


DEEP LEARNING

구글 트렌드: 딥러닝



10 BREAKTHROUGH TECHNOLOGIES 2013

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.



Temporary Social Media

Messages that quickly self-destruct could enhance the privacy of online communications and make people freer to be spontaneous.



Prenatal DNA Sequencing

Reading the DNA of fetuses will be the next frontier of the genomic revolution. But do you really want to know about the genetic problems or musical aptitude of your unborn child?



Additive Manufacturing

Skeptical about 3-D printing? GE, the world's largest manufacturer, is on the verge of using the technology to make jet parts.



Baxter: The Blue-Collar Robot

Rodney Brooks's newest creation is easy to interact with, but the complex innovations behind the robot show just how hard it is to get along with people.



Memory Implants

A maverick neuroscientist believes he has deciphered the code by which the brain forms long-term memories. Next: testing a prosthetic implant for people suffering from long-term memory loss.



Smart Watches

The designers of the Pebble watch realized that a mobile phone is more useful if you don't have to take it out of your pocket.



Ultra-Efficient Solar Power

Doubling the efficiency of a solar cell would completely change the economics of renewable energy. Nanotechnology just might make it possible.



Big Data from Cheap Phones

Collecting and analyzing information from simple cell phones can provide surprising insights into how people move about and behave – and even help us understand the spread of diseases.




Supergrids

A new high-power circuit breaker could finally make highly efficient DC power grids practical.



음성 인식 성능



- NAVER 음성 인식기: 20% 성능 개선

Scientists See Promise in Deep-Learning Programs



Hao Zhang/The New York Times

A voice recognition program translated a speech given by Richard F. Rashid, Microsoft's top scientist, into Mandarin Chinese.

By JOHN MARKOFF
Published: November 23, 2012

Using an artificial intelligence technique inspired by theories about how the brain recognizes patterns, technology companies are reporting startling gains in fields as diverse as computer vision, speech recognition and the identification of promising new molecules for designing drugs.

- FACEBOOK
- TWITTER
- GOOGLE+
- SAVE
- EMAIL

MOST EMAILED

MOST VIEWED



1. THOMAS L. FRIEDMAN
How to Get a Job at Google, Part 2

2. OP-ED CONTRIBUTOR
The Public Health Crisis Hiding in Our Food



3. 50 Years Into the War on Poverty, Hardship Hits Back



4. Taking On Adam Smith (and Karl Marx)

사회 '딥 러닝 SW' 교통표지판 인지도 인가 인간 뇌 신경망 모사로 획기적 발전... 인공지능 연구·실용화 현주소

Tweet 0

8+1 0

목록 | 메일 | 인쇄 | 글씨크기

● 1400만장 사진 2만개로 분류면 정확도 15.8%

딥 러닝의 성과는 다양한 분야에서 활용이 가능하다. 무엇이 사용자의 기호에 맞는 것 인지를 고려한 애플리케이션의 정렬 프로그램뿐 아니라 마케팅이나 치안에도 쓸 수 있다. 예를 들어 사람들의 구매 습관에 대한 데이터베이스를 활용해 구매고객에 따라 스스로 진열을 바꾸는 상점을 만들거나, 얼굴인식의 기능을 획기적으로 개선한 통합 범죄 감시 시스템을 구축할 수도 있다.

딥 러닝이 구현한 인식기술은 이제 사람과의 경쟁에 나서고 있다. 문서는 정형화된 구조로 돼 있어 검색이 쉽지만, 이미지나 비디오는 약간의 변형이나 각도 전환만으로도 완벽하게 일치하지 않기 때문에 일정 수준 이상의 '판단'이 필요하다. 결국 이미지와 비디오 검색이 가능하다는 것은 인공지능이 제대로 작동하고 있다는 것을 의미한다. 페이스북에는 2000억장의 사진이 게재돼 있고, 매분마다 72시간 분량의 새로운 비디오가 올라온다. 이 같은 데이터가 용량만 차지하는 쓰레기가 되지 않기 위해서는 분류가 필요하지만, 기존의 기술로는 이를 자동화할 방법이 없었다.

지난해 스위스 인공지능연구소 연구진은 자체적으로 개발한 딥 러닝 기술로 사람들과 교통표지판 인식 대결을 벌였다. 연구진의 딥 러닝 프로그램은 총 5만장의 표지판 그림 중 99.46%를 정확하게 인지해내 32명의 사람들로 구성된 인간팀의 99.22%를 앞섰다. 오랜 운전경력을 가진 인간팀의 개인당 정확도는 98.84%였다.

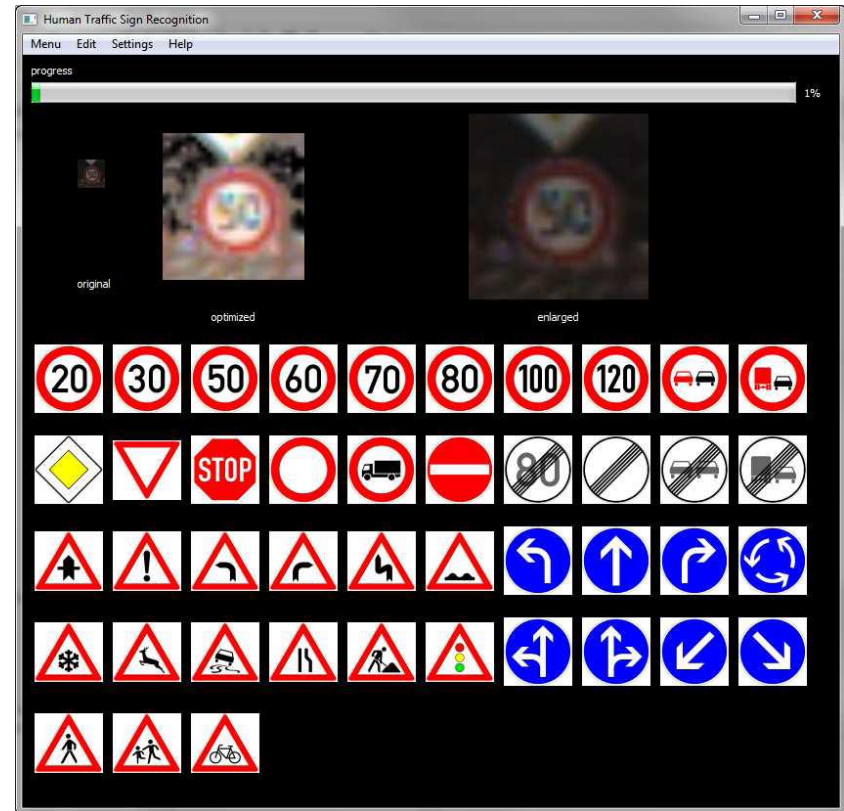
하지만 교통표지판처럼 명확한 한계를 설정하지 않으면, 아직 갈 길이 멀다. 올여름 구글과 스탠퍼드대 연구팀은 1만 6000대의 컴퓨터를 사람의 뇌 신경처럼 연결해 1400만장에 이르는 사진을 2만개의 카테고리에 자동으로 분류하는 작업을 진행했다. 최종 분류의 정확도는 15.8%에 불과했다.

프로젝트 책임자인 제프 딘은 "이전에 최고 수준으로 평가받던 이미지 인식 기술보다 70% 이상 향상된 수치인 만큼 아직까지 무궁무진한 개선의 여지가 있다."면서 "아무리 숙련된 사람이라도 5분에 250개가량의 이미지만 분류할 수 있고, 이는 인간의 힘으로는 정보처리가 불가능하다는 것을 의미한다."고 지적했다.

German Traffic Sign Recognition benchmark

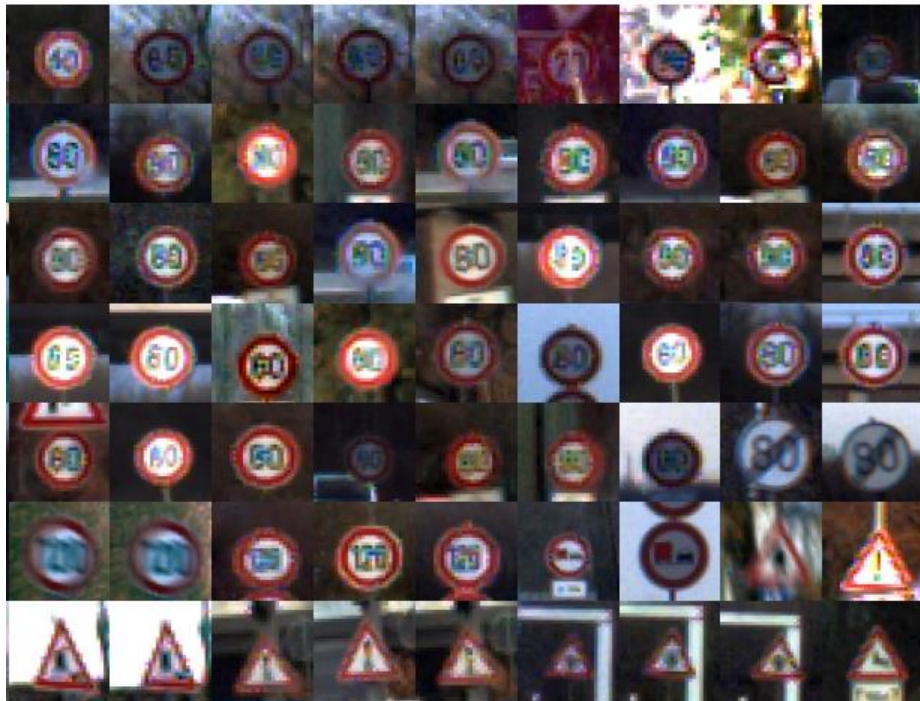


샘플 이미지



사용자 인터페이스

GTSRB 결과



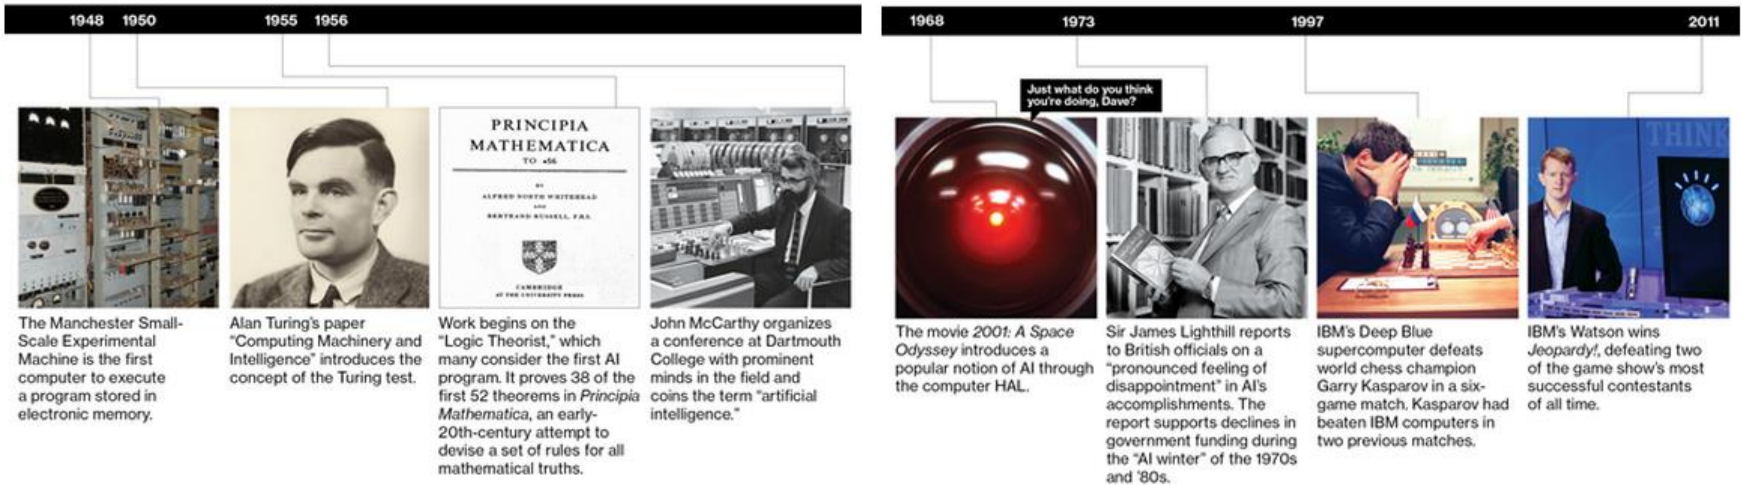
실패한 경우

CCR (%)	Team	Method
99.46	IDSIA	Committee of CNNs
99.22	INI-RTCV	Human (best individual)
98.84	INI-RTCV	Human (average)
98.31	Sermanet	Multi-Scale CNN
96.14	CAOR	Random Forests
95.68	INI-RTCV	LDA (HOG 2)
93.18	INI-RTCV	LDA (HOG 1)
92.34	INI-RTCV	LDA (HOG 3)

AI 역사




AI's Evolution



The Manchester Small-Scale Experimental Machine is the first computer to execute a program stored in electronic memory.



Alan Turing's paper "Computing Machinery and Intelligence" introduces the concept of the Turing test.



Work begins on the "Logic Theorist," which many consider the first AI program. It proves 38 of the first 52 theorems in *Principia Mathematica*, an early-20th-century attempt to devise a set of rules for all mathematical truths.



John McCarthy organizes a conference at Dartmouth College with prominent minds in the field and coins the term "artificial intelligence."



The movie *2001: A Space Odyssey* introduces a popular notion of AI through the computer HAL.



Sir James Lighthill reports to British officials on a "pronounced feeling of disappointment" in AI's accomplishments. The report supports declines in government funding during the "AI winter" of the 1970s and '80s.




IBM's Deep Blue supercomputer defeats world chess champion Garry Kasparov in a six-game match. Kasparov had beaten IBM computers in two previous matches.



IBM's Watson wins *Jeopardy!*, defeating two of the game show's most successful contestants of all time.

Turing test

- 지능의 기준
 - 대화를 통한 컴퓨터의 사고력 측정

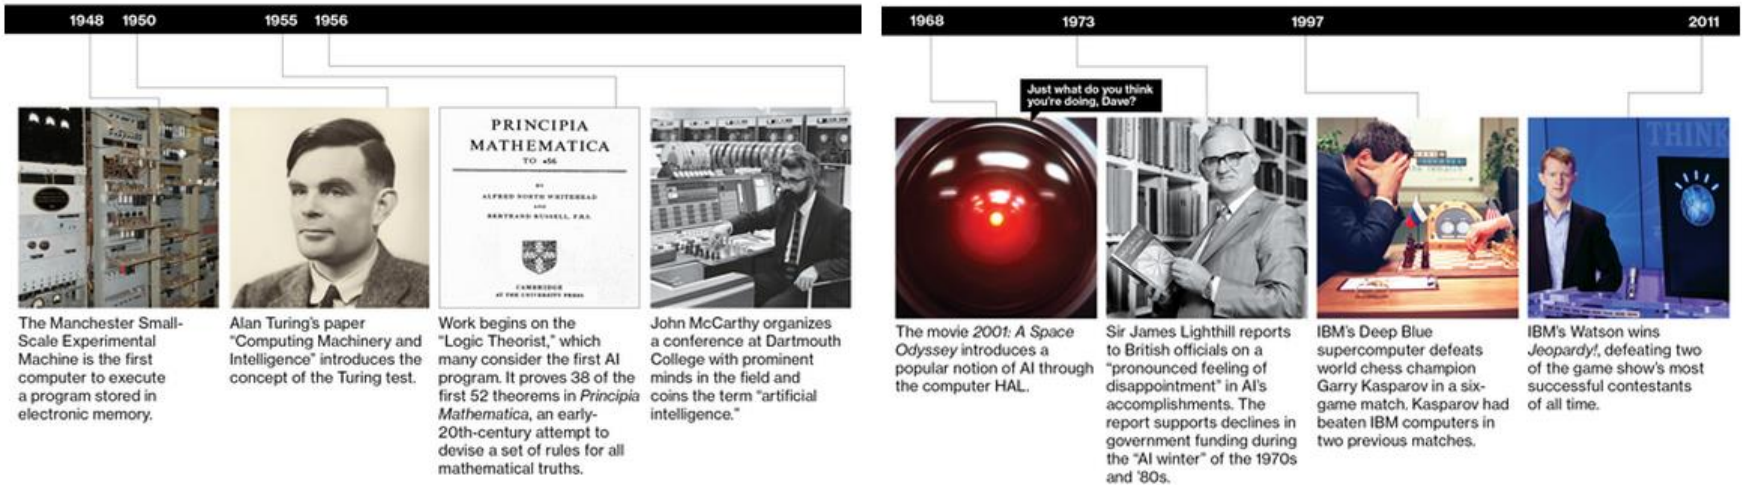



튜링테스트 진행 과정 개요도 (삽화 제공 : 안치영)

AI 역사




AI's Evolution





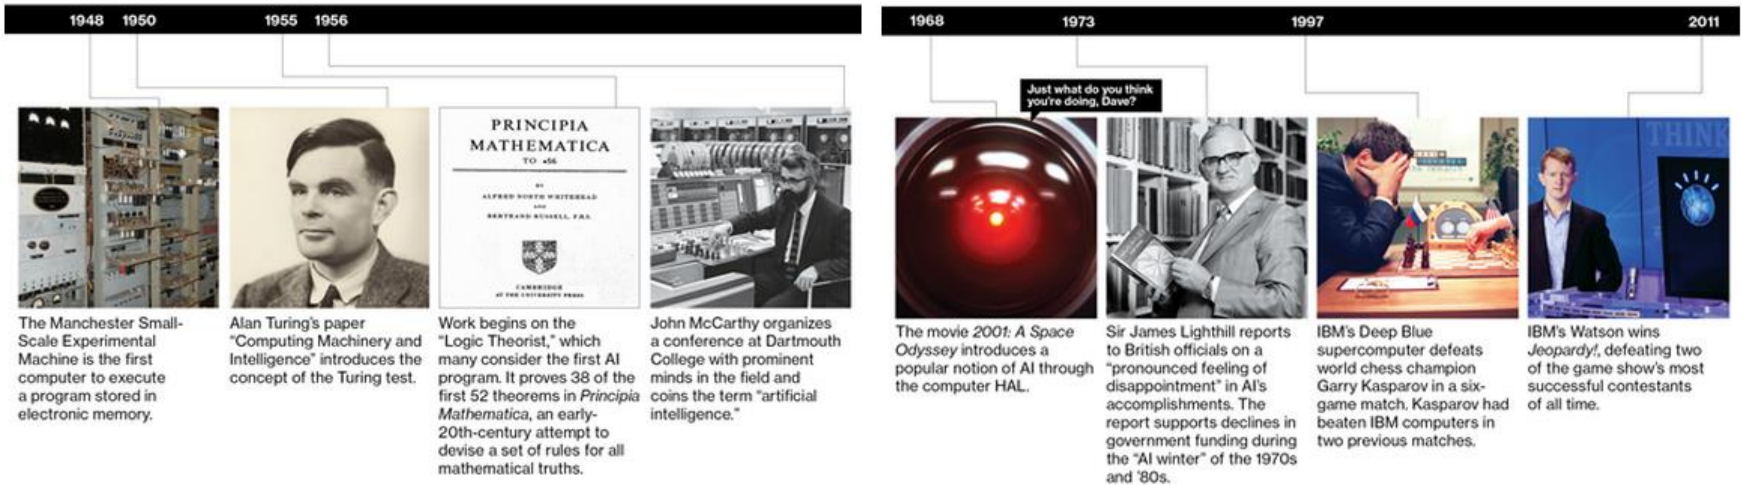
I'm sorry Dave.



AI 역사




AI's Evolution



The Manchester Small-Scale Experimental Machine is the first computer to execute a program stored in electronic memory.



Alan Turing's paper "Computing Machinery and Intelligence" introduces the concept of the Turing test.



Work begins on the "Logic Theorist," which many consider the first AI program. It proves 38 of the first 52 theorems in *Principia Mathematica*, an early-20th-century attempt to devise a set of rules for all mathematical truths.



John McCarthy organizes a conference at Dartmouth College with prominent minds in the field and coins the term "artificial intelligence."



The movie *2001: A Space Odyssey* introduces a popular notion of AI through the computer HAL.



Sir James Lighthill reports to British officials on a "pronounced feeling of disappointment" in AI's accomplishments. The report supports declines in government funding during the "AI winter" of the 1970s and '80s.



IBM's Deep Blue supercomputer defeats world chess champion Garry Kasparov in a six-game match. Kasparov had beaten IBM computers in two previous matches.



IBM's Watson wins *Jeopardy!*, defeating two of the game show's most successful contestants of all time.

딥블루


- **딥블루 vs 게리 카스파로프, 1997**

- Deep Blue vs Kasparov

- $3\frac{1}{2}$ vs $2\frac{1}{2}$

- Brute-force search power

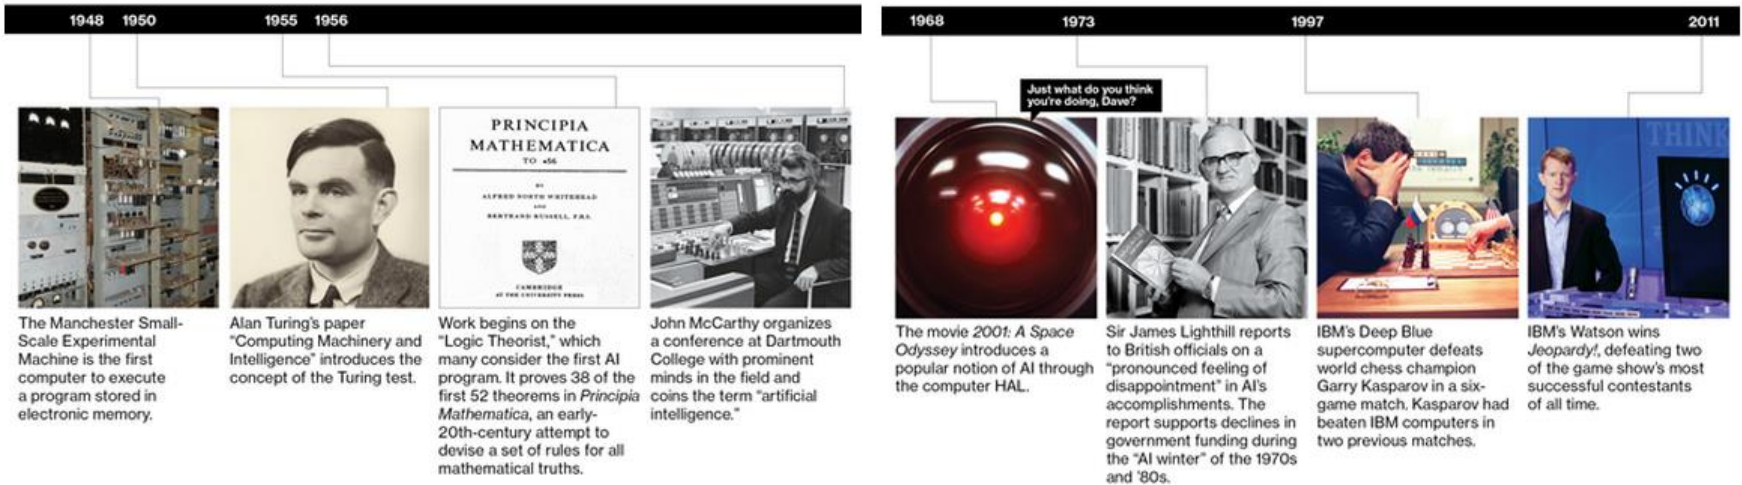
- 6~8 수를 내다봄



AI 역사



AI's Evolution



IBM Watson 슈퍼컴퓨터

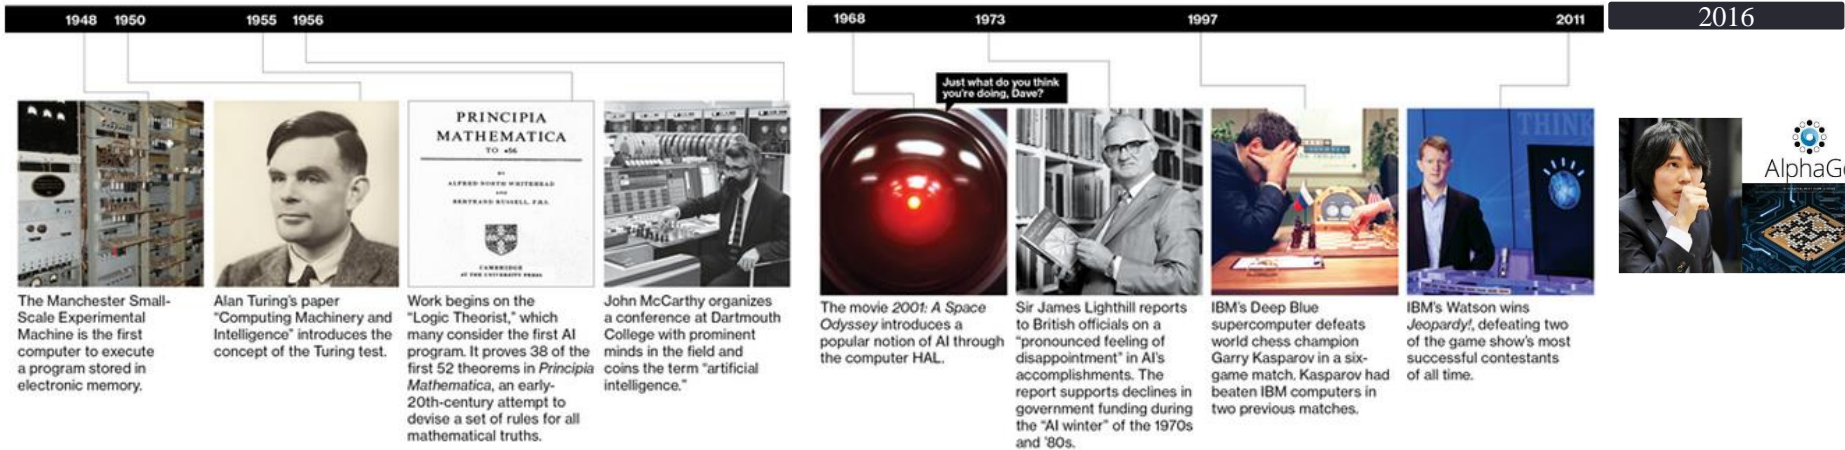
- 질문예시)
 - **KATHLEEN KENYON'S EXCAVATION OF THIS CITY MENTIONED IN JOSHUA SHOWED THE WALL HAD BEEN REPAIRED 17 TIMES**
 - **WHAT is "Jericho"**
 - **THIS CHILD STAR GOT HIS FIRST ON-SCREEN KISS IN "MY GIRL"**
 - **WHO is "Macaulay Culkin"**



AI 역사





AI's Evolution



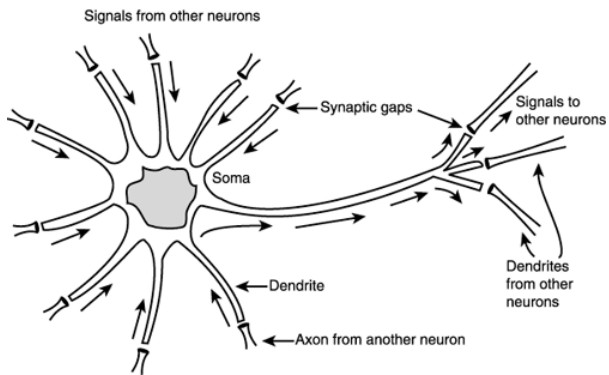
MACHINE LEARNING

인공신경망을 중심으로...


뉴런: 신경망의 기본 단위



인공 뉴런(Artificial Neuron)





실제 뉴런



뉴런의 수학적 모델

예시: 연어와 농어의 구별



예시: 연어와 농어의 구별







FIGURE 1.4. The two features of lightness and width for sea bass and salmon. The dark line could serve as a decision boundary of our classifier. Overall classification error on the data shown is lower than if we use only one feature as in Fig. 1.3, but there will still be some errors. From: Richard O. Duda, Peter E. Hart, and David G. Stork, *Pattern Classification*. Copyright © 2001 by John Wiley & Sons, Inc.




예시: 연어와 농어의 구별






Fully connected neural network

- Example
 - 1000x1000 image
 - 1M hidden units
- 10^{12} ($= 10^6 \times 10^6$) parameters!



역전파 알고리즘

- 역전파 알고리즘(Back-propagation algorithm)
 - D. E. Rumhart, G. E. Hinton, and R. J. Williams, Learning representation by back-propagation errors, Nature, 1986
- 역전파 알고리즘의 한계
 - 최적값 찾기에 실패
- MNIST 데이터 실험 (400회)
 - 2층 구조가 최고의 성능을 보임



Neural network
Back propagation,
Nature



1986

- 장점
 - 일반적인 문제에 적용할 수 있는 학습법
 - Biological 시스템과 관련이 깊음
- 문제점
 - Training 이 쉽지 않음
 - 현실적인 문제에 잘 동작하지 않음

Neural network
Back propagation,
Nature


Non-linear SVM



- 다양한 시도들


- Flat structure
 - SVM, Boosting, ...
- Biological 시스템과 거리가 생김
- 특정한 문제를 해결하는 특정한 방법 (SIFT, LBP, HOG, GMM-HMM)

Flat Processing Scheme



Neural network
Back propagation,
Nature

Deep belief network,
Science




- 비지도 학습을 이용한 pre-training
- Training 방법 향상
 - Dropout, RectLinear, Normalization, ...
- 컴퓨터 구조의 발달
 - GPU
 - Multi-core computer 시스템
- 빅데이터

Neural network
Back propagation,
Nature

Deep belief network,
Science

Speech



task	hours of training data	DNN-HMM	GMM-HMM with same data	GMM-HMM with more data
Switchboard (test set 1)	309	18.5	27.4	18.6 (2000 hrs)
Switchboard (test set 2)	309	16.1	23.6	17.1 (2000 hrs)
English Broadcast News	50	17.5	18.8	
Bing Voice Search (Sentence error rates)	24	30.4	36.2	
Google Voice Input	5,870	12.3		16.0 (>>5,870hrs)
Youtube	1,400	47.6	52.3	


Neural network
Back propagation,
Nature

Deep belief network,
Science

Speech

Object recognition

IMAGENET



Submission	Method	Error rate
Supervision	Deep CNN	0.16422
ISI	FV: SIFT, LBP, GIST, CSIFT	0.26172
XRCE/INRIA	FV: SIFT and color 1M-dim features	0.27058
OXFORD_VGG	FV: SIFT and color 270K-dim features	0.27302

ImageNet Large Scale Visual Recognition Competition (ILSVRC)


Steel drum



Output:
 Scale
 T-shirt
Steel drum
 Drumstick
 Mud turtle



Output:
 Scale
 T-shirt
 Giant panda
 Drumstick
 Mud turtle




Neural network
Back propagation,
Nature

Deep belief network,
Science

Speech

Object recognition

IMAGENET



• IMAGENET 2013: 영상 인식

RANK	Name	Error rate	Description
1	NYU	0.11197	Deep Learning
2	NUS	0.12535	Deep Learning
3	OXFORD	0.13555	Deep Learning


Neural network
Back propagation,
Nature

Deep belief network,
Science

Speech

Object recognition

IMAGENET



- IMAGENET 2013: 영상 인식

RANK	Name	Error rate	Description
1	Google	0.06656	Deep Learning
2	Oxford	0.07325	Deep Learning
3	MSRA	0.08062	Deep Learning


Neural network
Back propagation,
Nature

Deep belief network,
Science

Speech

Object recognition

IMAGENET



- Google 과 Baidu는 2013년에 Deep Learning기반 영상 검색 엔진을 발표함
- NAVER N-Drive에 사진 검색 서비스 실시



Neural network
Back propagation,
Nature

Deep belief network,
Science

Speech

Object recognition

The game of GO



The AI race is on

