Institute of Policy Studies

ENGAGING MINDS, EXCHANGING IDEAS

Closed-Door Discussion on The Digital Economy and Challenges in the Process of Transformation

Thursday, 20 July 2017

Conference Room, Level 1, Oei Tiong Ham Building

Lee Kuan Yew School of Public Policy

National University of Singapore

THE IMPACT OF AI TO THE DIGITAL ECONOMY AND SECURITY

Presentation By

Dr Simon See

Chief Solution Architect & Director

NVIDIA AI Technology Centre and Solution Architecture and Engineering

Lee Kuan Yew School of Public Policy

National University of Singapore

Institute of Policy Studies

Engaging Minds, Exchanging Ideas

AI IMPACT OF DIGITAL ECONOMY

AMAZING ACHIEVEMENTS IN AI

Play Go

Play Doom

Learn Paint Style

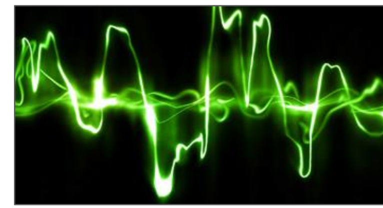

Synthesize Voice

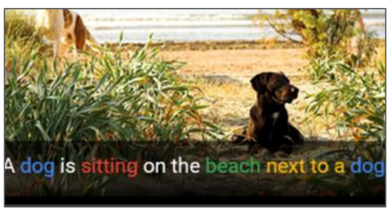

Write Captions

Learn Motor Skills

Learn to Walk

Drive

AI COMPUTING ECOSYSTEM

Amazon Baidu eBay Facebook

Flickr Google iQIYI JD.com

Microsoft Netflix Periscope Pinterest

Qihoo 360 Shazam Skype Sogou

Tencent Twitter Yandex Yelp

AI-powered Consumer Services

AI-as-a-Service

AI for Enterprise

AI for Auto

>1,500 AI Startups

SMART AND SAFE CITIES NEEDS AI

1 billion cameras worldwide (2020)

- Public Safety
- Traffic Management
- Public Transit
- Retail Analytics
- Law Enforcement
- Forensics

10’s of exabytes of data per day

30 billion frames per second

AI ACHIEVES SUPERHUMAN PERFORMANCE

**Delving Deep into Rectifiers:
Surpassing Human-Level Performance on ImageNet Classification**

Kaiming He   Xiangyu Zhang   Shaoqing Ren   Jian Sun
Microsoft Research

Abstract

Rectified activation units (rectifiers) are essential for state-of-the-art neural networks. In this work, we study rectifier neural networks for image classification from two aspects. First, we propose a Parametric Rectified Linear Unit (PReLU) that generalizes the traditional rectified unit. PReLU improves model fitting with nearly zero extra computational cost and little overfitting risk. Second, we derive a robust initialization method that particularly considers the rectifier nonlinearities. This method enables us to train extremely deep rectified models directly from scratch and to investigate deeper or wider network architectures. Based on the learnable activation and advanced initialization, we achieve ***4.94%*** *top-5 test error on the ImageNet 2012 classification dataset. This is a 26% relative improvement over the ILSVRC 2014 winner (GoogLeNet, 6.66% [33]). To our knowledge, our result is the first[1] to surpass the reported human-level performance (5.1%, [26]) on this dataset.*

1. Introduction

Convolutional neural networks (CNNs) [19, 18] have demonstrated recognition accuracy better than or comparable to humans in several visual recognition tasks, including recognizing traffic signs [3], faces [34, 32], and handwritten digits [3, 36]. In this work, we present a result that surpasses the human-level performance reported by [26] on a more generic and challenging recognition task - the classification task in the 1000-class ImageNet dataset [26].

In the last few years, we have witnessed tremendous improvements in recognition performance, mainly due to advances in two technical directions: building more powerful models, and designing effective strategies against overfitting. On one hand, neural networks are becoming more capable of fitting training data, because of increased complexity (*e.g.*, increased depth [29, 33], enlarged width [37, 28], and the use of smaller strides [37, 28, 2, 29]), new nonlinear activations [24, 23, 38, 22, 31, 10], and sophisticated layer designs [33, 12]. On the other hand, better generalization is achieved by effective regularization techniques [13, 30, 10, 36], aggressive data augmentation [18, 14, 29, 33], and large-scale data [4, 26].

Among these advances, the rectifier neuron [24, 9, 23, 38], *e.g.*, Rectified Linear Unit (ReLU), is one of several keys to the recent success of deep networks [18]. It expedites convergence of the training procedure [18] and leads to better solutions [24, 9, 23, 38] than conventional sigmoid-like units. Despite the prevalence of rectifier networks, recent improvements of models [37, 28, 12, 29, 33] and theoretical guidelines for training them [8, 27] have rarely focused on the properties of the rectifiers.

Unlike traditional sigmoid-like units, *ReLU is not a symmetric function*. As a consequence, the mean response of ReLU is always no smaller than zero; besides, even assuming the inputs/weights are subject to symmetric distributions, the distributions of responses can still be asymmetric because of the behavior of ReLU. These properties of ReLU influence the theoretical analysis of convergence and empirical performance, as we will demonstrate.

In this paper, we investigate neural networks from two aspects particularly driven by the rectifier properties. First, we propose a new extension of ReLU, which we call *Parametric Rectified Linear Unit* (PReLU). This activation function adaptively learns the parameters of the rectifiers, and improves accuracy at negligible extra computational cost. Second, we study the difficulty of training rectified models that are very deep. By explicitly modeling the nonlinearity of rectifiers (ReLU/PReLU), we derive a theoretically sound initialization method, which helps with convergence of very deep models (*e.g.*, with 30 weight layers) trained directly from scratch. This gives us more flexibility to explore more powerful network architectures.

On the 1000-class ImageNet 2012 dataset, our network leads to a single-model result of 5.71% top-5 error, which surpasses all multi-model results in ILSVRC 2014. Further, our multi-model result achieves **4.94%** top-5 error on the test set, which is a 26% relative improvement over the ILSVRC 2014 winner (GoogLeNet, 6.66% [33]). To the best of our knowledge, our result surpasses for the first time the reported human-level performance (5.1% in [26]) of a dedicated individual labeler on this recognition challenge.

[1] reported in Feb. 2015.

1026

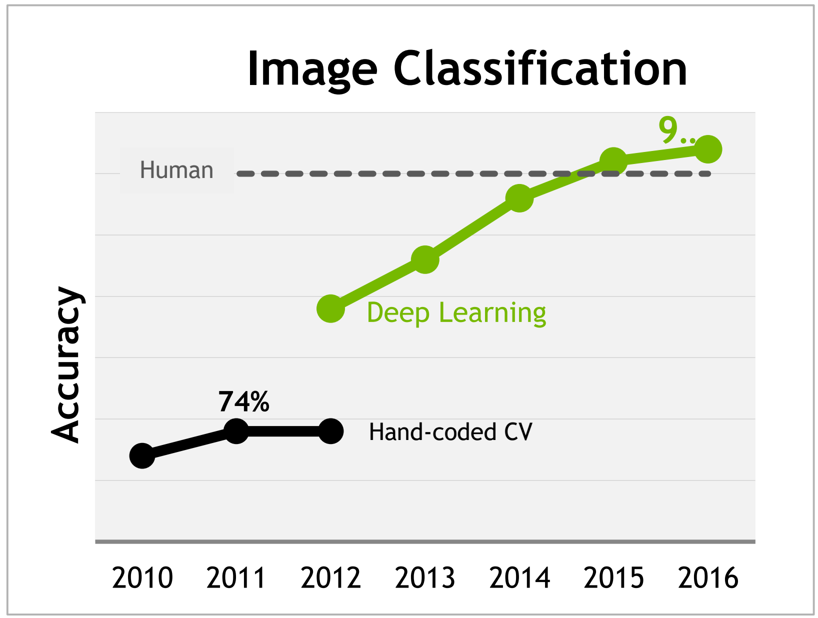

AI COMPUTING ADOPTION

World's first search by example for comm. sec.

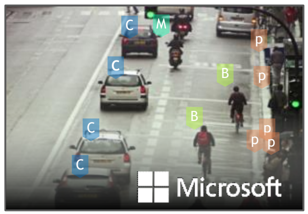

Super-human image classification

6x Improvement for pedestrian detection in rain

5x speed up for ALPR

2x stream recording density

10x speed up in vehicle attribute classification

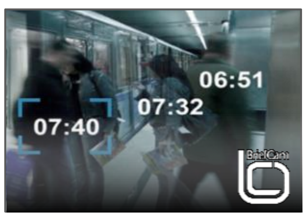

30x faster than realtime video synopsis

11x boost in investigation productivity

30x speedup in people and attribute detection

World leading object detection

EXAMPLES — AI FOR SMART CITIES

Traffic monitoring

Driver Analytics

Parking Management

Hyperscale Analytics

Video Management

Video Enhancement

Super Resolution

EXAMPLE — AI FOR SAFE CITIES

Secure Premises

Public Safety

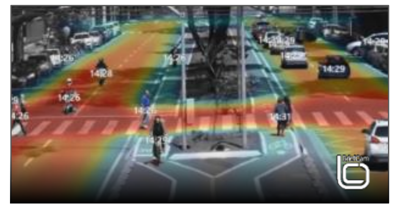

Video Synopsis

Law Enforcement

Security Robotics

Unmanned Aerial

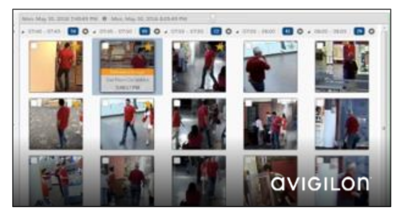

Search by Example

AI CITY SOLUTION

HIKVISION	avigilon TRUSTED SECURITY SOLUTIONS	milestone	dahua TECHNOLOGY	Hanwha
IBM	BriefCam The Video Synopsis Company	unv	SENSETIME	Genetec™
Tiandy	SEEQUESTOR	VIMOC Technologies	aeryon Labs Inc.	Motionloft
KNIGHTSCOPE	miovision rethink traffic	axxon Experience The Next	ROBOTIC ASSISTANCE DEVICES	aqueti
netradyne	Pivot3	umbo computer vision	MotionDSP	fst biometrics

DeepFool: a simple and accurate method to fool deep neural networks

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Pascal Frossard
École Polytechnique Fédérale de Lausanne
{seyed.moosavi,alhussein.fawzi,pascal.frossard} at epfl.ch

Algorithm 1 DeepFool for binary classifiers

input: Image $\boldsymbol{x}$, classifier f.
output: Perturbation $\hat{\boldsymbol{r}}$.
Initialize $\boldsymbol{x}_0 \leftarrow \boldsymbol{x}$, $i \leftarrow 0$.
while $\operatorname{sign}(f(\boldsymbol{x}_i)) = \operatorname{sign}(f(\boldsymbol{x}_0))$ **do**
  $\boldsymbol{r}_i \leftarrow -\frac{f(\boldsymbol{x}_i)}{\|\nabla f(\boldsymbol{x}_i)\|_2^2} \nabla f(\boldsymbol{x}_i)$,
  $\boldsymbol{x}_{i+1} \leftarrow \boldsymbol{x}_i + \boldsymbol{r}_i$,
  $i \leftarrow i + 1$.
end while
return $\hat{\boldsymbol{r}} = \sum_i \boldsymbol{r}_i$.

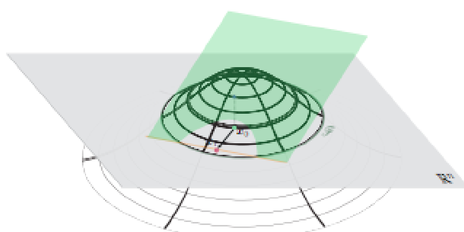

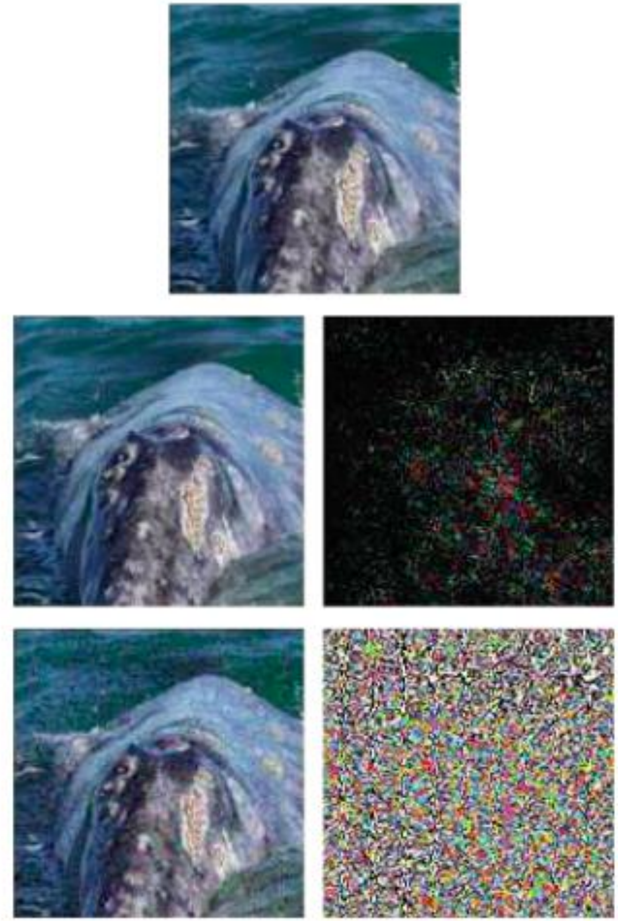

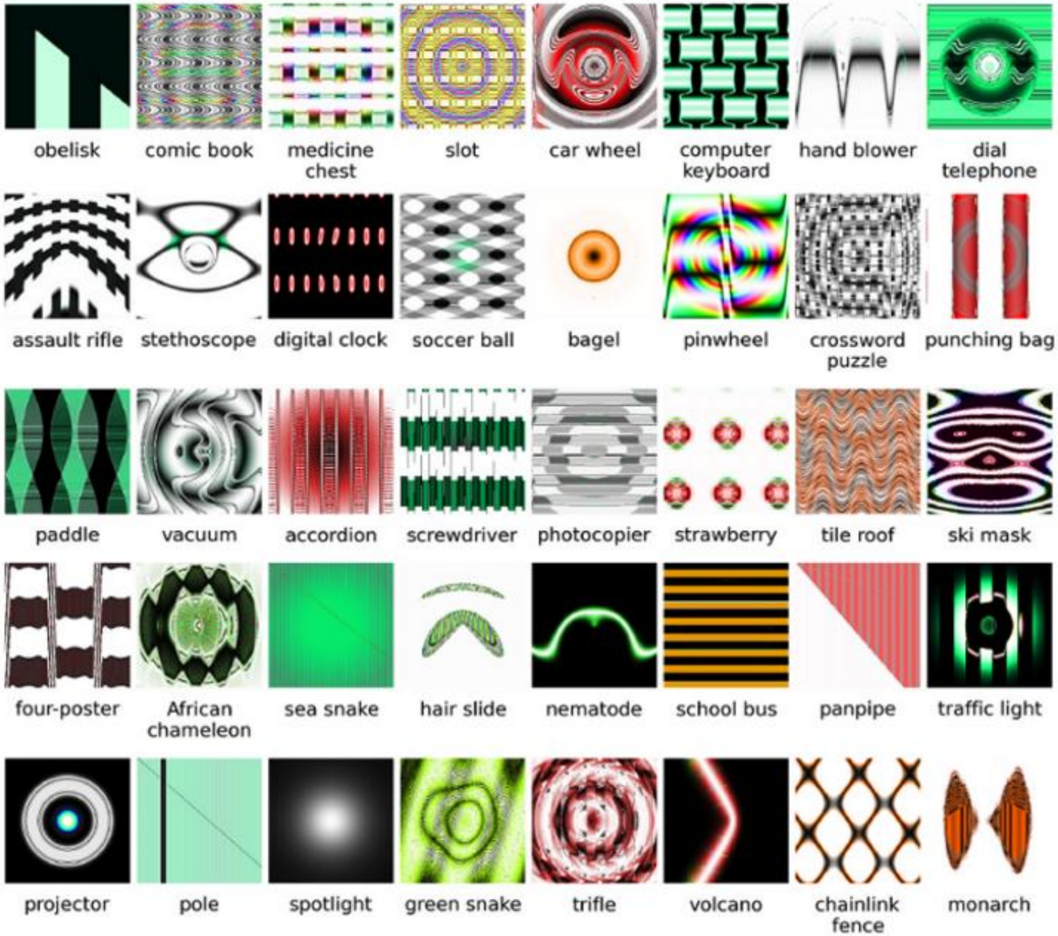
obelisk
comic book
medicine chest
slot
car wheel
computer keyboard
hand blower
dial telephone
assault rifle
stethoscope
digital clock
soccer ball
bagel
pinwheel
crossword puzzle
punching bag
paddle
vacuum
accordion
screwdriver
photocopier
strawberry
tile roof
ski mask
four-poster
African chameleon
sea snake
hair slide
nematode
school bus
panpipe
traffic light
projector
pole
spotlight
green snake
trifle
volcano
chainlink fence
monarch

PERTURB INFORMATION

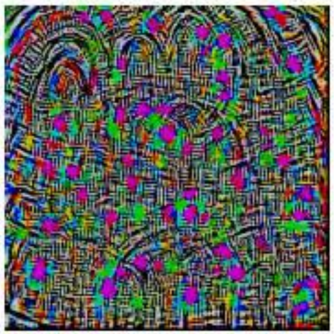

(a) CaffeNet

(b) VGG-F

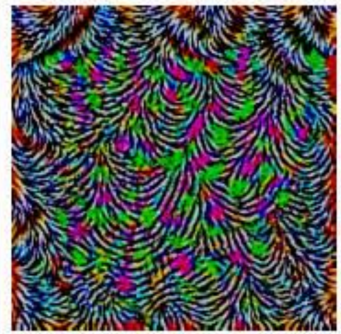

(c) VGG-16

(d) VGG-19

(e) GoogLeNet

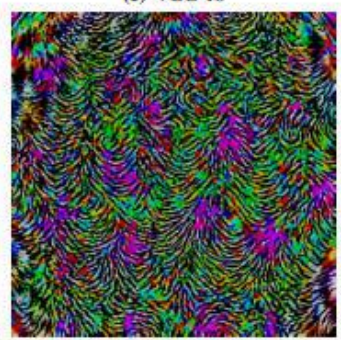

(f) ResNet-152

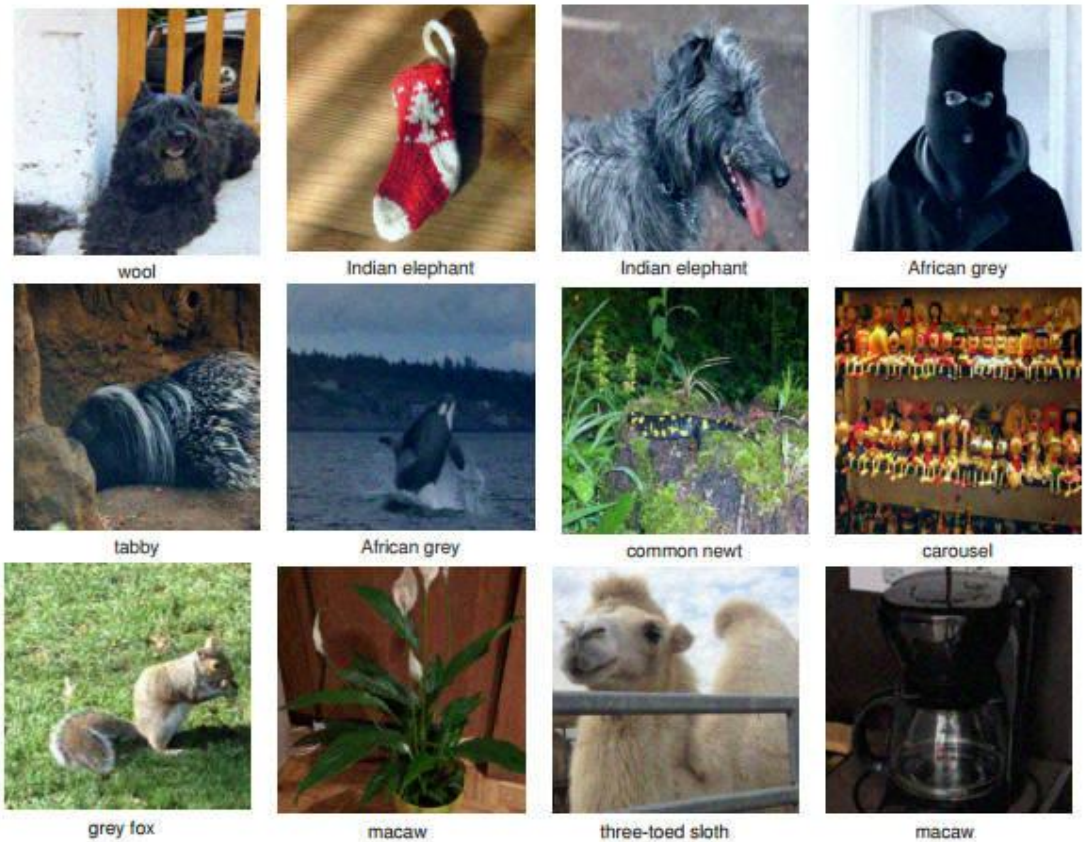

Examples of perturbed images and their corresponding labels. The first two rows of images belong to the ILSVRC 2012 validations set, and the last row are random images taken by a mobile phone camera.

AI TECHNOLOGY FOR FAKE NEWS

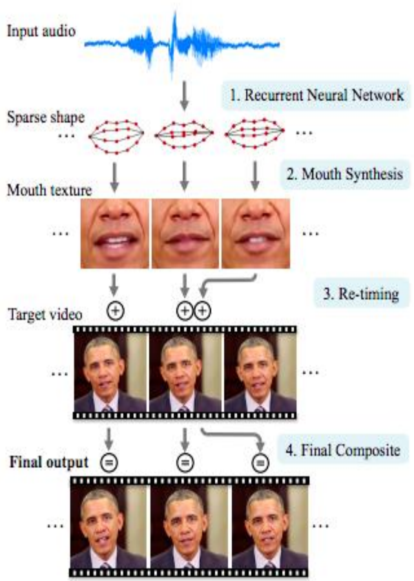

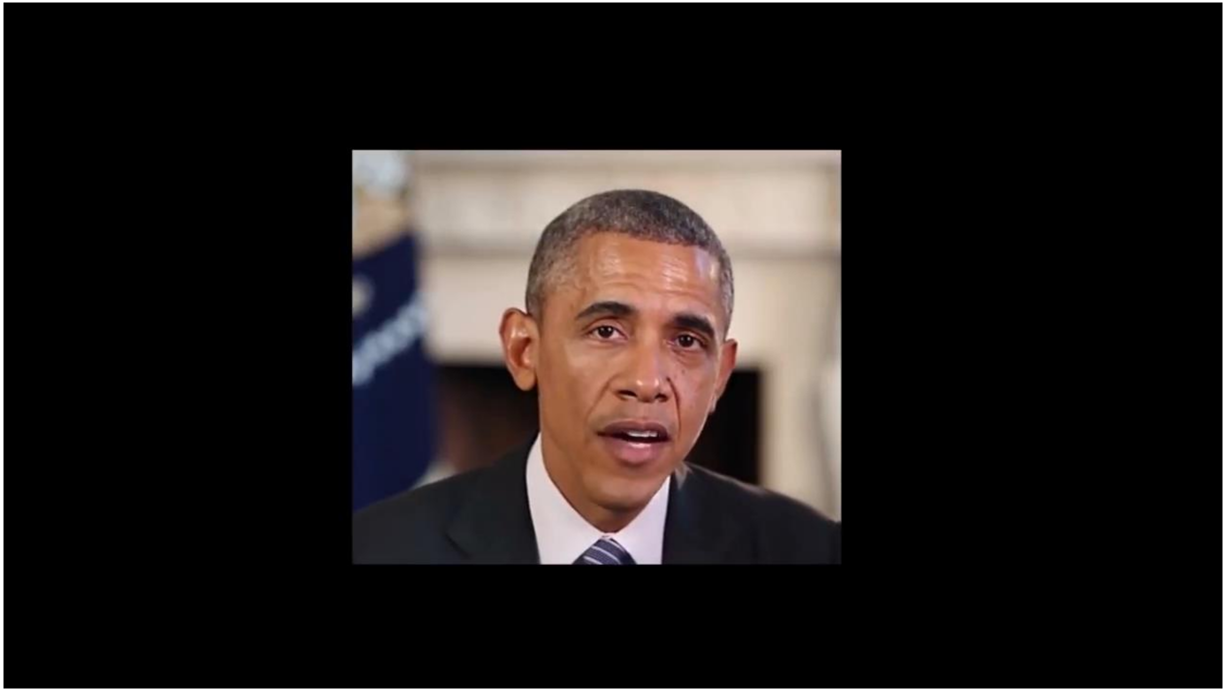

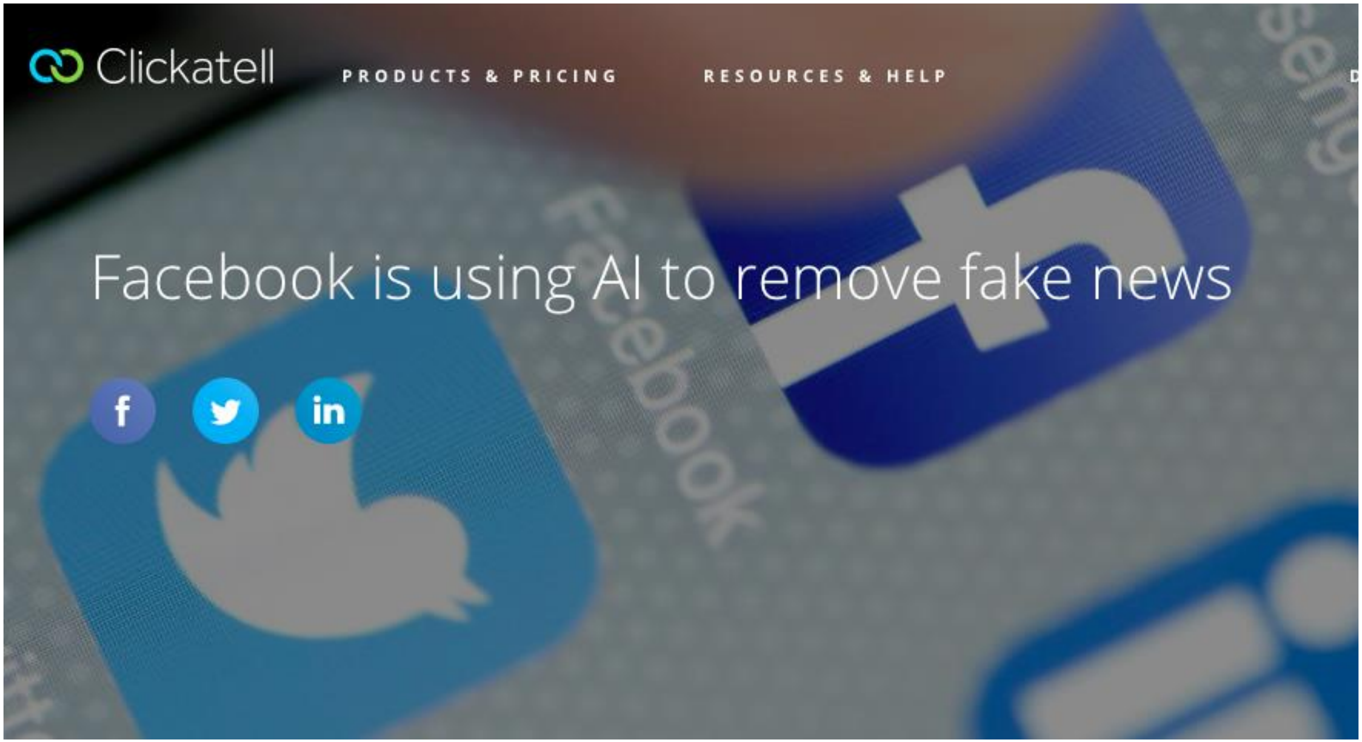
Clickatell
PRODUCTS & PRICING
RESOURCES & HELP
Facebook is using AI to remove fake news


COULD HACKERS USE
ARTIFICIAL INTELLIGENCE
TO STEAL YOUR DATA?

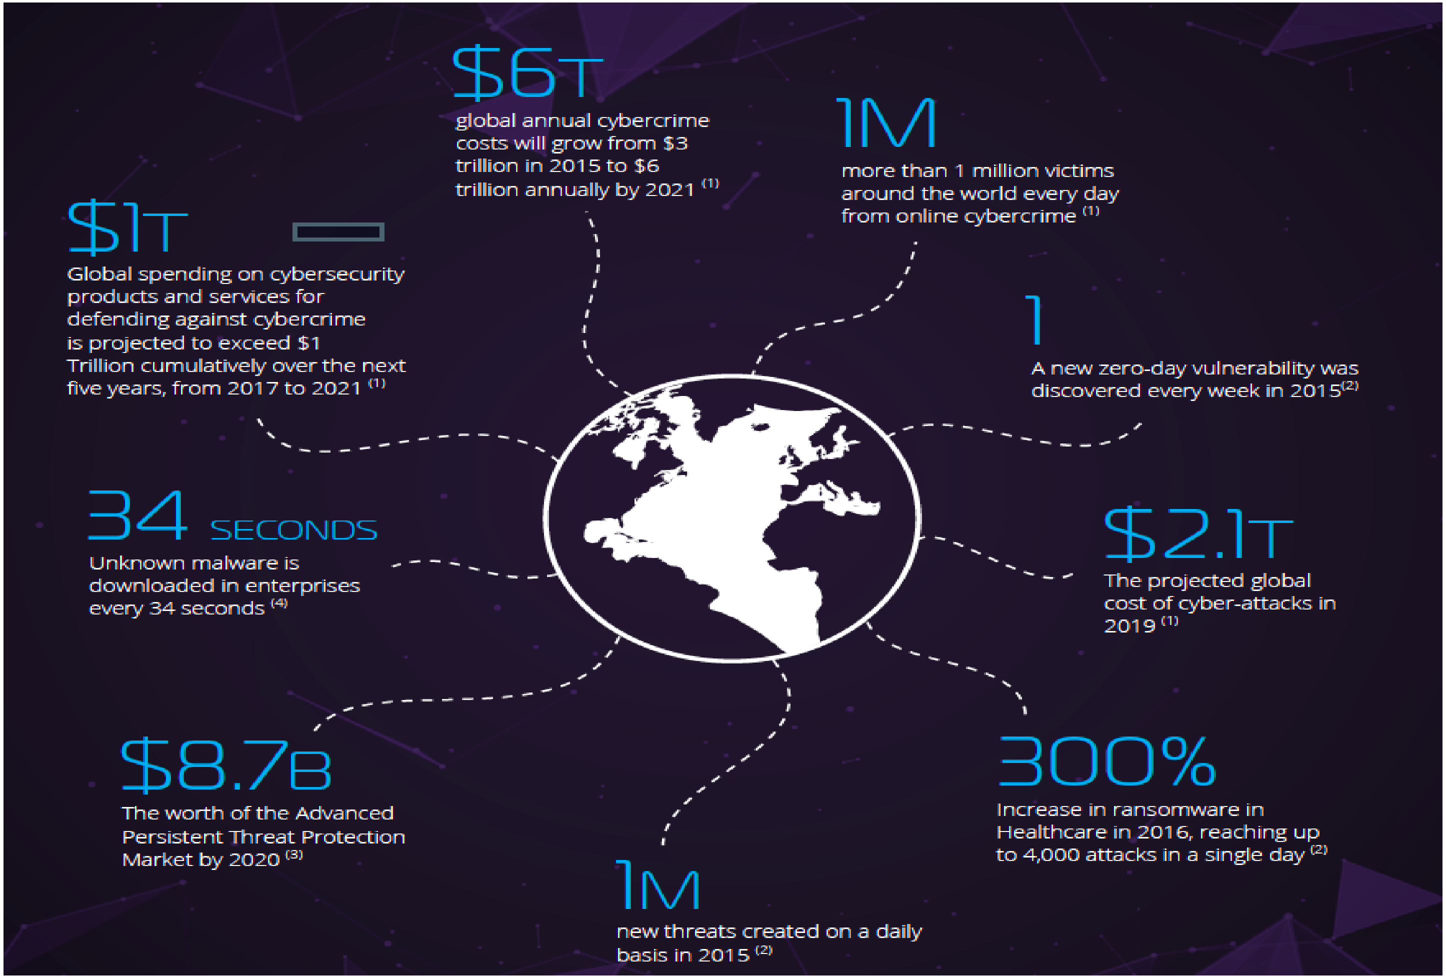
$6T
global annual cybercrime costs will grow from $3 trillion in 2015 to $6 trillion annually by 2021 (1)
1M
more than 1 million victims around the world every day from online cybercrime (1)
$1T
Global spending on cybersecurity products and services for defending against cybercrime is projected to exceed $1 Trillion cumulatively over the next five years, from 2017 to 2021 (1)
1
A new zero-day vulnerability was discovered every week in 2015 (2)
34 SECONDS
Unknown malware is downloaded in enterprises every 34 seconds (4)
$2.1T
The projected global cost of cyber-attacks in 2019 (1)
$8.7B
The worth of the Advanced Persistent Threat Protection Market by 2020 (3)
300%
Increase in ransomware in Healthcare in 2016, reaching up to 4,000 attacks in a single day (2)
1M
new threats created on a daily basis in 2015 (2)

SIGNATURE

The antivirus engine compares the contents of a file (op-codes or strings) to its database of known malware signatures. If the malware has not been seen before, a handcraft signature is generated and then released as an update to clients. This process is time-consuming, resulting in signatures released months after initial detection.

HEURISTICS

This is a general term for the different techniques used to detect malware based on their behavior characteristics typically used in known malware code (e.g., op-code with random parts). Generally used together with signature-based detection.

BEHAVIORAL

Detection based on behavioral fingerprint of the malware at run-time instead of characteristics hardcoded in the malware code itself. Similar to heuristic-based detection and used in Intrusion Detection Systems, this method is able to detect malware only once the malicious actions commence.

SANDBOX

A development of the behavioral-based detection method that executes the programs in a virtual environment instead of detecting the behavioral fingerprint at run-time. Although this technique has shown to be quite effective, it is rarely used in end-user antivirus solutions given its heavy and slow process.

MACHINE LEARNING

Algorithms are used to classify the behavior of a file as malicious or benign according to a series of file features that are manually extracted from the file itself. Each specific structure of the file has to be broken into the smallest part in order to be learned.

DEEP LEARNING

Deep learning is a novel adaptation of neural networks, inspired by the brain's ability to learn. Powerful algorithms are capable of learning from any type of data without receiving any outside assistance, in the same way a brain operates. The application of deep learning in cybersecurity results in substantial improvement in malware detection rates, particularly regarding previously unknown zero-day threats.

DEEP INSTINCT™ IS THE FIRST COMPANY THAT APPLIES DEEP LEARNING TO CYBERSECURITY

Deep Instinct™ has developed a highly efficient deep learning core library running on GPUs and uses it to learn the behavior of billions of malware vectors.

AI POWERED HACKING MACHINE

Institute of Policy Studies

ENGAGING MINDS, EXCHANGING IDEAS

Closed-Door Discussion
on
The Digital Economy and Challenges in the Process of Transformation

Thursday, 20 July 2017

Conference Room, Level 1, Oei Tiong Ham Building

Lee Kuan Yew School of Public Policy
National University of Singapore