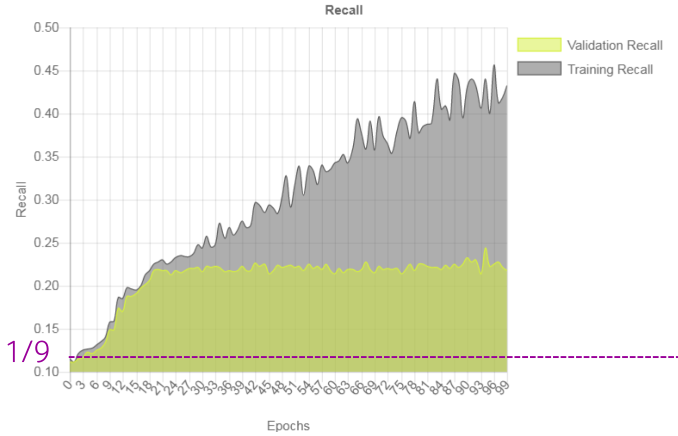

Correct key byte candidate
(good generalization)

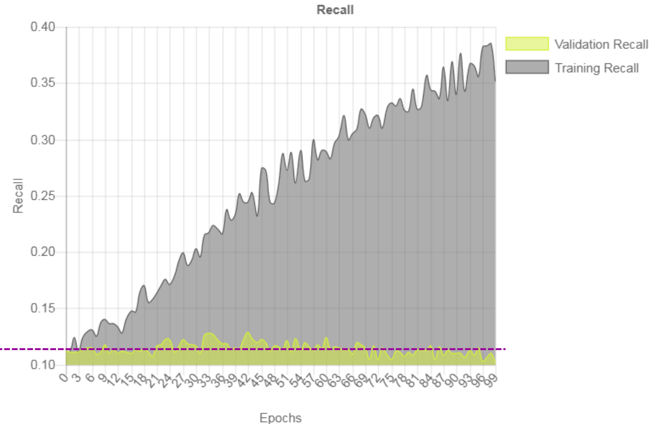

Wrong key byte candidate
(poor generalization)

Breaking AES with First-Order Masking

Overfitting can be verified by checking where the NN is learning

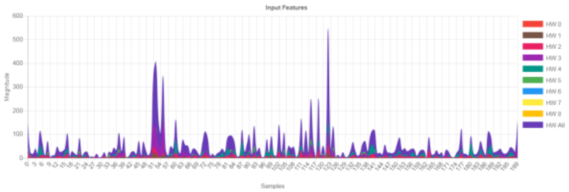

Correct key byte candidate
(CNN learns from specific
and leaky samples)

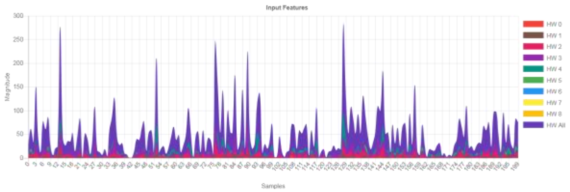

Wrong key byte candidate
(CNN overfits because it
can't distinguish leaky
samples from noise)

Breaking AES with First-Order Masking

Neural Network: Input Layer > ConvLayer > 50 > 50 > 50 > Output Layer

Training/validation/test sets: 36000/2000/2000 traces

Leakage Model: HW of S-Box Out (Round 1) → 9 classes

Results for key byte 0:

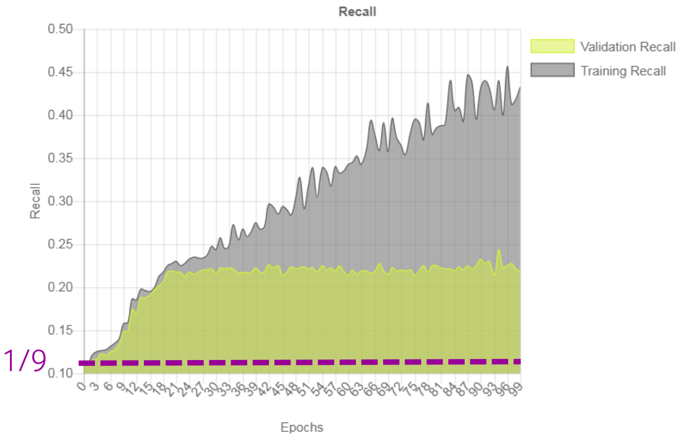

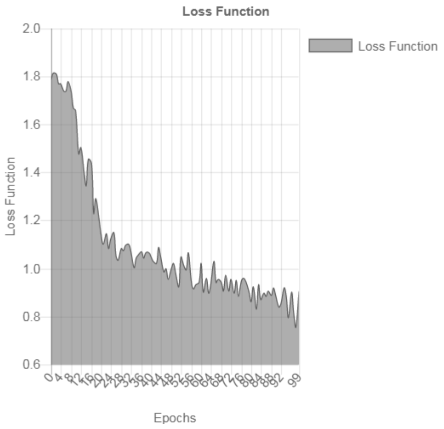

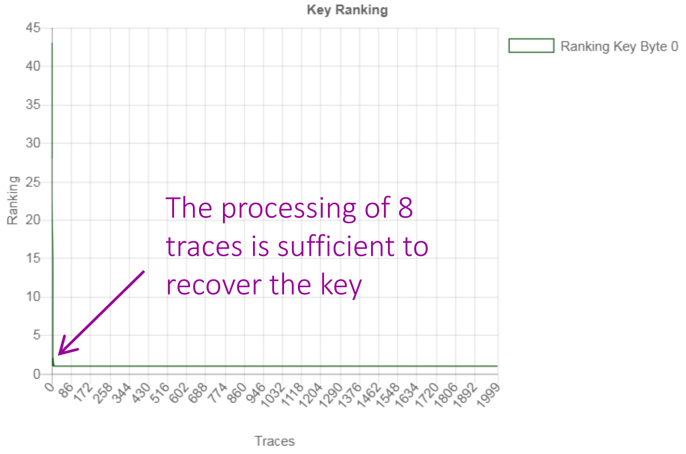

1st cool thing

This shouldn't work

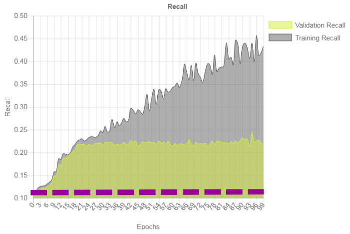

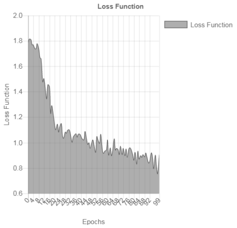

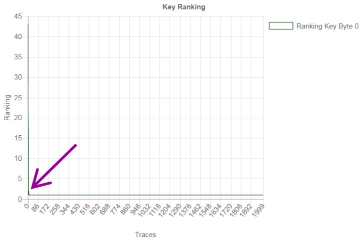

2nd cool thing

DL is up there with dozens of SCA research teams

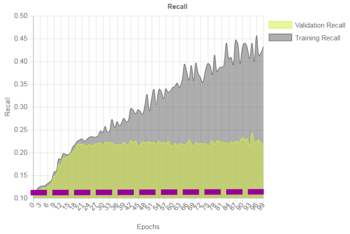

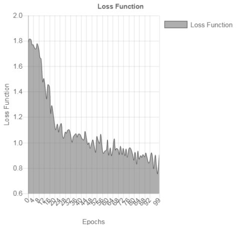

Wrapping up

I want to learn more!

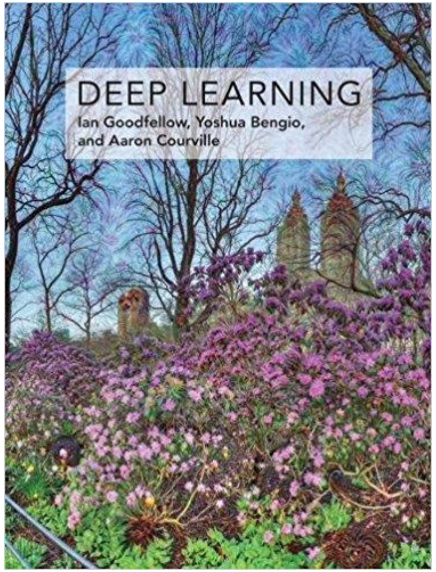

Deeplearningbook.org

introtodeeplearning.com

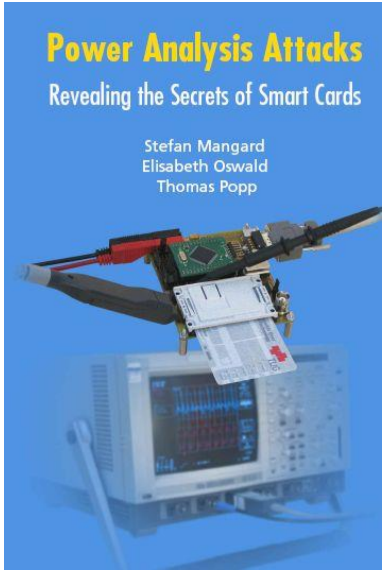

bookstores

?

By Colin & Jasper

nostarch

Key takeaways

- DL does SCA art + science and scales
- DL requires network fiddling, the bar is low, not yet at 0
- Automation needed to put a dent in embedded insecurity

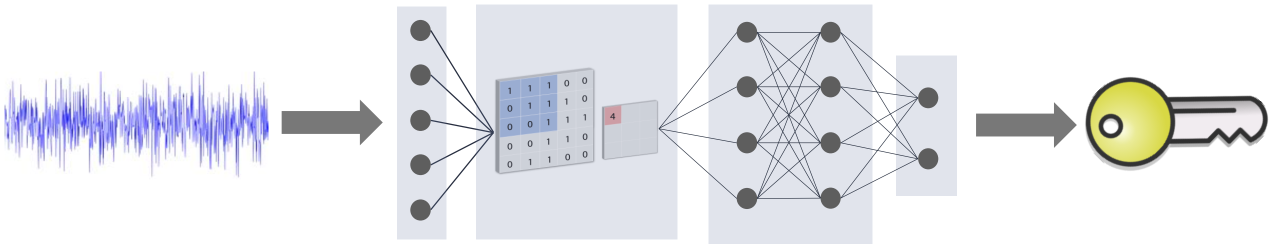

References

- http://www.deeplearningbook.org/
- S. Haykin, “Neural Networks and Learning Machines”.
- E. Cagli et al, “Breaking Cryptographic Implementations Using Deep Learning” https://eprint.iacr.org/2016/921.pdf
- Benadjila et al, “Study of Deep Learning Techniques for Side-Channel Analysis and Introduction to ADCAD Database” https://eprint.iacr.org/2018/053.pdf
- H. Maghrebi et al, “Convolutional Neural Networks with Data Augmentation Against Jitter-Based Countermeasures” https://eprint.iacr.org/2017/740
- Zhang et al, “Understanding deep learning requires re-thinking generalization” https://arxiv.org/abs/1611.03530
- Keskar et al, “On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima”, https://arxiv.org/pdf/1609.04836.pdf
- Shwartz and Tishby, “Opening the black-box of Deep Learning via Information”, https://arxiv.org/abs/1703.00810

Riscure B.V.
Frontier Building, Delftechpark 49
2628 XJ Delft
The Netherlands
Phone: +31 15 251 40 90

www.riscure.com

Riscure North America
550 Kearny St., Suite 330
San Francisco, CA 94108 USA
Phone: +1 650 646 99 79

inforequest@riscure.com

Riscure China
Room 2030-31, No. 989, Changle Road, Shanghai 200031
China
Phone: +86 21 5117 5435

inforcn@riscure.com

riscure

Challenge your security

jasper@riscure.com
@jzvw