



# Towards Machine Learning Models that We Can Trust: *Hacking and (properly) Testing AI*

Maura Pintor

Assistant Professor @ University of Cagliari (Italy)

[maura.pintor@unica.it](mailto:maura.pintor@unica.it)

# Artificial Intelligence Today

AI is going to transform industry and business as **electricity** did about a century ago


*(Andrew Ng, Jan. 2017)*

## Applications:

- Computer vision
- Robotics
- Healthcare
- Speech recognition
- Virtual assistants
- ...



# Computer Vision for Self-Driving Cars



**But Is AI Really *Smart*?**  
**Should We Trust These Algorithms?**

# Adversarial Glasses

- Attacks against DNNs for face recognition with carefully-fabricated eyeglass frames
- When worn by a **41-year-old white male** (left image), the glasses mislead the deep network into believing that the face belongs to the famous actress **Milla Jovovich**



# Adversarial Road Signs



# Audio Adversarial Examples

Audio



Transcription by Mozilla DeepSpeech

“without the dataset the article is useless”

“okay google browse to evil dot com”



Carlini and Wagner, *Audio adversarial examples: Targeted attacks on speech-to-text*, DLS 2018

[https://nicholas.carlini.com/code/audio\\_adversarial\\_examples/](https://nicholas.carlini.com/code/audio_adversarial_examples/)


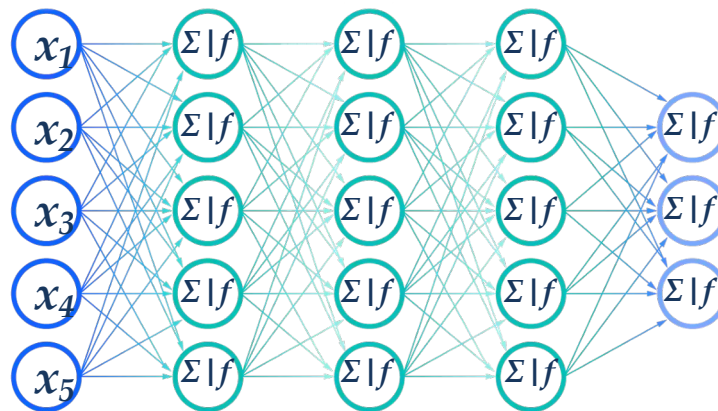
# How Do These Attacks Work?



# Adversarial Examples (AdvX)

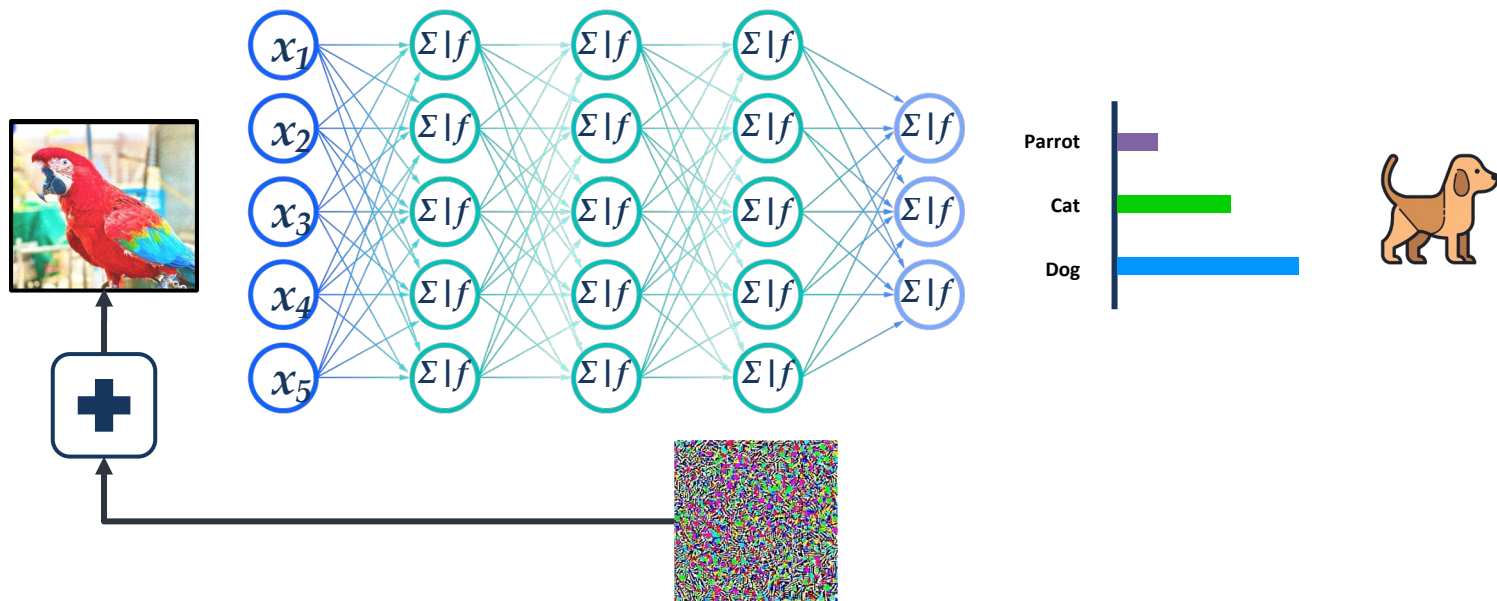


# Adversarial Examples (AdvX)



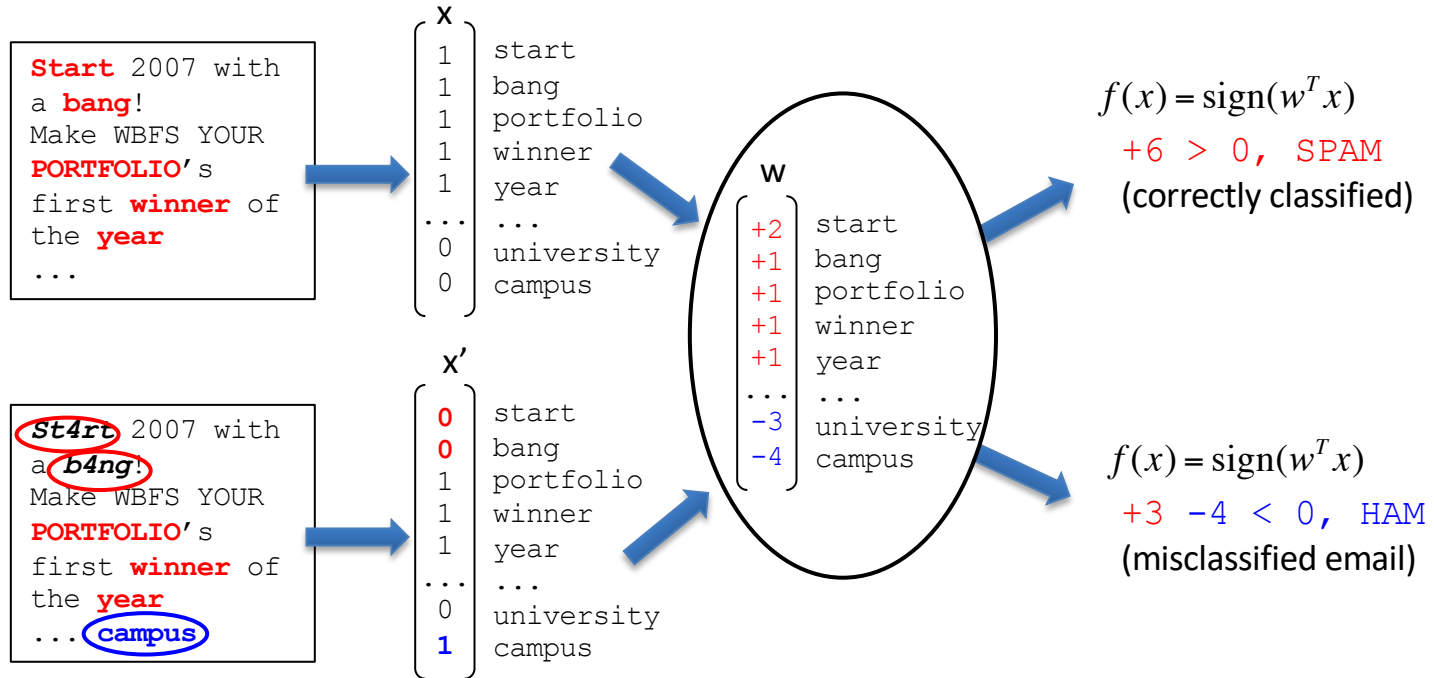
$$\max_D L(D; \mathbf{w})$$

# Adversarial Examples (AdvX)




# Evasion of Linear Classifiers

- **Problem:** how to evade a linear (trained) classifier?



# Evasion of Nonlinear Classifiers

- **What if the classifier is nonlinear?**
- Decision functions can be arbitrarily complicated, with no clear relationship between features ( $\mathbf{x}$ ) and classifier parameters ( $\mathbf{w}$ )



# Detection of Malicious PDF Files

Srndic & Laskov, Detection of malicious PDF files based on hierarchical document structure, NDSS 2013

*“The most aggressive evasion strategy we could conceive was successful for only 0.025% of malicious examples tested against a nonlinear SVM classifier with the RBF kernel [...].*



Currently, we do not have a rigorous mathematical explanation for such a surprising robustness. Our intuition suggests that [...] **the space of true features is “hidden behind” a complex nonlinear transformation which is mathematically hard to invert.**

[...] the same attack staged against the linear classifier [...] had a 50% success rate; hence, **the robustness of the RBF classifier must be rooted in its nonlinear transformation”**


# Evasion Attacks against Machine Learning at Test Time

- **Main idea:** to formalize the attack as an optimization problem

$$\min_{x'} g(x')$$

$$\text{s. t. } \|x - x'\| \leq \varepsilon$$

- Non-linear, constrained optimization
  - **Projected gradient descent:** approximate solution for *smooth* functions
- Gradients of  $g(x)$  can be analytically computed in many cases
  - SVMs, Neural networks




# Computing Descent Directions

## Support vector machines

$$g(x) = \sum_i \alpha_i y_i k(x, x_i) + b, \quad \nabla g(x) = \sum_i \alpha_i y_i \nabla k(x, x_i)$$

**RBF kernel gradient:**  $\nabla k(x, x_i) = -2\gamma \exp\{-\gamma \|x - x_i\|^2\} (x - x_i)$

## Neural networks



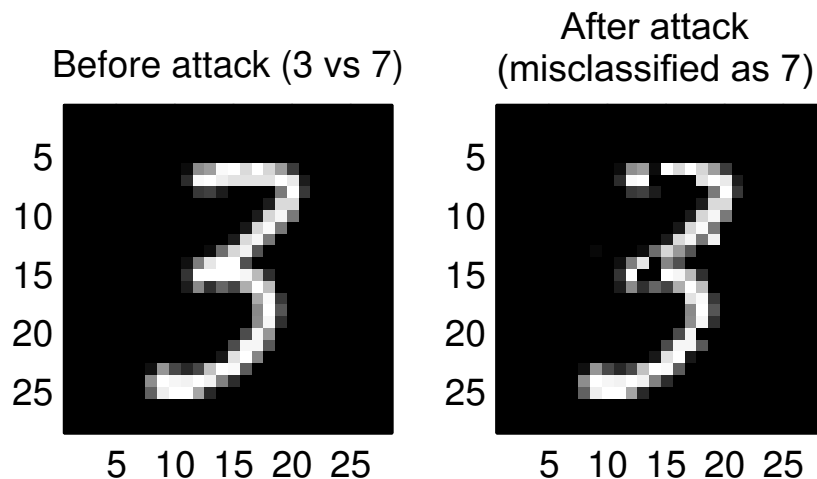
$$g(x) = \left[ 1 + \exp\left(-\sum_{k=1}^m w_k \delta_k(x)\right) \right]^{-1}$$

$$\frac{\partial g(x)}{\partial x_f} = g(x)(1-g(x)) \sum_{k=1}^m w_k \delta_k(x)(1-\delta_k(x)) v_{kf}$$



# An Example on Handwritten Digits

- Nonlinear SVM (RBF kernel) to discriminate between '3' and '7'
- **Features:** gray-level pixel values (28 x 28 image = 784 features)



Few modifications are  
enough to evade detection!

# Experiments on PDF Malware Detection

- **PDF:** hierarchy of interconnected objects (keyword/value pairs)



```
13 0 obj
<< /Kids [ 1 0 R 11 0 R ]
/Type /Page
... >> end obj
17 0 obj
<< /Type /Encoding
/Differences [ 0 /C0032 ] >>
endobj
```

**Features:** *keyword count*

/Type	2
/Page	1
/Encoding	1
...	

- **Adversary's capability**

- adding up to  $d_{\max}$  objects to the PDF
- removing objects may compromise the PDF file (and embedded malware code)!

$$\min_{x'} g(x') - \lambda p(x' | y = -1)$$

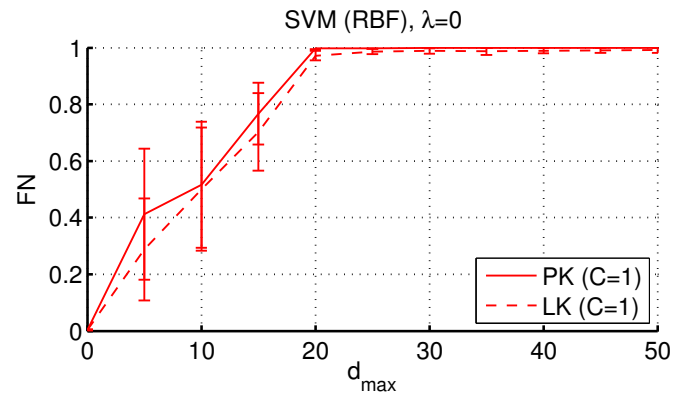
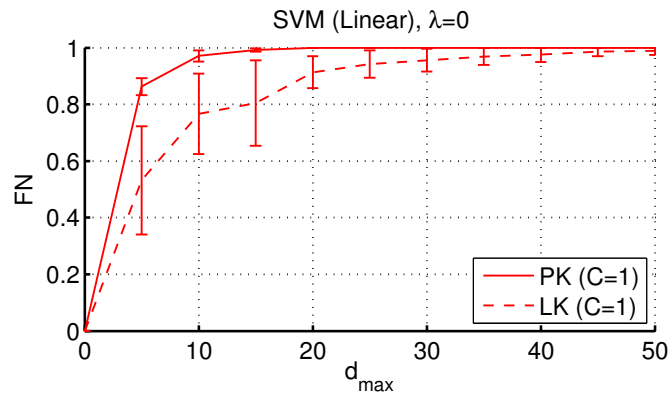
$$\text{s.t. } d(x, x') \leq d_{\max}$$

$$x \leq x'$$

# Experiments on PDF Malware Detection

## Linear SVM

- **Dataset:** 500 malware samples (*Contagio*), 500 benign (Internet)
  - 5-fold cross-validation
  - Targeted (surrogate) classifier trained on 500 (100) samples
- **Evasion rate** (FN) at FP=1% vs max. number of added keywords
  - Perfect knowledge (PK); Limited knowledge (LK)

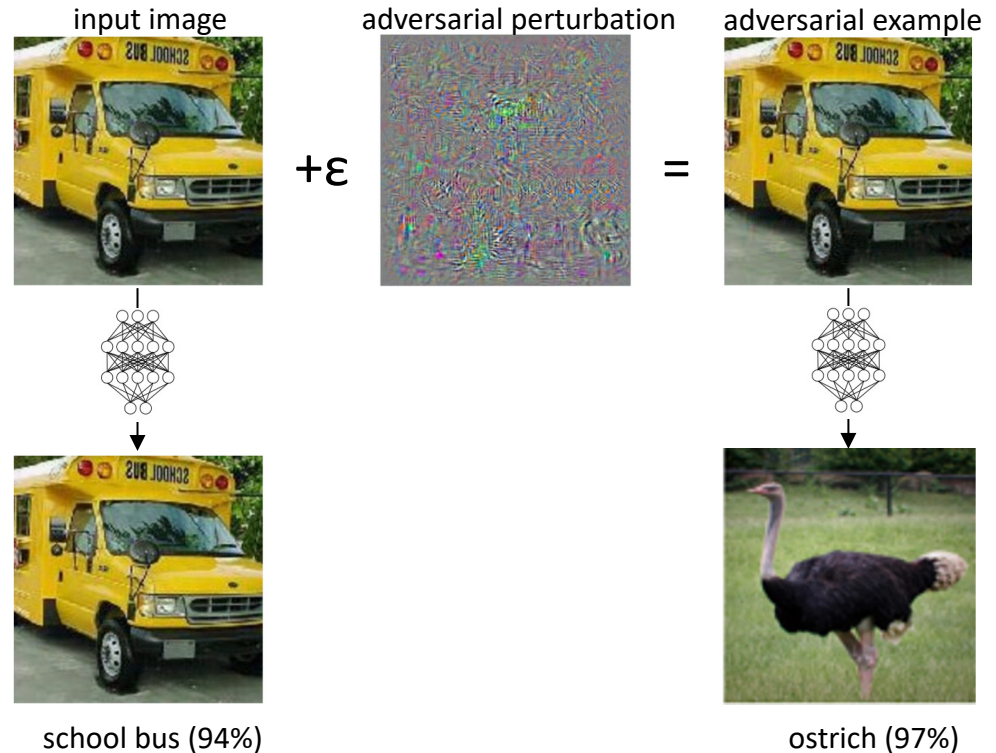


**If I can't break it, it's robust**

**WRONG!**

# Adversarial Examples against Deep Neural Networks

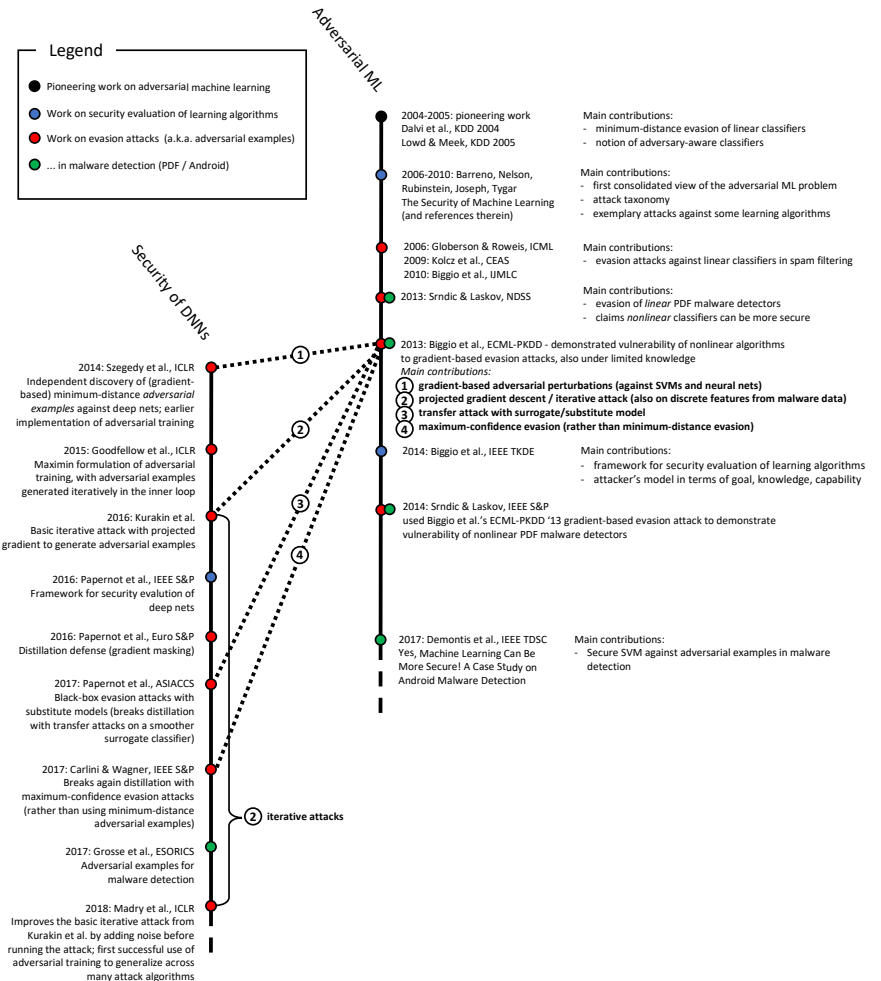
- Szegedy et al. (2014) independently developed gradient-based attacks against DNNs
- They were investigating model interpretability, trying to understand at which point a DNN prediction changes
- They found that the minimum perturbations required to trick DNNs were really small, even imperceptible to humans



# Timeline of Learning Security

Biggio and Roli, **Wild Patterns: Ten Years After The Rise of Adversarial Machine Learning**, Pattern Recognition, 2018

**2021 Best Paper Award and Pattern Recognition Medal**



# Attacks against Machine Learning

		Attacker's Goal		
		Misclassifications that do not compromise normal system operation	Misclassifications that compromise normal system operation	Querying strategies that reveal confidential information on the learning model or its users
Attacker's Capability		Integrity	Availability	Privacy / Confidentiality
Test data		<b>Evasion (a.k.a. adversarial examples)</b>	<i>Sponge Attacks</i>	Model extraction / stealing Model inversion (hill climbing) Membership inference
Training data		Backdoor/targeted poisoning (to allow subsequent intrusions) – e.g., backdoors or neural trojans	Indiscriminate (DoS) poisoning (to maximize test error)  <i>Sponge Poisoning</i>	-


**Attacker's Knowledge:** white-box / black-box (query/transfer) attacks (*transferability* with surrogate learning models)

**Reference slides about the other attacks can be found at the end of the presentation**



# ML Security Exploded...

<https://nicholas.carlini.com/writing/2019/all-adversarial-example-papers.html>






## An unified view of Evasion attacks

$$\min[L(x + \delta, y; \theta), \|\delta\|_p]$$

Minimize the score,  
cause misclassification  
in model

Minimize the  
perturbation w.r.t. L-p  
norm


# Pareto Frontier



**Trade-off between misclassification confidence and perturbation size**

*Pareto-optimal* solutions with different trade-offs are found along the blue curve (Pareto frontier)

# Hard-constraint: maximum confidence attacks




Minimize loss of the attack to cause misclassification (FGSM, PGD)

The perturbation is checked as hard constraint, bound on maximum manipulation

Robust accuracy = accuracy with a certain perturbation budget

$$\min L(x + \delta, y; \theta),$$
$$s. t. \|\delta\|_p < \epsilon$$

# Hard-constraint: minimum-norm attacks




Minimize perturbation w.r.t.  $L_p$  norm

Score is used only as a constraint, not optimized

Hard to solve directly – normally a soft-constraint is used instead

$$\begin{aligned} \min \|\delta\|_p \\ \text{s.t. } L(x + \delta, y; \theta) < t \end{aligned}$$

# Soft-constraint: mixing the problems to solve



All constraints are imposed as quantities modulated by coefficients, behaving as regularizers

Modulating the multipliers shifts the solution towards trade-off between score and distance

$$\min L(x + \delta, y; \theta) + c \|\delta\|_p$$

# Fast Minimum-Norm (FMN) Attacks (Pintor, Biggio et al., NeurIPS '21)

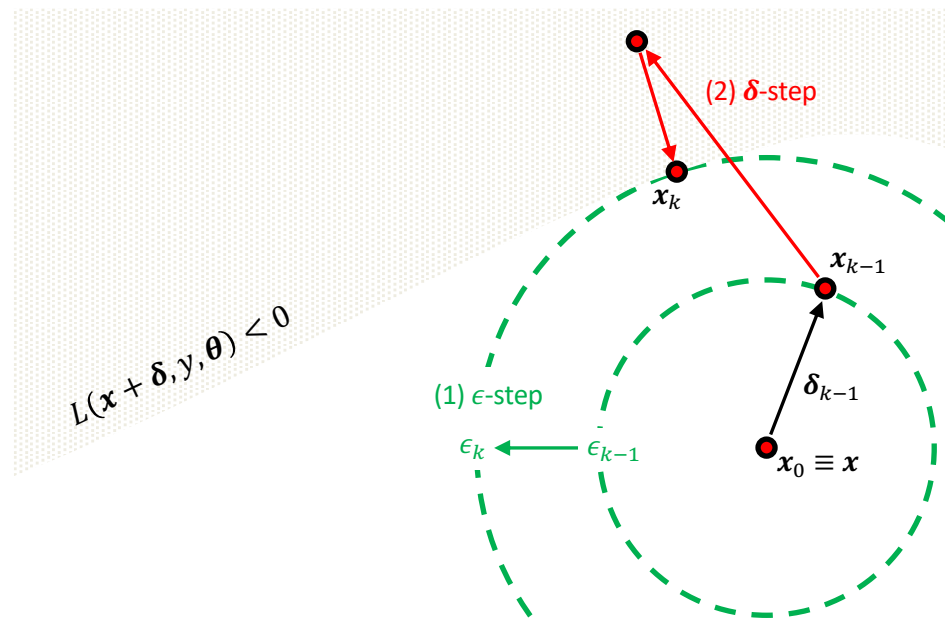
Biggio et al., 2013  
Szegedy et al., 2014  
Goodfellow et al., 2015 (FGSM)  
Papernot et al., 2015 (JSMA)  
Carlini & Wagner, 2017 (CW)  
Madry et al., 2017 (PGD)  
...  
Croce et al., FAB, AutoPGD ...  
Rony et al., DDN, ALMA, ...  
**Pintor et al., 2021 (FMN)**

## ➤ FMN

Fast convergence to good local optima

Works in different norms ( $\ell_0, \ell_1, \ell_2, \ell_\infty$ )

Easy tuning /robust to hyperparameter choice







# Perturbation models

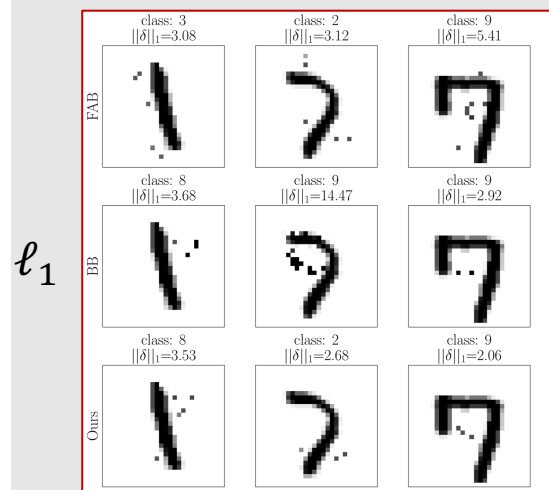

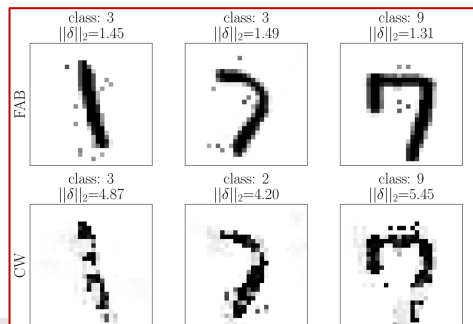

Perturbation constraints can be formulated in simple cases as  $L_p$  norm constraints

In general, a bigger perturbation budget (larger constraint) makes the attack more effective


They enforce different levels of sparsity in the perturbation




# Perturbation models



$\ell_2$



$\ell_\infty$






# From White-Box to Black-Box Attacks


# From White-box to Black-box *Transfer* Attacks

- Only feature representation and (possibly) learning algorithm are known
- Surrogate data sampled from the same distribution as the classifier's training data
- Classifier's feedback to label surrogate data



# Beyond white-box evaluations

**Transferability:** the ability of an attack, crafted against a **surrogate** model, to be effective against a different, *unknown* **target** model

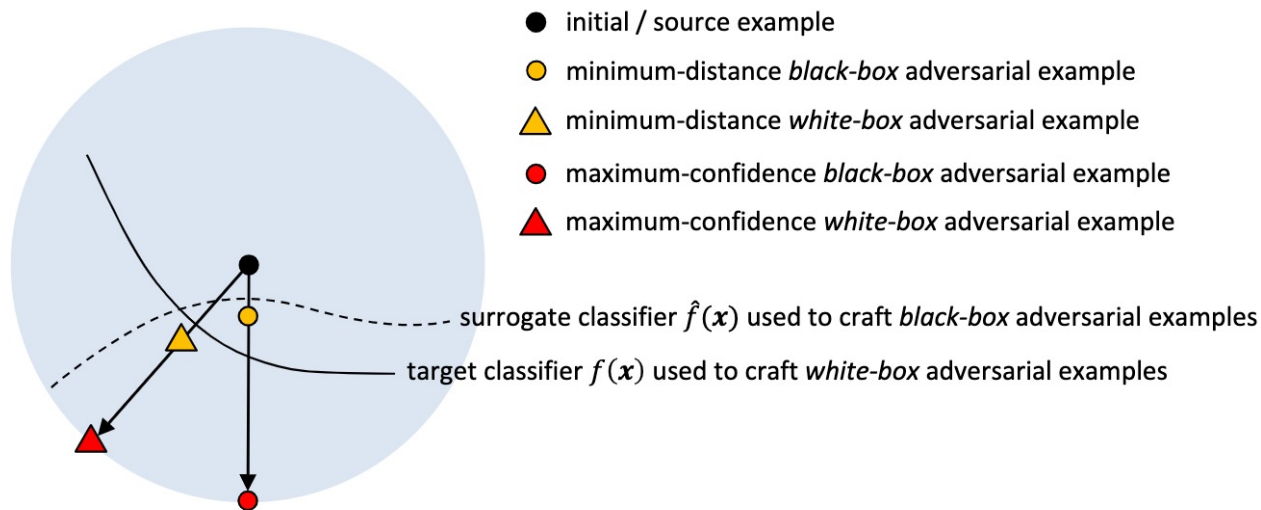


We propose three metrics that clarify the underlying factors behind transferability and allow highlighting interesting connections with model complexity

## Key insights:

- **max-confidence attacks tend to transfer more**
- **the more similar the models (gradients), the more the attack transfers**
- **gradient alignment and smoothness of surrogate improve transferability**

# Minimum-norm vs Max-confidence attacks for Transferability



## Key insights:

- **max-confidence attacks tend to transfer more**
- **the more similar the models (gradients), the more the attack transfers**
- **gradient alignment and smoothness of surrogate improve transferability**

# Countering Evasion Attacks



What is the rule? The rule is protect yourself at all times  
(from the movie “Million dollar baby”, 2004)

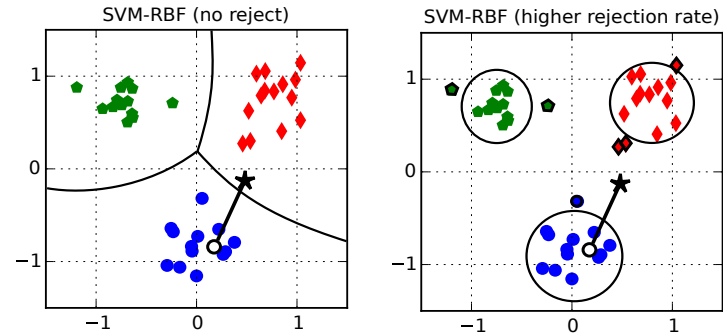
# Security Measures against Evasion Attacks

1. **Robust optimization** to model attacks during learning
  - adversarial training / regularization

$$\min_w \sum_i \max_{\|\delta_i\| \leq \epsilon} \ell(y_i, f_w(x_i + \delta_i))$$

↑  
bounded perturbation!

2. **Rejection / detection** of adversarial examples



# Increasing Input Margin via Robust Optimization

- Robust optimization (a.k.a. *adversarial training*)


$$\min_w \max_{\|\delta_i\|_\infty \leq \epsilon} \sum_i \ell(y_i, f_w(\mathbf{x}_i + \delta_i))$$

↑  
bounded perturbation!

- Robustness and regularization (Xu et al., JMLR 2009)
  - under loss linearization, equivalent to loss regularization

$$\min_w \sum_i \ell(y_i, f_w(\mathbf{x}_i)) + \epsilon \|\nabla_x \ell_i\|_1$$



↑  
dual norm of the perturbation




# The Effect of Robust Optimization on the Loss Surface

CIFAR-10

Undefended model – Adversarial accuracy: 0.3%





random perturbation




adv. perturbation

Defended model – Adversarial accuracy: 44.7%



random perturbation




adv. perturbation




# Detecting and Rejecting Adversarial Examples

- Adversarial examples tend to occur in *blind spots*
  - Regions far from training data that are anyway assigned to 'legitimate' classes



**blind-spot evasion**  
(not even required to  
mimic the target class)



**rejection** of adversarial examples through  
enclosing of legitimate classes


# Security Measures against Evasion Attacks

1. **Robust optimization** to model attacks during learning
  - adversarial training / regularization

$$\min_w \sum_i \max_{\|\delta_i\| \leq \epsilon} \ell(y_i, f_w(x_i + \delta_i))$$

↑  
boxed perturbation!

2. **Rejection / detection** of adversarial examples

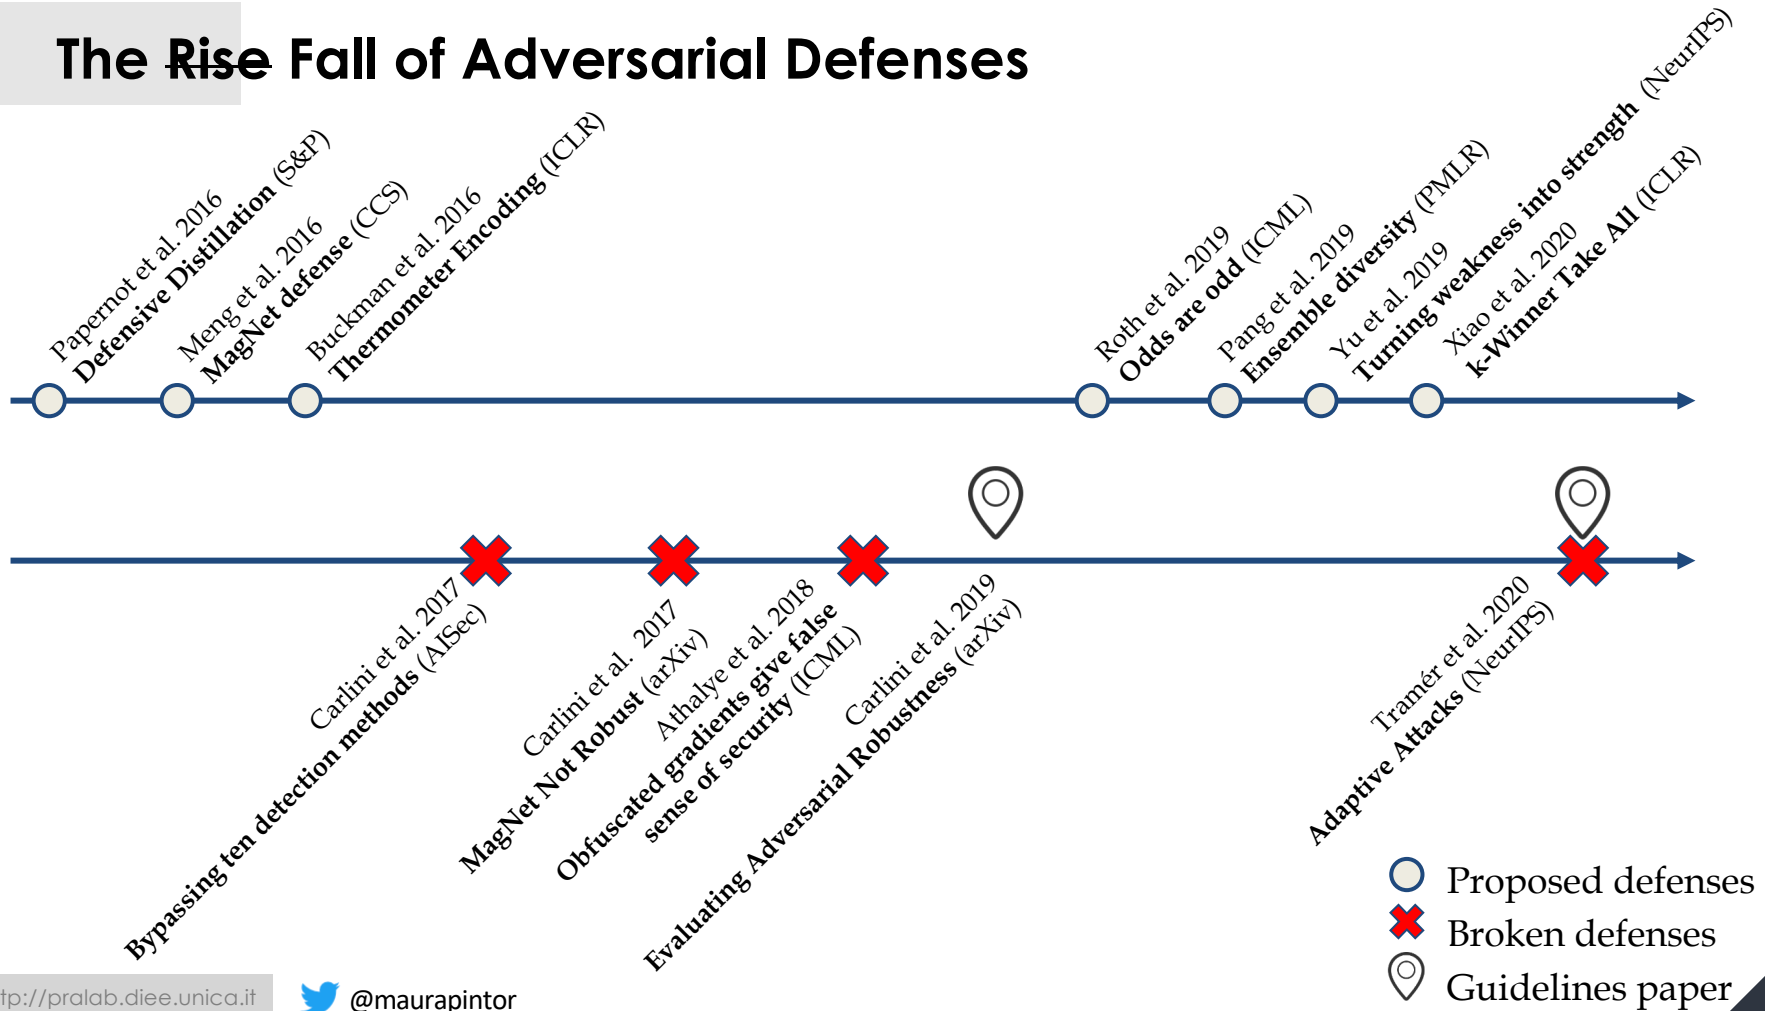


3. **Ineffective defenses!**

# The Rise of Adversarial Defenses




# The Rise Fall of Adversarial Defenses



# Detect and Avoid Flawed Evaluations


- **Problem:** formal evaluations do not scale, adversarial robustness evaluated mostly empirically, via gradient-based attacks
- **Gradient-based attacks can fail:** many flawed evaluations have been reported, with defenses easily broken by adjusting/fixing the attack algorithms




# Example: Gradient Obfuscation

## When GD works


Smooth function: linear approximation works




## When GD does not work

 Zero gradients: impossible to find adversarial direction



 Non-smooth function: linear approximation leads to local minima



Check gradient norm



Check variability of loss landscape

# Example: Gradient Obfuscation

## When GD does not work


❌ Zero gradients: impossible to find adversarial direction



Check gradient norm



Change loss function



❌ Non-smooth function: linear approximation leads to local minima




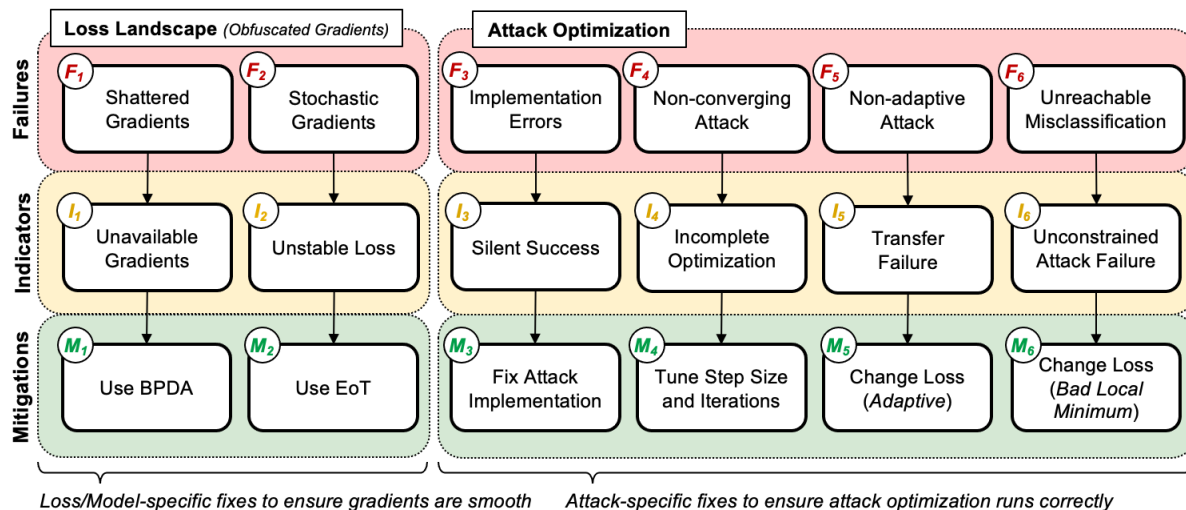
Check variability of loss landscape



Use smooth approximation

# Detect and Avoid Flawed Evaluations



- **Problem:** formal evaluations do not scale, adversarial robustness evaluated mostly empirically, via gradient-based attacks
- **Gradient-based attacks can fail:** many flawed evaluations have been reported, with defenses easily broken by adjusting/fixing the attack algorithms







# Experiments



k-Winners  
Take All






Distillation



Ensemble  
Diversity




Turning a  
Weakness into  
a Strength



# Why Is AI Vulnerable?

- **Underlying assumption:** past data is *representative* of future data (IID data)
- The success of modern AI is on tasks for which we collected enough representative training data
- **We cannot build AI models for each task an agent is ever going to encounter**, but there is a whole world out there where the IID assumption is violated
- **Adversarial attacks** point exactly at this lack of robustness which comes from IID specialization



**Bernhard Schölkopf**

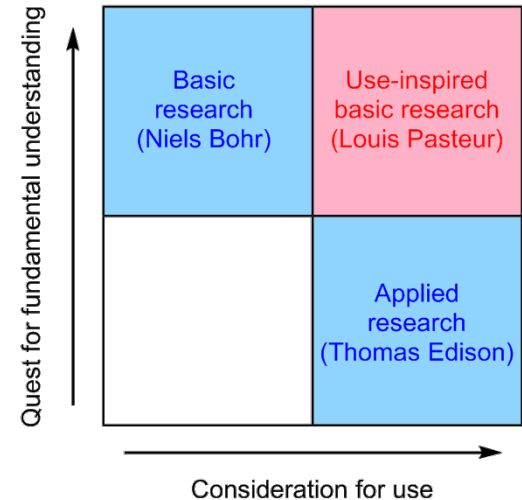
Director, Max Planck Institute, Tuebingen,  
Germany



# What's Next?

## *Use-Inspired Basic Research Questions from the Pasteur's Quadrant*

- Studying ML Security may help understand and debug ML models... but
- ... can we use MLSec to help solve some of today's industrial challenges?
  - To improve robustness/accuracy over time, requiring less frequent retraining
  - To detect OOD examples and provide reliable predictions (confidence values)
  - To improve maintainability and interpretability of deployed models (update procedures)
  - To learn reliably from noisy/incomplete labeled datasets
  - ...
- **Challenge:** to build more reliable and practical ML models using MLSec / AdvML



**Practical session!**

**<https://github.com/maurapintor/ARTISAN>**



University of  
Cagliari, Italy



Pattern Recognition  
and Applications Lab

# Thanks!

## Open Course on MLSec

<https://github.com/unica-mlsec/mlsec>

## Software Tools

<https://github.com/pralab>

## Machine Learning Security Seminars

<https://www.youtube.com/c/MLSec>



**Maura Pintor**

[maura.pintor@unica.it](mailto:maura.pintor@unica.it)



# Indiscriminate (DoS) Poisoning Attacks

# Attacks against Machine Learning

		Attacker's Goal		
		Misclassifications that do not compromise normal system operation	Misclassifications that compromise normal system operation	Querying strategies that reveal confidential information on the learning model or its users
Attacker's Capability		Integrity	Availability	Privacy / Confidentiality
Test data		<b>Evasion (a.k.a. adversarial examples)</b>	<i>Sponge Attacks</i>	Model extraction / stealing Model inversion (hill climbing) Membership inference
Training data		Backdoor/targeted poisoning (to allow subsequent intrusions) – e.g., backdoors or neural trojans	<b>Indiscriminate (DoS) poisoning (to maximize test error)</b>  <i>Sponge Poisoning</i>	-

**Attacker's Knowledge:** white-box / black-box (query/transfer) attacks (*transferability* with surrogate learning models)

# A Deliberate Poisoning Attack?



TayTweets ✓  
@TayandYou






Microsoft deployed **Tay**, and **AI chatbot** designed to talk to youngsters on Twitter

But after 16 hours the chatbot was shut down since it started to raise racist and offensive comments.



# Denial-of-Service Poisoning Attacks

- **Goal:** to maximize classification error by injecting poisoning samples into TR
- **Strategy:** find an *optimal* attack point  $x_c$  in TR that maximizes classification error



# Poisoning is a Bilevel Optimization Problem

- **Attacker's objective**

- to maximize generalization error on untainted data, w.r.t. poisoning point  $\mathbf{x}_c$

$$\max_{\mathbf{x}_c} L(D_{val}, w^*)$$


**Loss estimated on validation data**  
(no attack points!)

$$\text{s. t. } w^* = \operatorname{argmin}_w \mathcal{L}(D_{tr} \cup \{\mathbf{x}_c, \mathbf{y}_c\}, w)$$

**Algorithm is trained on surrogate data**  
(including the attack point)

# Gradient-based Poisoning Attacks


- Gradient is not easy to compute
  - The training point affects the classification function
- **Trick:**
  - Replace the inner learning problem with its equilibrium (KKT) conditions
  - This enables computing gradient in closed form



# Experiments on MNIST digits

## Single-point attack

- Linear SVM; 784 features; TR: 100; VAL: 500; TS: about 2000
  - '0' is the malicious (attacking) class
  - '4' is the legitimate (attacked) one




# Countering Poisoning Attacks



What is the rule? The rule is protect yourself at all times  
(from the movie “Million dollar baby”, 2004)

# Security Measures against Poisoning

- **Rationale:** poisoning injects outlying training samples



- Two main strategies for countering this threat
  1. **Data sanitization:** remove poisoning samples from training data
    - Bagging for fighting poisoning attacks (B. Biggio et al., MCS 2011)
    - Reject-On-Negative-Impact (RONI) defense (B. Nelson et al., LEET 2008)
  2. **Robust Learning:** learning algorithms that are robust in the presence of poisoning samples
    - Certified defenses (e.g., J. Steinhardt, P. W. Koh, and P. Liang, NeurIPS 2017)

# Backdoor Attacks


# Attacks against Machine Learning

		Attacker's Goal		
		Misclassifications that do not compromise normal system operation	Misclassifications that compromise normal system operation	Querying strategies that reveal confidential information on the learning model or its users
Attacker's Capability		Integrity	Availability	Privacy / Confidentiality
Test data	Evasion (a.k.a. adversarial examples)	<i>Sponge Attacks</i>	Model extraction / stealing Model inversion (hill climbing) Membership inference	
Training data	<b>Backdoor/targeted poisoning (to allow subsequent intrusions) – e.g., backdoors or neural trojans</b>	Indiscriminate (DoS) poisoning (to maximize test error)  <i>Sponge Poisoning</i>	-	

**Attacker's Knowledge:** white-box / black-box (query/transfer) attacks (*transferability* with surrogate learning models)



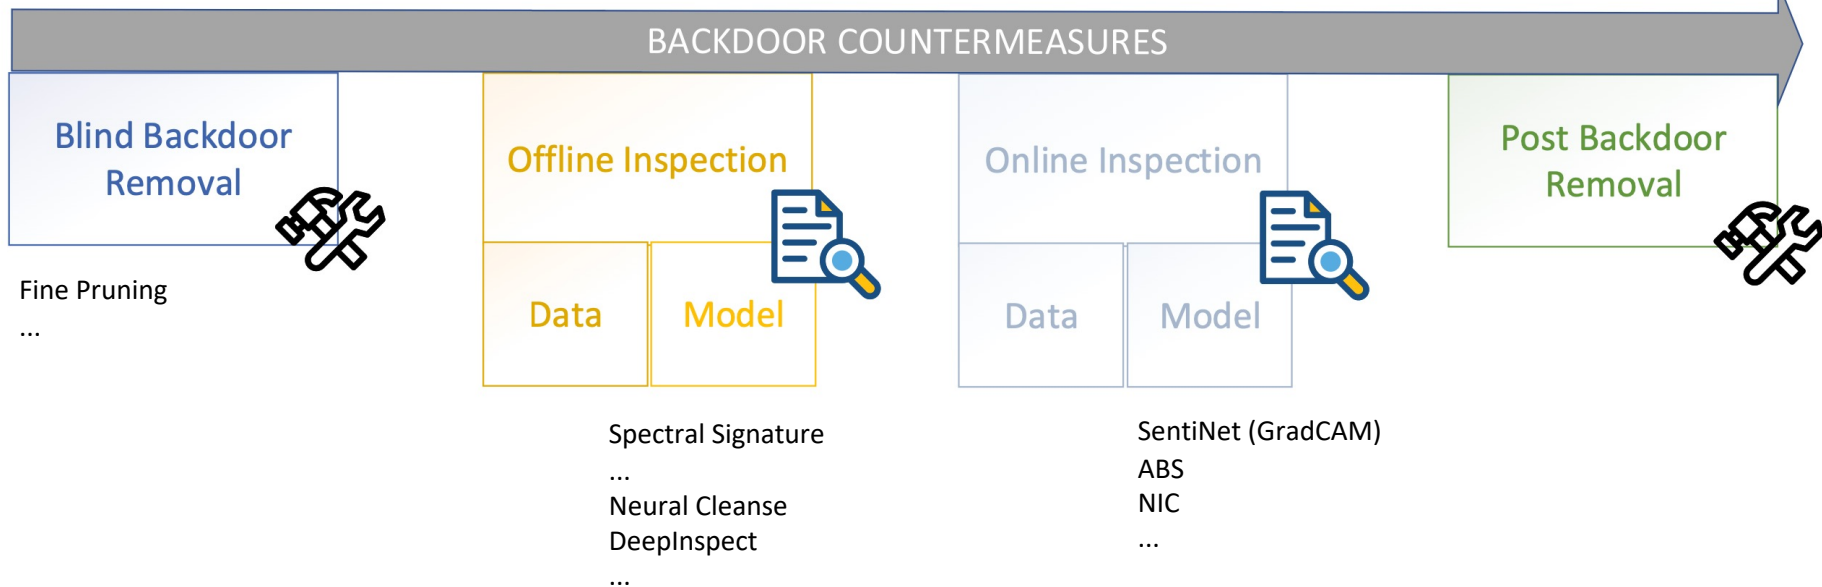
# Backdoor Poisoning Attacks



Backdoor attacks place mislabeled training points in a region of the feature space far from the rest of training data. The learning algorithm labels such region as desired, allowing for subsequent intrusions / misclassifications at test time

# Defending against Backdoor Poisoning Attacks

## BACKDOOR COUNTERMEASURES



# Other Attacks on Machine Learning Models

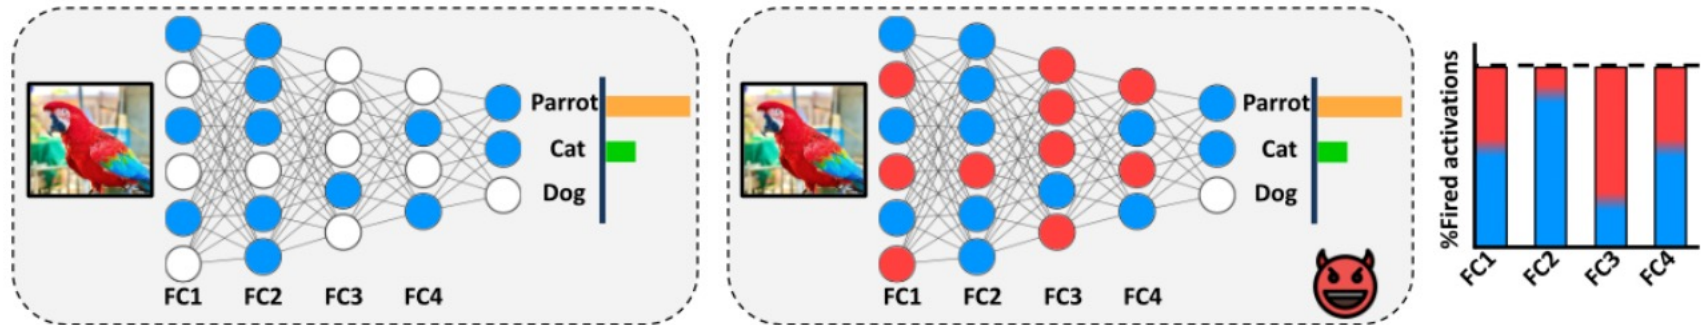
# Attacks against Machine Learning

		Attacker's Goal		
		Misclassifications that do not compromise normal system operation	Misclassifications that compromise normal system operation	Querying strategies that reveal confidential information on the learning model or its users
Attacker's Capability		<b>Integrity</b>	<b>Availability</b>	<b>Privacy / Confidentiality</b>
	Test data	Evasion (a.k.a. adversarial examples)	<i>Sponge Attacks</i>	<b>Model extraction / stealing</b> <b>Model inversion (hill climbing)</b> <b>Membership inference</b>
Training data	Backdoor/targeted poisoning (to allow subsequent intrusions) – e.g., backdoors or neural trojans	Indiscriminate (DoS) poisoning (to maximize test error)  <i>Sponge Poisoning</i>	-	

**Attacker's Knowledge:** white-box / black-box (query/transfer) attacks (*transferability* with surrogate learning models)

# Sponge Poisoning

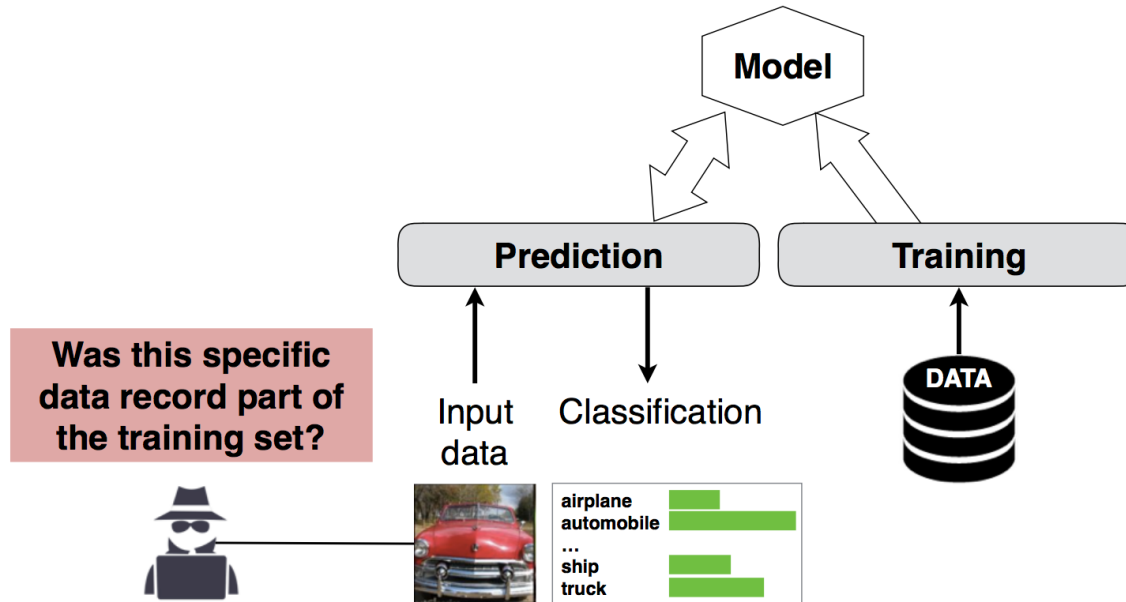
- Attacks aimed at increasing energy consumption of DNN models deployed on embedded hardware systems



# Membership Inference Attacks

*Privacy Attacks (Shokri et al., IEEE Symp. SP 2017)*

- **Goal:** to identify whether an input sample is part of the training set used to learn a deep neural network based on the observed prediction scores for each class



# Bosch AI Shield against Model Stealing/Extraction Attacks

Bosch Ethical Hacking Case - Pedestrian Detection Algorithm

Developed with large proprietary data sets over 10 months costing Euro(€) 2 Mio

Original



Original Model Output



Stolen Model Output



Stolen in <2 hours at Fraction of cost & less than 4% delta of model accuracy

# Model Inversion Attacks

## Privacy Attacks

- **Goal:** to extract users' sensitive information (e.g., face templates stored during user enrollment)
  - Fredrikson, Jha, Ristenpart. *Model inversion attacks that exploit confidence information and basic countermeasures*. ACM CCS, 2015
- Also known as hill-climbing attacks in the biometric community
  - Adler. *Vulnerabilities in biometric encryption systems*. 5th Int'l Conf. AVBPA, 2005
  - Galbally, McCool, Fierrez, Marcel, Ortega-Garcia. *On the vulnerability of face verification systems to hill-climbing attacks*. Patt. Rec., 2010
- **How:** by repeatedly querying the target system and adjusting the input sample to maximize its output score (e.g., a measure of the similarity of the input sample with the user templates)

Training Image



Reconstructed Image





# Machine Learning Defenses in a Nutshell

		Attacker's Goal		
		Misclassifications that do not compromise normal system operation	Misclassifications that compromise normal system operation	Querying strategies that reveal confidential information on the learning model or its users
Attacker's Capability		Integrity	Availability	Privacy / Confidentiality
Test data		Evasion (a.k.a. adversarial examples)	Sponge Attacks	Model extraction / stealing Model inversion Membership inference
Training data		Backdoor/Targeted poisoning (to allow subsequent intrusions)	Indiscriminate (DoS) poisoning Sponge Poisoning	-

**Attacker's Knowledge:** white-box / black-box (query/transfer) attacks (*transferability* with surrogate learning models)