


Deep Neural Nets concepts

Dept. Of computer Engineering
Hanbat National University
Yoon Joong Kim

High-level Analogy with Astrophysics and coffee

High-Level Analogy with Astrophysics and Coffee





Likewise for coffee:




[based on Sean Carroll lecture]


Similar for iterative algorithms and DNN?




Ultimate dream: thinking machine




Ultimate dream : thinking machine



Dendrite : 수상돌기
Axon : 축색돌기
synapse : axon terminal + dendrite
Myelin sheath : 축색돌기를 감싸는 막





Schematic of a biological neuron.





Schematic of a biological neuron.

Neural model(Perceptron)




A diagram showing how the Perceptron works. (Source) Frank Rosenblatt, 1969

Logistic regression units




<http://www.andreykurenkov.com/writing/a-brief-history-of-neural-nets-and-deep-learning/>

Hardware implementations



Frank Rosenblatt, ~1957: Perceptron



Widrow and Hoff, ~1960: Adaline/Madaline



False Promises

- The Navy revealed the embryo of an electronic computer today that *it expects will be able to walk, talk, see, write, reproduce itself* and be conscious of its existence ... Dr. Frank Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said perceptrons might be fired to the planets as mechanical space explorers” July 08, 1958


The New York Times

<http://query.nytimes.com/gst/abstract.html?res=9D01E4D8173DE53BBC4053DFB1668383649EDE>



(Simple) AND/OR problem: linearly separable?



(Simple) AND/OR problem: linearly separable?




x1	x2	y
0	0	0
0	1	1
1	0	1
1	1	0



Perceptrons (1969)

by Marvin Minsky, founder of the MIT AI Lab


- In his book
- We need to use MLP, multilayer perceptrons (multilayer neural nets)



- No one on earth had found a viable way to train MLPs(W, b) good enough to learn such simple functions.


Frank Rosenblatt, ~1957: Perceptron

Marvin Minsky, MLP (1969), unable to train w, b of MLP



“No one on earth had found a viable way to train*”

*Marvin Minsky, 1969



<http://cs231n.github.io/convolutional-networks/>



Backpropagation

(1974, 1982 by Paul Werbos, 1986 by Jeffrey Hinton)





Jeffrey Hinton,
University of Toronto

Training



Inference



Frank Rosenblatt, ~1957: Perceptron


Marvin Minsky, MLP (1969), unable to train W, b of MLP

Paul 1974/1982, Hinton 1986, Error Backpropagation

<https://devblogs.nvidia.com/parallelforall/inference-next-step-gpu-accelerated-deep-learning/>

Convolutional Neural Networks


- Visual System by Hubel & Wiesel, 1959,62,58,...





[Yann LeCun](#), Paris
Jeffrey Hinton's lab
1996, AT&T Bell Labs-Research
as head of Image Processing
Research, which was part of
Lawrence Labinar's S&IPRL
2003, NYU
2013, 1st director of Facebook
AI Research
....

Convolutional Neural Networks, [LeNet-5, LeCun 1980]



“At some point in the late 1990s, one of these systems was reading 10 to 20% of all the checks in the US.”

[LeNet-5, LeCun 1980]

[Fei-Fei Li & Andrej Karpathy & Justin Johnson Lecture 7 - 6 27 Jan 2016](#)

Project NavLab 1984-1994, CMU



<https://en.wikipedia.org/wiki/Navlab>

Neural-net in movie

Terminator 2 (1991)



JOHN: Can you learn? So you can be... you know. More human. Not such a dork all the time.

TERMINATOR: My CPU is a **neural-net** processor... a learning computer. But **Skynet** presets the switch to "read-only" when we are sent out alone.

...

We'll learn how to **set** the neural net

TERMINATOR Basically. (starting the engine, backing out) The **Skynet** funding bill is passed. The system goes on-line August 4th, 1997. Human decisions are removed from strategic defense. **Skynet** begins to learn, at a geometric rate. It becomes **self-aware** at 2:14 a.m. eastern time, August 29. In a panic, they try to pull the plug.

SARAH: And **Skynet** fights back.

TERMINATOR: Yes. It launches its ICBMs against their targets in Russia.


SARAH: Why attack Russia?

TERMINATOR: Because **Skynet** knows the Russian counter-strike will remove its enemies here.

A BIG problem

- number of layers-vanishing weights

- Backpropagation just did not work well for normal neural nets with many layers
- Other rising machine learning algorithms: SVM, RandomForest, etc.
- **1995** “Comparison of Learning Algorithms For Handwritten Digit Recognition” by LeCun et al. found that this new approach worked better



<http://neuralnetworksanddeeplearning.com/chap6.html>

CIFAR's contribution

- Canadian Institute for Advanced Research (CIFAR)
- CIFAR encourages basic research without direct application, was what motivated Hinton to move to Canada in 1987, and funded his work afterward.



CIFAR

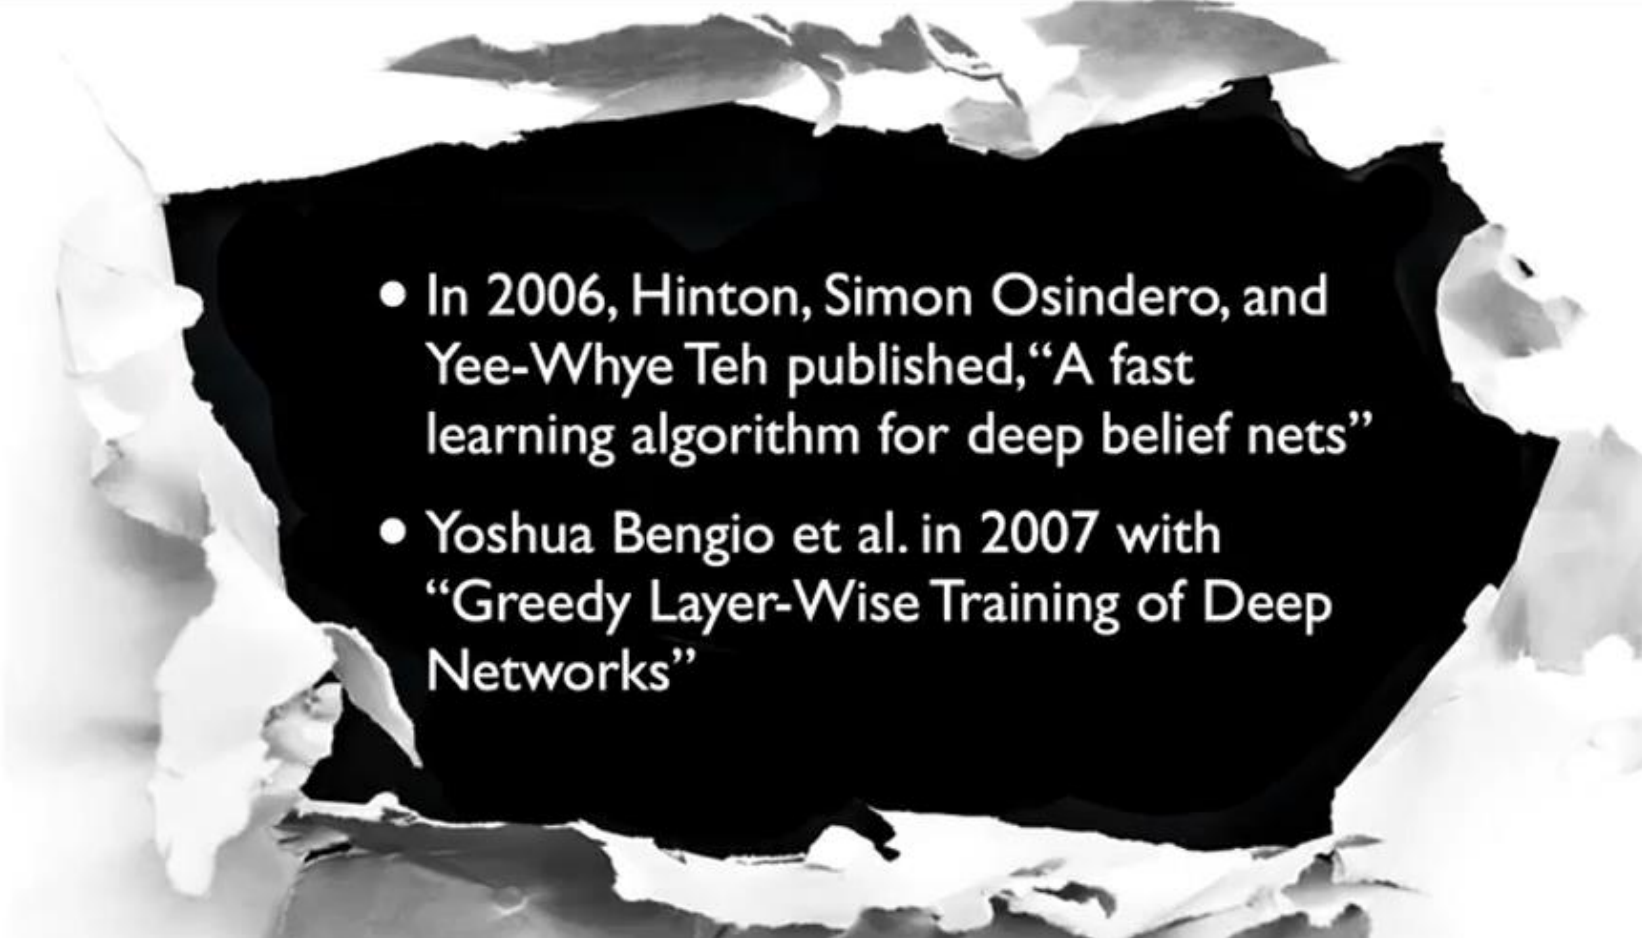
CANADIAN INSTITUTE
for ADVANCED RESEARCH

“Everyone else was doing something different”

- “It was the *worst possible time*”, says Bengio, a professor at universite de Montreal and co-director of the CIFAR program since it was rewarded last year. “Everyone else was doing something different. Somehow, Geoff convinced them.”
- “We should give(CIFAR) a lot of credit for making that gamble.
- CIFAR had huge impact in forming a community around deep learning



Hinton and Bengio's two breakthrough papers


- 
- In 2006, Hinton, Simon Osindero, and Yee-Whye Teh published, “A fast learning algorithm for deep belief nets”
 - Yoshua Bengio et al. in 2007 with “Greedy Layer-Wise Training of Deep Networks”

Breakthrough in 2006 and 2007 by Hinton and Bengio


- In 2006, initializing weights
Neural networks with many layers really could be trained well, if the weights are initialized in a clever way rather than randomly.
- In 2007, DNN learning
Deep learning methods are more efficient for difficult problems than shallow methods.
- Rebranding to *Deep Nets, Deep Learning*

<https://chatbotslife.com/a-brief-history-of-neural-nets-and-deep-learning-part-4-61be90639182>

Large Scale Visual Recognition Challenge in IMAGENET



The Image Classification Challenge:
1,000 object classes
1,431,167 images

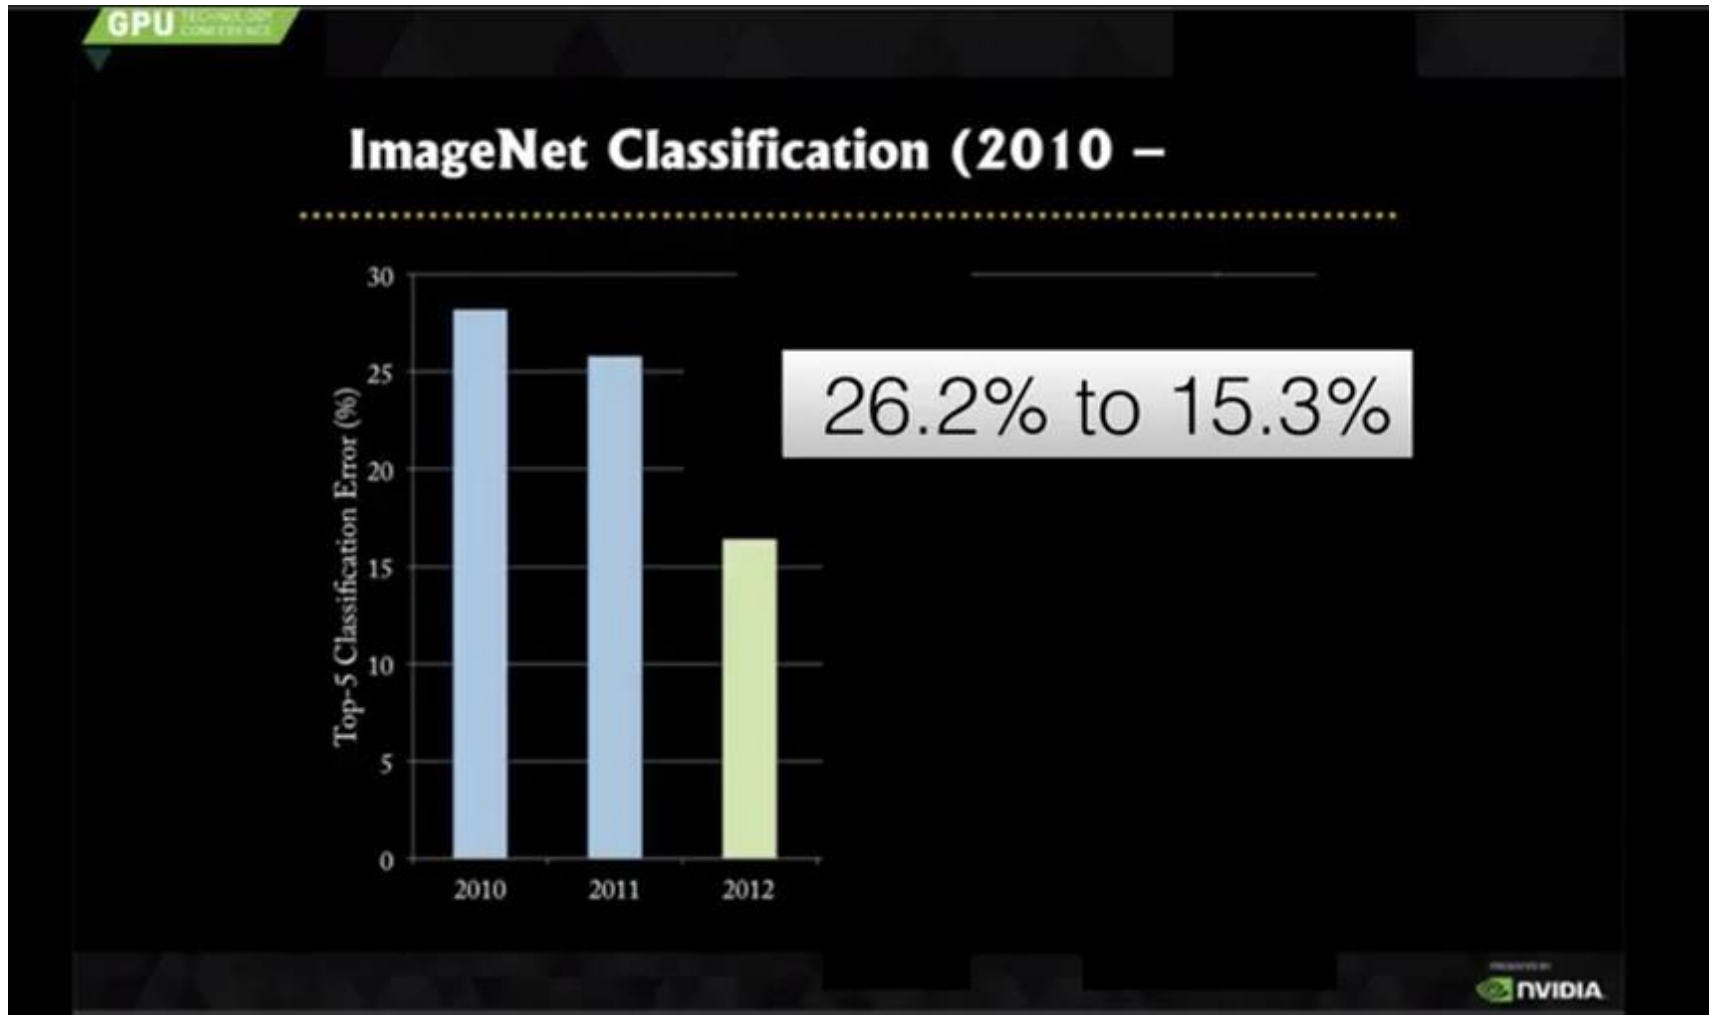


Output: Scale T-shirt <u>Steel drum</u> Drumstick Mud turtle	✓	Output: Scale T-shirt Giant panda Drumstick Mud turtle	✗
--	---	--	---

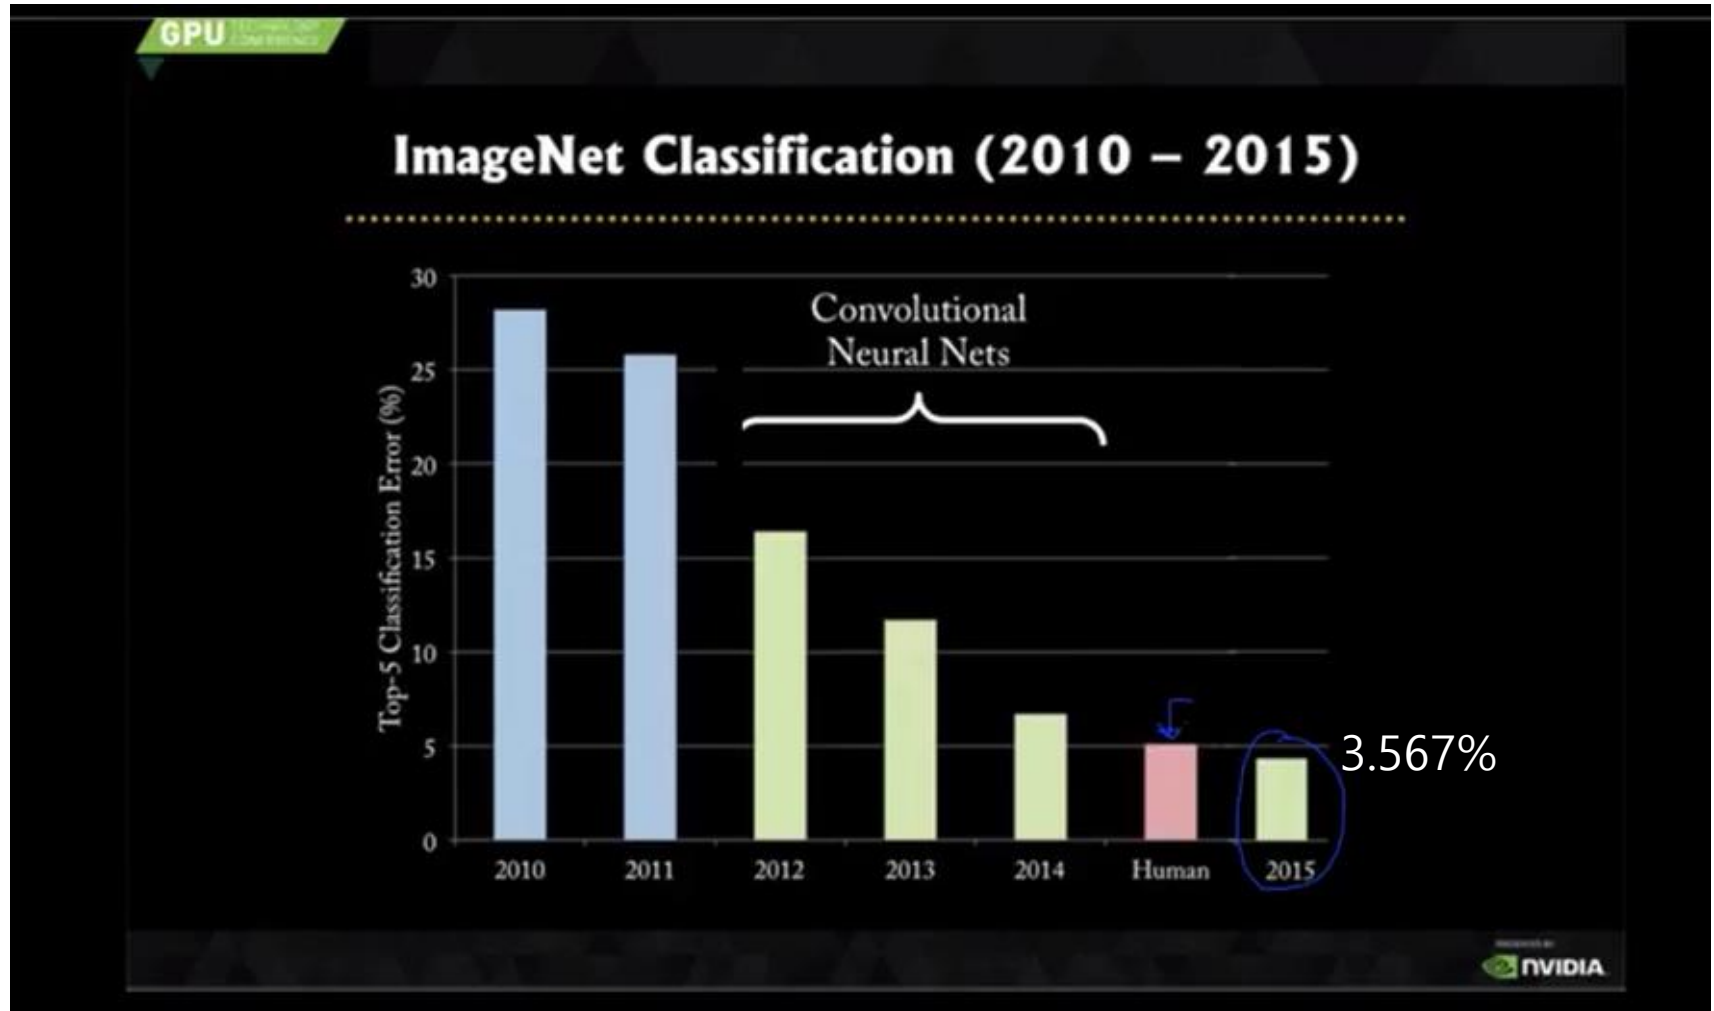
Russakovsky et al. arXiv, 2014

Fei-Fei LI & Andrej Karpathy & Justin Johnson Lecture 1 - 23 4-Jan-16

AlexNet of Doctoral research by Alex in Hinton's lab, 2010



System based on Deep learning, MSRA team 2015




Ensemble 2 by Trimps-Soushen(2016)

- Jie Shao, Xiaoteng Zhang, Zhengyan Ding, Yixin Zhao, Yanjun Chen, Jianying Zhou, Wenfei Wang, Lin Mei, Chuanping Hu

The Third Research Institute of the Ministry of Public Security, P.R. China.

- Object classification/localization (CLS-LOC)
Based on image classification models like Inception, Inception-Resnet, ResNet and Wide Residual Network (WRN), we predict the class labels of the image. Then we refer to the framework of "Faster R-CNN" to predict bounding boxes based on the labels. Results from multiple models are fused in different ways, using the model accuracy as weights.
- **classification error : 2.99%**

Neural networks that can explain photos



A group of people shopping at an outdoor market.

There are many vegetables at the fruit stand.

Deep API Learning

- Explain how to use API for a question



Figure 3: The Overall Workflow of DEEPAPI

66 copy a file and save it to
-your destination path 95

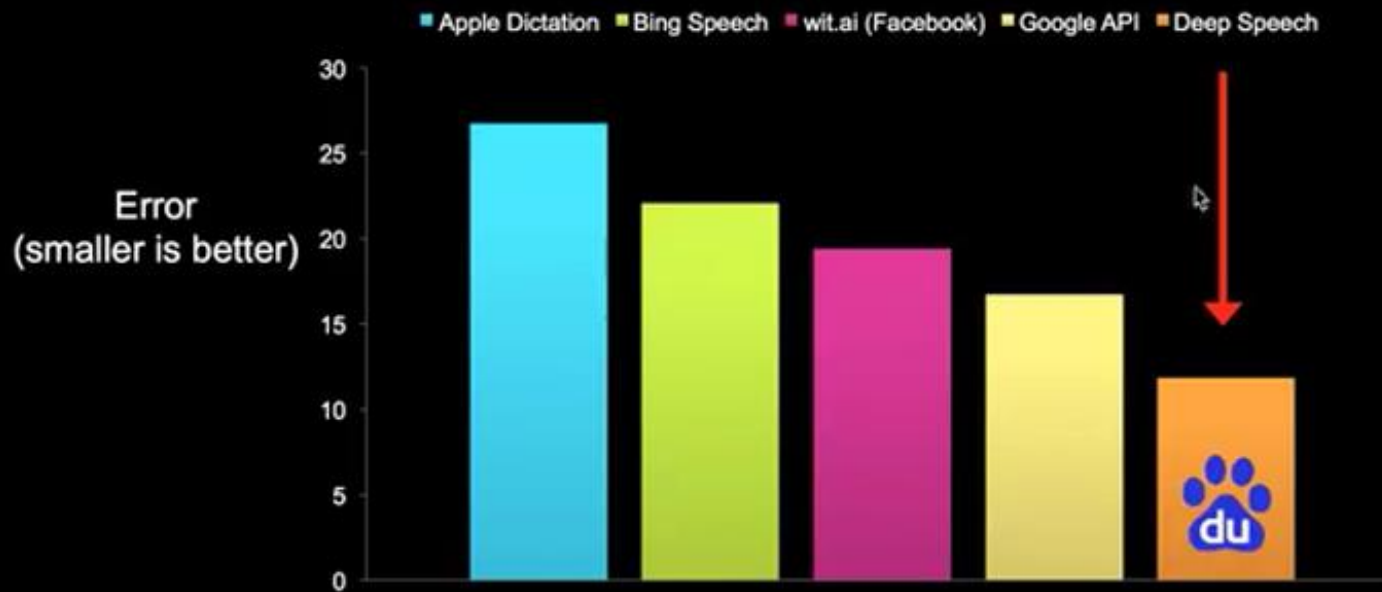


```
FileInputStream.new FileOutputStream.new FileInputStream.getChannel File-  
OutputStream.getChannel FileChannel.size FileChannel.transferTo FileInput-  
Stream.close FileOutputStream.close FileChannel.close FileChannel.close
```


*GU et al. at HKUST with MSR/

Speech Recognition in noise environment

Speech recognition errors



Game



AlphaGo Lee (2016)



AlphaGo Versions

Configuration and strength^[61]

Versions	Hardware	Elo rating	Matches
AlphaGo Fan	176 GPUs, ^[52] distributed	3,144 ^[51]	5:0 against Fan Hui
AlphaGo Lee	48 TPUs, ^[52] distributed	3,739 ^[51]	4:1 against Lee Sedol
AlphaGo Master	4 TPUs, ^[52] single machine	4,858 ^[51]	60:0 against professional players; Future of Go Summit
AlphaGo Zero	4 TPUs, ^[52] single machine	5,185 ^[51]	100:0 against AlphaGo Lee 89:11 against AlphaGo Master
AlphaZero	4 TPUs, single machine	N/A	60:40 against AlphaGo Zero

Automatic Bird-Species Recognition using the Deep Learning and Web Data Mining ,ICTC2018




Fig. 1. Flow chart of the Automatic Bird-Species Recognition.




Fig. 2. Left, header error due to data loss. Right, the white-background images.



Fig. 3. Left, outlier-removed images of birds. Right, outlier images of birds

ICTC 2018

- Bayesian Deep Learning-based Confidence-aware Solar Irradiance Forecasting System
 - HyunYong Lee and Byung Tak Lee (ETRI, Korea)
- Ensemble Classifier based on Decision-Fusion of Multiple Models for Speech Emotion Recognition
 - Kyoung-Ju Noh (ETRI, Korea)
- Distributed Deep Learning Framework based on Shared Memory for Fast Deep Neural Network Training
 - Eun-Ji Lim, Shinyoung Ahn, Wan Choi and Yoo-mi Park (ETRI, Korea)
- Automatic Bird-Species Recognition using the Deep Learning and Web Data Mining
- A development of a speech data transcription tool for building a spoken corpus
 - Hanbat National University
- [Samples in ICTC 2018](#)

Geoffrey Hinton's summary of findings up to today

- Our labeled databases were thousands of times too small
- Our computers were millions of times too small
- We initialized the weights in a stupid way
- We used the wrong-type of non-linearity

Why should I care?

- I am a researcher, not a computer scientist!
- Do you have a idea?
- Do you sell something?
- Are doing any business?

Yube subtitle(자막)

TensorFlow implementation

```
dropout_rate = tf.placeholder("float")  
_L1 = tf.nn.relu(tf.add(tf.matmul(X, W1), B1))  
L1 = tf.nn.dropout(_L1, dropout_rate)
```

TRAIN:
sess.run(optimizer, feed_dict={X: batch_xs, Y: batch_ys,
dropout_rate: 0.7})


EVALUATION:
accuracy.eval({X: mnist.test.images, Y:
dropout_rate: 1})

자막

Facebook

The screenshot displays a Facebook profile for 'Sung Kim'. The top navigation bar includes the search function and the user's name 'Sung' with a 'Home' button. The left sidebar lists various navigation options: 'Edit Profile', 'News Feed', 'Ad Manager', 'Messages' (15), 'Events', 'Photos', '소프트웨어스토리', 'ISSTA' (2), 'I'm a runner', '누가 이법안을 발의 했...' (1), '시와 날씨', 'Top Photos' (1), 'TensorFlow KR', 'Saved' (11), 'PAGES', 'IEEE Transactions...' (20+), 'Like Pages' (14), 'Pages Feed' (20+), 'Create Page', and 'Create Ad'. The main content area shows a status update from 'ISSTA' published by 'Andreas Zeller' 45 minutes ago. The update text reads: 'Doing a PhD in Software Testing and Analysis? Submit to the ISSTA 2016 Doctoral Symposium by April 22! Featuring a keynote by Alex Orso! Details: https://issta2016.cispa.saarland/doctoralsymposium/'. Below the text is a call for doctoral symposium submissions, stating that ISSTA is a leading research symposium in software testing and analysis, bringing together academics, industrial researchers, and practitioners. The event page link is 'ISSTA2016.CISPA.SAARLAND'. The update shows '38 people reached' and a 'Boost Post' button. The right sidebar contains 'YOUR ADS' with an ad for '누가 이법안을 발의 했나?' (1), 'This Week' statistics for 'Post Reach' (3) and 'People Engaged' (1), 'Recent Posts' with a 'Boost Post' button, and '2 event invites' including 'Sungjin Kim's birthday is today'. At the bottom of the right sidebar, there are 'INVITE FRIENDS TO LIKE PAGES' for 'HKUST Water Sports Center' and 'Amazon Web Services'.

Google search engine



Amazon

NEW & INTERESTING FINDS ON AMAZON **EXPLORE**

amazon **Try Prime**

Departments Your Amazon

1-16 of over 1,000,000 results for "galaxy note"

Show results for

- Cell Phones & Accessories
 - Unlocked Cell Phones
 - Carrier Cell Phones
- Computers & Accessories
 - Computer Tablets
 - Styluses
 - Tablet Cases

See All 34 Departments

Refine by

International Shipping (What's this?)

Ship to Korea, Republic of

Amazon Prime

prime

Eligible for Free Shipping

Free Shipping by Amazon

Brand

- Samsung
- BoxWave
- Galaxy Not
- Sansung
- Fosmon
- Generic

galaxy note

- galaxy note 8 glass screen protector
 - in Cell Phones & Accessories
- galaxy note 8 case
- galaxy note 8 phone case
- galaxy note 8
- galaxy note 4 case
- galaxy note 2 case
- galaxy note 3
- galaxy note 8 wireless charger
- galaxy note 3 battery

Unlocked Cellphone, Black Sapphire

★★★★☆ 165

Trade-in eligible for an Amazon gift card

Product Description

... Bigger has always been better. *Galaxy Note* proved that time and again. ...

Samsung Galaxy Note Edge, Charcoal Black 32GB

by Samsung

\$289⁰⁰

FREE Shipping on eligible orders

More Buying Choices

\$183.12

★★★★☆ 181

Trade-in eligible for an Amazon gift card

Product Description

... edge of the display, the *Galaxy Note Edge* redefines smartphone ...

Samsung Galaxy Note8 (US Version) Factory Unlocked Phone - 6.3" Screen - 64GB - Orchid Gray (U.S. Warranty)

by Samsung

The Halloween Shop

EN Hello, Sign in Account & Lists Orders Try Prime Cart

Sort by Relevance

Windows [선택]으로 이

Display



Why now ?

- Students/Researchers

- Not too late to be a world experts
- Not too complicated (mathematically)

- Practitioner

- Accurate enough to be used in practice
- Many ready-to-use tools such TensorFlow
- Many easy/simple programming language such as Python

- After all, it is fun !!

References

- Courses, Artificial Intelligence Laboratory, Stanford
 - <http://ai.stanford.edu/courses/>
- 김성훈 교수, 모두를 위한 머신러닝/딥러닝 강의
 - <https://hunkim.github.io/ml/>
- Deep Learning
 - https://ko.m.wikipedia.org/wiki/딥_러닝
- 딥러닝 스터디 자료 모음
 - <https://bbongcol.github.io/deep-learning-bookmarks/>