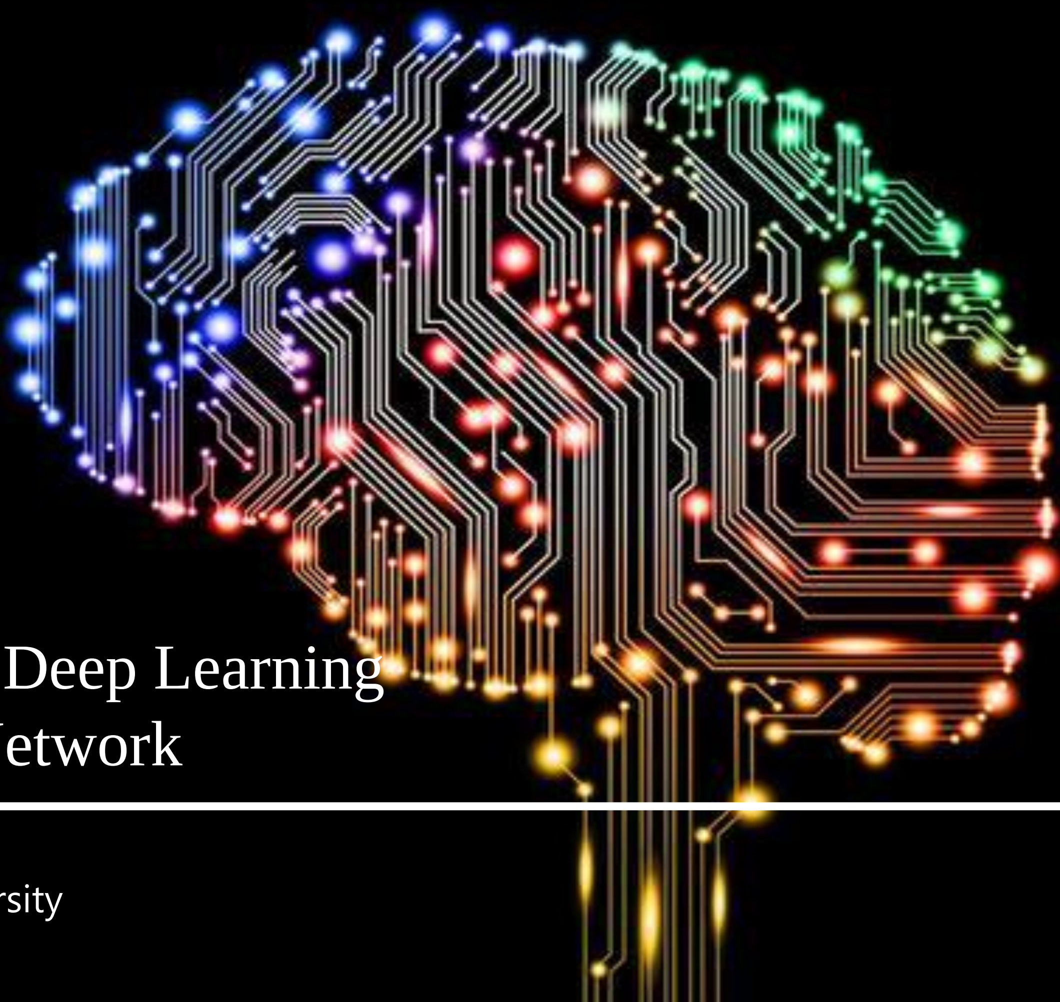

Introduction of Deep Learning
Deep Neural Network

Yoon Joong Kim,
Hanbat National University

Introduction to Deep Learning

Yoon-Joong Kim
Department of Computer Engineering
Hanbat National University

yjkim@hanbat.ac.kr

contents

1. Perceptron

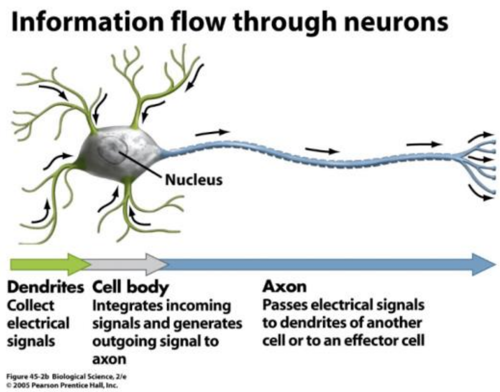

simple functions

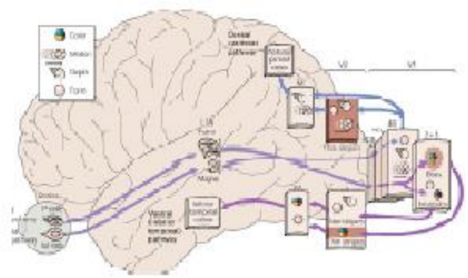

(Van Essen&Gallant, 1994)

Multi-layered

- 연결주의(connectionism)는 길고 멋진 역사를 가지고 있습니다.
- 뉴런은 간단합니다. 그러나 다층 네트워크에서의 배열은 매우 강력합니다.
- 그들은 스스로 조직합니다. 효과적으로 학습하는 것은 조직의 변화 (또는 연결 강도)입니다.

- 인간은 패턴을 잘 인식합니다. 뇌는 어떻게 이러한 일을 하는 걸까요?

1. Perceptron(cont.)

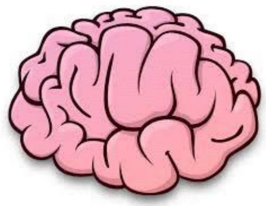

Dendrite : 수상돌기
Axon : 축색돌기
synapse : axon terminal + dendrite
Myelin sheath : 축색돌기를 감싸는 막

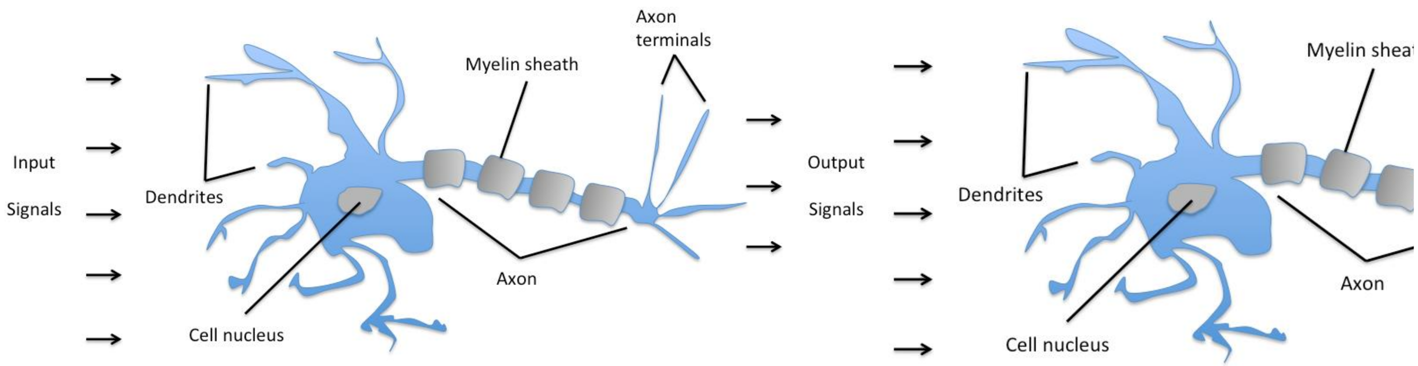

Schematic of a biological neuron.

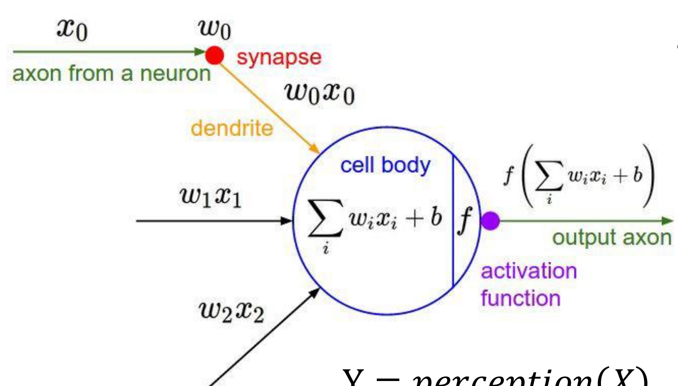

$$Y = perception(X)$$
$$= f(\sum w_i x_i + b))$$

1.Perceptron(cont.)

- The perceptron algorithm
  - was invented in 1958 at the Cornell Aeronautical Laboratory by Frank Rosenblatt, funded by the United States Office of Naval Research. (퍼셉트론 알고리즘은 1958 년 미국 해군 연구소에 의해 자금을 조달 한 Frank Rosenblatt의 "Cornell Aeronautical Laboratory"에서 발명되었습니다.)

- 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

1. Perceptron(cont.)

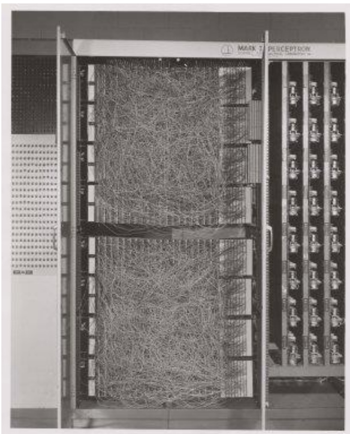

Frank Rosenblatt, ~1957: Perceptron

- Hardware implementations of perceptron
  - The Mark I perceptron machine was the first implementation of the perceptron algorithm.
  - The machine was connected to a camera that used 2020 cadmium sulfide photocells to produce a 400-pixel image. The main visible feature is a patch panel that allowed experimentation with different combinations of input features.
  - To the right of that are arrays of potentiometers that implemented the adaptive weights

1. Perceptron(cont.)

- Adaptive Neuron: Perceptron, linearly separable?

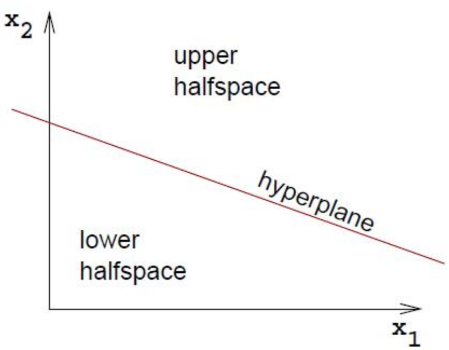

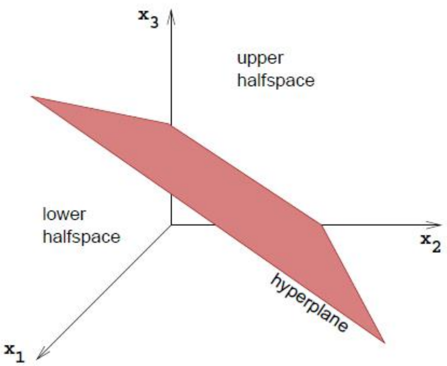

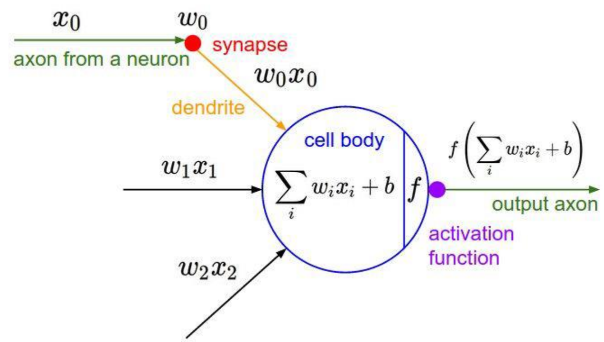

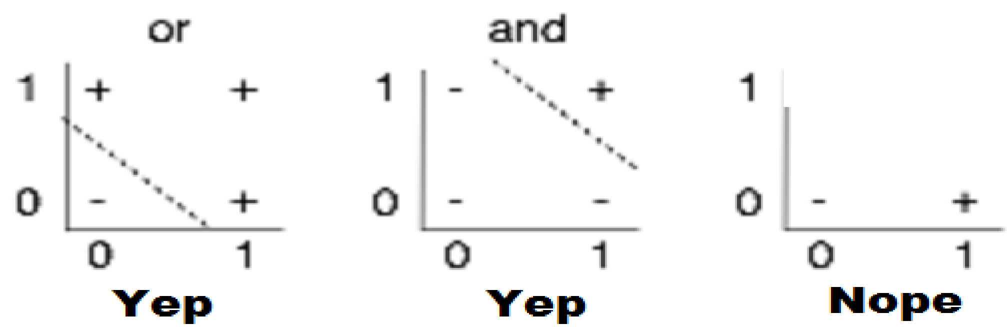

x1	x2	y
0	0	0
0	1	1
1	0	1
1	1	1

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

1. Perceptron(cont.)

- Adaptive Neuron: Perceptron, linearly separable?

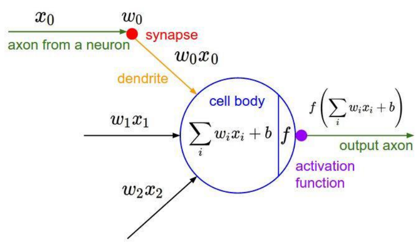

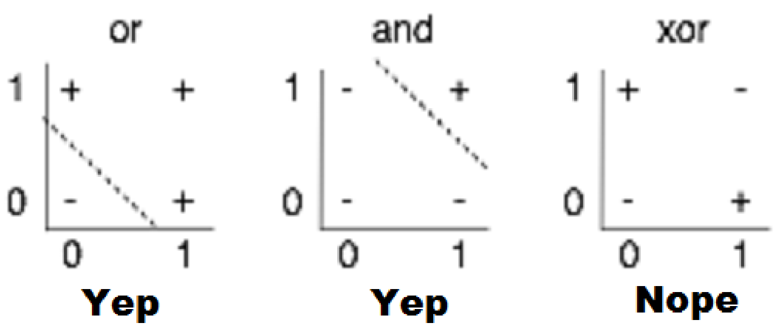

x1	x2	y
0	0	0
0	1	1
1	0	1
1	1	1

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

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

2. Multi-layer Perceptron

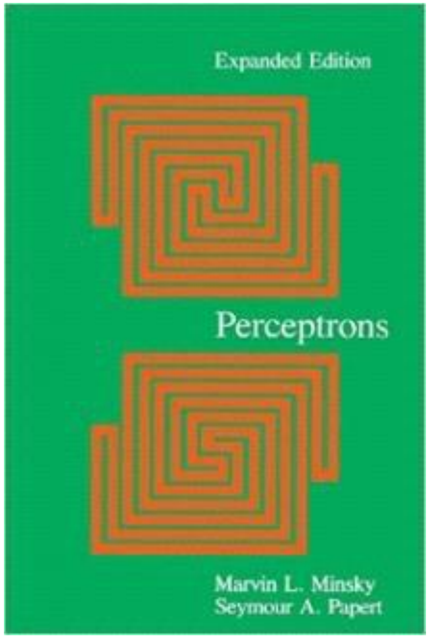

- Multi-layer Perceptron (1969)
  - by Marvin Minsky, founder of the MIT AI Lab
- We need to use MLP, multilayer perceptron (multilayer neural nets)
- Famous examples: XOR, Group Invariance Theorems (Minsky, Papert, 1969)
  - MLP로 XOR 문제 증명
- No one on earth had found a viable way to train MLPs(W, b) good enough to learn such simple functions.

Frank Rosenblatt, ~1957: Perceptron
Marvin Minsky, MLP (1969), unable to train w,b of MLP

2. Multi-layer Perceptron(cont.)

*Marvin Minsky, 1969

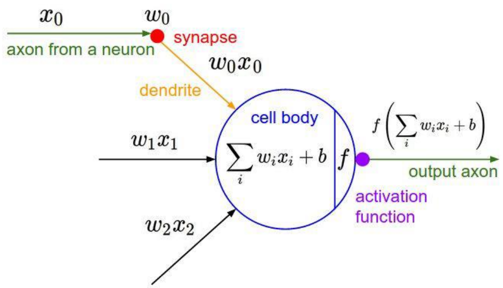

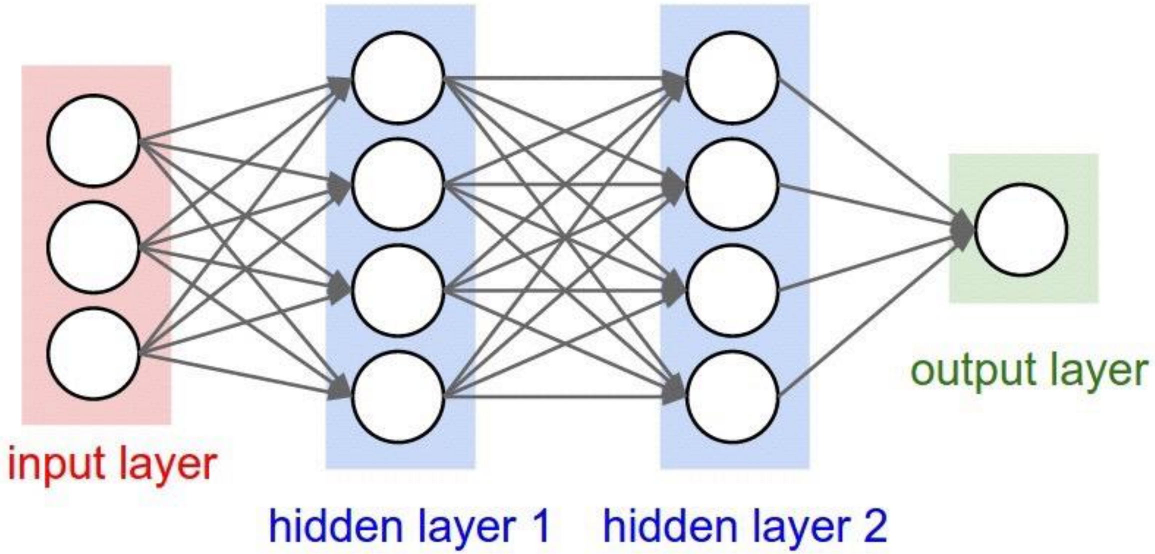

http://cs231n.github.io/convolutional-networks/

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

3. Backpropagation algorithm

Jeoffrey Hinton,
University of Toronto

- Backpropagation algorithm of feature learning
  (1974, 1982 by Paul Werbos, 1986 by Jeoffrey Hinton)

W_i, b_i

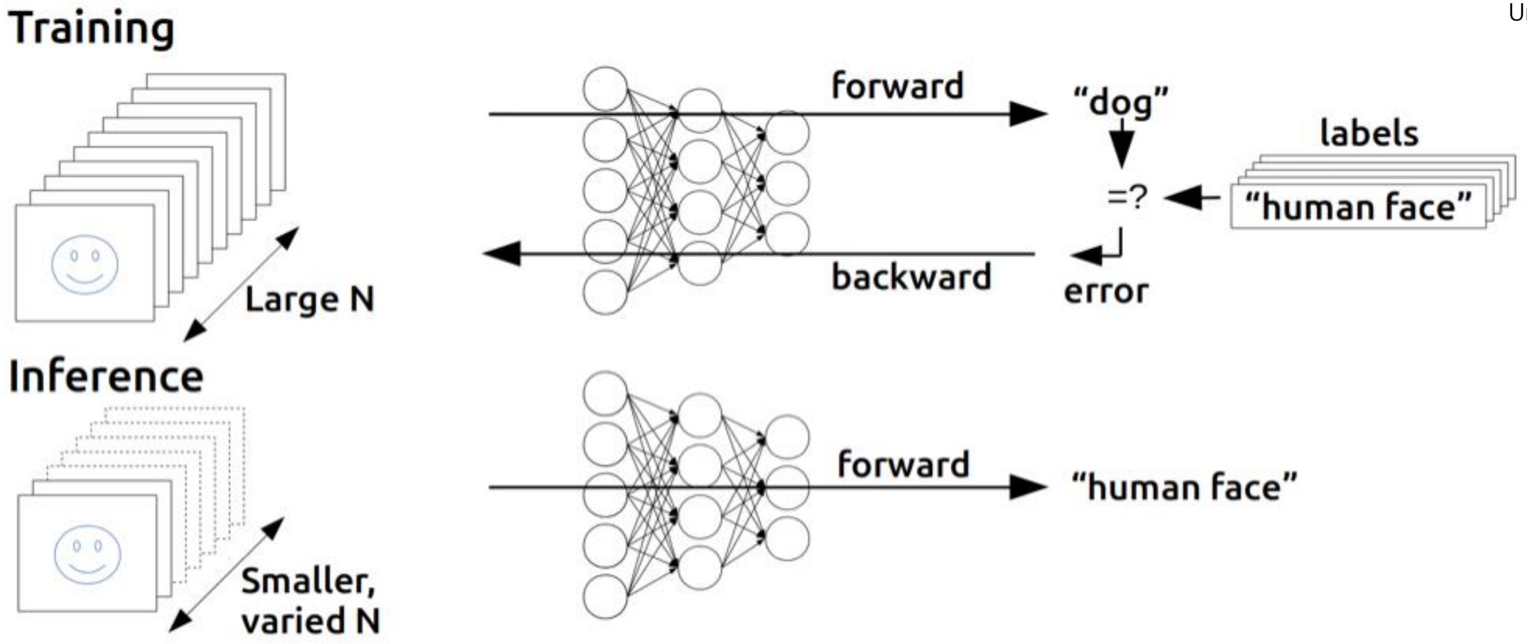

Frank Rosenblatt, ~1957: Perceptron
Marvin Minsky, MLP (1969), unable to train W,b of MLP
Paul 1974/1982, Hinton 1986, Error Backpropagation

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

3. Backpropagation algorithm(cont.)

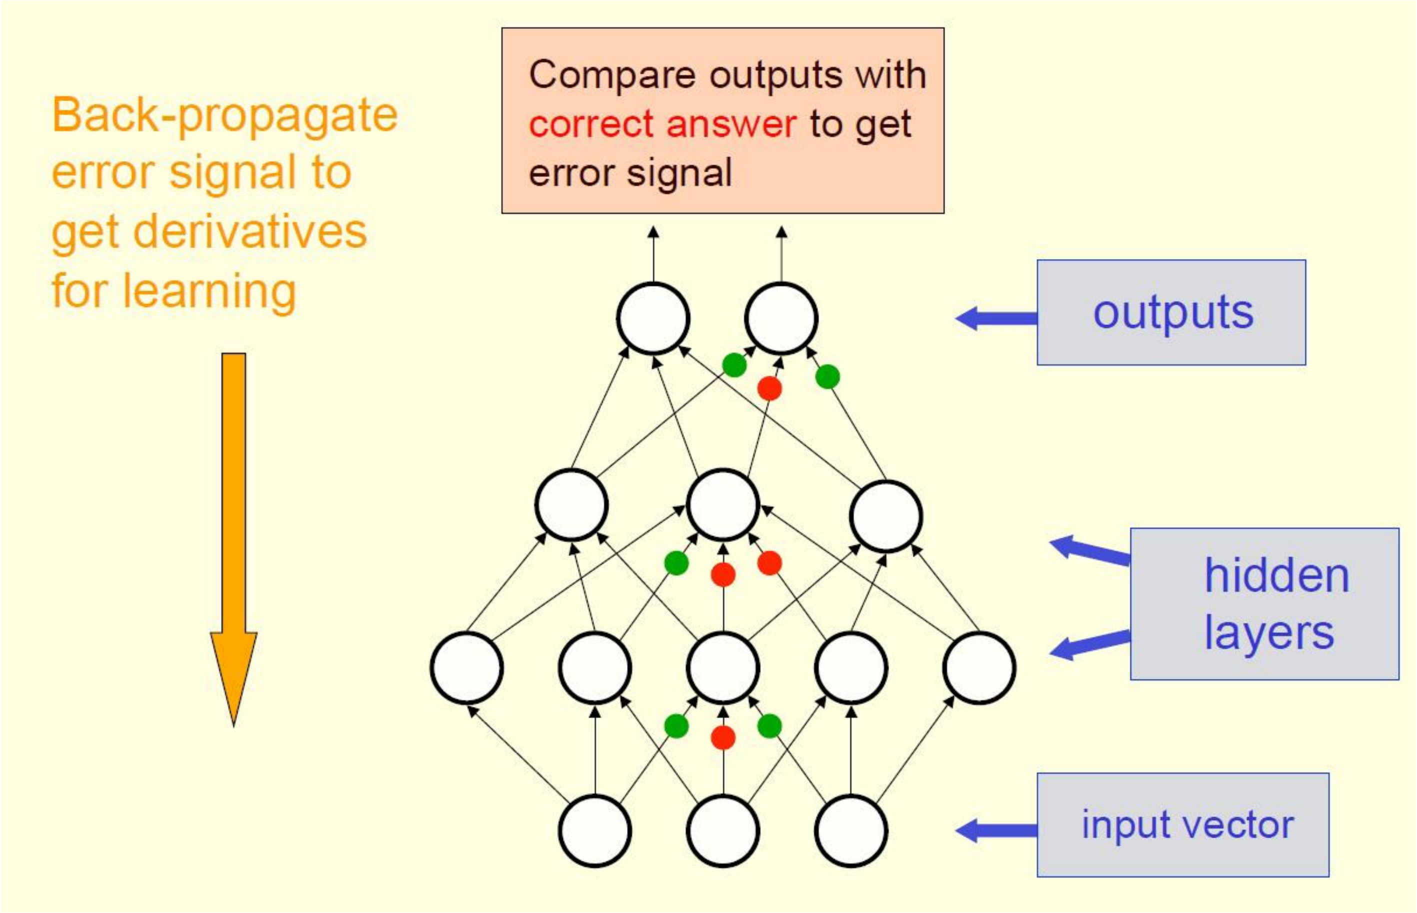

[Slide: G. E. Hinton]

3. Backpropagation algorithm(cont.)

- Summary of Deep learning
  - 1958, Perception
    - The perceptron algorithm was invented in 1958
  - 1969, Multi Layer Perceptron
    - Marvin Minsky , MIT
  - **1986,** Backpropagation algorithm of feature learning
    - **Hinton, Toronto University**
    - The more number of hidden layers, the better...…
  - 1990-2006 ,
    - **MLP-Vanishing weights** :In practice deeper neural networks would need a lot of labeled data and could be **not trained easily**
    - Neural Networks and Backpropagation (with the exception of use in Convolutional Networks) went out of fashion between 1990-2006
  - **In 2006**
    - Hinton and colleagues found a way to pre-train feedforward networks using a **Deep Belief Network** trained greedily
    - This allowed larger networks to be trained by simply using backpropagation for tuning the pre-trained network (easier!)
  - **Since 2010**
    - pre-training of large feedforward networks in this sense also out
    - Availability of large datasets and **fast GPU implementations** have made backpropagation from scratch almost unbeatable

3. Backpropagation algorithm(cont.)

- 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

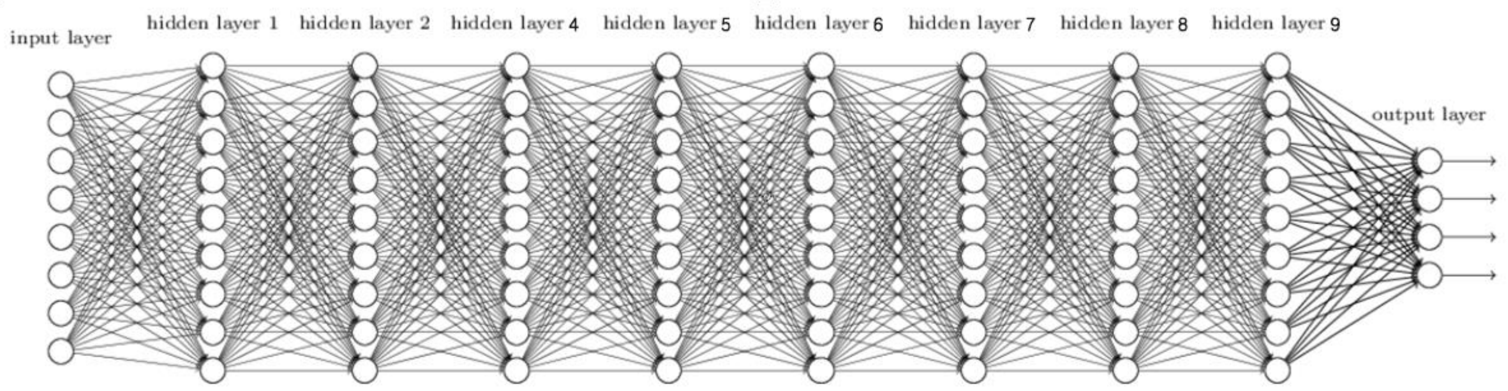

http://neuralnetworksanddeeplearning.com/chap6.html

4. Breakthrough by Hinton and Bengio

- 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.

https://hunkim.github.io/ml/

4. Breakthrough (cont.)

- "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

https://hunkim.github.io/ml/

4. Breakthrough (cont.)

- In 2006, initializing weights
  - Neural networks with many layers really could trained well, if *the weights are initialized in a clever way* rather than randomly.
  - "A fast learning algorithm for deep belief nets"

- In 2007,
  - Deep learning methods are more efficient for difficult problems than shallow methods.
  - "Greedy Layer-wise Learning of Deep Networks"

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

- Rebranding to ***Deep Nets, Deep Learning***

4. Breakthrough (cont.)

- Geoffrey Hitton’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

https://hunkim.github.io/ml/

5. Why use Deep Multi Layered Models?

- Argument 1:
  - 시각적 장면은 계층 적으로 구성됩니다 (언어도 마찬가지입니다). (Visual scenes are hierarchically organized (so is language!))
- Argument 2:
  - 생물학적 비전은 계층 적으로 구성되어 있으며, 우리는 거기에서 몇 가지 아이디어를 수집하고 싶습니다.(Biological vision is hierarchically organized, and we want to glean some ideas from there)
- Argument 3:
  - 얕은 표현은 매우 다양한 기능을 표현하는 데 비효율적입니다 (Shallow representations are inefficient at representing highly varying functions)

5. Why use Deep Multi Layered Models?(cont.)

- Argument 1: 시각적 장면은 계층 적으로 구성됩니다 (언어도 마찬가지입니다). (Visual scenes are hierarchically organized (so is language!))

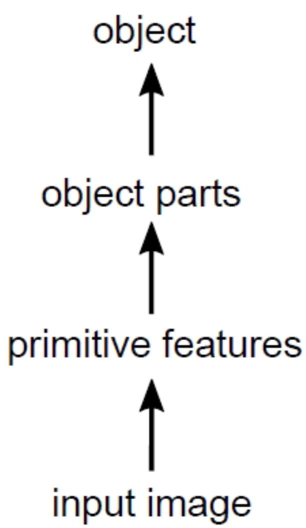

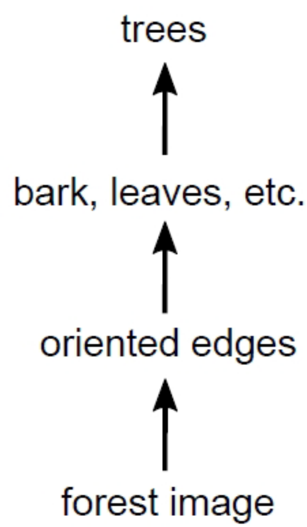

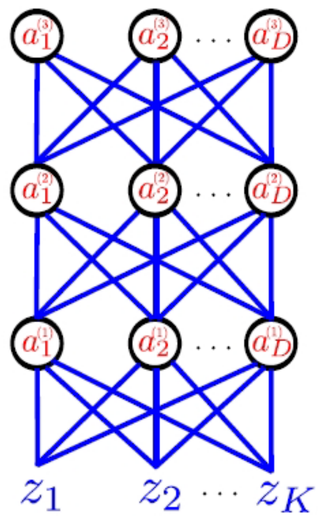

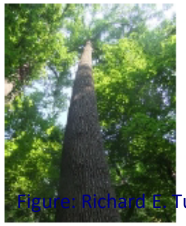

Figure: Richard E. Turner

5. Why use Deep Multi Layered Models?(cont.)

- Argument 2: 생물학적 비전은 계층 적으로 구성되어 있으며, 우리는 거기에서 몇 가지 아이디어를 수집하고 싶습니다.(Biological vision is hierarchically organized, and we want to glean some ideas from there)

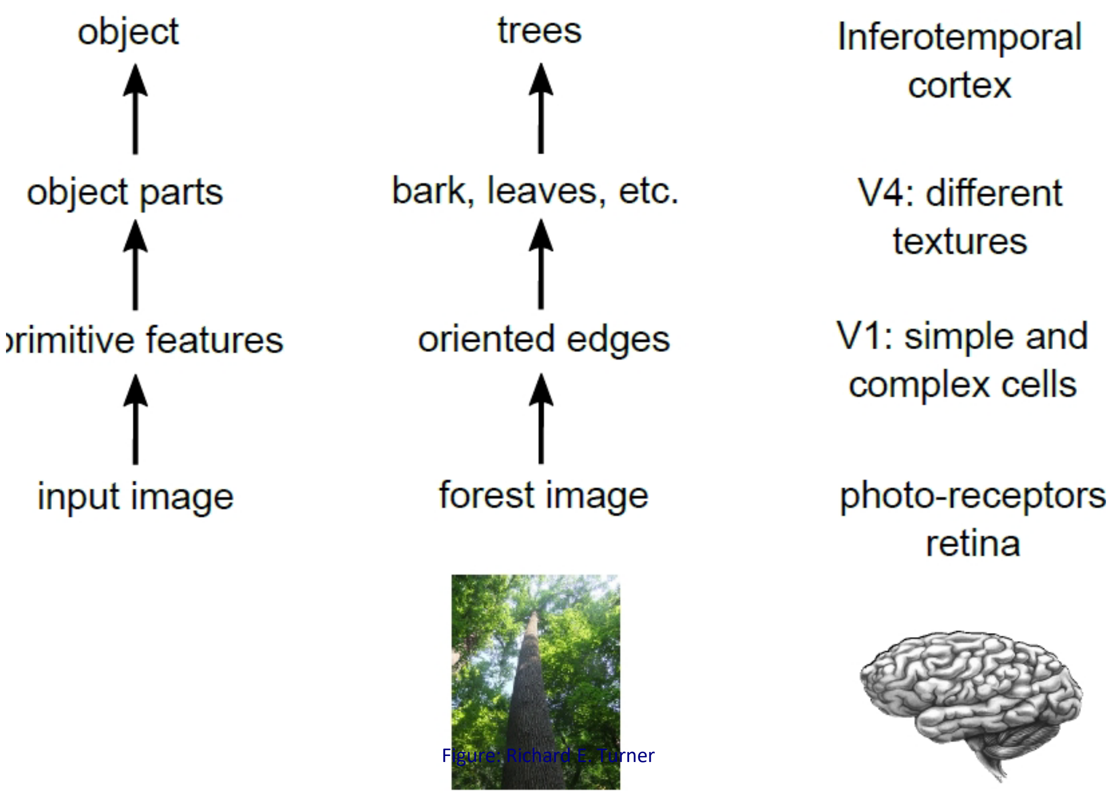

Figure: Richard E. Turner

5. Why use Deep Multi Layered Models?(cont.)

- In the **perceptual system**,
  - **neurons represent features of the sensory input**
  - **The brain learns to extract many layers of features**. Features in one layer represent more complex combinations of features in the layer below. (e.g. Hubel Weisel (Vision), 1959, 1962)

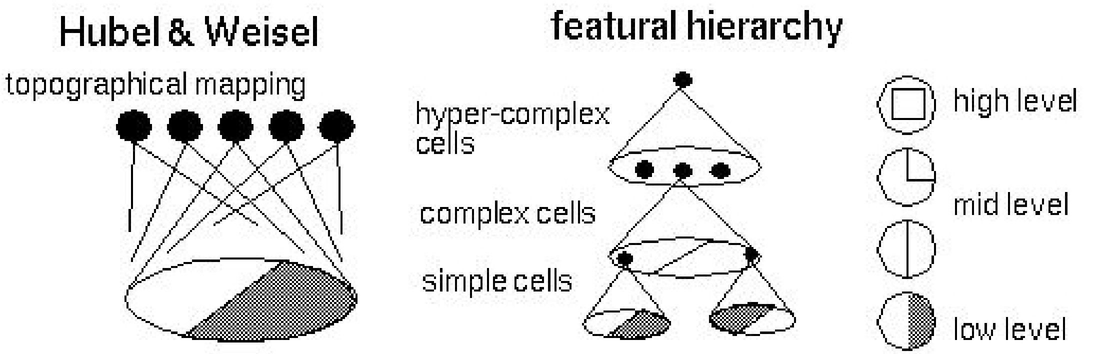

- How can we imitate such a process on a computer?

5. Why use Deep Multi Layered Models?(cont.)

- Argument 3: 얕은 표현은 매우 다양한 기능을 표현하는 데 비효율적입니다 (Shallow representations are inefficient at representing highly varying functions)
  - 기능이 간결한 심층 아키텍처로 표현 될 수 있는 경우,
    이러한(기능이 충분하지 않은) 심층 아키텍처로 표현하려면 매우 큰 아키텍처가 필요할 수 있습니다.

6. Convolutional Neural Networks

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

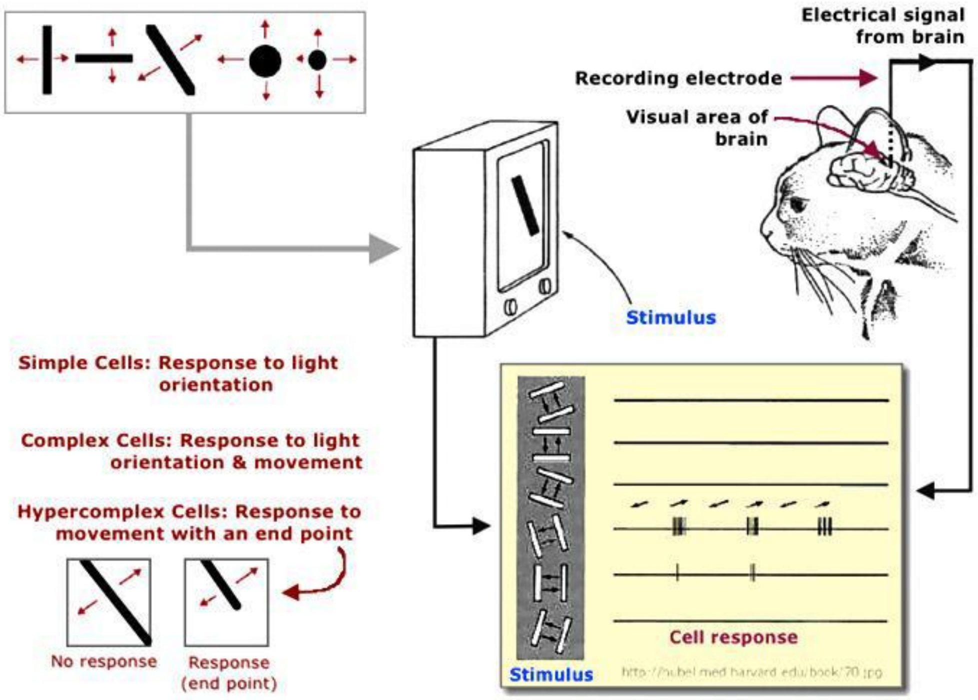

Fei-Fei Li & Andrej Karpathy & Justin Johnson Lecture 7 - 6 27 Jan 2016

6. Convolutional Neural Networks(cont.)

[Yann LeCun](), Paris
Jeoffrey 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
....

- 10 digit recognizer

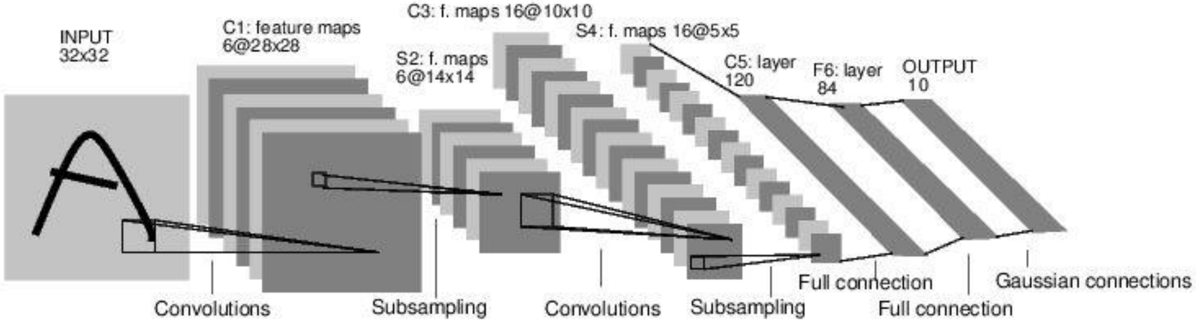

"At some point in the late 1990s, one of these systems was reading 10 to 20% of all the checks in the US."

[LeNet-5, LeCun 1980]

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

6. Convolutional Neural Networks(cont.)

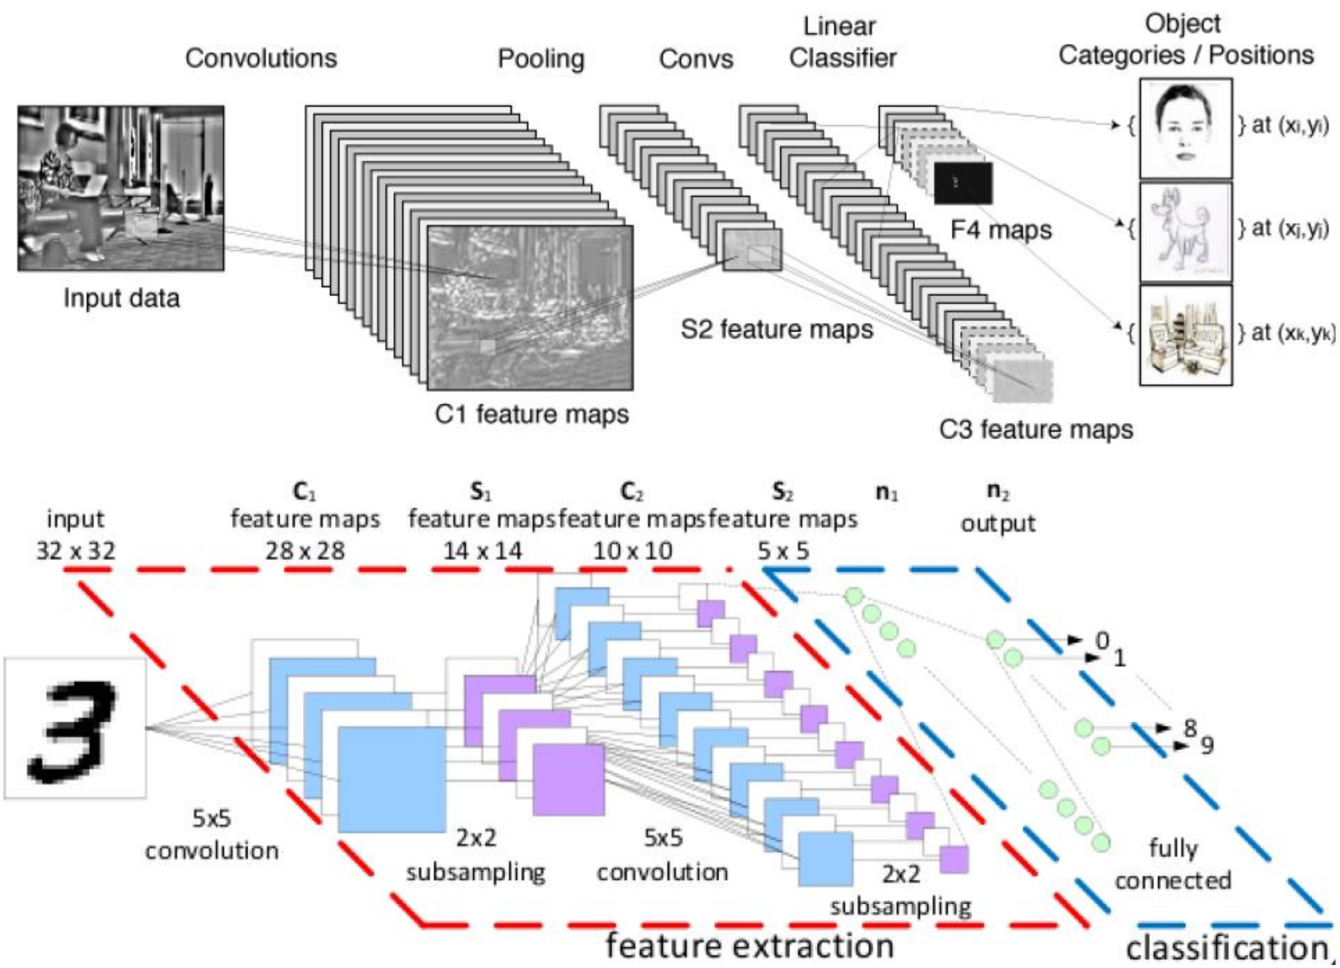

Figure: Yann LeCun

7. Deep Learning Applications

- Project NavLab 1984-1994, CMU

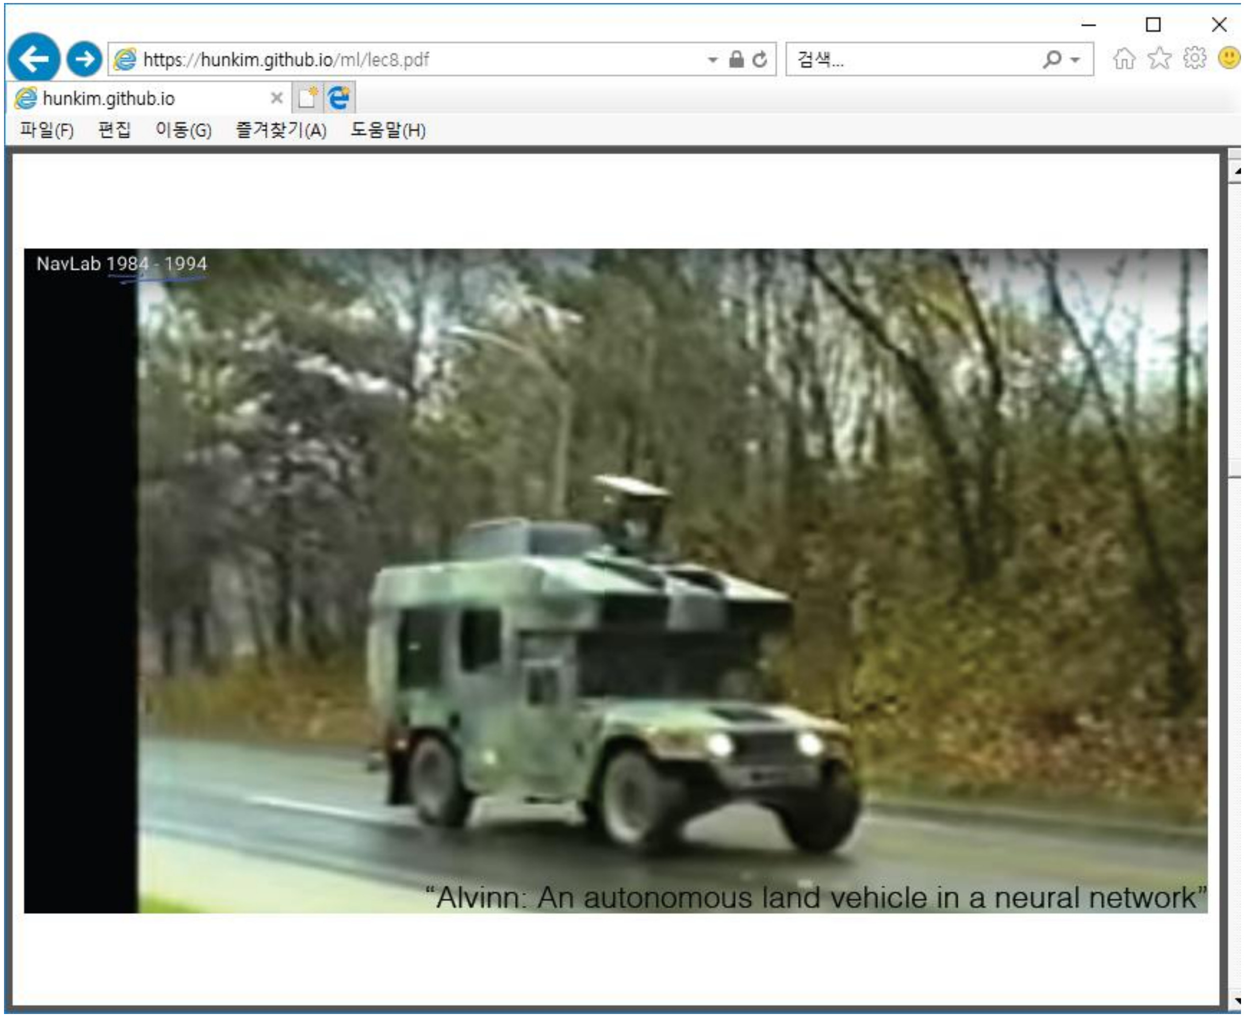

7. Deep Learning Applications(cont.)

- Neural-net in movie

Terminator 2 (1991)

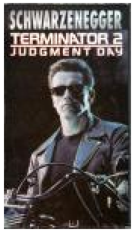

JOHN: Can you learn? So you can be... you know. More human. Not such a dork all the time.

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.

http://pages.cs.wisc.edu/~jerryzhu/cs540/handouts/neural.pdf

7. Deep Learning Applications(cont.)

- Large Scale Visual Recognition Challenge in IMAGENET

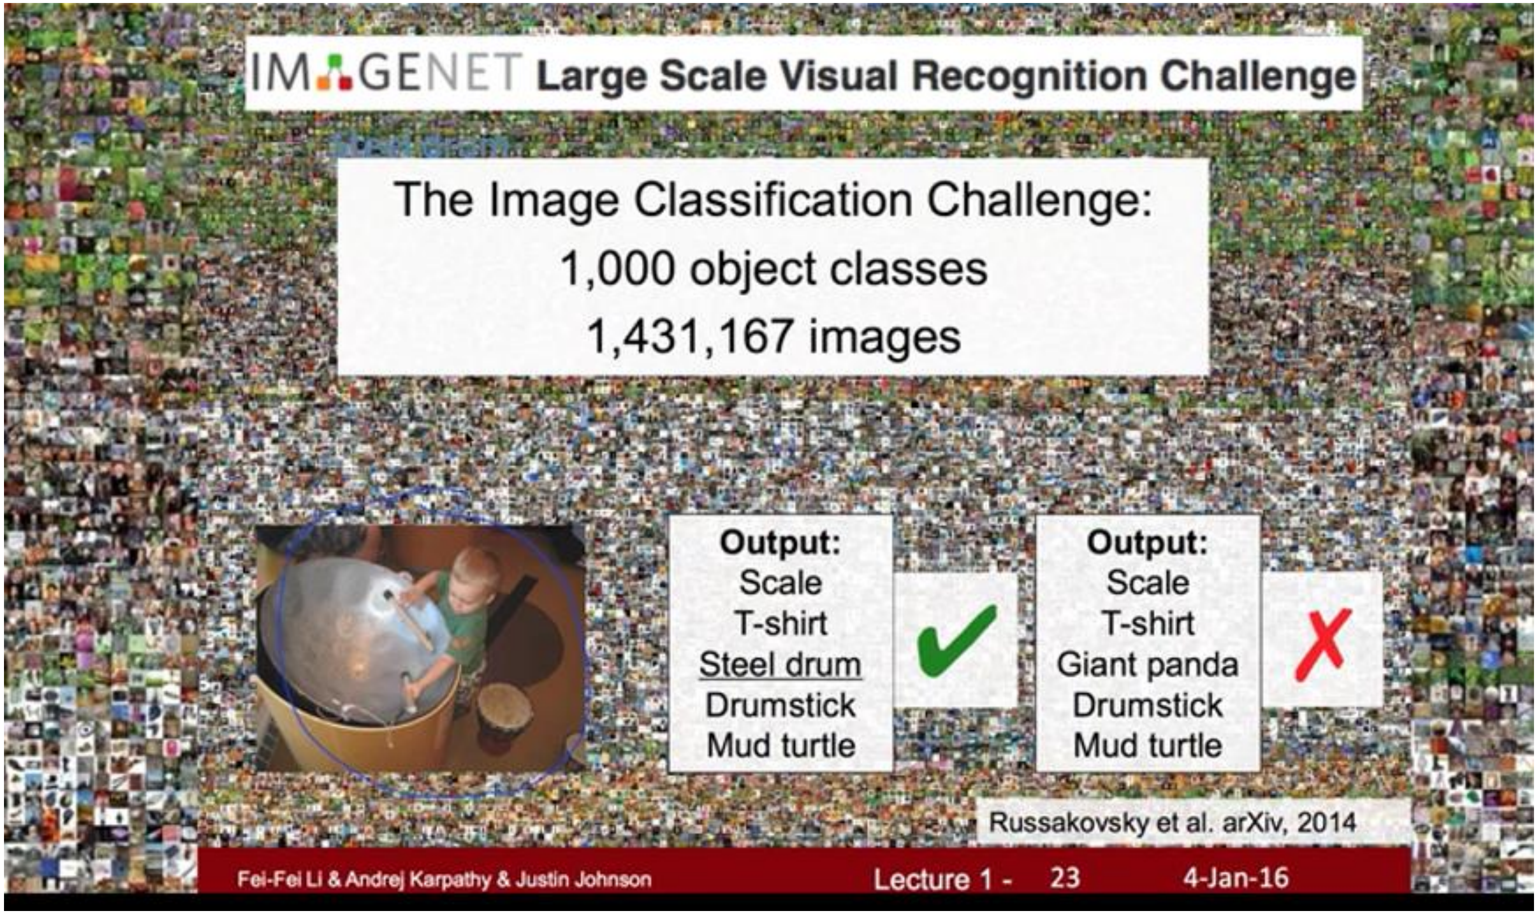

7. Deep Learning Applications(cont.)

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

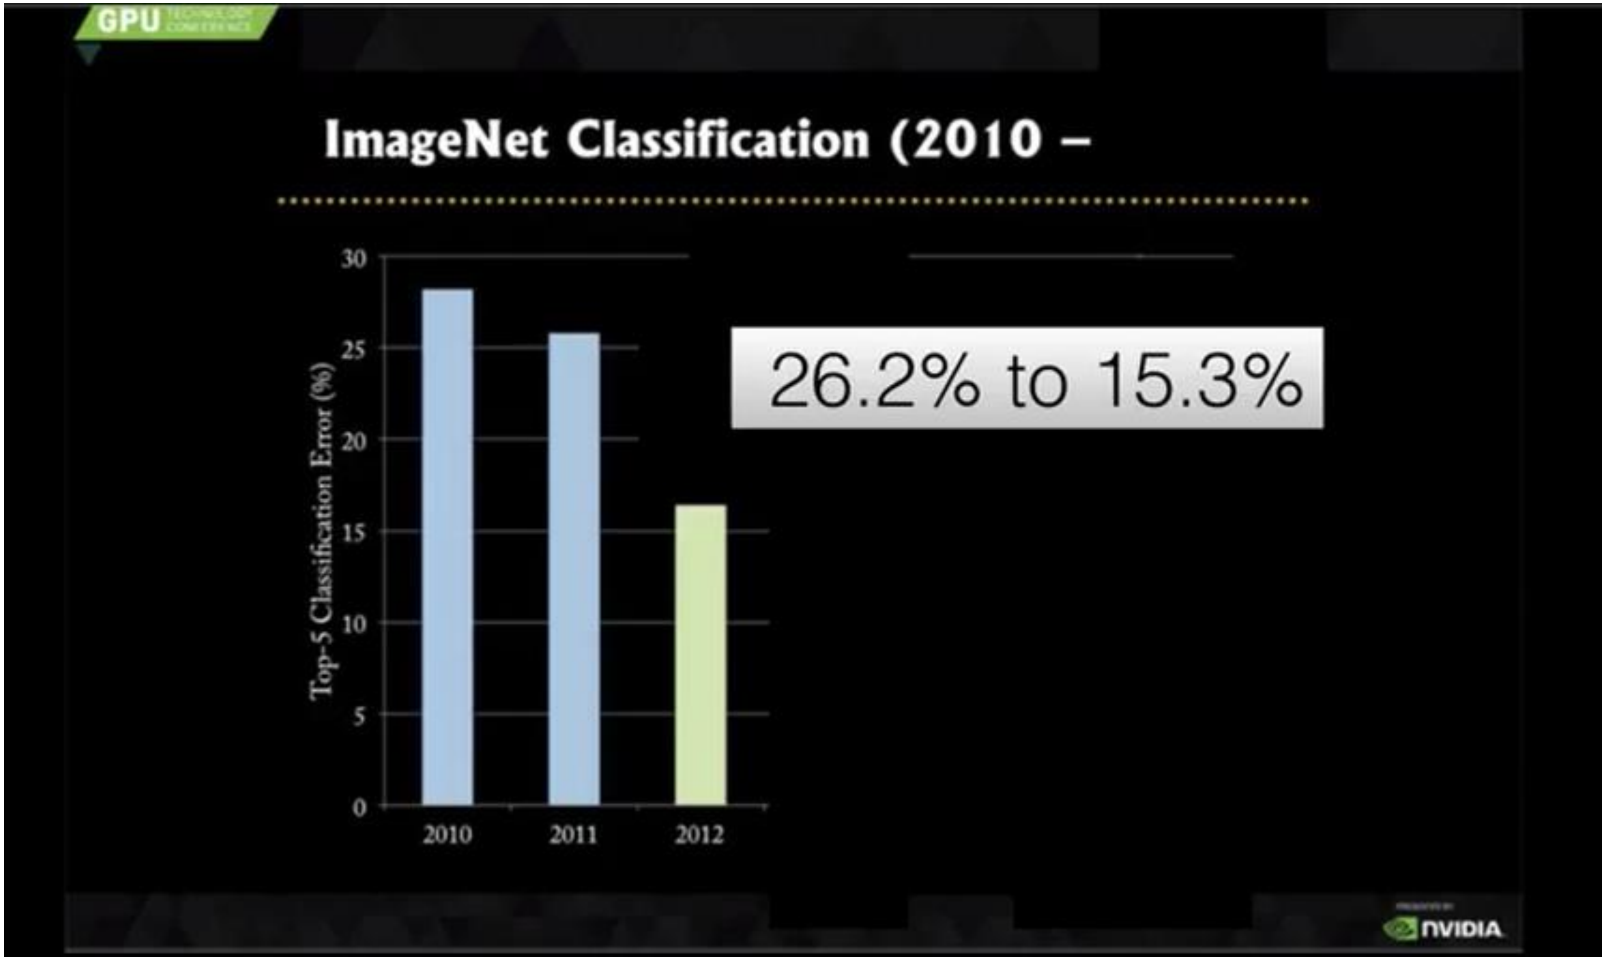

7. Deep Learning Applications(cont.)

- System based on Deep learning, MSRA team 2015
  in IMAGENET

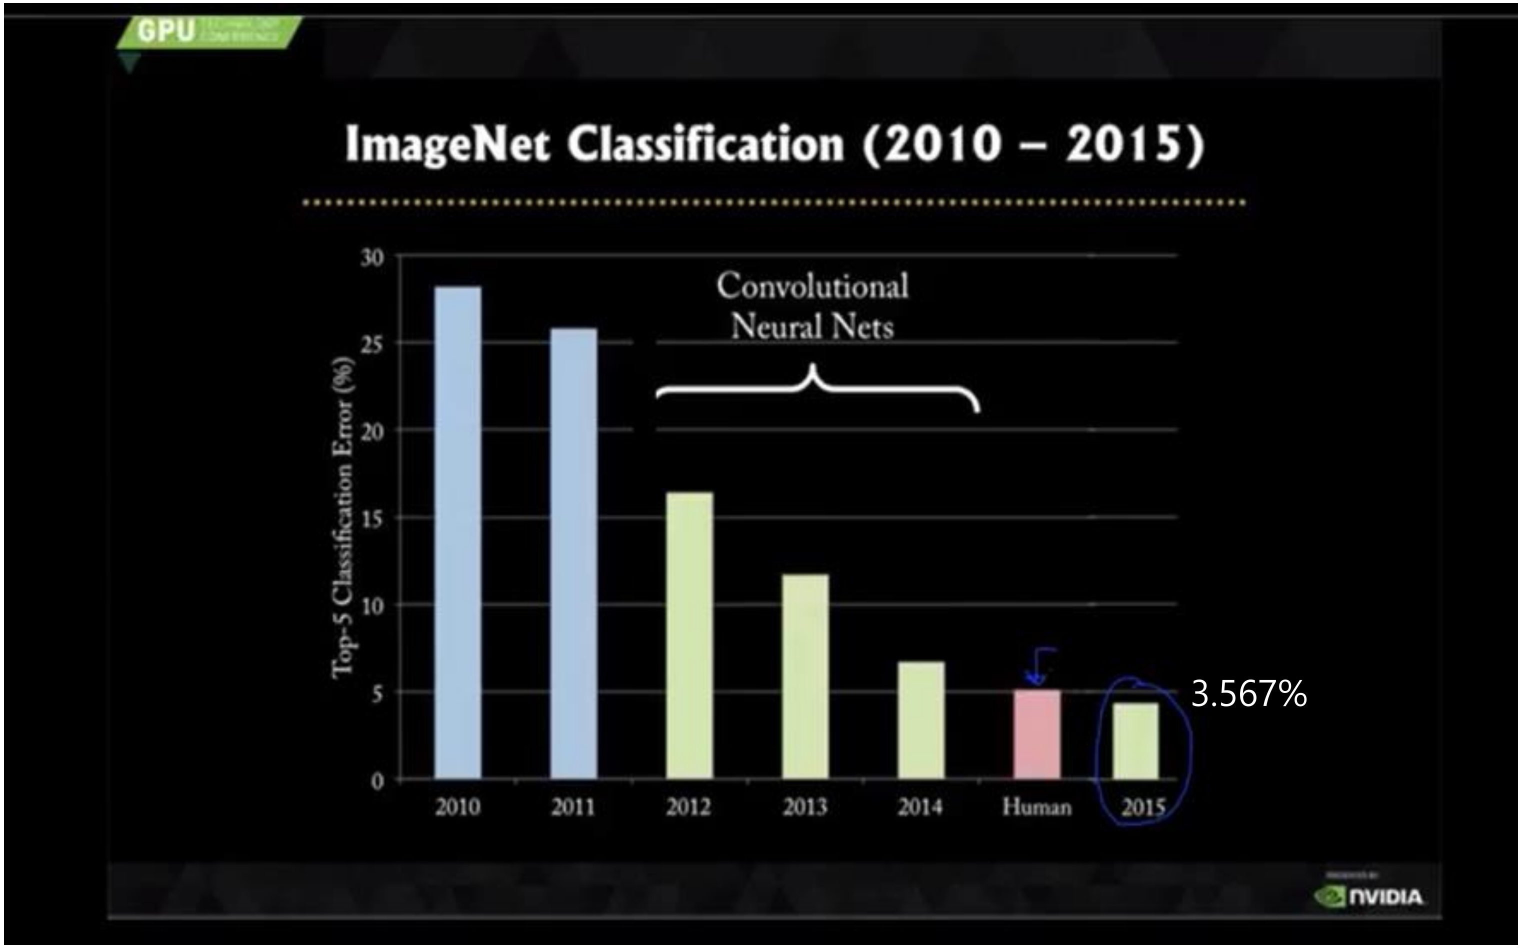

7. Deep Learning Applications(cont.)

- Ensemble 2 by Trimps-Soushen(2016) in IMAGENET
  - 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%

7. Deep Learning Applications(cont.)

- Neural networks that can explain photos

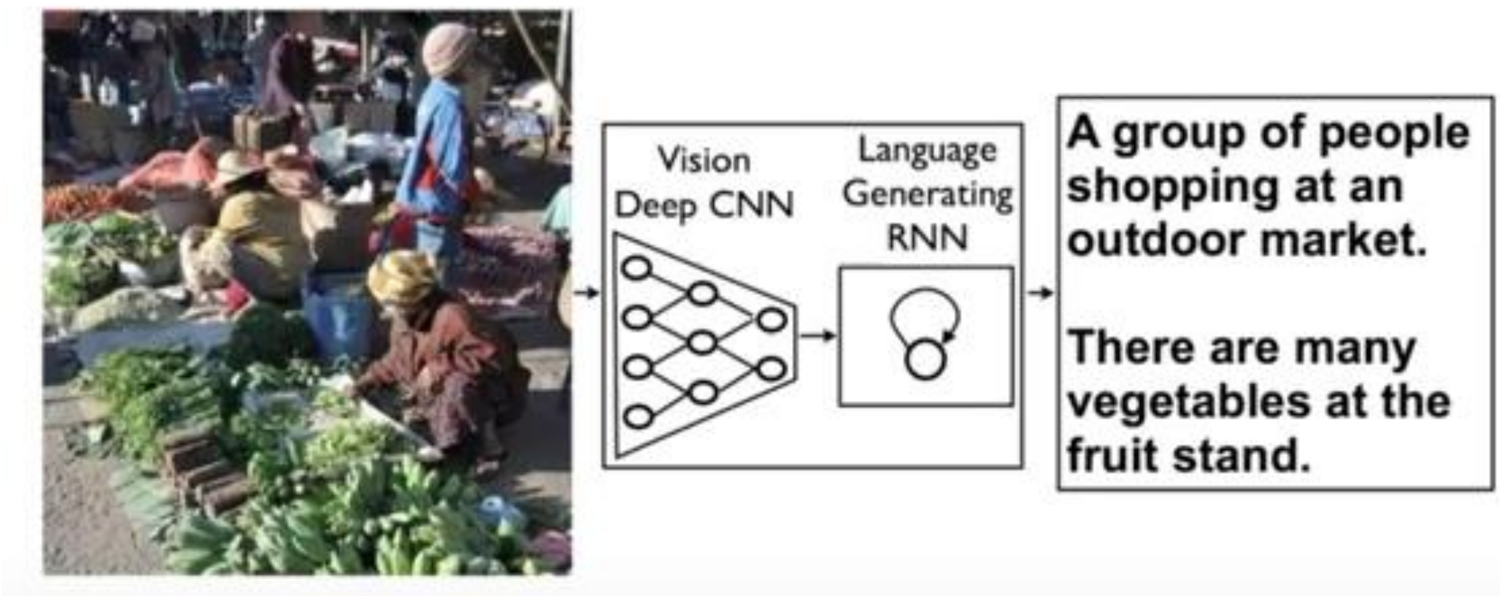

https://hunkim.github.io/ml/

7. Deep Learning Applications(cont.)

- Explain how to use API for a question

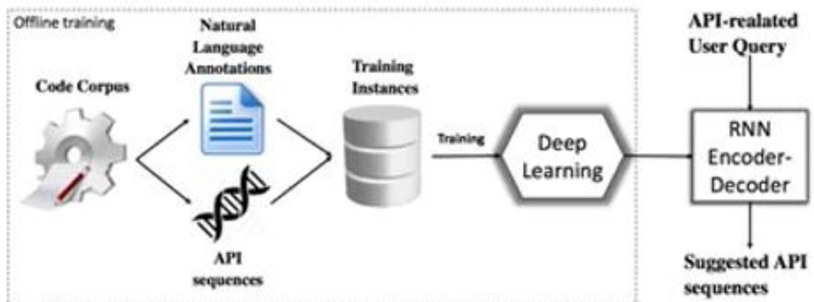

Figure 3: The Overall Workflow of DEEPAPI

“copy a file and save it to -your destination path” → 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 MSRA

7. Deep Learning Applications(cont.)

- Speech Recognition in noise environment

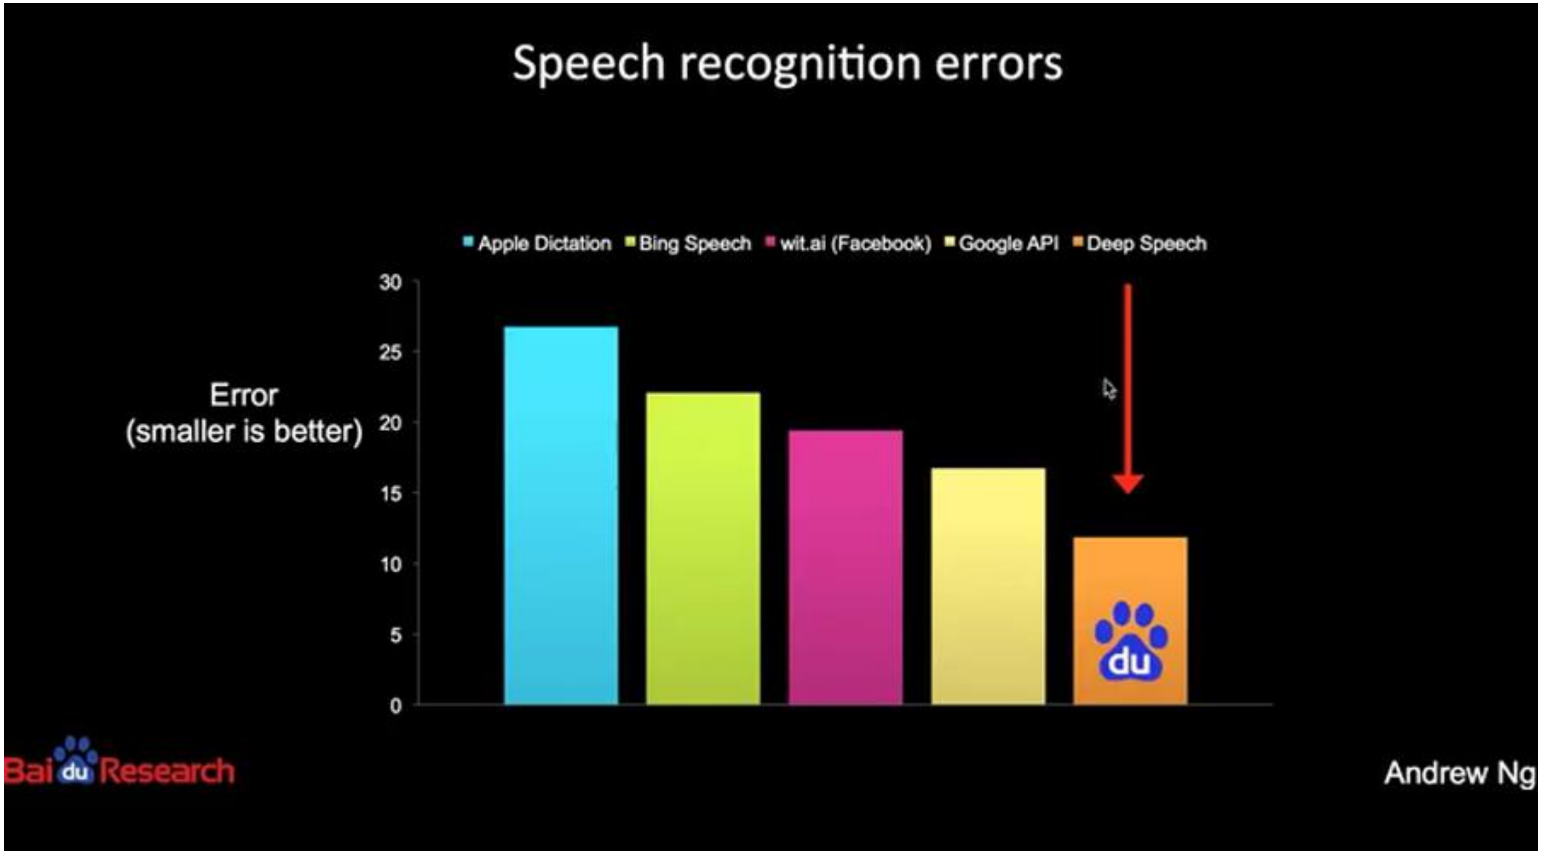

7. Deep Learning Applications(cont.)

- 프레임단위특징과 발음단위특징을 통합하는 Attention Mechanism을 이용한 음성감성인식 시스템

MSER(Merged Sepeech Emotion Recognition)

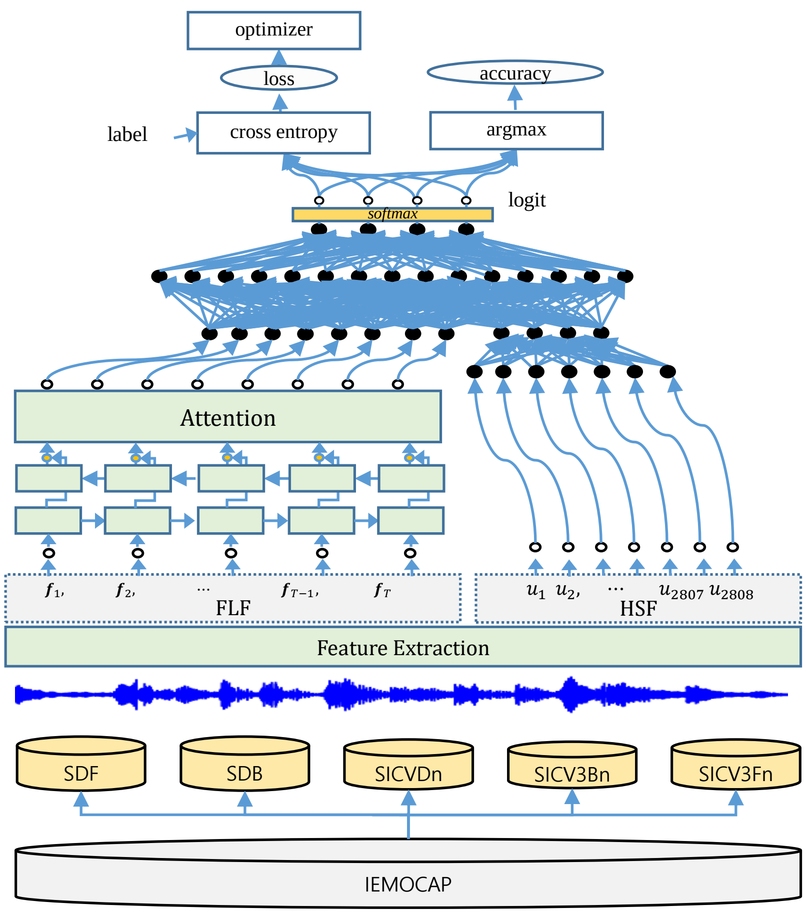

7. Deep Learning Applications(cont.)

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

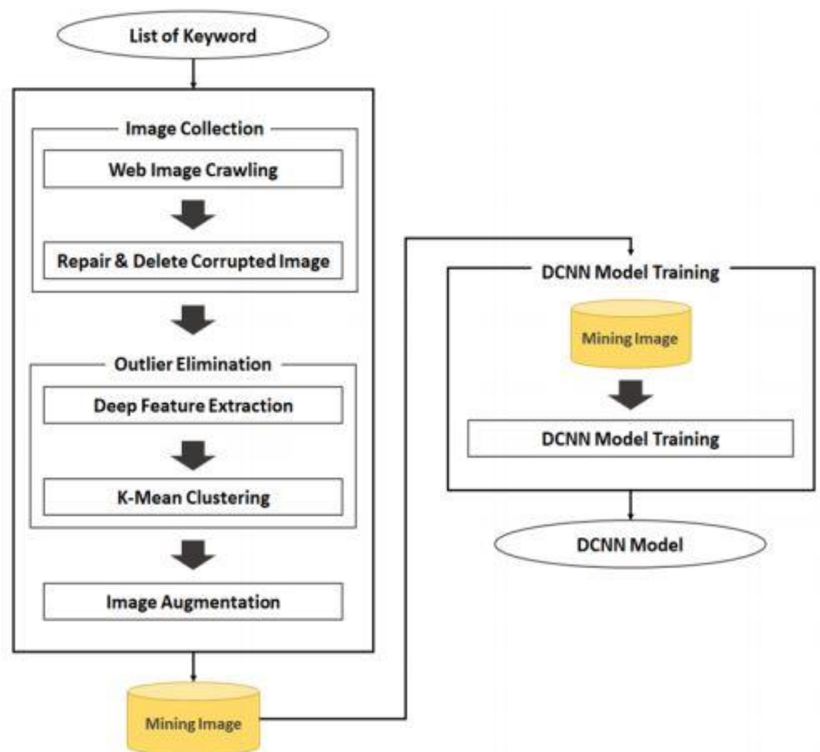

Fig. 1. Flow chart of the Automatic Bird-Species Recognition.

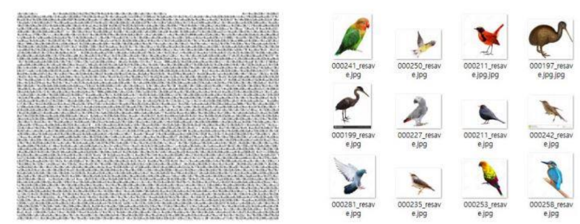

Fig. 2. Left, header error due to data loss. Right, the white-background images.

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

7. Deep Learning Applications(cont.)

- GAN을 이용한 음성 감정 인식 모델의 성능 개선

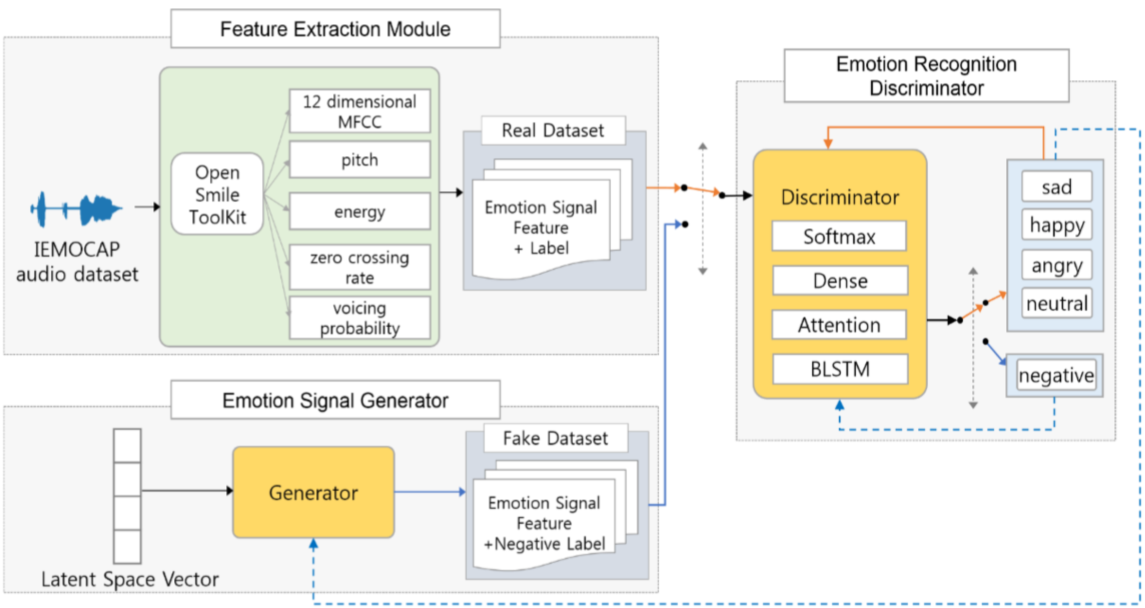

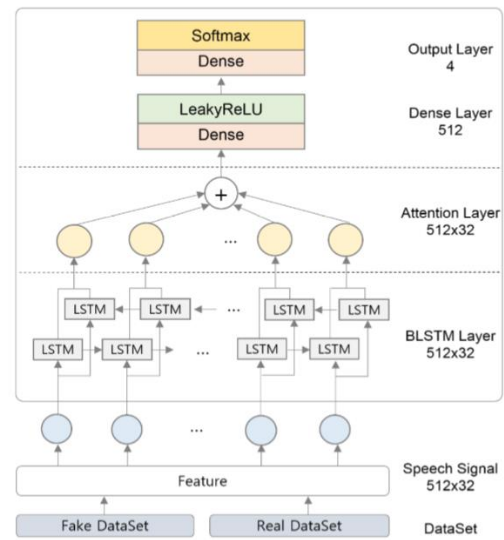

그림 1. 제안한 GAN을 이용한 음성 감정 인식 시스템 구조
Fig. 1. Proposed speech emotion recognition system architecture using GAN

http://www.wins.or.kr/papers/view.aspx?idx=213

그림 2. Attention과 BLSTM을 이용한 판별기
Fig. 2. Discriminator using attention and BLSTM

7. Deep Learning Applications(cont.)

- Two-Dimensional Attention-Based LSTM Model for Stock Index Prediction(ESCI)

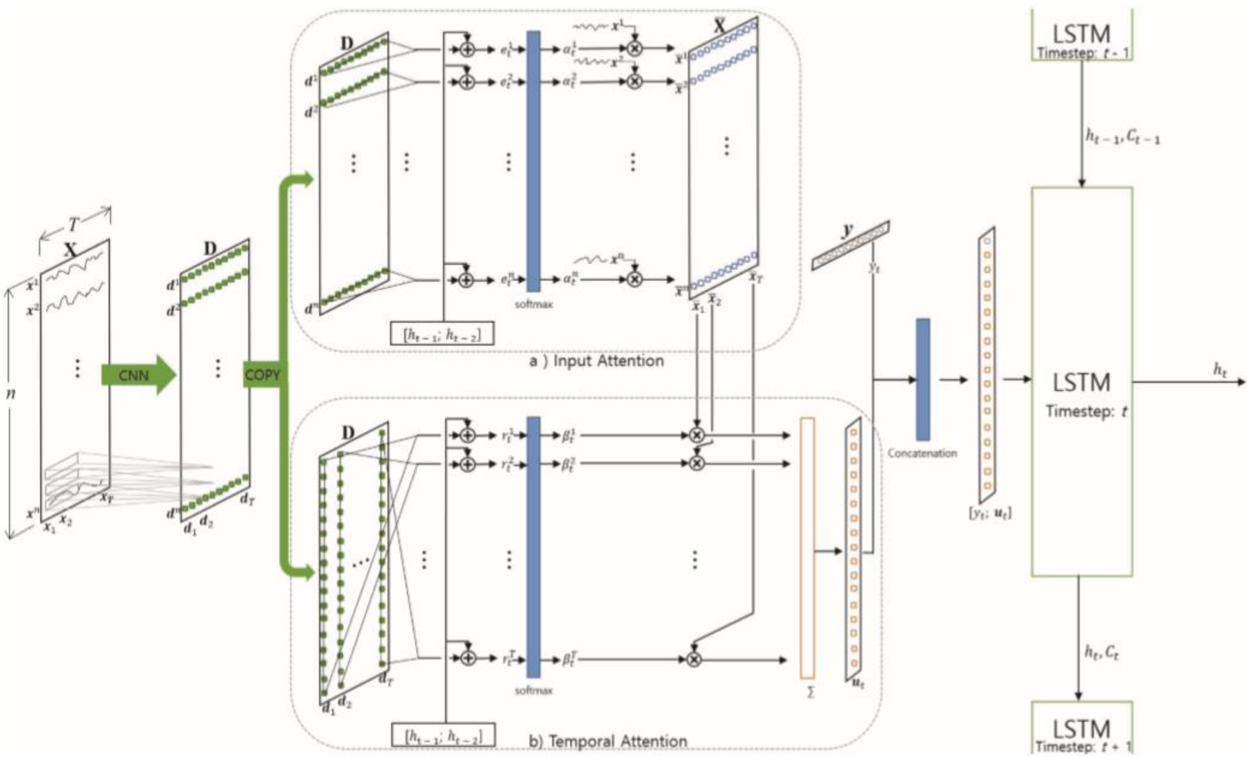

Fig. 1. Conceptual diagram of the 2D-ALSTM model composed of a CNN layer, an input and temporal attention layer, and an LSTM RNN layer.

http://www.wins.or.kr/papers/view.aspx?idx=193

7. Deep Learning Applications(cont.)

- A Voice Activity Detection Model composed of Bidirectional Long-Short Term Memory and Attention Mechanism.(IEEE Best Paper Award)

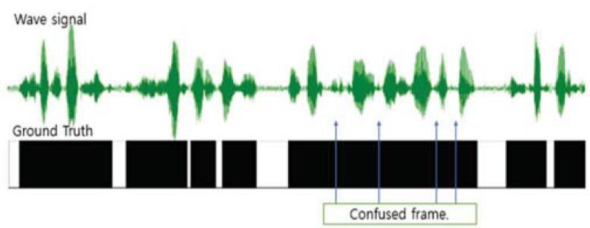

Fig.1. 10 seconds long noise added wave signal (SNR=10) and its ground truth for voice activity. Black indicates that the frame is labeled as voice.

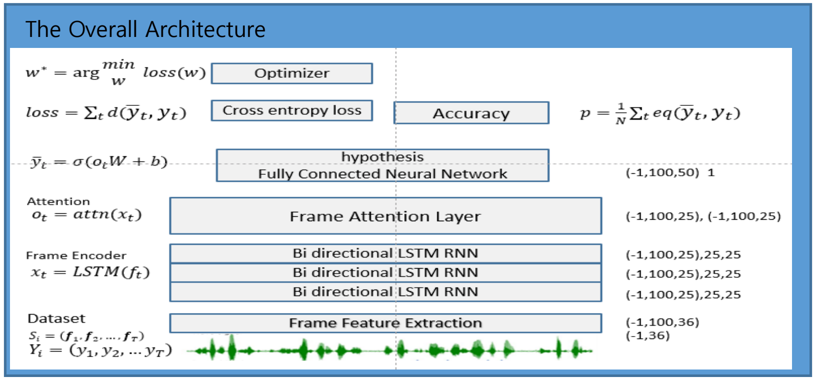

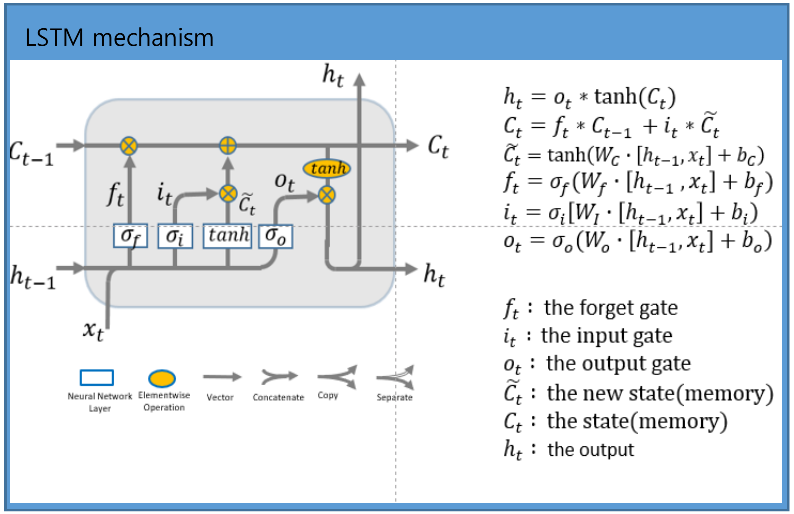

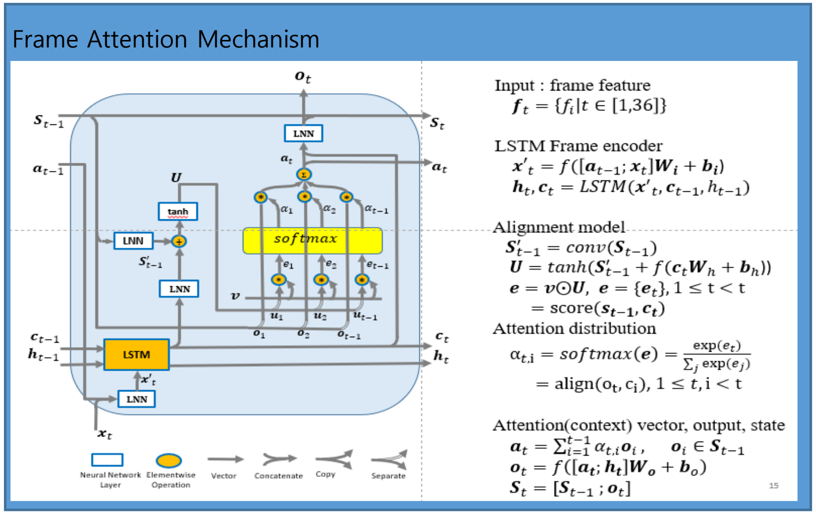

http://www.wins.or.kr/papers/view.aspx?idx=189

7. Deep Learning Applications(cont.)

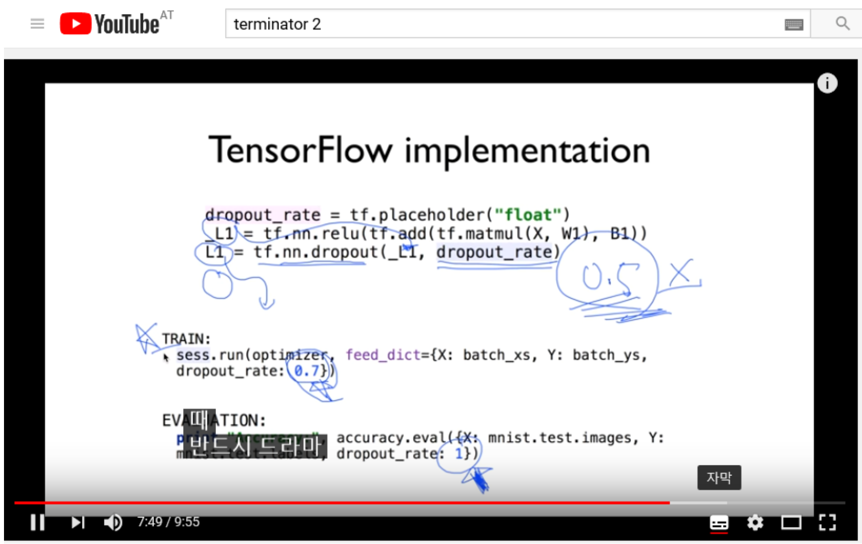

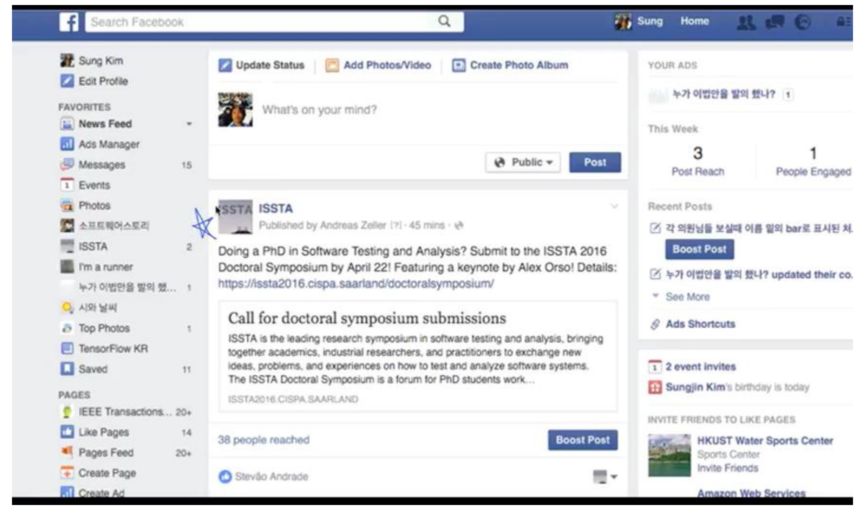

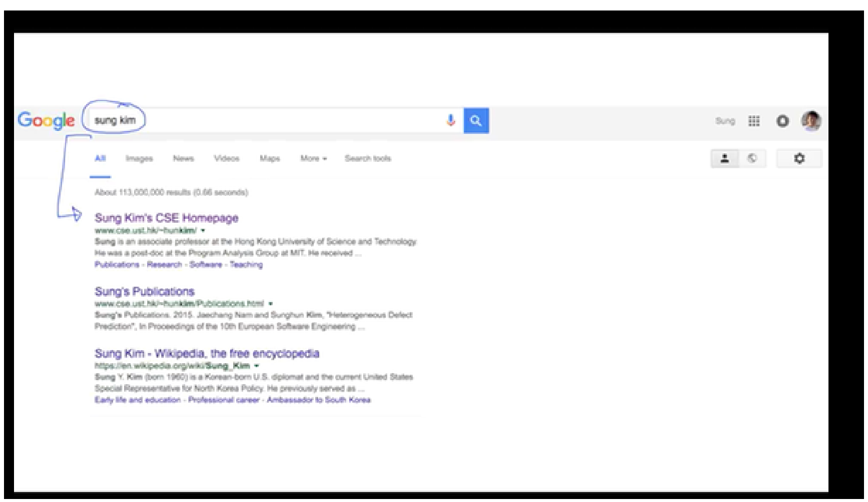

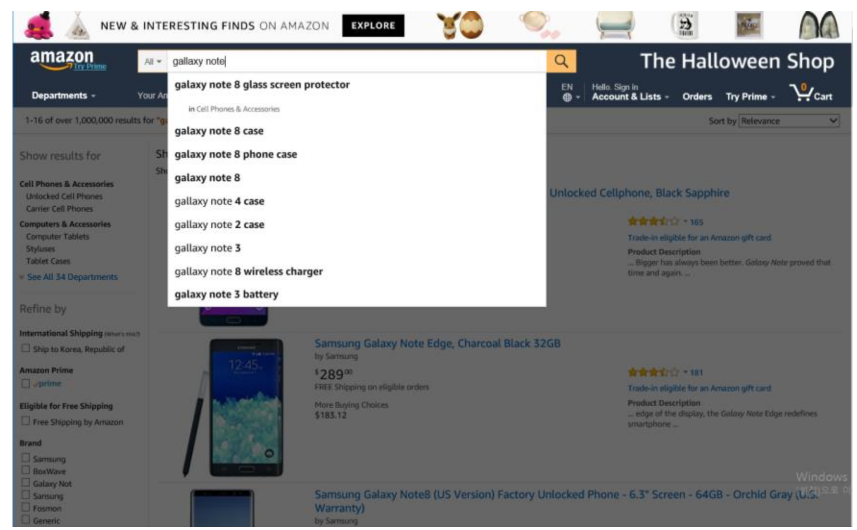

8. Recent Deep Learning Successes and Research Areas

8. Recent Deep Learning Successes and Research Areas

- 2016: Year of Deep Learning

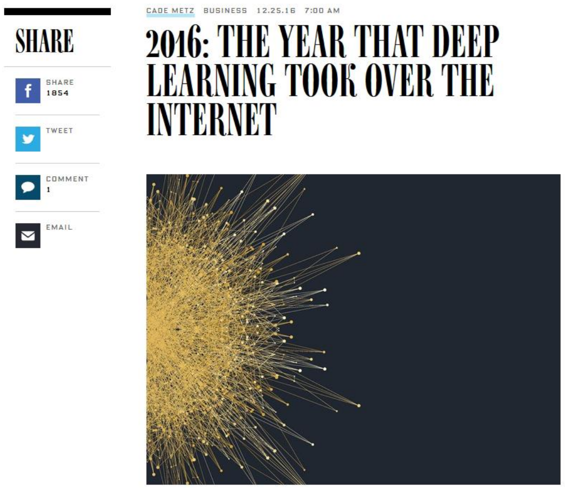

- 아만다 호지 슨 (Amanda Hodgson)은 호주 서해안에서 인도양을 향해 드론을 발사하여 물 위에서 사진을 찍을 수 있습니다. 사진은 멸종 위기에 처한 해양 포유류의 멸종을 막기위한 노력의 일환으로 퍼스 근처의 만에서 듀공(dugong) 또는 해우(sea cow)를 찾는 방법입니다. 문제는 Hodgson과 그녀의 팀이 모든 항공 사진을 조사 할 시간이 없다는 것입니다. 그들 중 너무 많은 수가 약 45,000 명이며, 듀공을 발견하는 것은 훈련되지 않은 눈으로는 너무 어렵습니다. 그래서 그녀는 작업을 깊은 신경망(DNN)에 제공하고 있습니다.

- 신경망은 Facebook 뉴스 피드에 게시 된 사진에서 얼굴을 식별하는 기계 학습 모델입니다. 또한 Android 휴대 전화에 질문 한 내용을 인식하고 Google 검색 엔진을 실행하는 데 도움이됩니다. 인간 두뇌의 뉴런 네트워크에서 느슨하게 모델링 된이 포괄적 인 수학적 모델은 방대한 디지털 데이터를 분석하여이 모든 것을 배웁니다. 현재 Perth에있는 Murdoch University의 해양 생물학자인 Hodgson은 동일한 기술을 사용하여 수천 개의 오픈 워터 사진에서 듀공을 찾고 동일한 오픈 소스 소프트웨어 인 TensorFlow에서 신경 네트워크를 실행하여 머신 러닝 서비스를 지원합니다.

8. Recent Deep Learning Successes and Research Areas(cont.)

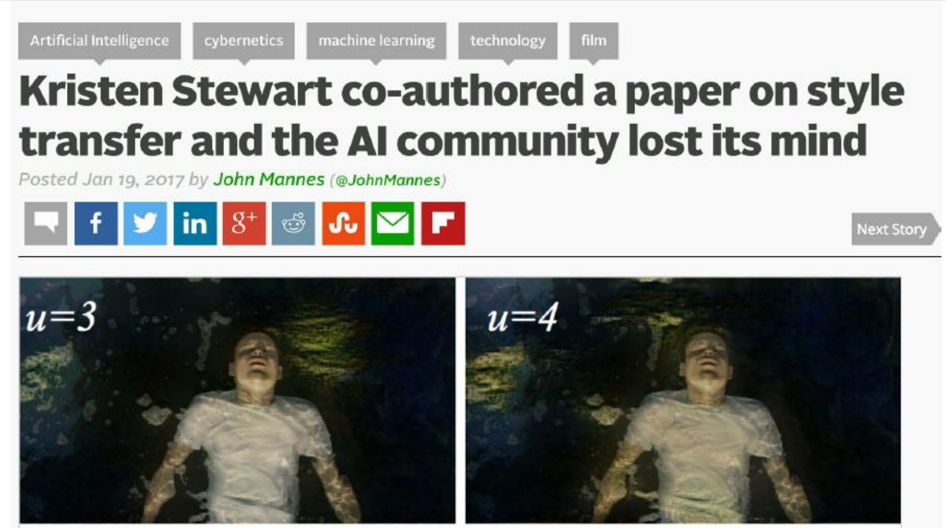

- Kristen Stewart는 기계 학습에 관한 논문을 발표함으로써 인공 지능 커뮤니티를 놀라게 했습니다.
- Twilight 여배우는 최근 단편 영화 Come Swim으로 감독 데뷔했으며,“스타일 전송”(일부 이미지 나 비디오의 미학이 다른 스타일에 적용되는)이라는 기계 학습 기술을 사용하여 인상적인 비주얼 스타일을 만들었습니다. 특수 효과 엔지니어인 Bhautik J Joshi 및 프로듀서 David Shapiro와 함께 Stewart는이 작품에 대한 논문을 공동 논문으로 저술했으며 비 피어 리뷰 작업을 위해 인기있는 온라인 저장소에 게시했습니다.

8. Recent Deep Learning Successes and Research Areas(cont.)

MIT Technology Review

Sign in | Subscribe

Topics | Magazine | Newsletters | Events

Intelligent Machines

Deep Learning

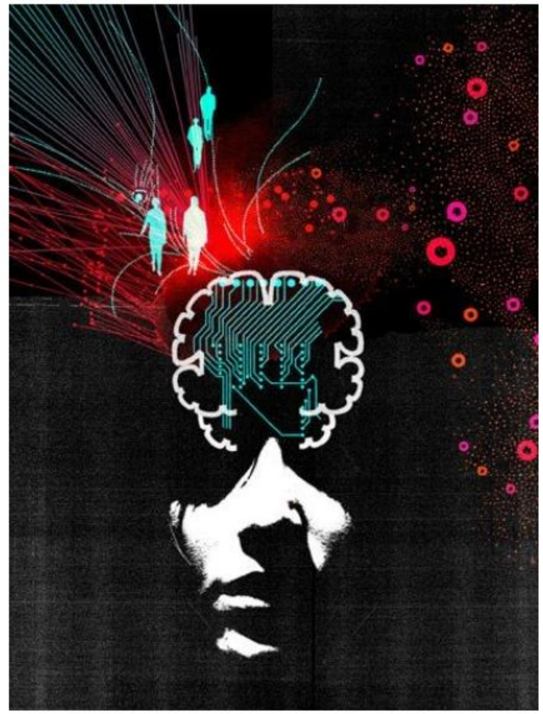

With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart.

by **Robert D. Hof** Apr 23, 2013

- 대량의 계산 능력으로 이제 기계는 물체를 인식하고 실시간으로 음성을 번역 할 수 있습니다. 인공 지능이 마침내 똑똑해 지고 있습니다.

8. Recent Deep Learning Successes and Research Areas(cont.)

- Machine Translation

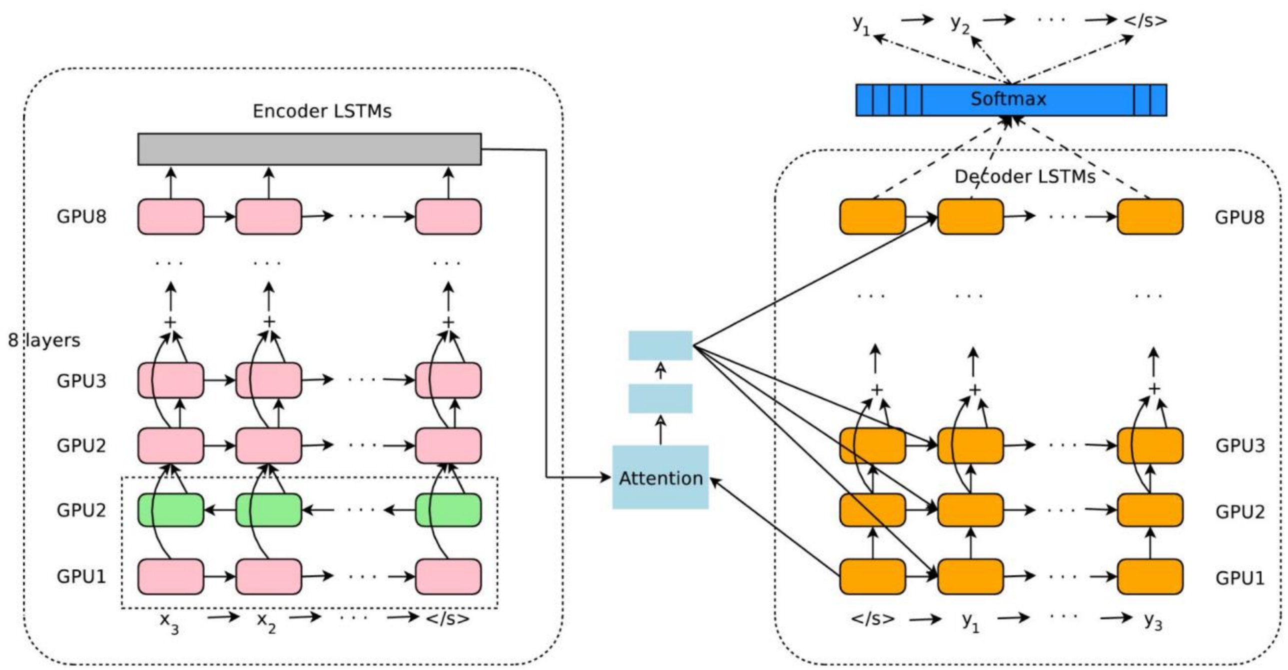

- Your Google Translate usage will now be powered by an **8 layer Long Short Term Memory Network with residual connections and attention**

Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation; Wu et al.

8. Recent Deep Learning Successes and Research Areas(cont.)

- Artistic Style

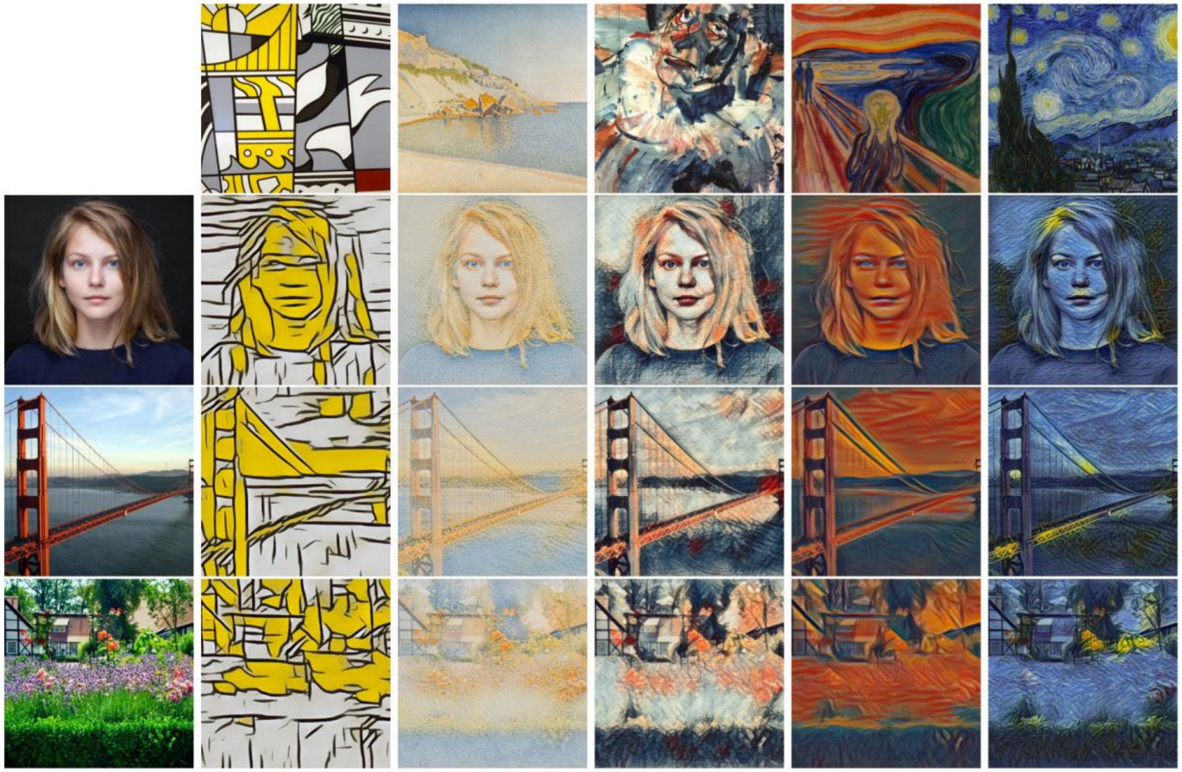

(a) With conditional instance normalization, a single style transfer network can capture 32 styles at the same time, five of which are shown here. All 32 styles in this single model are in the Appendix. Golden Gate Bridge photograph by Rich Niewiroski Jr.

A Learned Representation for Artistic Style; Dumoulin, Shlens, Kudlur; ICLR 2017

8. Recent Deep Learning Successes and Research Areas(cont.)

- Speech Synthesis

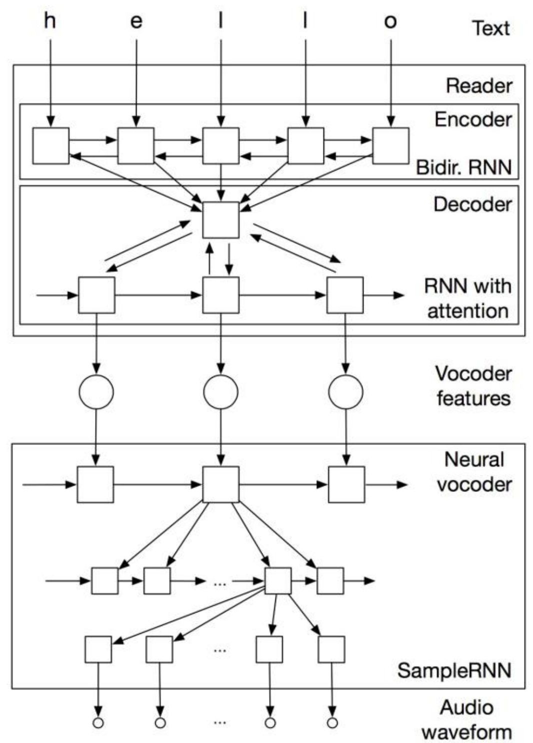

Figure 1: Char2Wav: An end-to-end speech synthesis model.

Char2Wav: End-to-End Speech Synthesis; Sotelo et al., ICLR 2017; http://josesotelo.com/speechsynthesis/

8. Recent Deep Learning Successes and Research Areas(cont.)

- Neuroevolution of Architectures

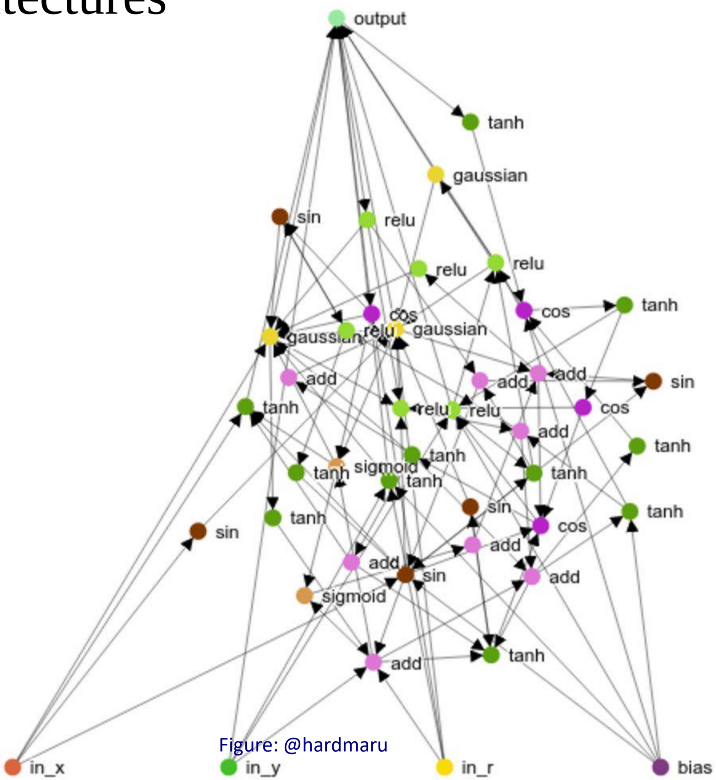

Figure: @hardmaru

- Recent large scale studies by Google show that evolutionary methods are catching with intelligently designed architectures

8. Recent Deep Learning Successes and Research Areas(cont.)

9th and 15th of March 2016

- Game Playing

Mastering the game of Go with deep neural networks and tree search; Silver et al., Nature; 2016

8. Recent Deep Learning Successes and Research Areas(cont.)

- 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

16

8. Recent Deep Learning Successes and Research Areas(cont.)

- Protein Folding(단백질 폴딩)
- Drug discovery(약물발견)
- Particle Physics(입자물리학)
- Energy Management(에너지관리)
- ...

Next time

- 기계학습
- 지도학습
- …

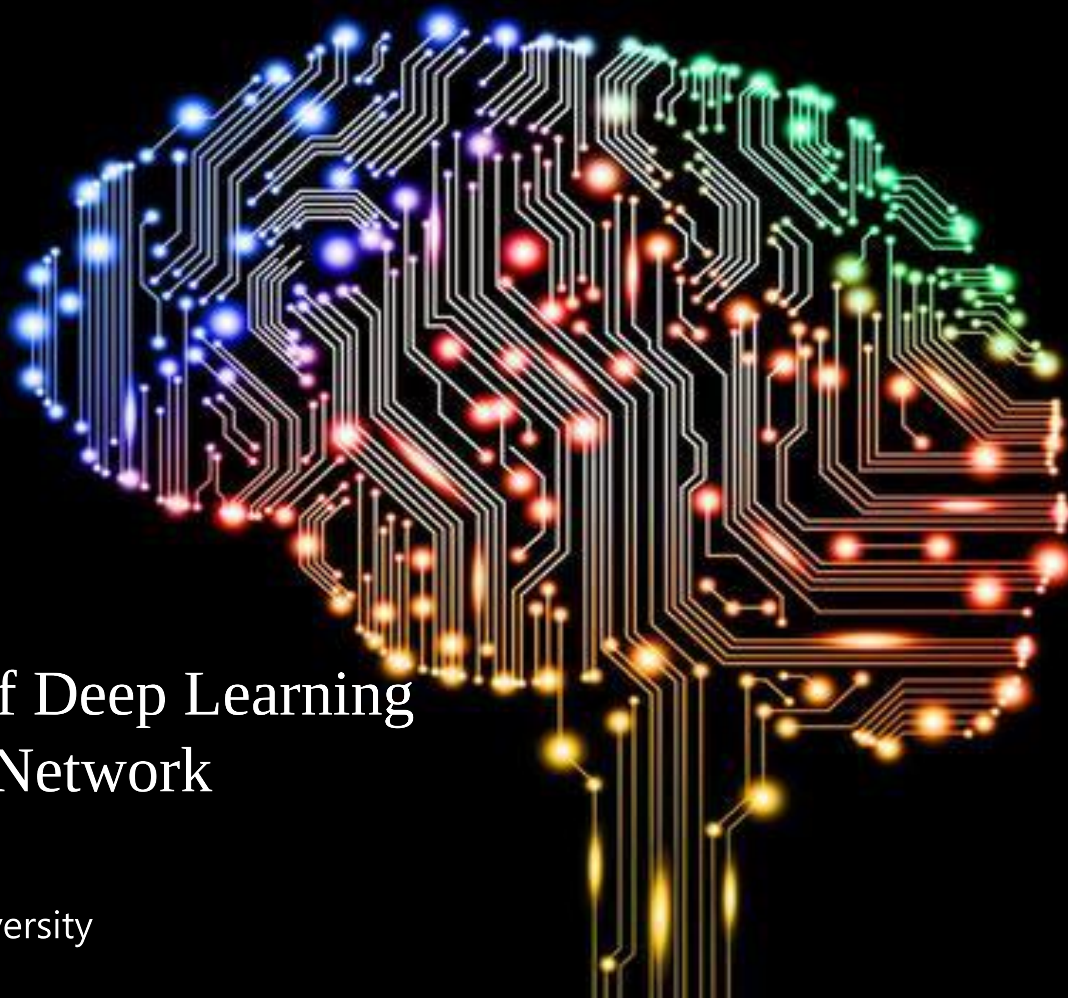

Introduction of Deep Learning
Deep Neural Network

Yoon Joong Kim,
Hanbat National University