Reliable and Interpretable Artificial Intelligence

Lecture 2: Adversarial Attacks

Martin Vechev ETH Zurich

Fall 2020



http://www.sri.inf.ethz.ch



Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Pointing the need for broader view on AI

Human-Level Intelligence or Animal-Like Abilities? Communications of the ACM, Oct 2018 Adnan Darwiche

https://cacm.acm.org/magazines/2018/10/231373-human-level-intelligence-or-animal-like-abilities/fulltext

"...We need a new generation of AI researchers who are well versed in and appreciate **classical AI**, **machine learning**, and computer science more broadly while also being informed about AI history..."

Today: Adversarial Examples

What are these? More examples in various domains

Why do they exist?

How to generate adversarial examples?

(techniques also used for adversarial training, logic, etc...discussed in later lectures)

"Adversarial examples are inputs to machine learning models that an attacker has intentionally designed to cause the model to make a mistake"

(Goodfellow et al 2017)

Adversarial Examples [from lecture 1]

=

Noisy attack: vision system thinks we now have a gibbon...



x "panda" 57.7% confidence



 $\operatorname{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{y}))$

"nematode" 8.2% confidence



$$\begin{array}{c} x + \\ \epsilon \text{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{y})) \\ \text{"gibbon"} \\ 99.3 \% \text{ confidence} \end{array}$$

Explaining and Harnessing Adversarial Examples, ICLR '15

Tape pieces make network predict a 45mph sign



Robust Physical-World Attacks on Deep Learning Visual Classification, CVPR'18

Self-driving car: in each picture one of the 3 networks makes a mistake...







DRV_C1: right DRV_C2: right

DRV_C3: right

DeepXplore: Automated Whitebox Testing of Deep Learning Systems, SOSP'17

Adversarial Geometric Perturbations



Adversarial Examples (more)



Reese Witherspoon





Russel Crowe

Accessorize to a Crime: Real and Stealthy Attacks on State-of-the-Art Face Recognition, CCS '16

Real World Impersonation/Dodging Attacks

Real glasses





Lujo Bauer



=

100% success



John Malkovich

Adversarial Examples in Reinforcement Learning



An agent (Deep Q Network) plays the game by selecting actions from a given state (image) that the game produces.

An attacker can perturb the image slightly so that the DQN agent chooses the wrong action: here, it wrongly picks noop (do nothing) in the right image, instead of moving the paddle down (left image).

Adversarial Examples in NLP

Article: Super Bowl 50

Paragraph: "Peyton Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver's Executive Vice President of Football Operations and General Manager. Quarterback Jeff Dean had jersey number 37 in Champ Bowl XXXIV."

Question: "What is the name of the quarterback who was 38 in Super Bowl XXXIII?" Original Prediction: John Elway Prediction under adversary: Jeff Dean

The Ensemble model is fooled by the addition of an adversarial distracting sentence in blue.

Adversarial Examples in Audio Processing: Speech to Text

An attack on DeepSpeech:



Adding small noise to the input audio makes the network transcribe any arbitrary phrase

Audio Adversarial Examples: Targeted Attacks on Speech-to-Text, ICML 2018 workshop

Adversarial Examples in Audio Processing: Text classification



Adding small noise to the input audio changes the classification of the output

Adversarial Examples: Some History

2006: Deep learning models gain renewed interest

2012: Multiple works showed that deep networks can achieve near-human performance (sometimes even better)

2013: Research in understanding neural networks behavior becomes critical with society implications beyond computer science

2014: While trying to understand decision making in neural networks, Szegedy et al. discovered adversarial examples

2015-on: Finding adversarial examples and proving their absence becomes an active research area...

Robustness

Robustness: A network is **robust** if it returns correct output on all inputs

Impractical: the input space is too large to be covered

Local Robustness (informal): A learning model is locally-robust if it returns the correct output on inputs *similar* to inputs in the training set

This was believed to be evident by having high accuracy on the test set

Why is High Accuracy Not Enough?

Inputs in the training and test set are taken from a given distribution

Neural networks aim to achieve high accuracy on test sets drawn from the given distribution

There are still many similar inputs that are never tested (and have low-probability for the given distribution)

The Story of Clever Hans



Why Do Adversarial Examples Exist?

Neural Networks are too linear



Why are they designed to be linear?

Linear functions are easy to optimize!

(Goodfellow 2016)

Neural Networks are too linear

Learned Model not powerful enough to fit data properly

Learned Model trained on adversarial examples now more powerful (less linear)

(Madry 2017)

Experimental Linearity of Perturbations



- We have 1 image and 10 classes (0 to 9). The correct classification is 4.
- x-axis is perturbation (chose direction and perturb the image by that direction).
- y-axis are logits (values before calling softmax), unnormalized probabilities.

We see that the function which computes the particular logit, is basically almost (piece wise) linear in the perturbation.

Only around 0 (no perturbation) does the function behave in a non-linear manner and where classification is correct (i.e. 4).

Explaining and harnessing adversarial examples, Ian J. Goodfellow, Jonathon Shlens & Christian Szegedy, ICLR'15

Generating Adversarial Examples (somewhat possible due to these mostly linear properties)

Targeted vs. Untargeted Attacks

Targeted Attack – aims to misclassify the input (e.g., image) to a specific label (e.g. panda to gibbon)

Untargeted Attack – aims to misclassify the input to any wrong label (e.g. panda to any other animal)

Formulated as a slightly different optimization problem

Targeted Attack: Problem Statement

Input:

- neural network $f: X \to C$
- input $x \in X$
- target label $t \in C$, such that $f(x) \neq t$

Output:

• A perturbation η such that $f(x + \eta) = t$

Adversarial example $x' = x + \eta$

Untargeted Attack: Problem Statement

Input:

- neural network $f: X \to C$
- input $x \in X$

Output:

• A perturbation η such that $f(x + \eta) \neq f(x)$

Adversarial example $x' = x + \eta$

Types of Attacks

White box attacker: the attacker knows the model, the parameters, and the architecture

Black box attacker: the attacker knows the architecture (e.g., the layers) but not its parameters (e.g., weights)

Note: it was found adversarial examples are **transferrable**, hence given the same training data as the original network, an attacker can train their own **mirror network** of the black box original network and then attack the mirror network with white-box techniques. If attack on mirror network succeeds, it will likely succeed on the original.

We will look at white box attacks

Targeted Fast Gradient Sign Method

1. Compute perturbation:

 $\eta = \epsilon \cdot \operatorname{sign}(\nabla_x \operatorname{loss}_t(x))$, where

$$\nabla_{x} \operatorname{loss}_{t} = \left(\frac{\partial \operatorname{loss}_{t}}{\partial x_{1}}, \dots, \frac{\partial \operatorname{loss}_{t}}{\partial x_{n}}\right) \quad \operatorname{sign}(g) = \begin{cases} -1, \text{ if } g < 0\\ 0, \text{ if } g = 0\\ 1, \text{ if } g > 0 \end{cases}$$

2. Perturb the input:

$$x' = x - \eta$$

3. Check if:

$$f(x') = t$$

- Here, each x_i is a pixel
- ϵ is a very small constant (e.g., 0.007)
- As FGSM is 1-step, x' is guaranteed to stay inside the box [x ε, x + ε], so no need to project.
- *t* is the target, bad label
- loss_t is the loss w.r.t target label
- FGSM was designed to be fast, not optimal (may not compute minimal perturbation)

Untargeted version of FGSM

1. Compute perturbation:

 $\eta = \epsilon \cdot \operatorname{sign}(\nabla_x \operatorname{loss}_s(x))$

2. Perturb the input:

 $x' = x + \eta$

- 3. Check if:
 - $f(x') \neq \underline{s}$

- With untargeted FGSM, we do not know what the target (bad) label is that we want.
- We just want some label different than the correct label s.
- So we try to "get away" from the correct label by maximizing the value of the loss

FGSM



Explaining and Harnessing Adversarial Examples, ICLR'15

Importance of Small Perturbations



We need some notion of distance....

Norm: Notion of Distance

Similarity of $\mathbf{x} \sim \mathbf{x}'$ is usually captured by an l_p norm:

$$x \sim x'$$
 iff $||x - x'||_p < \epsilon$,

where
$$||x - x'||_p = ((|x_1 - x'_1|)^p + \dots + (|x_n - x'_n|)^p)^{\frac{1}{p}}$$

 l_0 (when $0^0 = 0$ and we get rid of 1/p root) captures the number of changed pixels.

 l_2 captures the Euclidian distance between x and x'. It can remain small if there are many small changes to many pixels.

 l_{∞} captures **maximum noise (change)** added to any coordinate. It is the maximum of the absolute values of the entries:

$$||x - x'||_{\infty} = max(|x_1 - x'_1|, \dots, |x_n - x'_n|)$$

This is the most common norm used for adversarial example generation and it is argued that it most naturally captures human vision.

To derive the max equation, see:

https://math.stackexchange.com/questions/3099179/proving-the-infinity-norm-is-equal-to-the-maximum-value-of-the-vector&sa=D&source=hangouts&ust=1600861806197000&usg=AFQjCNFLvo4hpAfiNBK06EwAZVFkuRiDNw

Targeted Attack with Small Changes

Input:

- neural network $f: X \to C$
- input $x \in X$
- target label $t \in C$, such that $f(x) \neq t$

Output:

- A perturbation η such that $f(x + \eta) = t$
- $\|\boldsymbol{\eta}\|_p$ is minimized

The problem of generating small perturbations can be phrased as an optimization problem:

 $p \in \{0, 2, \infty\}$

find η minimize $\|\eta\|_p$ such that $f(x + \eta) = t$ $x + \eta \in [0,1]^n$ This is a hard discrete constraint which is difficult to optimize for with gradient methods.

Note: η can have negative components.

Key insight: Relaxation of the hard constraint

Carlini et al., Towards Evaluating the Robustness of Neural Networks, 2017

Two steps:

Step 1: Define an objective function obj_t such that: if $obj_t(x + \eta) \le 0$ then $f(x + \eta) = t$

Step 2: Solve the following optimization problem:

find	η
minimize	$\ \boldsymbol{\eta}\ _p + c \cdot \boldsymbol{obj}_t(\boldsymbol{x} + \boldsymbol{\eta})$
such that	$x+\eta\in [0,1]^n$

Two steps:

Step 1: Define an objective function obj_t such that: if $obj_t(x + \eta) \le 0$ then $f(x + \eta) = t$

What are examples of functions for *obj* with the property of Step 1?

Choice I:	Lets take cross
$obj_t(x') = loss_t(x') - 1$	entropy loss for loss _t
Choice II: $obj_t(x') = \max(0, 0.5 - \mathbf{p}_f(x')_t)$	$\mathbf{p_f}(\mathbf{x}')_t$ returns the probability of class \mathbf{t} for input \mathbf{x}' on network f

Choice I: $obj(x) = loss_t(x) - 1$

Choice I:

$$obj_t(x) = loss_t(x) - 1$$

= $-log_2(\mathbf{p}(t)) - 1$



Plug in cross entropy loss for $loss_t$ with logarithm base 2

Here, we use $\mathbf{p}(t)$ as a shortcut for $\mathbf{p}_{\mathbf{f}}(\mathbf{x})_{t}$ so to avoid clutter

What we see here is that if the obj_t function is 0 or negative, then the probability $\mathbf{p}(t)$ is ≥ 0.5 (50%).

But if $\mathbf{p}(t)$ is ≥ 0.5 for the input x, then f will return t as a classification for x because this is the highest probability class. Hence, the desired property of Step 1 holds.

Choice II: $\max(0, 0.5 - \mathbf{p}_{f}(x)_{t})$

Choice II:

 $obj_t(x) = \max(0, 0.5 - \mathbf{p}(t))$



Again we use $\mathbf{p}(t)$ as a shortcut for $\mathbf{p}_{\mathbf{f}}(\mathbf{x})_{t}$ so to avoid clutter

What we see here is that the obj_t function is always 0 or greater.

It is only 0 when $\mathbf{p}(t)$ is ≥ 0.5 for the input \mathbf{x} .

Again, then f will return t as a classification for x because this is the highest probability class.

Hence, the desired property holds for Step 1.

Two steps:

Step 1: Define an objective function obj_t such that: if $obj_t(x + \eta) \le 0$ then $f(x + \eta) = t$

Step 2: Solve the following optimization problem:

find	η
minimize	$\ \boldsymbol{\eta}\ _p + c \cdot obj_t(\boldsymbol{x} + \boldsymbol{\eta})$
such that	$x+\eta\in [0,1]^n$

Two steps:

Step 1: Define an objective function obj_t such that: if $obj_t(x + \eta) \le 0$ then $f(x + \eta) = t$

Step 2: Solve the following optimization problem:



Lets take a closer look at $\|\eta\|_{\infty}$

 $\|\eta\|_{\infty}$ computes the maximum change: it takes the absolute value of every coordinate in η and returns the maximum value.

$$\boldsymbol{\eta} = (0.5, 0.49, 0.48) \qquad \qquad \frac{\partial \|\boldsymbol{\eta}\|_{\infty}}{\partial \boldsymbol{\eta}_1} = 1 \qquad \frac{\partial \|\boldsymbol{\eta}\|_{\infty}}{\partial \boldsymbol{\eta}_2} = 0 \qquad \frac{\partial \|\boldsymbol{\eta}\|_{\infty}}{\partial \boldsymbol{\eta}_3} = 0$$

After one step we get: $\eta = \eta - \gamma \cdot (1,0,0) = \eta - 0.03 \cdot (1,0,0) = (0.47, 0.49, 0.48)$ After two steps we get: $\eta = \eta - \gamma \cdot (0,1,0) = \eta - 0.03 \cdot (0,1,0) = (0.47, 0.46, 0.48)$ After three steps we get: $\eta = \eta - \gamma \cdot (0,0,1) = \eta - 0.03 \cdot (0,0,1) = (0.47, 0.46, 0.45)$

What we see is that because the gradient is 0 at all non-max locations, the gradient does not impose a penalty on the optimizer increasing a little bit those locations (due to the $obj_t(x + \eta)$ term in the optimization). Also, only one entry is changed at a time.

Lets take a closer look at $\|\eta\|_{\infty}$

Going back to the full optimization problem which also includes $obj(x + \eta)$

After one step we get: $\eta = \eta - \gamma \cdot (1,0,0) = \eta - 0.03 \cdot (1,0,0) = (0.47, 0.49, 0.48)$

Now the optimizer can slightly bump up the second location:

After one full step of optimizer, it may also bump up the 2 nd location:	(<mark>0.47, 0.5</mark> , 0.48)
After second full step of optimizer, we may get:	(0.5 , 0.47 , 0.48)

Well, we are just oscillating now and bouncing around...turns out SGD may not be a good way to optimize $\|\eta\|_{\infty}$ especially with other terms

One approach to solving the issue

Replace $\|\boldsymbol{\eta}\|_{\infty}$ with other proxy functions that reflect the distance One idea is to penalize large values in $\boldsymbol{\eta}$ via a term τ :

Replace $\|\boldsymbol{\eta}\|_{\infty}$ with $\sum_{i} \max(0, (|\boldsymbol{\eta}_{i}| - \tau))$

- τ is basically intended to capture the $\|\eta\|_{\infty}$ bound when optimization finishes.
- τ will be continuously minimized. Initially, τ starts at 1
- τ is decreased with some factor (say 0.9) at every iteration if all $|\eta_i|$ are less than τ (then, entire expression will be 0).

Note: an iteration consist of K small steps.

Note: when τ is large, gradient of $\sum_{i} \max(0, (|\boldsymbol{\eta}_{i}| - \tau))$ is similar to gradient of $\|\boldsymbol{\eta}\|_{\infty}$.

Example & Notes on Optimization

Let $L(\eta) = \sum_{i} \max(0, (|\eta_{i}| - \tau))$ and $\eta = (0.47, 0.49, 0.48)$

Start with	τ = 1,	then $L(\boldsymbol{\eta}) = 0$
Next iteration:	τ = 0.9,	then $L(\boldsymbol{\eta}) = 0$
Next iteration:	$\tau = 0.81,$	then $L(\boldsymbol{\eta}) = 0$

Here we only show the optimization of one term, namely $L(\eta)$, to illustrate the relationship with optimizing $\|\eta\|_{\infty}$

At some iteration: $\tau = 0.478$, now via one step, we get:

 $L(\boldsymbol{\eta}) = \sum_{i} \max(0, (|\boldsymbol{\eta}_{i}| - 0.478)) = \sum_{i} (0, 0.012, 0.002) = 0.014$ $\nabla_{\boldsymbol{\eta}} L(\boldsymbol{\eta}) = (0, 1, 1)$

We can then update η as usual, complete this step, and continue with the next step

Notes on optimization:

...

- There are K steps within an iteration, each updating η .
- Entire optimization stops if after K steps $L(\eta) \neq 0$. Otherwise, if $L(\eta) = 0$, optimization continues with a new $\tau = 0.9$ * previous τ .
- Entire optimization stops if it also reaches some pre-defined value of τ (1/256 for Carlini & Wagner).
- When optimization stops, we return η at the iteration before the last one. This means $\|\eta\|_{\infty} \leq \tau$ where τ is the one used at iteration before last.
- If $\eta_{top2} < \tau < \eta_{top1}$ where η_{top1} is the largest element and η_{top2} is second largest element in η , then the gradient $\nabla_{\eta} L(\eta)$ will be the same as $\nabla_{\eta} \|\eta\|_{\infty}$

Summary: Optimization Problem

Two steps:

Step 1: Define an objective function obj_t such that: if $obj_t(x + \eta) \le 0$ then $f(x + \eta) = t$

Step 2: Solve the following optimization problem:

find	η
minimize	$\ \boldsymbol{\eta}\ _{\infty} + c \cdot obj_t(\boldsymbol{x} + \boldsymbol{\eta})$
such that	$x+\eta\in [0,1]^n$

Lecture Summary

Deep Learning is susceptible to adversarial examples in various domains



 $+.007 \times$

 \boldsymbol{x} "panda" 57.7% confidence



 $sign(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{y}))$ "nematode" 8.2% confidence



=

x + $\epsilon \operatorname{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{y}))$ "gibbon" 99.3 % confidence

Generating Adversarial examples (basically, an optimization problem)

- FGSM: targeted and untargeted
- small perturbation attacks
- Need suitable optimization problem

Next lecture: Dealing with Constraints and Adversarial defenses