areaperta
parlando di scienza al cnr di pisa

con il patrocinio di

Consiglio Nazionale delle Ricerche

REGIONE TOSCANA

PROVINCIA DI PISA

Comune di Pisa

AlphaGo e Reti Neurali

Fabrizio Falchi
ISTI-CNR

ISTI-CNR
Fabrizio Falchi

Algebra e Reti Neurali

Dalla Teoria alla Pratica

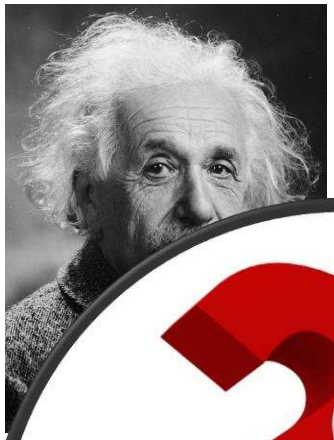

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 8\pi G T_{\mu\nu}$$

A cosa serve?

DALLA PRATICA ALLA TEORIA

AlphaGo

Google

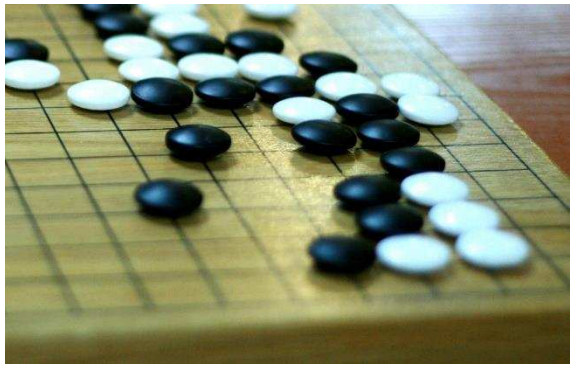

Perché funziona?

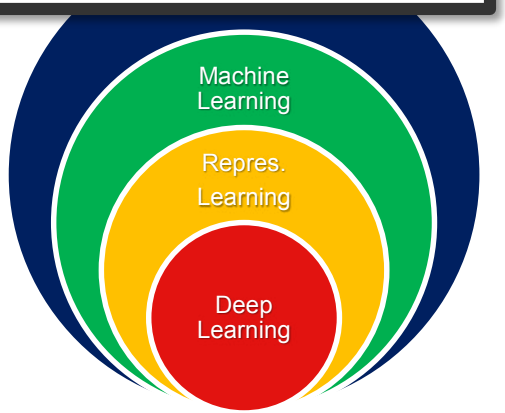

Nucleo

Estremità dell'assone

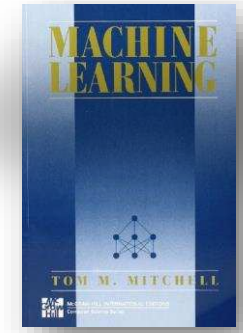

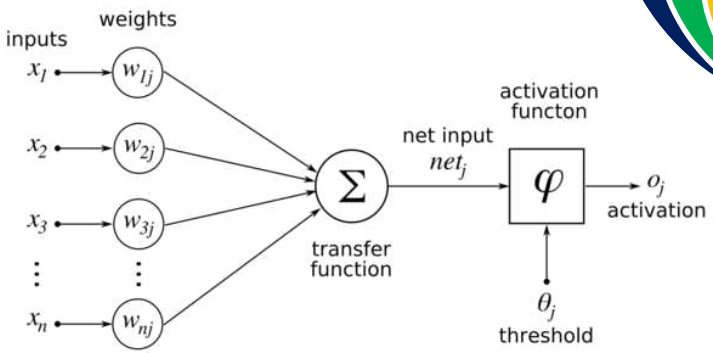

PERCHÉ FUNZIONA?

Quora

How far along are we in the understanding of why deep learning works?

Yoshua Bengio, My lab has been one of the three that started the deep learning approach, bac...

10.9k Views • Upvoted by Tao Xu, Built ML systems at Airbnb, Quora, Facebook and Microsoft., Nikhil Garg, I lead a team of Quora engineers working on ML/NLP problems, Alberto Bietti, MS in machine learning, Ecole Normale Supérieure

Yoshua is a Most Viewed Writer in Machine Learning.

Originally Answered: Yoshua Bengio: How far along are we in understanding why deep learning works?

Contrary to what some people think, I believe that we have already a good basic understanding of fundamentals about why deep learning works, e.g., * We understand that distributed representations,... (more)

SANDY
CARNEGIE ROAD
'Mingol Camp'
Mount Hoskin
Timperley Range
BM478
173
Kingston Pass
Mount Alexandra
Wellington Range
BM508
WONGAWOL
RD
335
Mount Wellesley
'Lorna Glen'
Calvert
'Windidda'
Tooloo Bluff
Windidda
Dingle Claypan
Waguin Hill
Bonython Bluff
Lake Jeffries
Carclew Range
Sachse Bluff
Banjo
Creek
'Yelma'
Little Greta Mine
BM494
BM488
Mount Eureka
BM527
Elder
BM499
84
BM498
WONGAWOL ROAD
BM509
Ironstone Well Mine
Red Lake Mine
Empire Mine
Lake Way
Wiluna
WILUNA
Desert Gold
Emu Farm
Lake Violet'
BM556
'Nimary Camp'
'Jundee Camp'
Mount Cleaver
Bare Granite Hill
Prominent Hill
Red Bluff
CANNING
STOCK
210
ROUTE
Bingah Pool
'Forbes Outcamp'
Private Track
Mooloonjoo Claypan
E
S
W

PENSIERI LENTI E VELOCI

Daniel Kahneman

- Psicologo
- Premio Nobel 2002 assieme a Vernon L. Smith per

avere integrato risultati della ricerca psicologica nella scienza economica, specialmente in merito al ***giudizio umano*** *e alla* ***teoria delle decisioni*** *in* ***condizioni d'incertezza***

- ***Thinking, Fast and Slow (2011)***
  riassume molte delle sue ricerche
  best seller internazionale

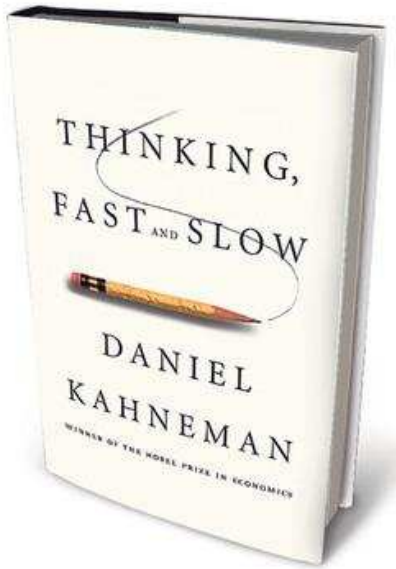


Test

PENSIERO VELOCE

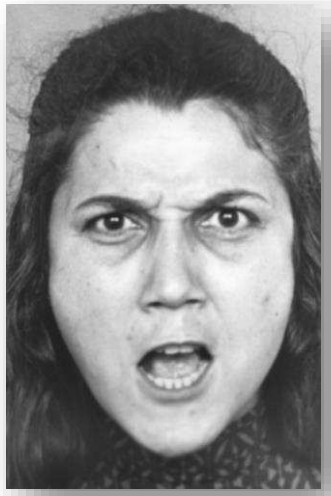

Avete probabilmente, intuitivamente, visto

- Donna, cappelli neri, arrabbiata

Quello che avete visto si estende nel **futuro**:

- sensazione che stesse per dire parole
  molto scortesi, magari a voce alta e con tono stridulo.

Non intendevate

- valutare il suo stato d'animo o
- prevedere che azione potesse compiere.

$$17 \times 24 = ?$$

17x24=?

- è una moltiplicazione
- sapete che potete risolverla
- riconoscete che 12,609 e 123 sono poco plausibili
- escludere 586 è difficile

Una soluzione precisa non vi è venuta in mente.

Vi siete chiesti se impegarvi o no nel calcolo.

17x24=?

Procedete lungo una **sequenza di stadi**.

recuperate dalla memoria
il **programma cognitivo** imparato a scuola

eseguire il calcolo è uno **sforzo**.

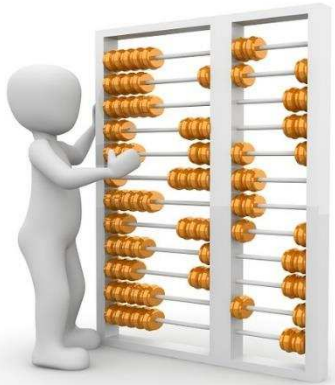

Sentite l'onere di conservare tanto materiale in **memoria**:

- il punto in cui eravate e
- quello in cui stavate andando
- il risultato intermedio.

THINKING, FAST AND SLOW

- Il processo è consistito in un lavoro mentale,
  un lavoro riflessivo, impegnativo e ordinato,
  il prototipo del pensiero lento.
- Era coinvolto anche il **corpo**.
  I muscoli si sono tesi, la pressione del sangue e la frequenza cardiaca sono aumentate.
  Le puppile si sono dilatate.

THINIKING, FAST AND SLOW

Kahneman descrive due modalità differenti di pensiero:

- **Sistema 1**:
  - veloce, intuitivo, stereotipico, involontario
- **Sistema 2**:
  - lento, conscio, infrequente, logico, calcolatore

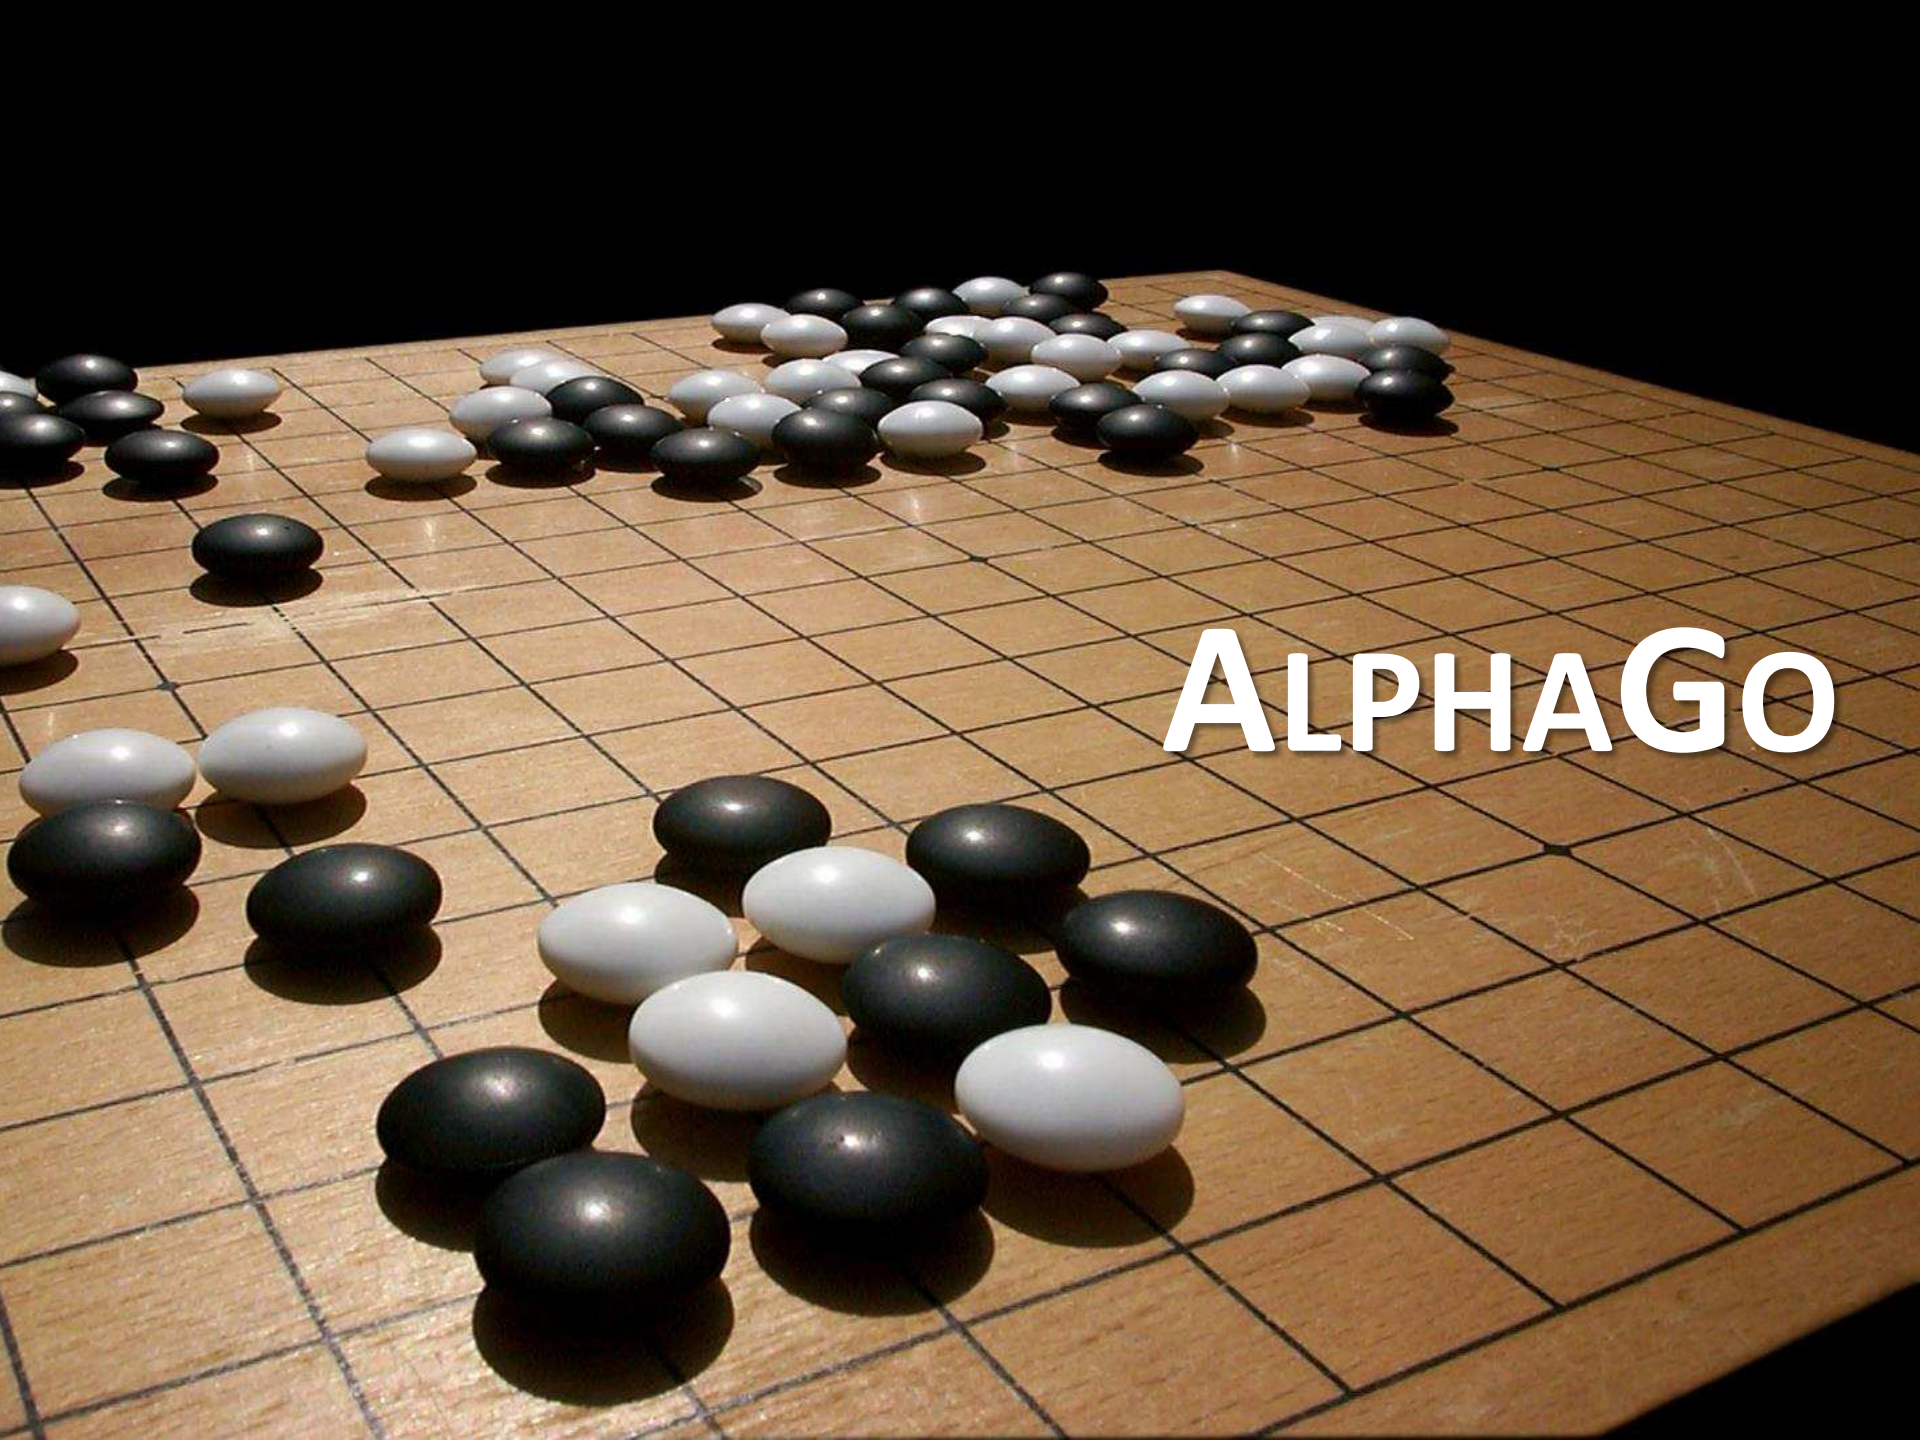
AlphaGo

ALBERO

Photo by: Stolz Gary M, U.S. Fish and Wildlife Service

Albero da Esplorare

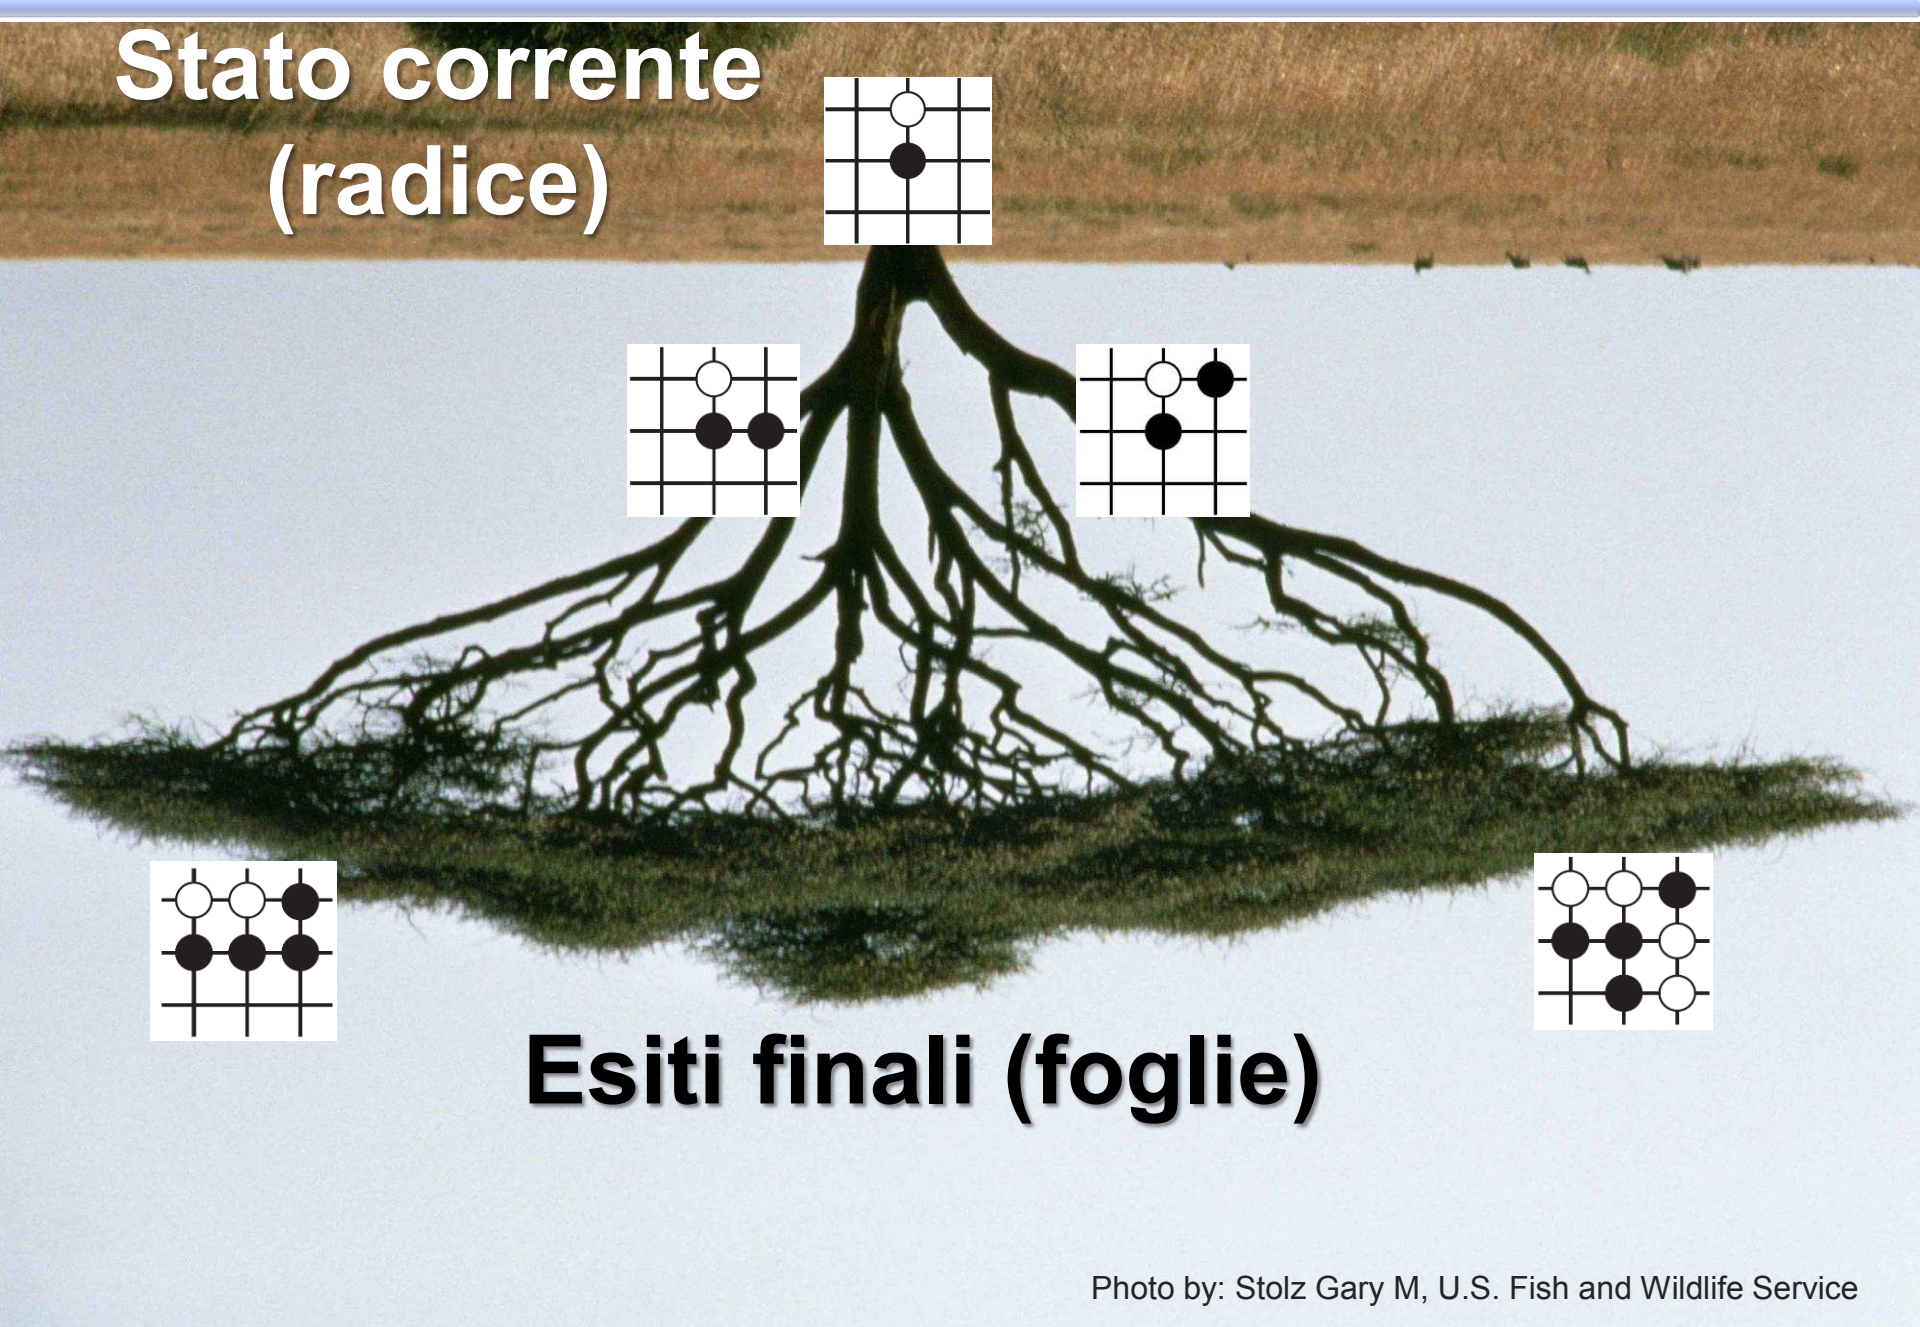

Photo by: Stolz Gary M, U.S. Fish and Wildlife Service

SIMULAZIONE: QUALE PARTE ESPLORIAMO?

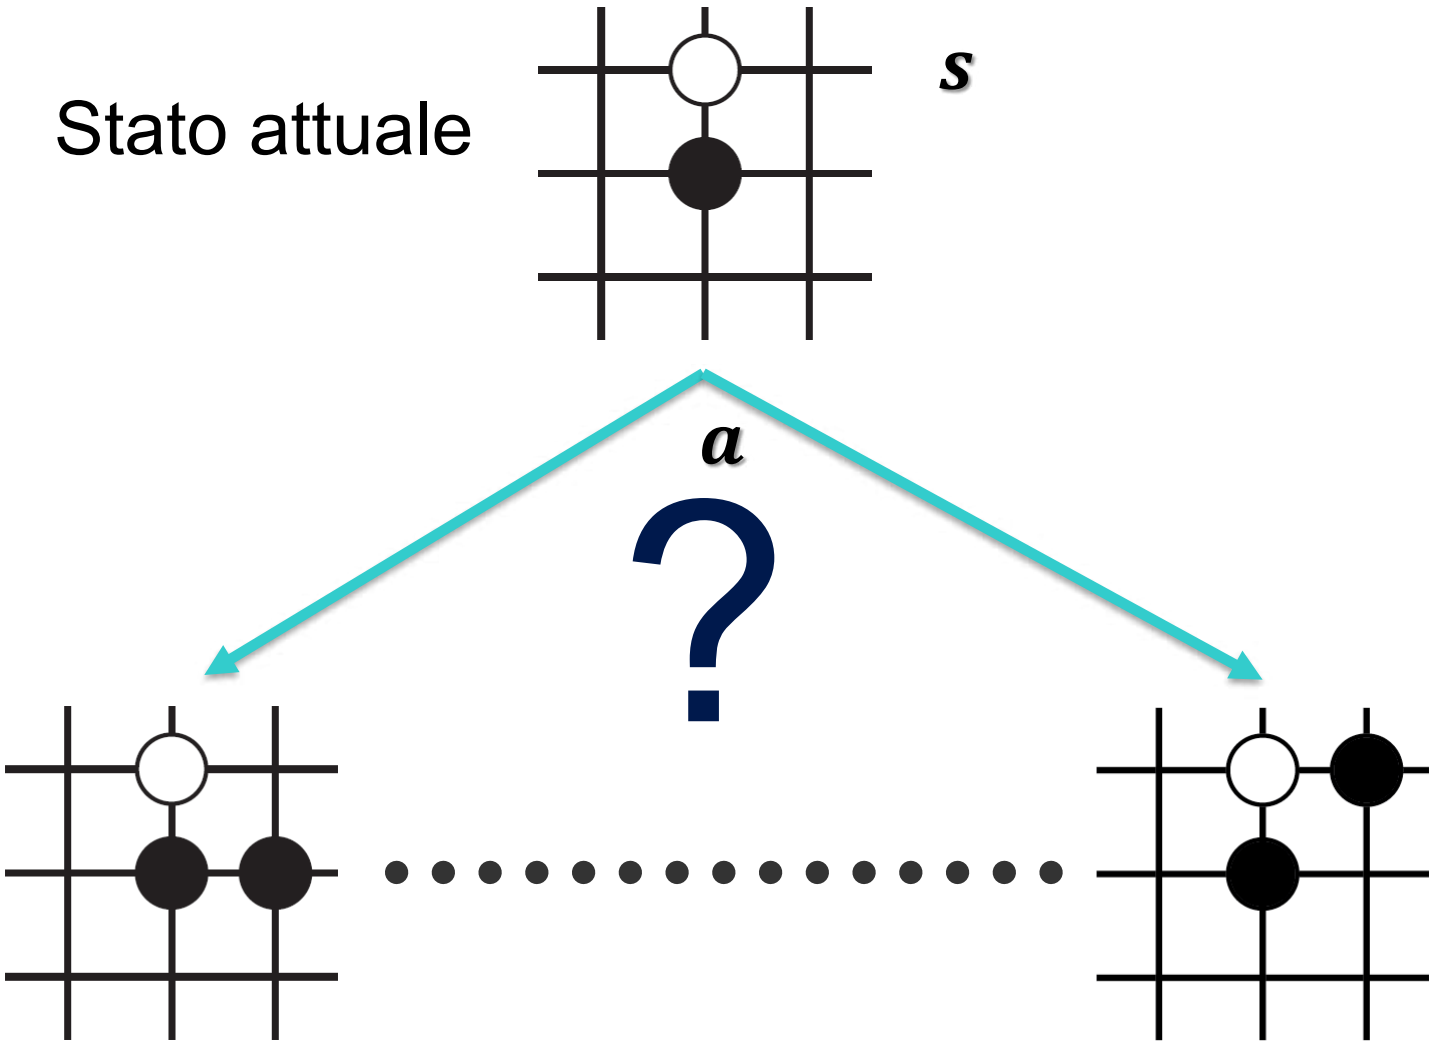

SIMULAZIONE: QUALE PARTE ESPLORIAMO?

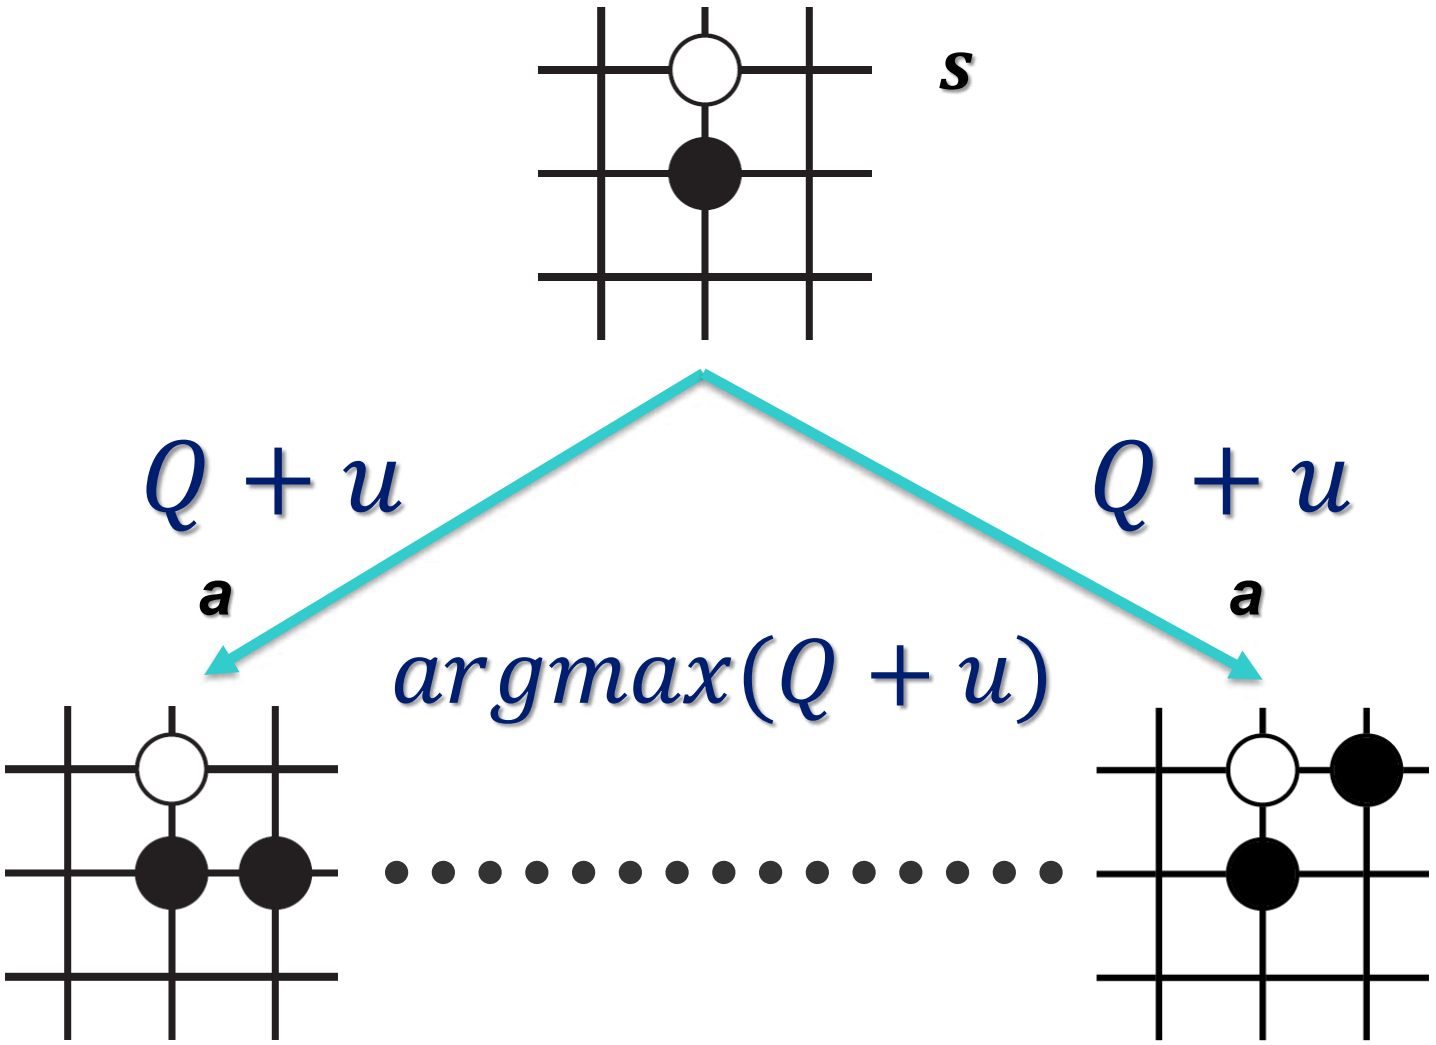

$Q + u$: BONUS ESPLORATIVO (u)

- u è la componente esplorativa

$$u(s,a) \propto \frac{P(s,a)}{1 + N(s,a)}$$

- $N(s,a)$ – n. passaggi sul ramo considerando tutte le simulazioni

- $P(s,a) = p(a|s)$
  - non varia da simulazione a simulazione
  - **Deep Convolutional Neural Network** (Policy Network)

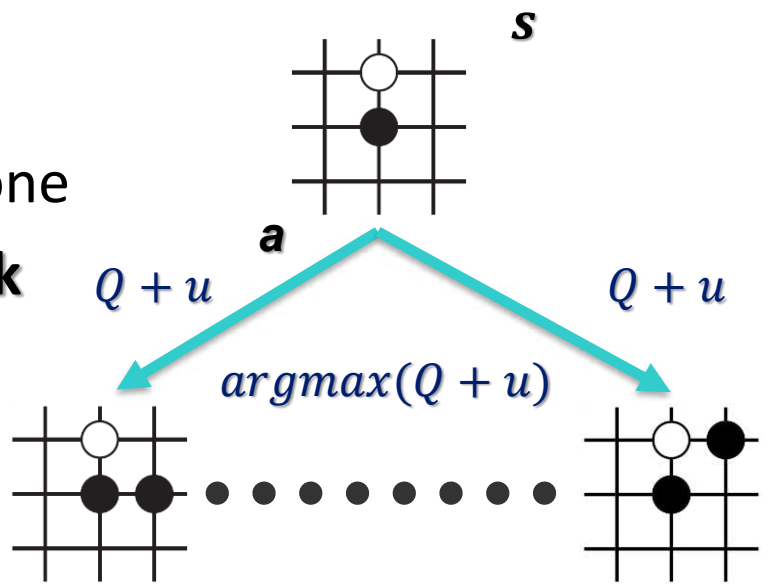

PROBABILITÀ A PRIORI (ESEMPIO) - *p(a|s)*

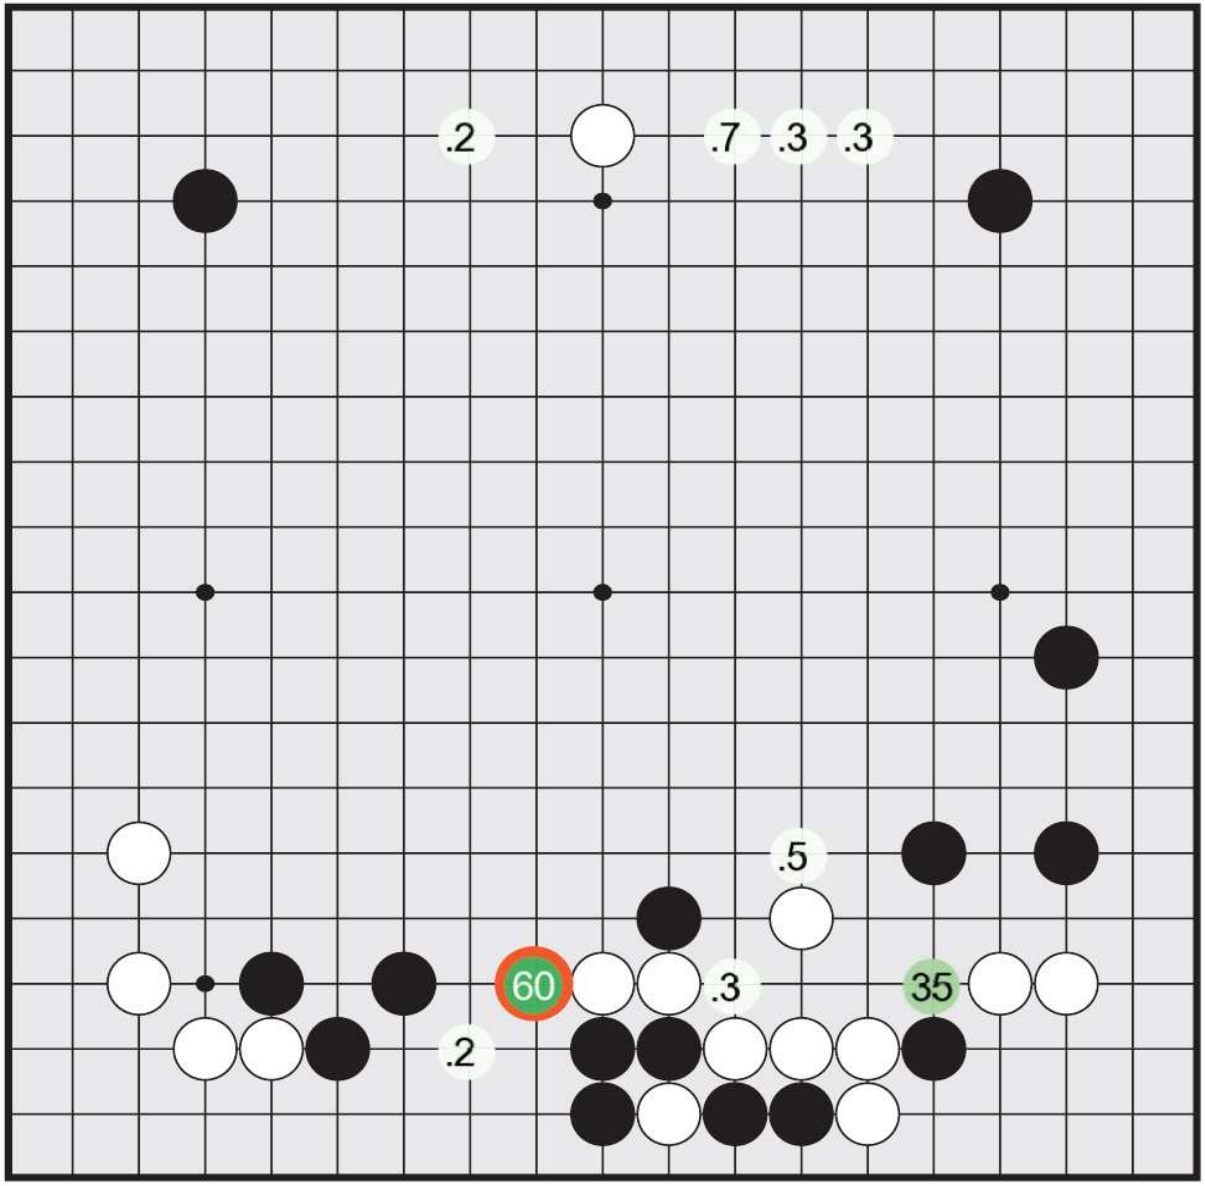

ALBERO

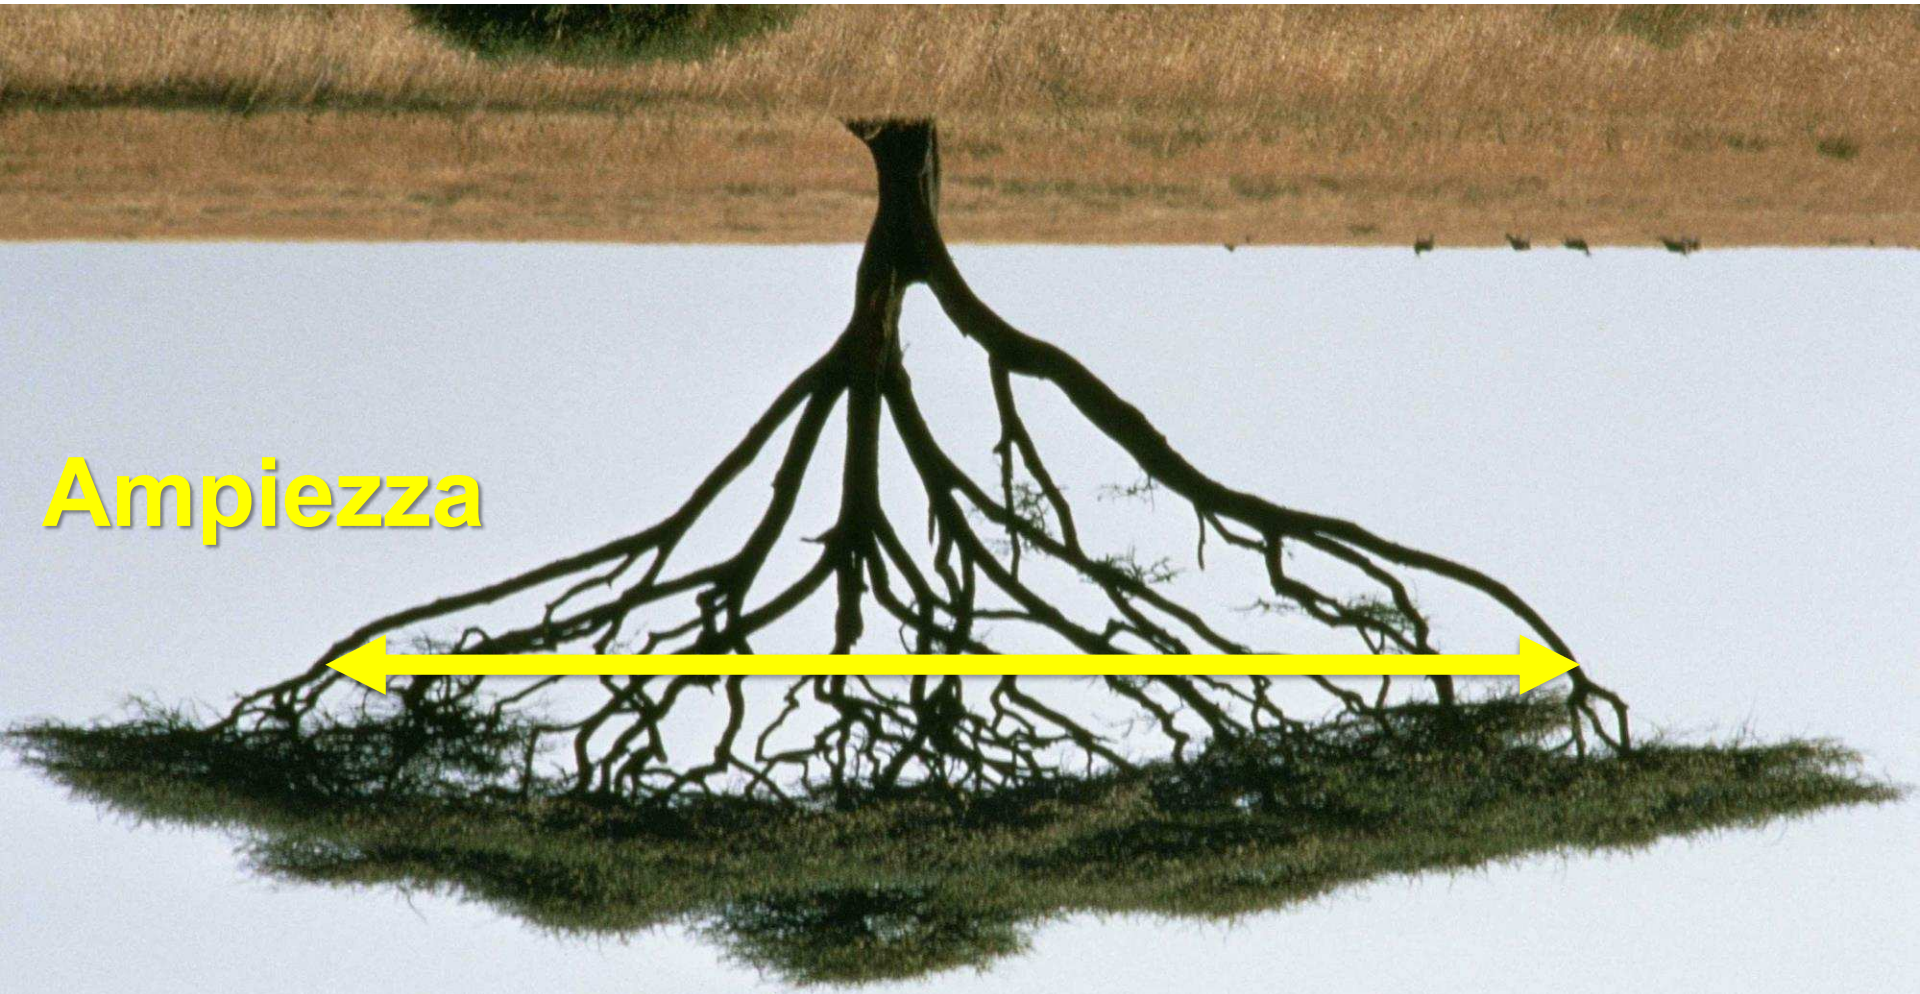

Photo by: Stolz Gary M, U.S. Fish and Wildlife Service

$Q + u$: Componente di Valutazione (Q)

- $\boldsymbol{Q(s,a)}$ è la componente di valutazione

$$Q(s,a) = \frac{1}{N(s,a)} \sum_{i=1}^{n} 1(s,a,i) V(s_L^i)$$

- media le valutazioni $\boldsymbol{V(s^i_L)}$
  - di tutti gli stati **foglia** $\boldsymbol{s^i_L}$
  - raggiunti nelle varie simulazioni $\boldsymbol{i}$

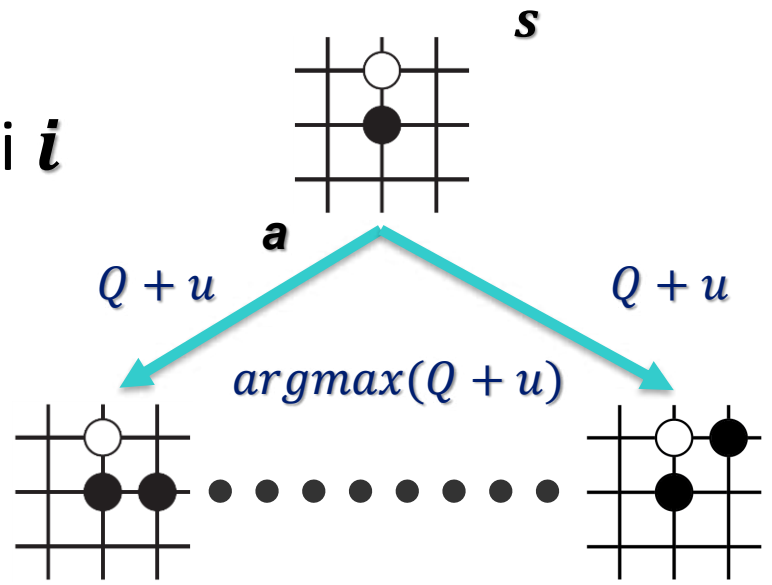

VALUTAZIONE DEGLI STATI FOGLIA

- Uno stato foglia ha valutazione

$$V(s_L) = (1 - \lambda)v_\theta(s_L) + \lambda z_L$$

- $\boldsymbol{v_\theta(s_L)}$
  - **Deep Convolutional Neural Network**

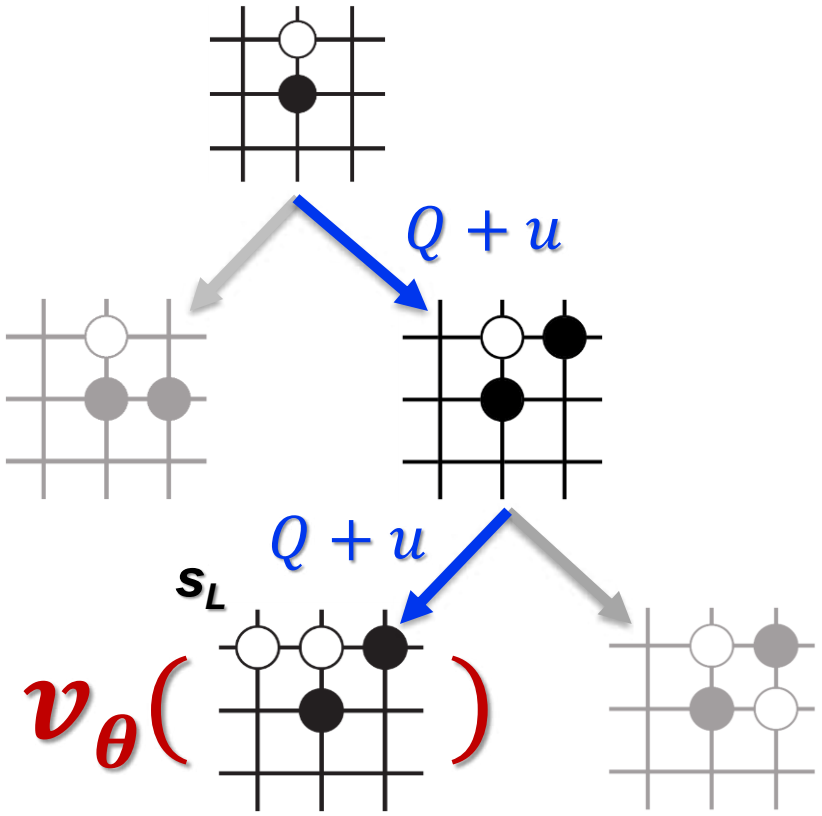

VALUTAZIONE DEGLI STATI CON $v_\theta(s_L)$

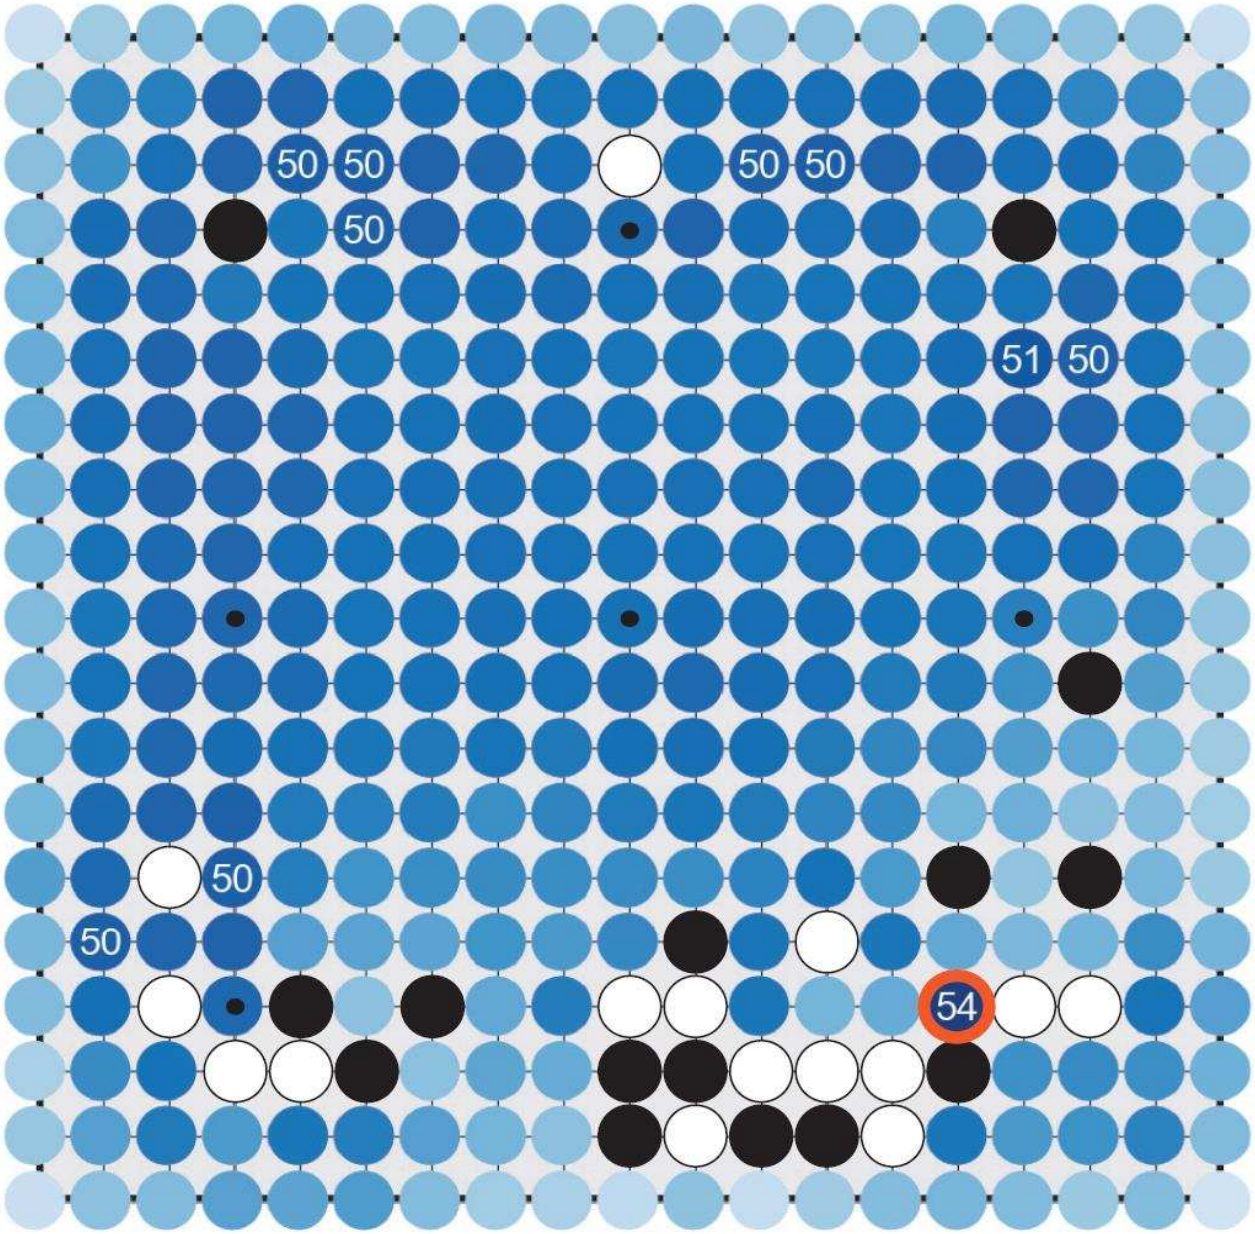

VALUTAZIONE DEGLI STATI FOGLIA

- Uno stato foglia ha valutazione

$$V(s_L) = (1-\lambda)v_\theta(s_L) + \lambda z_L$$

- $v_\theta(s_L)$
  - **Neural Network**

- z_L risultato finale
  - ***Fast rollout***
    - linear softmax

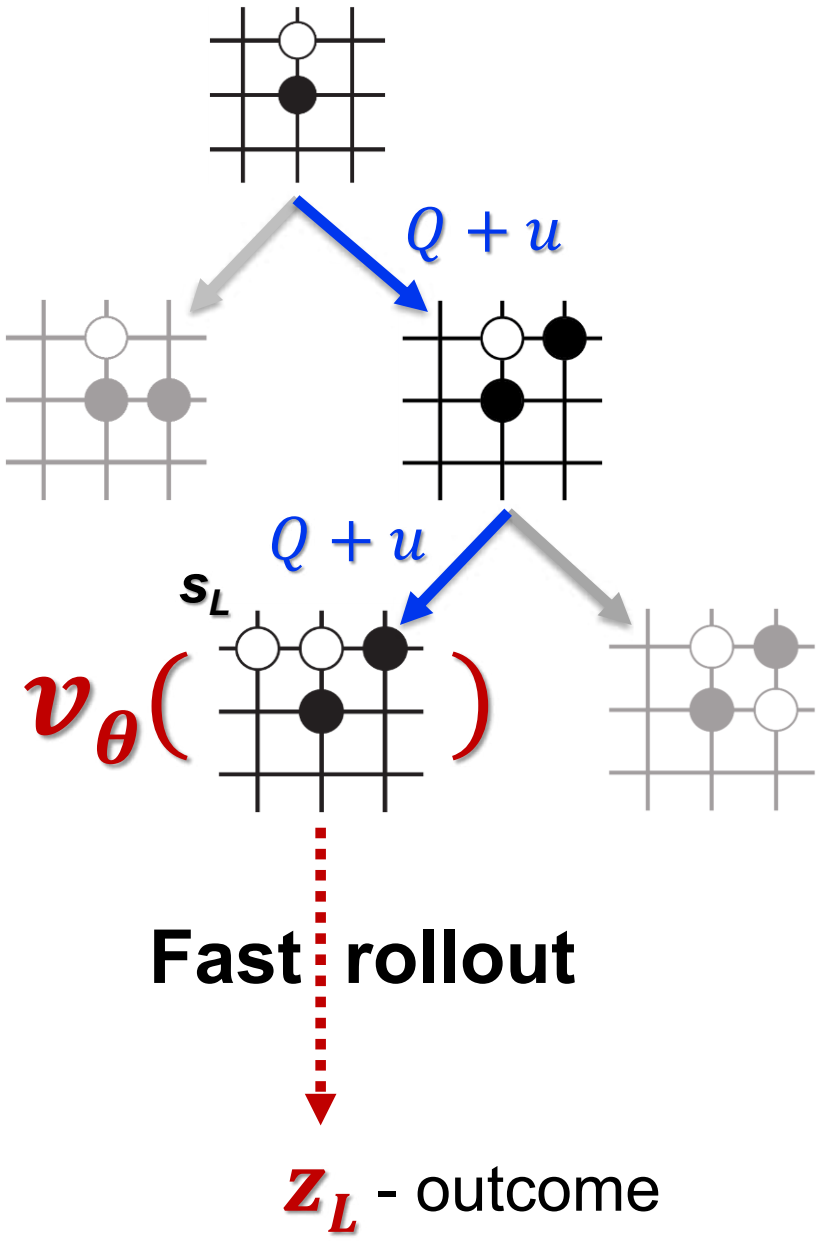

AGGIORNAMENTI

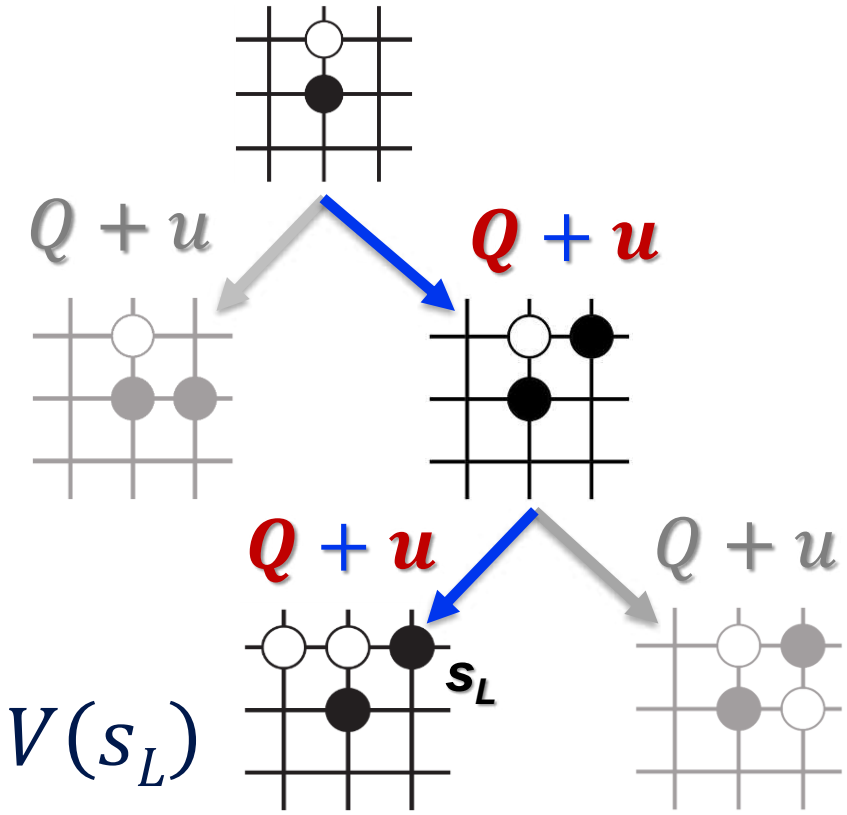

$$Q(s,a) = \frac{1}{N(s,a)} \sum_{i=1}^{n} 1(s,a,i) V(s_L^i)$$

$$u(s,a) \propto \frac{P(s,a)}{1 + N(s,a)}$$

$Q(s,a)$ USANDO SOLAMENTE $v_\theta(s_L)$

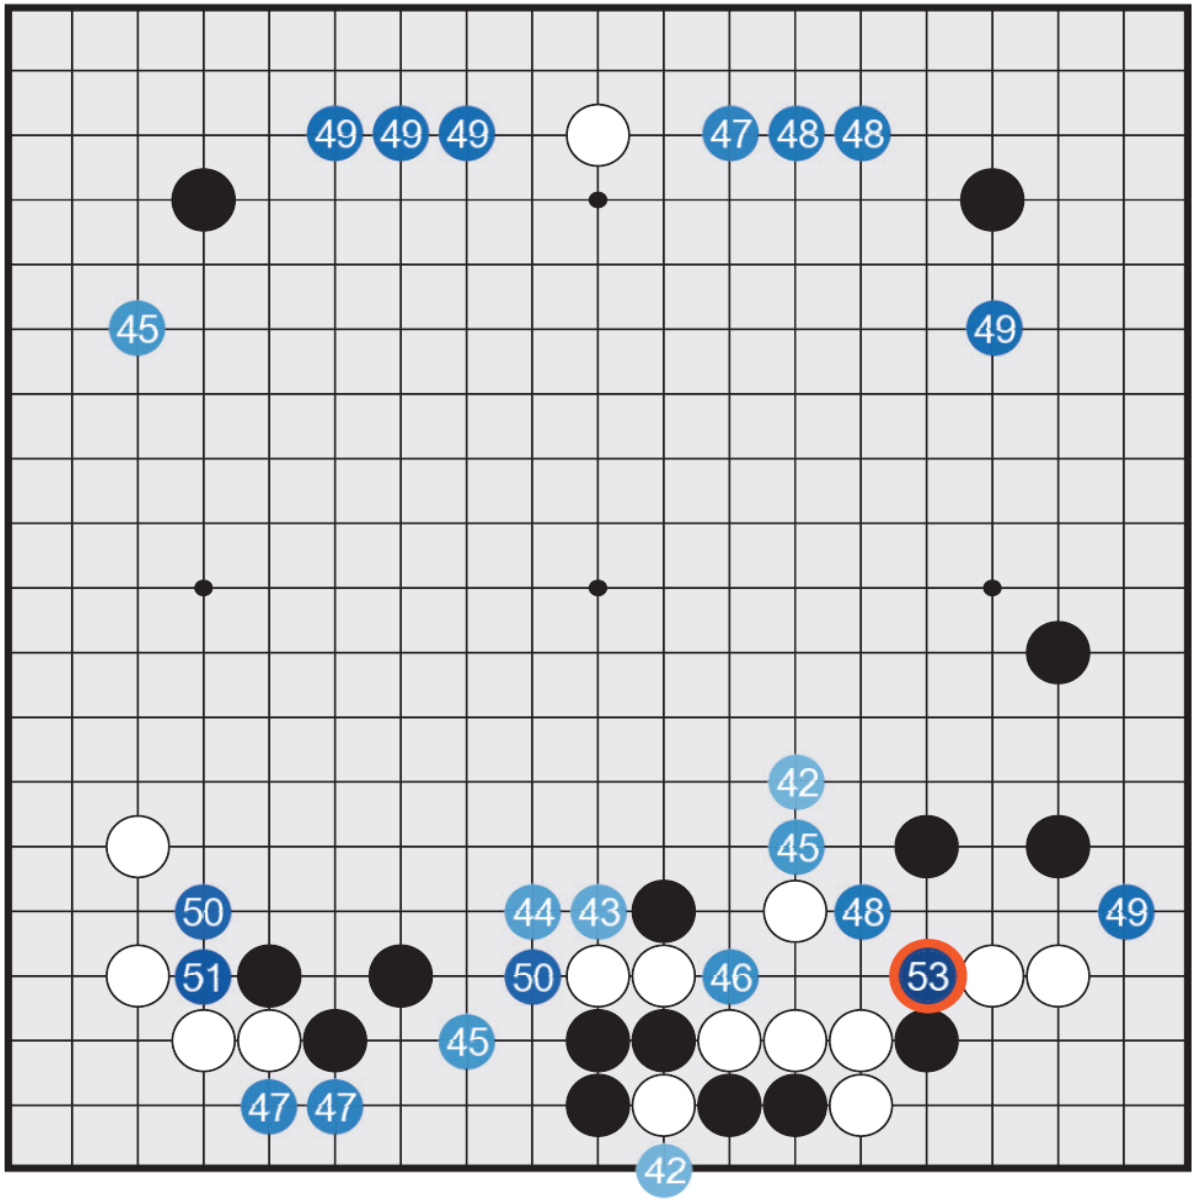

$Q(s, a)$ USANDO SOLAMENTE Z_L

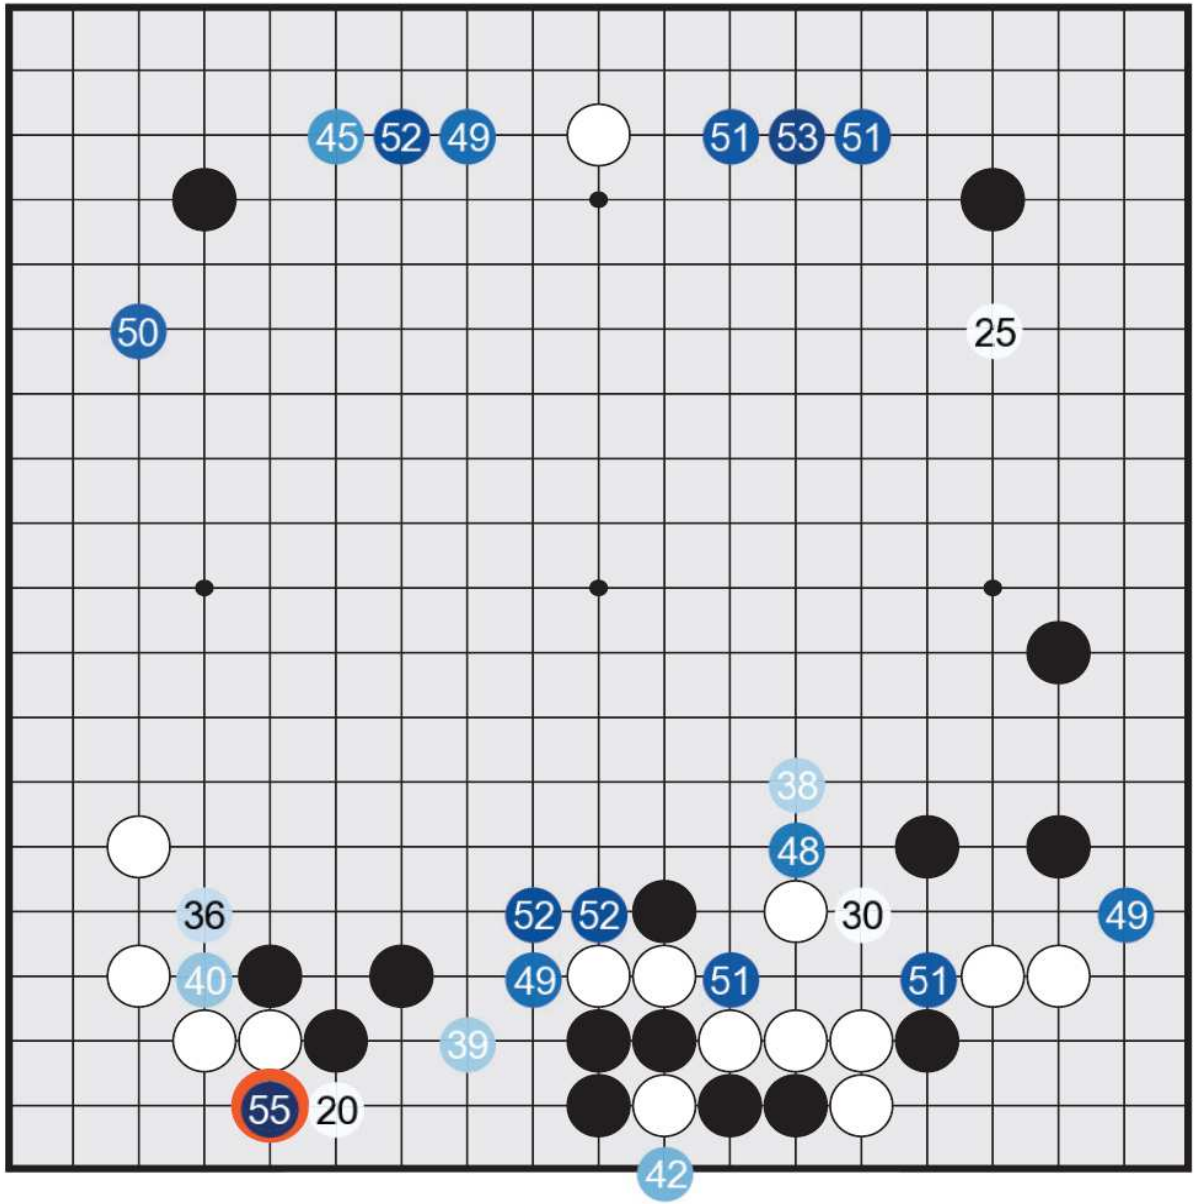

VALUE NETWORK VS ROLLOUTS

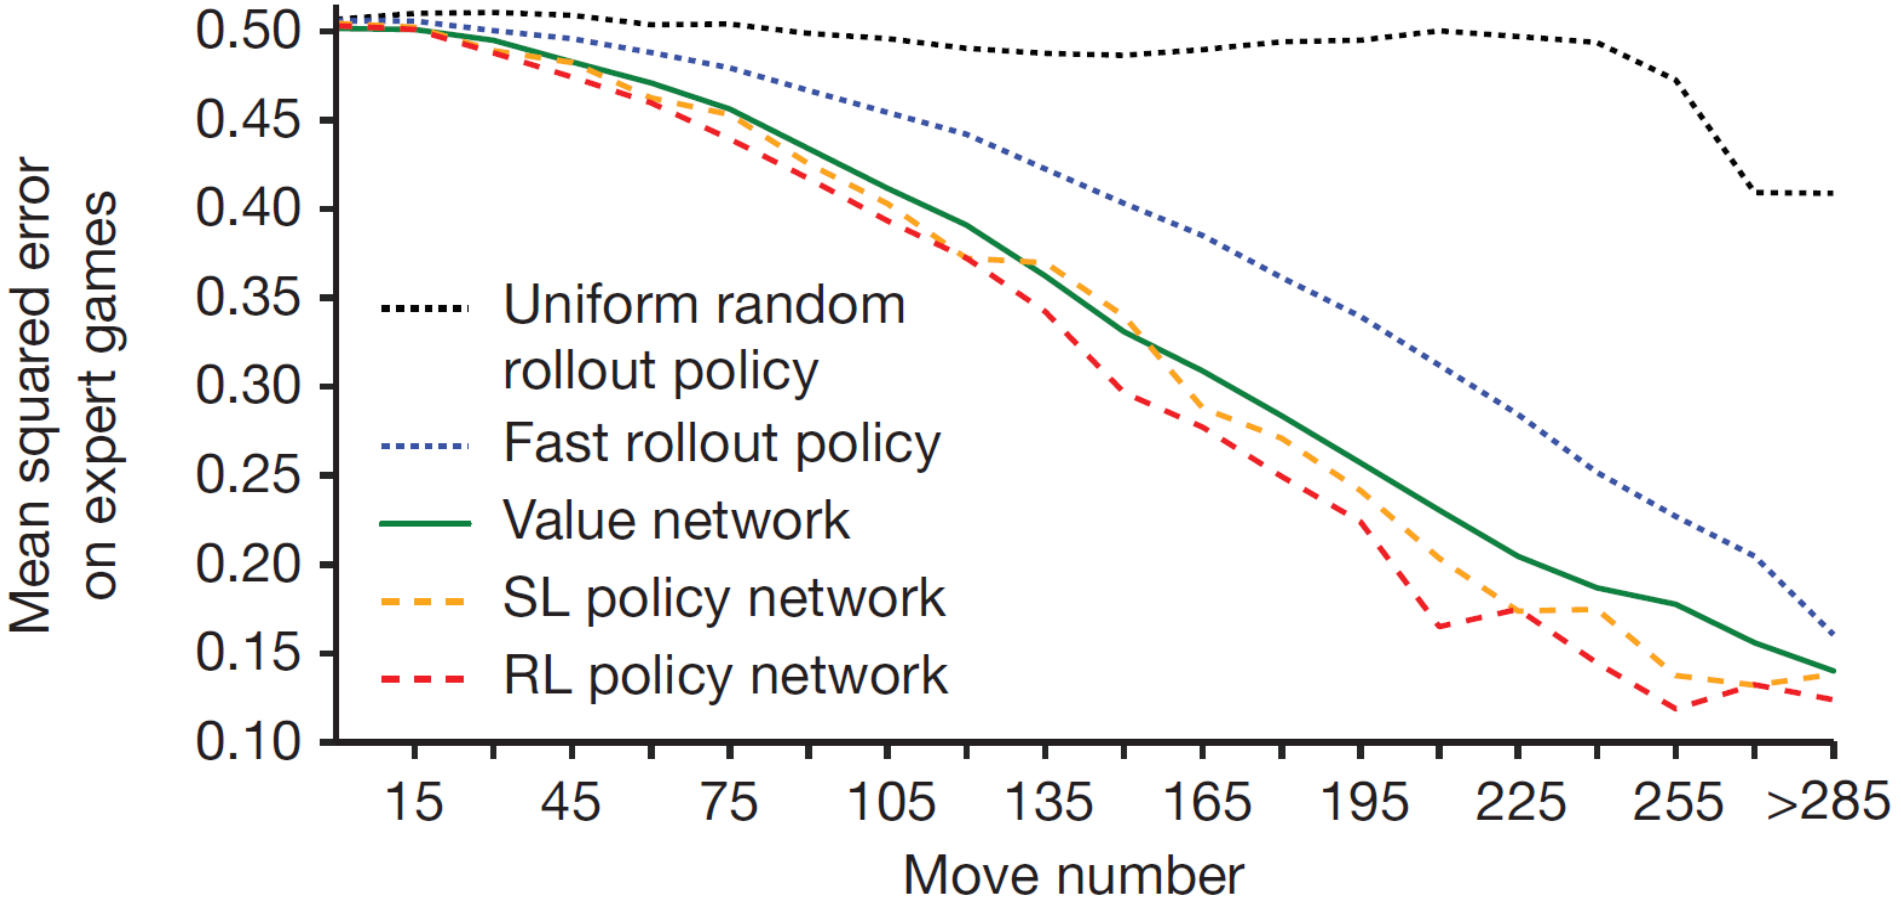

ALBERO

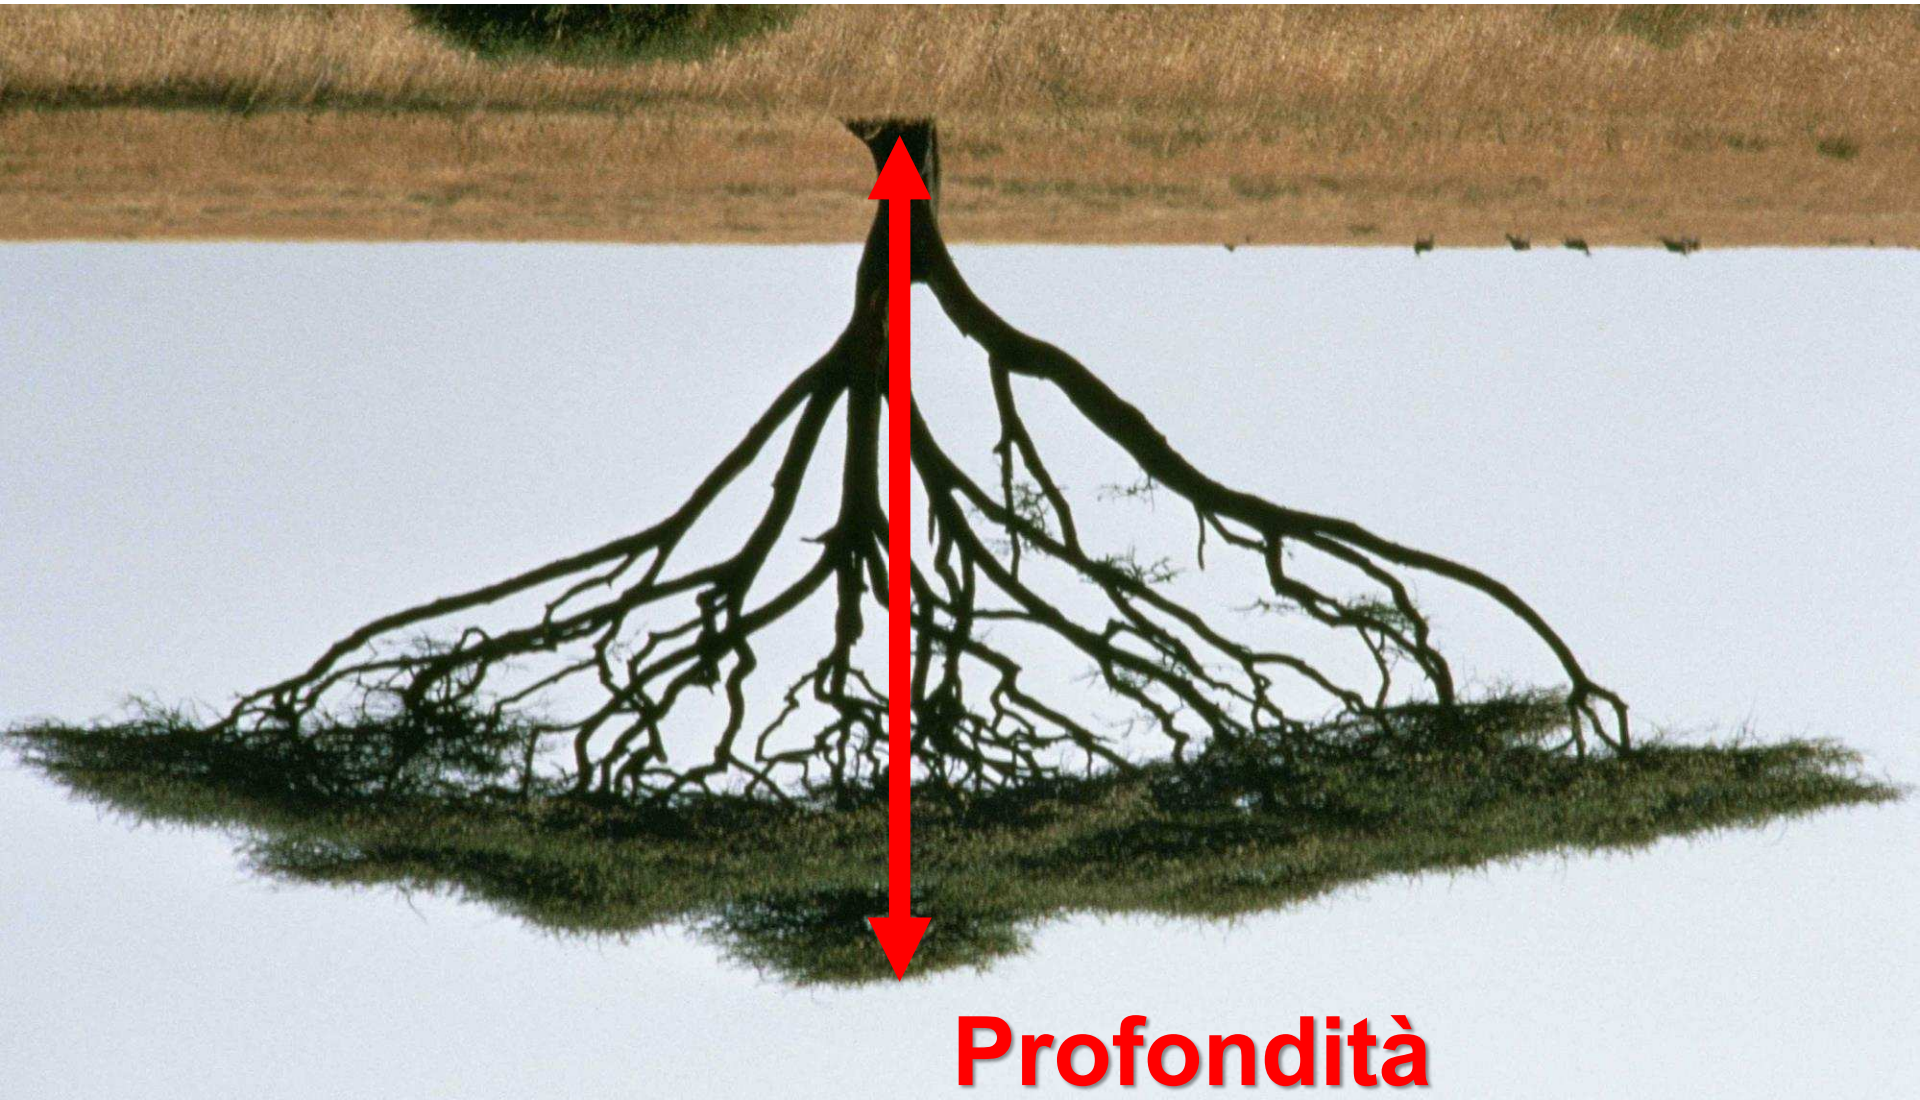

Photo by: Stolz Gary M, U.S. Fish and Wildlife Service

Percentuale di Simulazioni $N(s,a)/n$

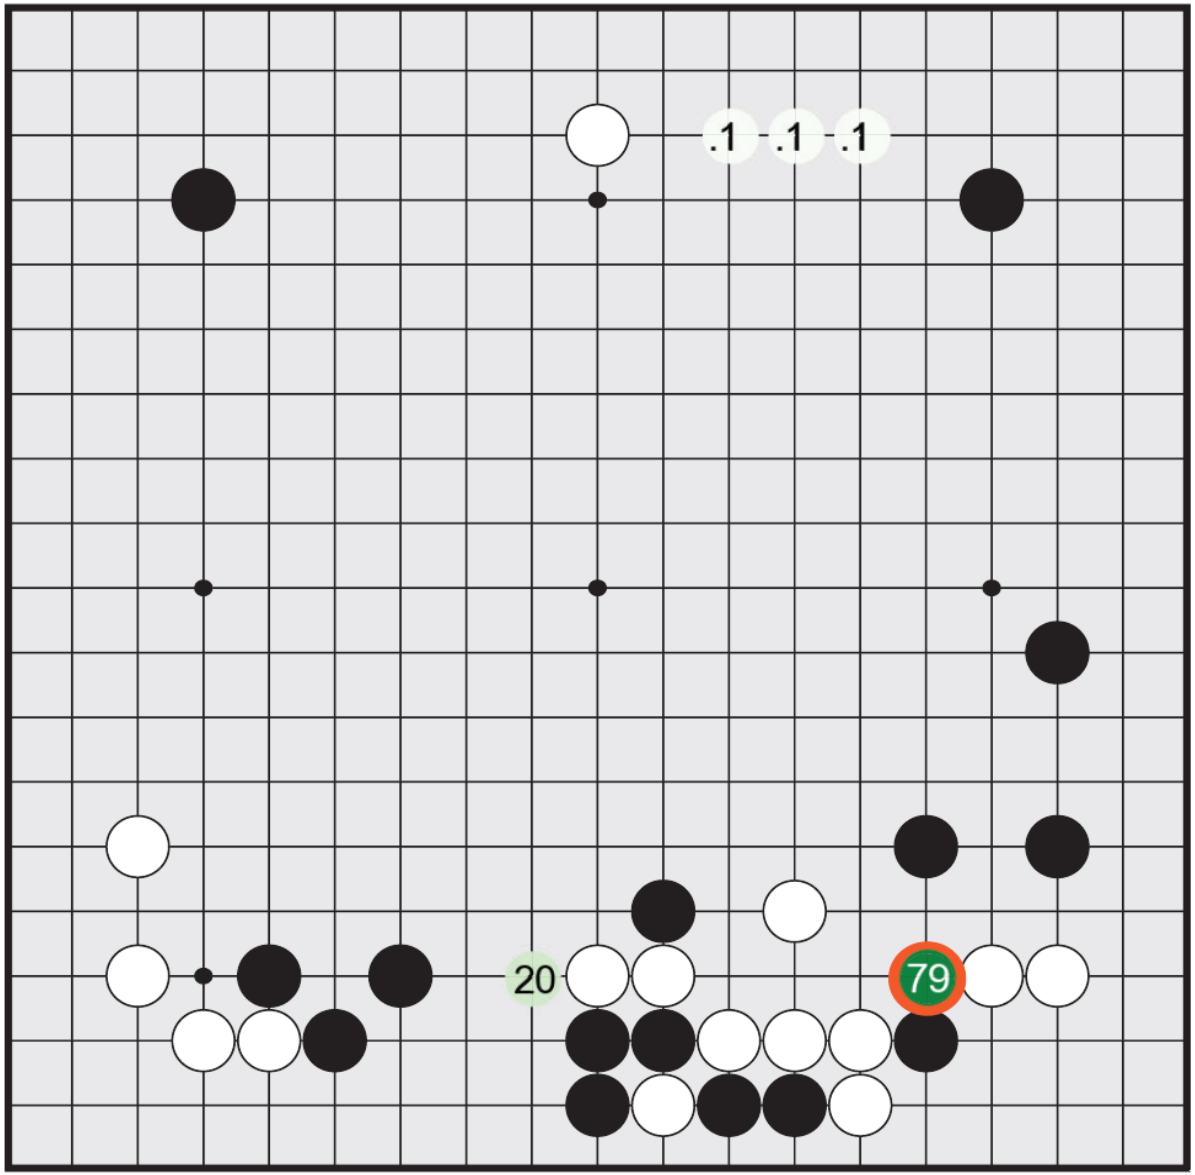

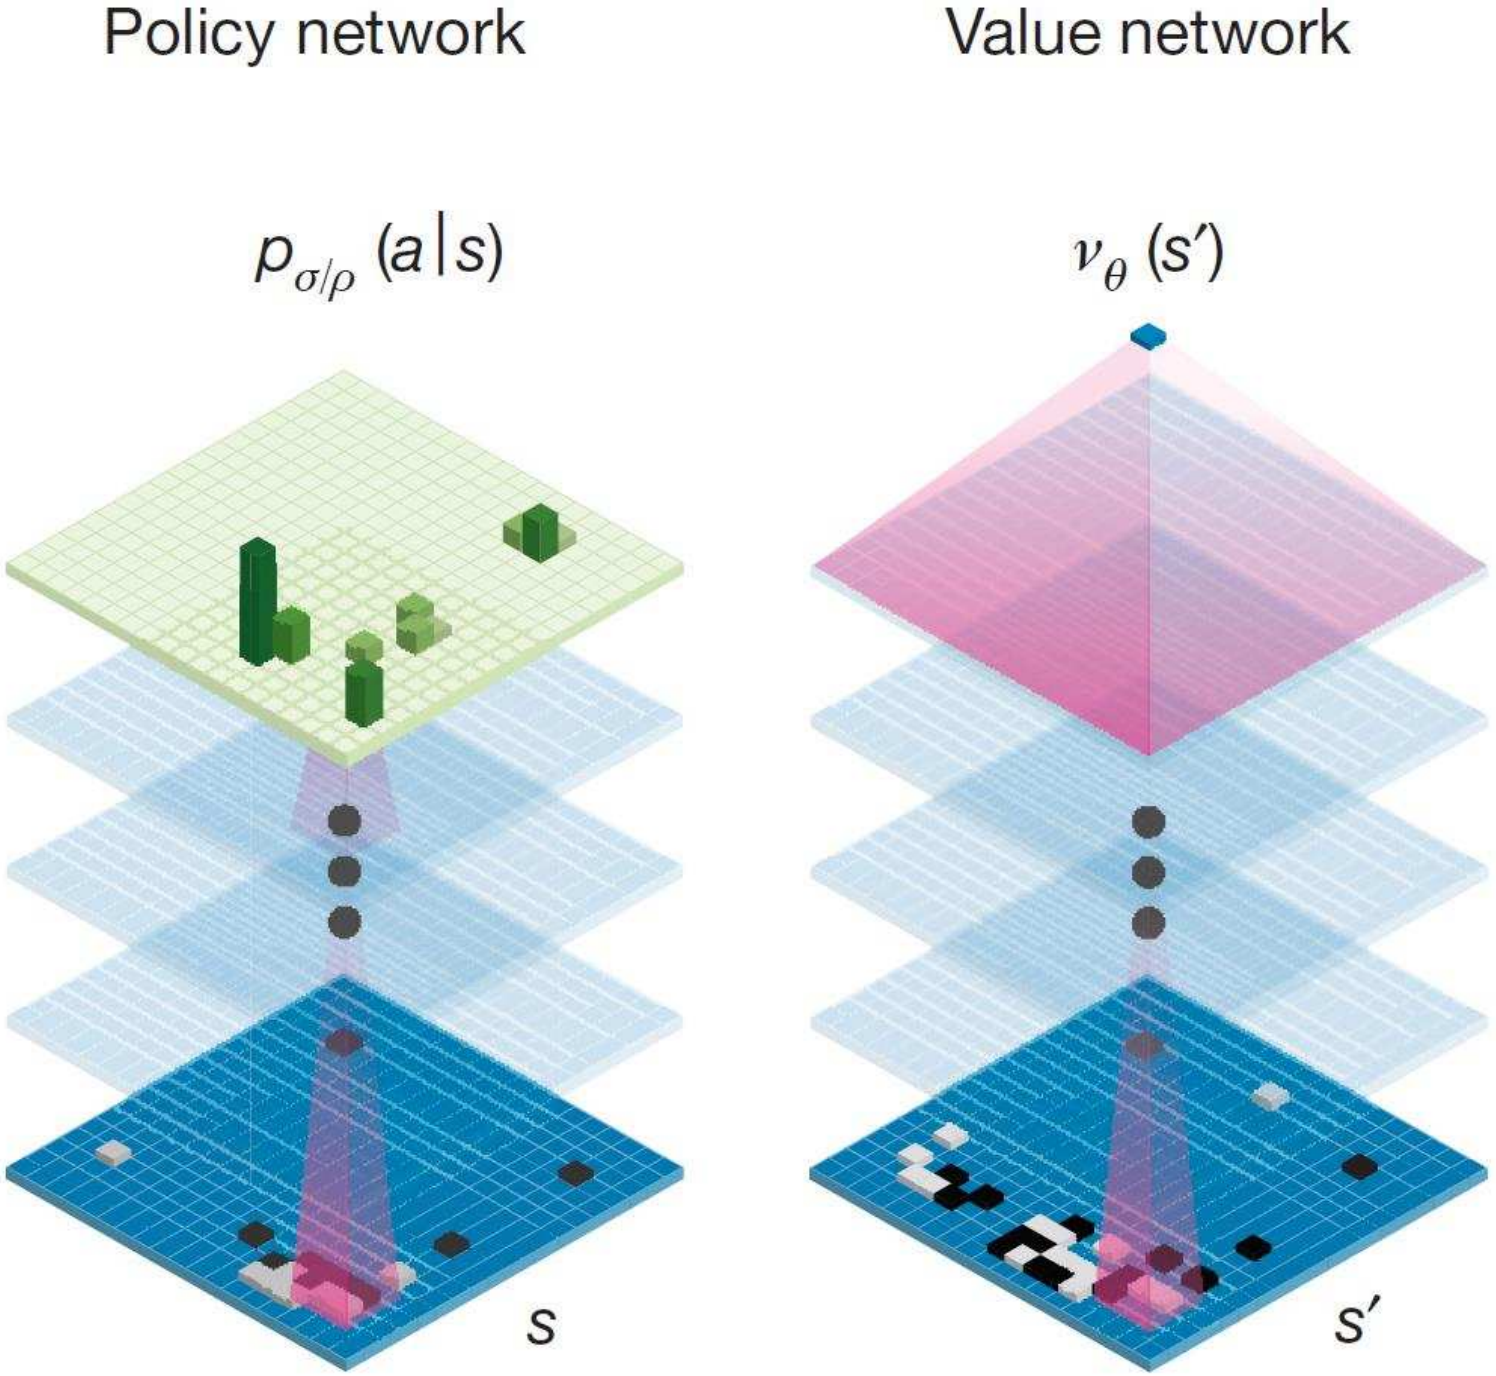
Policy network
Value network
$p_{\sigma/\rho}(a|s)$
$v_\theta(s')$
s
s'

DEMIS HASSABIS

Short Bio

Before embarking on my research career, I was previously a well-known UK videogames designer and AI programmer. I started my professional games career at the age of 16 working at the legendary Bullfrog Productions, and then after getting my degree at Cambridge, I joined the newly founded Lionhead Studios. In 1998 I founded and ran my own 60-person development studio, called Elixir Studios, producing games for Microsoft and Vivendi Universal, and have contributed to many best-selling and award-winning games over a successful 10-year career including: Syndicate (1993), Theme Park (1994), Black & White (2001), Republic (2003), and Evil Genius (2004).

I also play many games to an accomplished level including chess, shogi and poker and won the World Games Championships at the Mind Sports Olympiad a record 5-times before retiring from competitive play in 2003. You can find further information about my past from **here**.

Playing at the World Series of Poker

(hence the silly glasses)

Academic History

2005–2009

- PhD in Cognitive Neuroscience, **Wellcome Trust Centre for Neuroimaging** (WTCN), UCL, UK
- Thesis: "**The neural processes underlying episodic memory**" (supervisor: Prof **Eleanor Maguire**) (viva committee: Prof **John O'Keefe**, Prof **Lynn Nadel**)

1994–1997

- BA (First) in Computer Science, Queens' College, Cambridge University

Google DeepMind

Shane Legg

Sn. Staff Research Scientist
Google DeepMind
February 2014 – Present (1 year 4 months) | London, United Kingdom

Co-founder & Chief Science Officer
DeepMind Technologies
February 2011 – January 2014 (3 years) | London, United Kingdom

Research Associate
Gatsby Computational Neuroscience Unit, UCL
February 2009 – December 2010 (1 year 11 months)

Post Doc in Finance
University of Lugano
June 2007 – October 2008 (1 year 5 months)

Portfolio choice theory and behavioural finance models

PhD Student
IDSIA
April 2003 – May 2007 (4 years 2 months)

Senior Software Engineer
Adaptive Intelligence Inc.
January 2002 – October 2002 (10 months) | Wellington, New Zealand

Senior Software Engineer
Webmind
September 1999 – May 2001 (1 year 9 months)

Analyst Programmer
Freightways Information Services

Deep Learning
Convolutional Neural Network

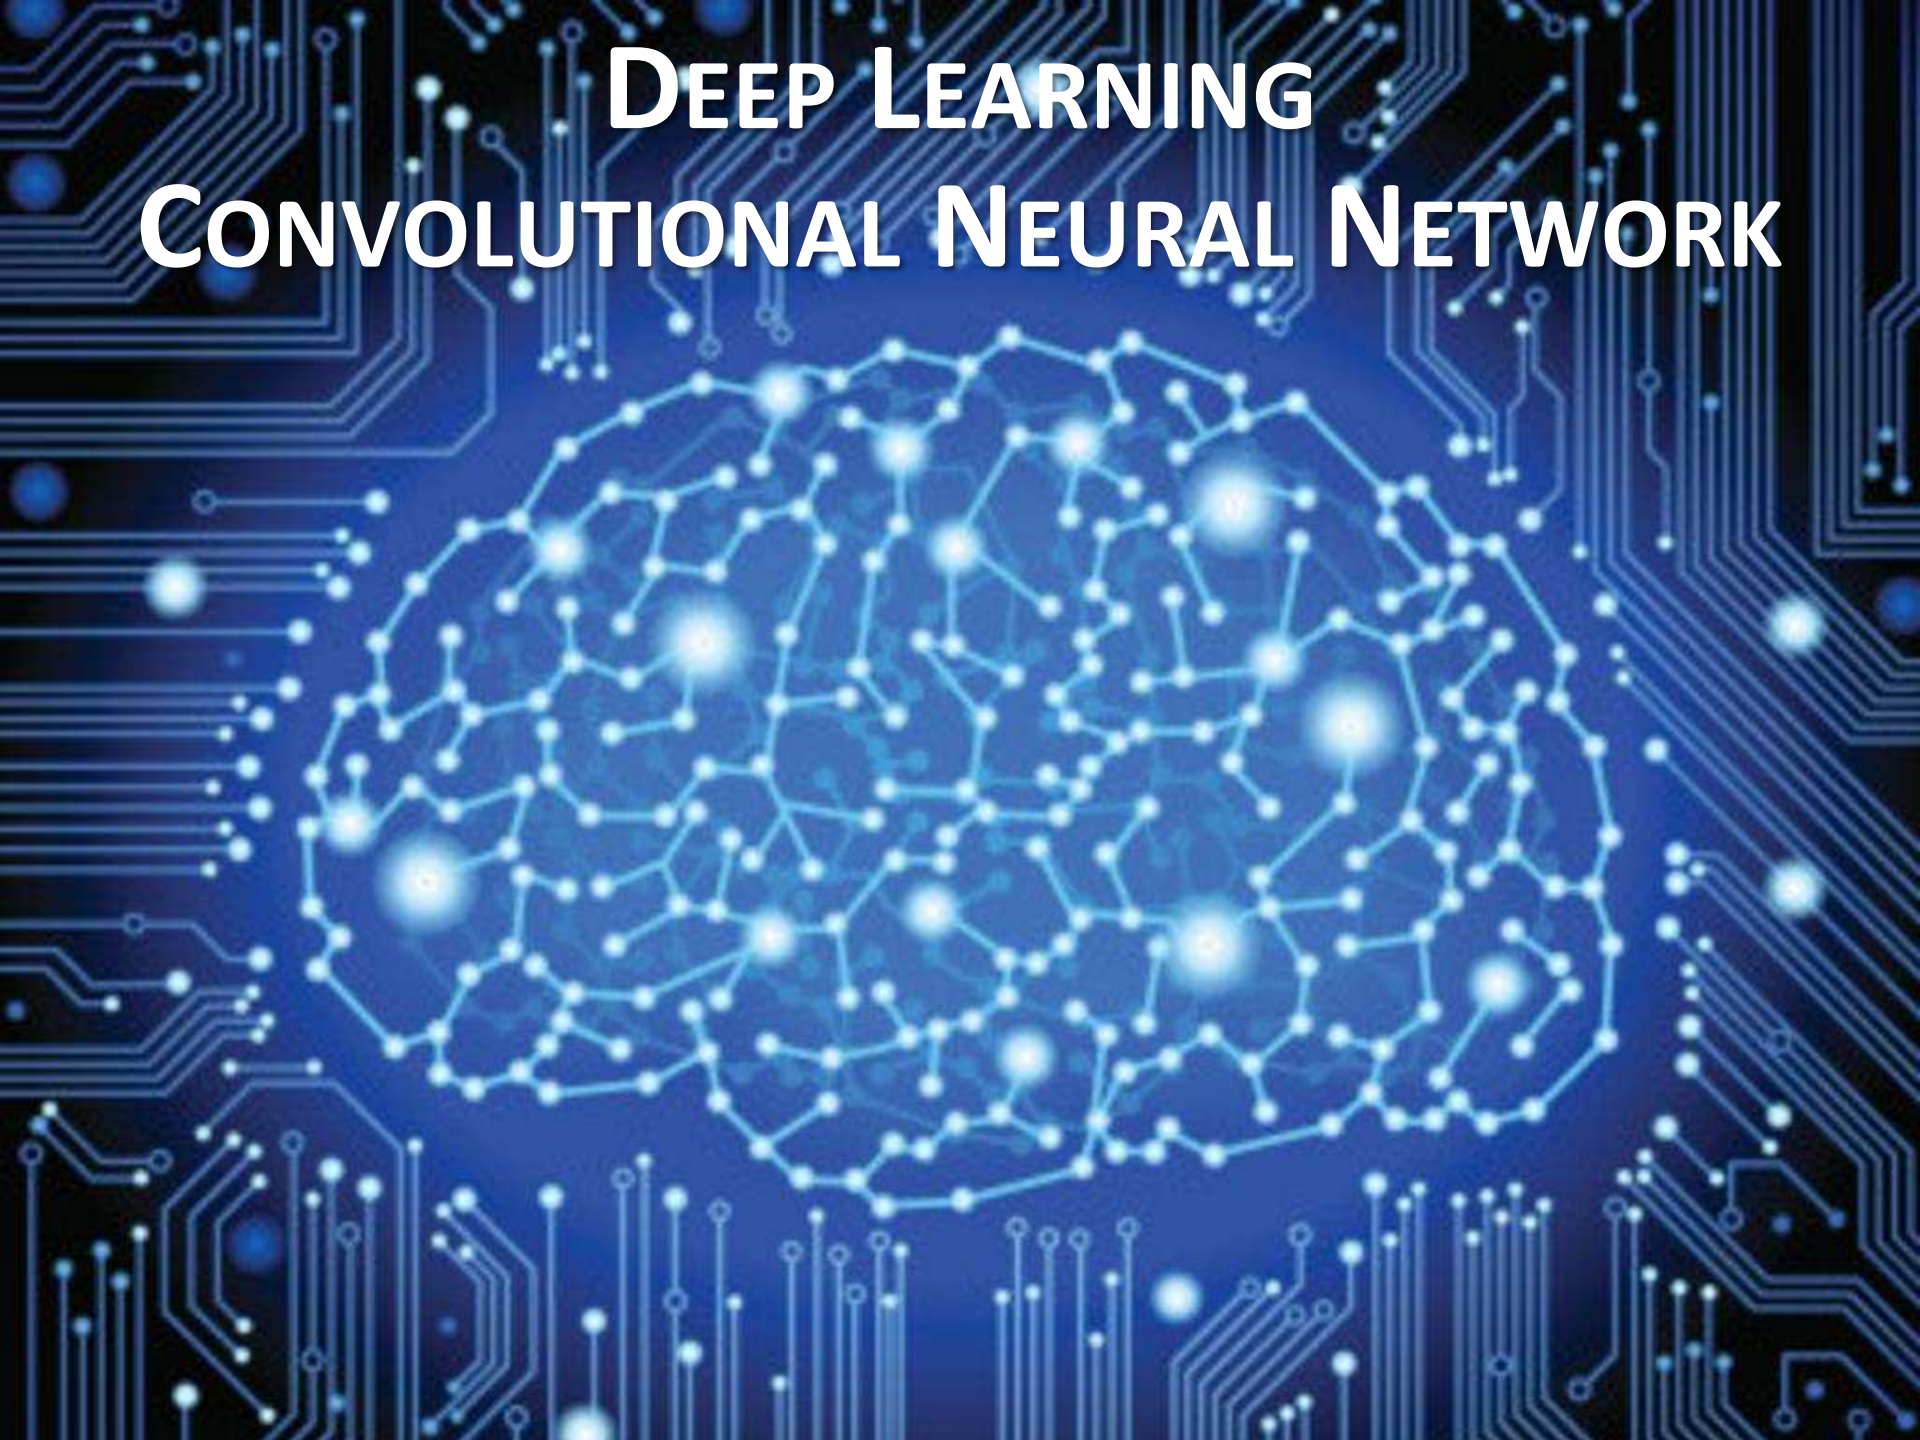

Deep learning (from Nature)

nature International weekly journal of science

Archive › Volume 521 › Issue 7553 › Insights › Reviews › Article

Yann LeCun, Yoshua Bengio & Geoffrey Hinton

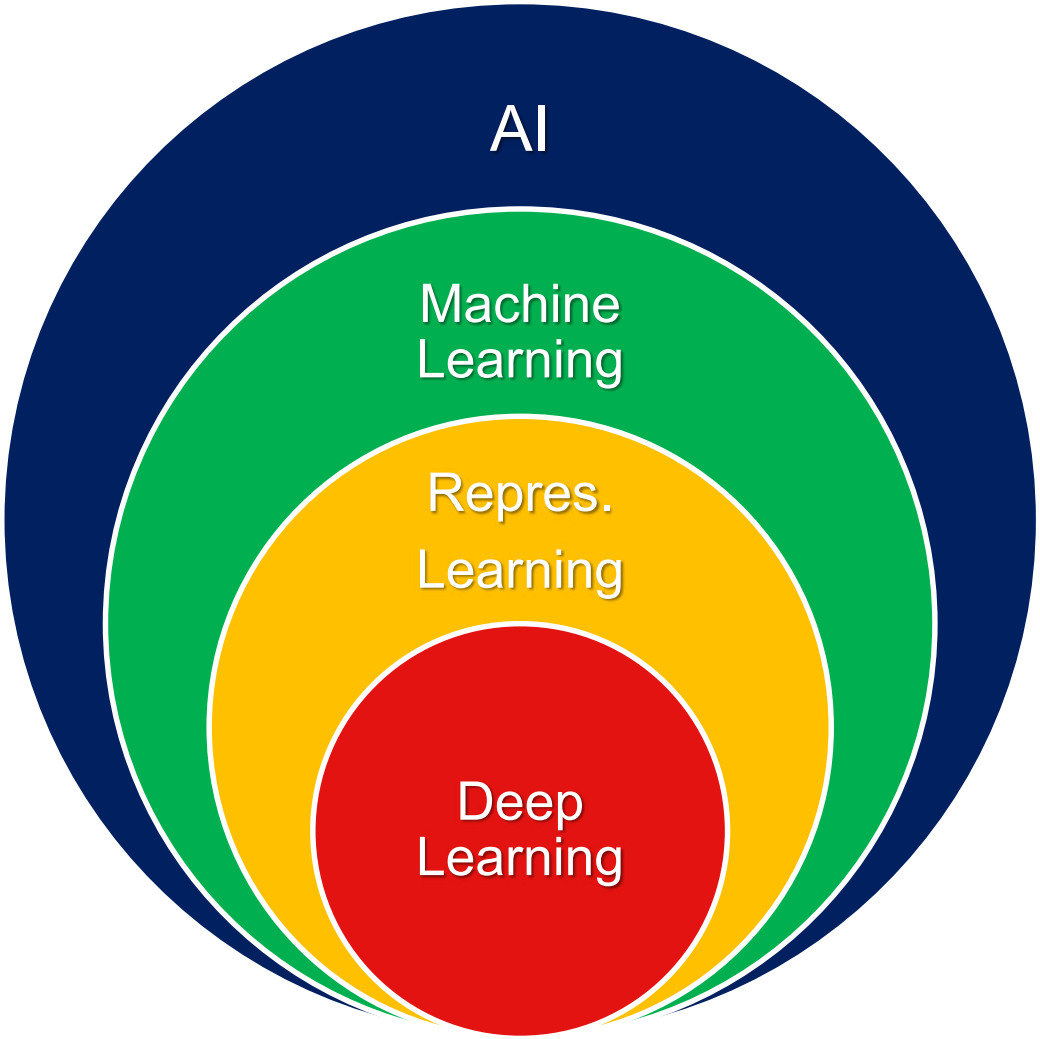

CNN and Deep Learning

C’è voluto tempo

Yoshua Bengio — Follow

Professor, U. Montreal, Computer Sc. & Op. Res., member/Fellow of CIFAR, CRM, REPARTI, GRSNC, CIRANO

Machine learning, deep learning, artificial intelligence

Verified email at umontreal.ca - Homepage

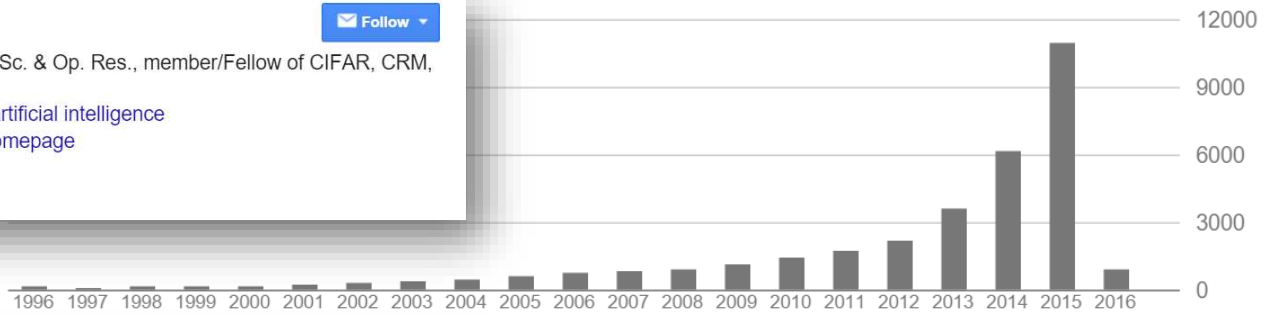

Yann LeCun — Follow

Director of AI Research at Facebook & Silver Professor at the Courant Institute, New York University

AI, machine learning, computer vision, robotics, image compression

Verified email at cs.nyu.edu - Homepage

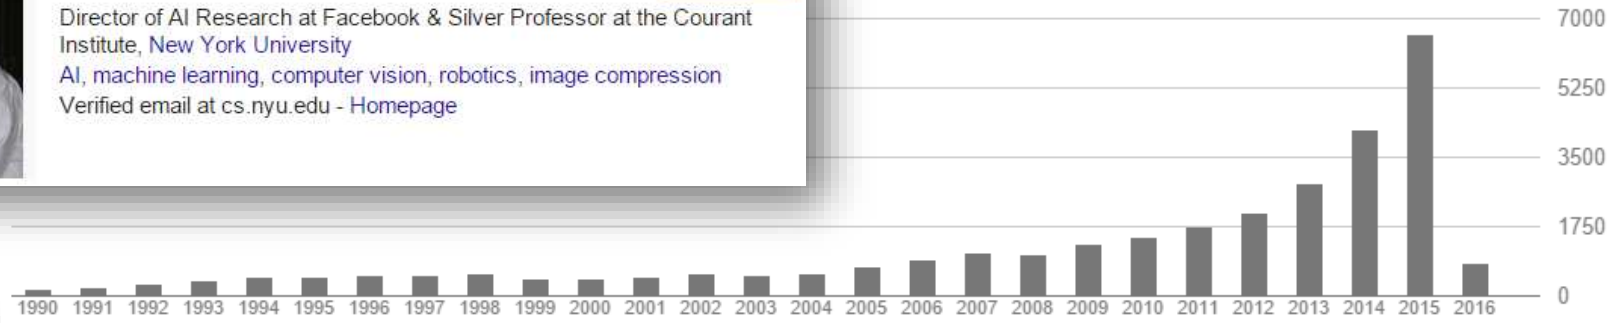

Geoffrey Hinton — Follow

Emeritus Professor of Computer Science, University of Toronto & Distinguished Researcher, Google Inc

machine learning, neural networks, artificial intelligence, cognitive science, object recognition

Verified email at cs.toronto.edu - Homepage

Machine Learning

Field of study that gives computers the ability to learn without being explicitly programmed.

Arthur Samuel (1959)

Some studies in machine learning using the game of checkers
IBM, J Res Dev 3:210–229

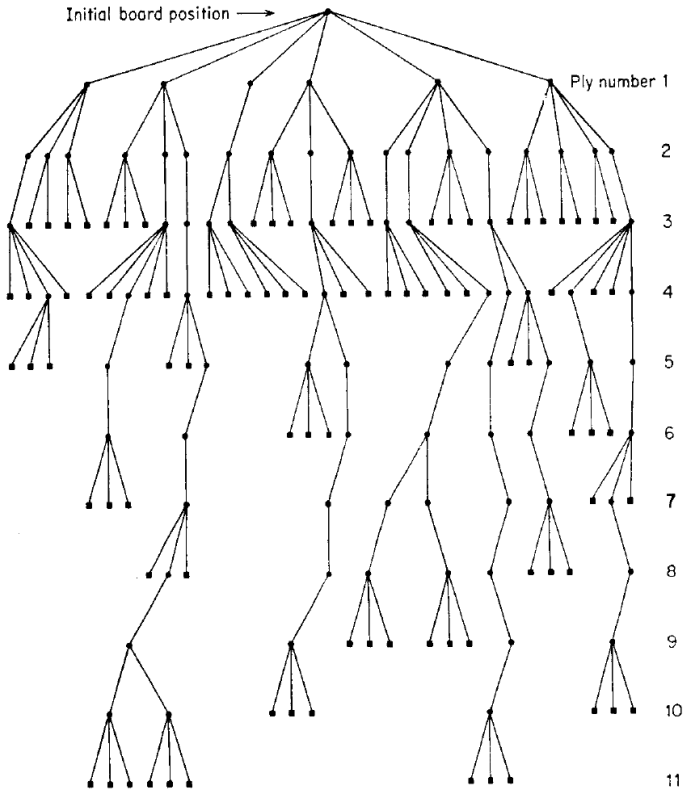

DeepMind Playing Atari

https://www.youtube.com/watch?v=Q70ulPJW3Gk

DeepMind Playing Atari

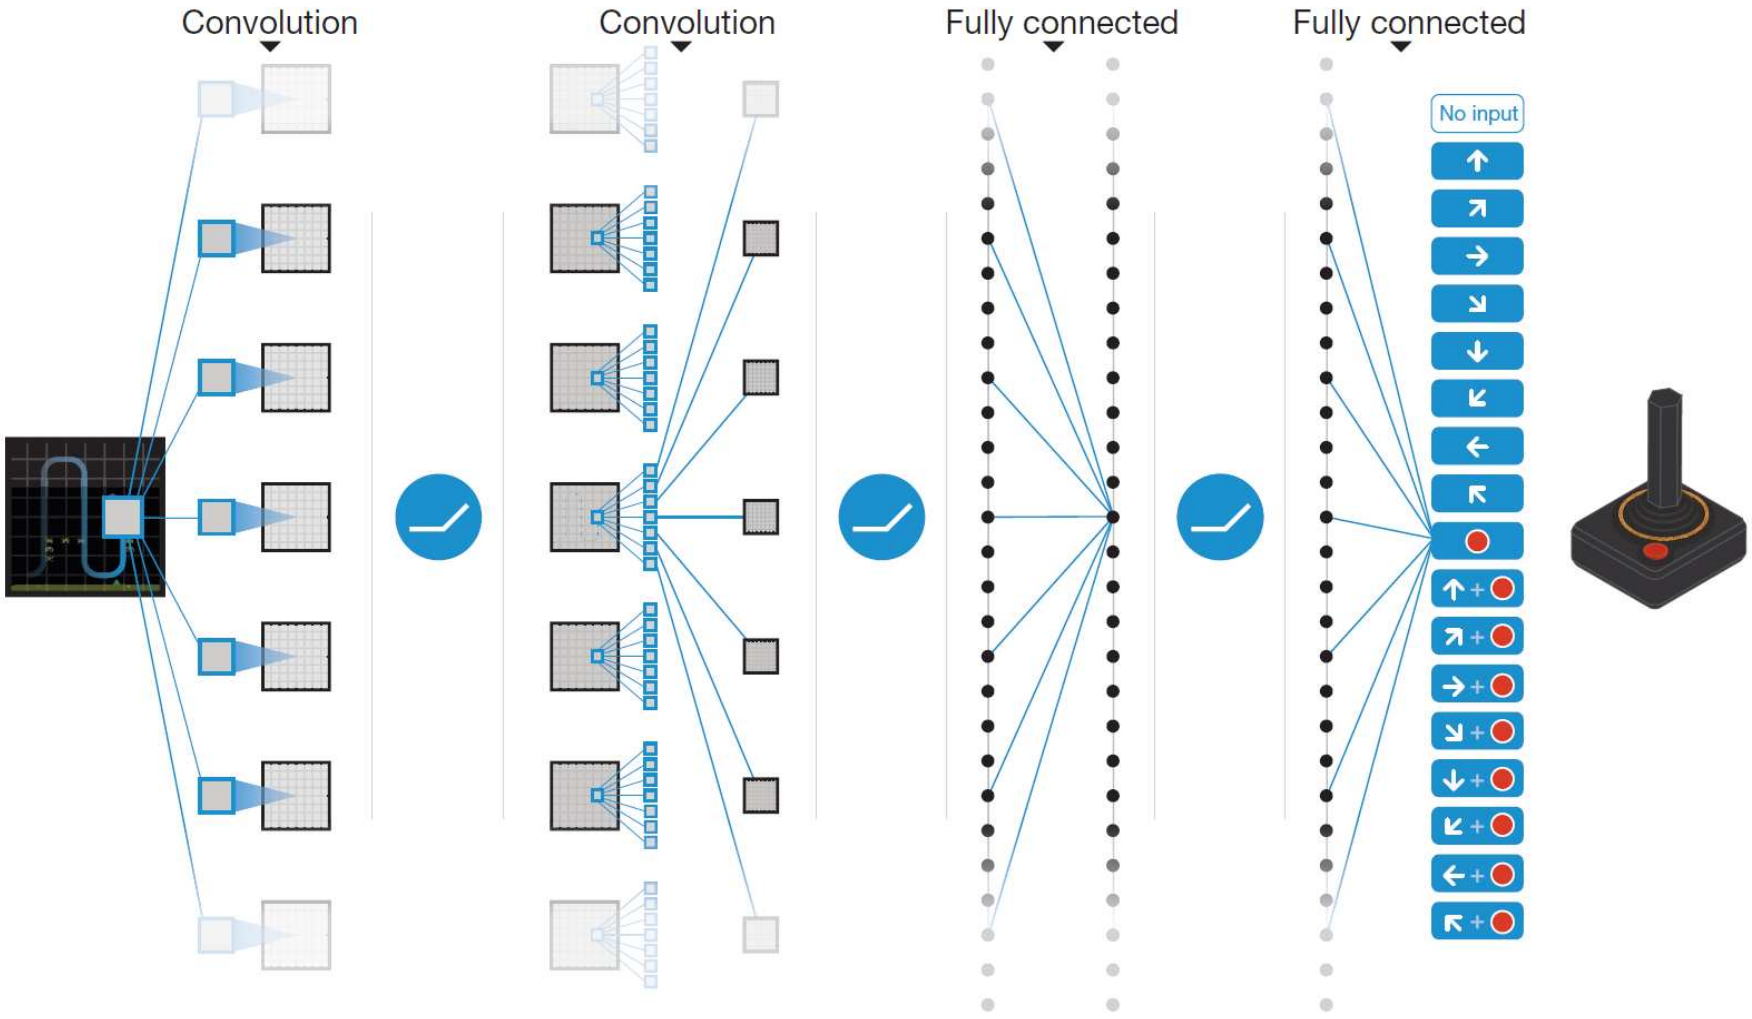

Machine Learning

The field of machine learning is concerned with the question of how to construct programs that automatically improve with experience.

*A computer program is said to **learn**:*

- *from experience* **E**
- *with respect to some task* **T**
- *and some performance measure* **P**

if its performance

- *on the task* **T**
- *as measured by* **P**
- ***improves** with experience* **E**

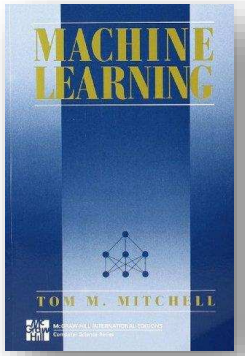

Tom M. Mitchelk
Machine Learning
McGraw-Hill, 1997

AlphaGo: Experience

1. 30 milioni di stati da **Kiseido Go Server**
   - Game server established in 2000 for people to play Go.

2. Policy Gradient **Reinforcement Learning**
   - Partite giocate dalla rete attuale
     con se stessa a uno dei passi precedenti (casuale)

3. 30 million posizioni distinte da partite giocate con se stesso

Deep learning (from Nature)

Yann LeCun, Yoshua Bengio & Geoffrey Hinton

Representation learning methods that

allow a machine to be fed with raw data and to automatically discover the representations needed for detection or classification.

Deep-learning are representation learning methods

- with **multiple levels** of representation, obtained by
- composing simple but **non-linear modules** that each
- transform the representation at one level into a representation at a higher, slightly more abstract level.

96 learned low-level filters

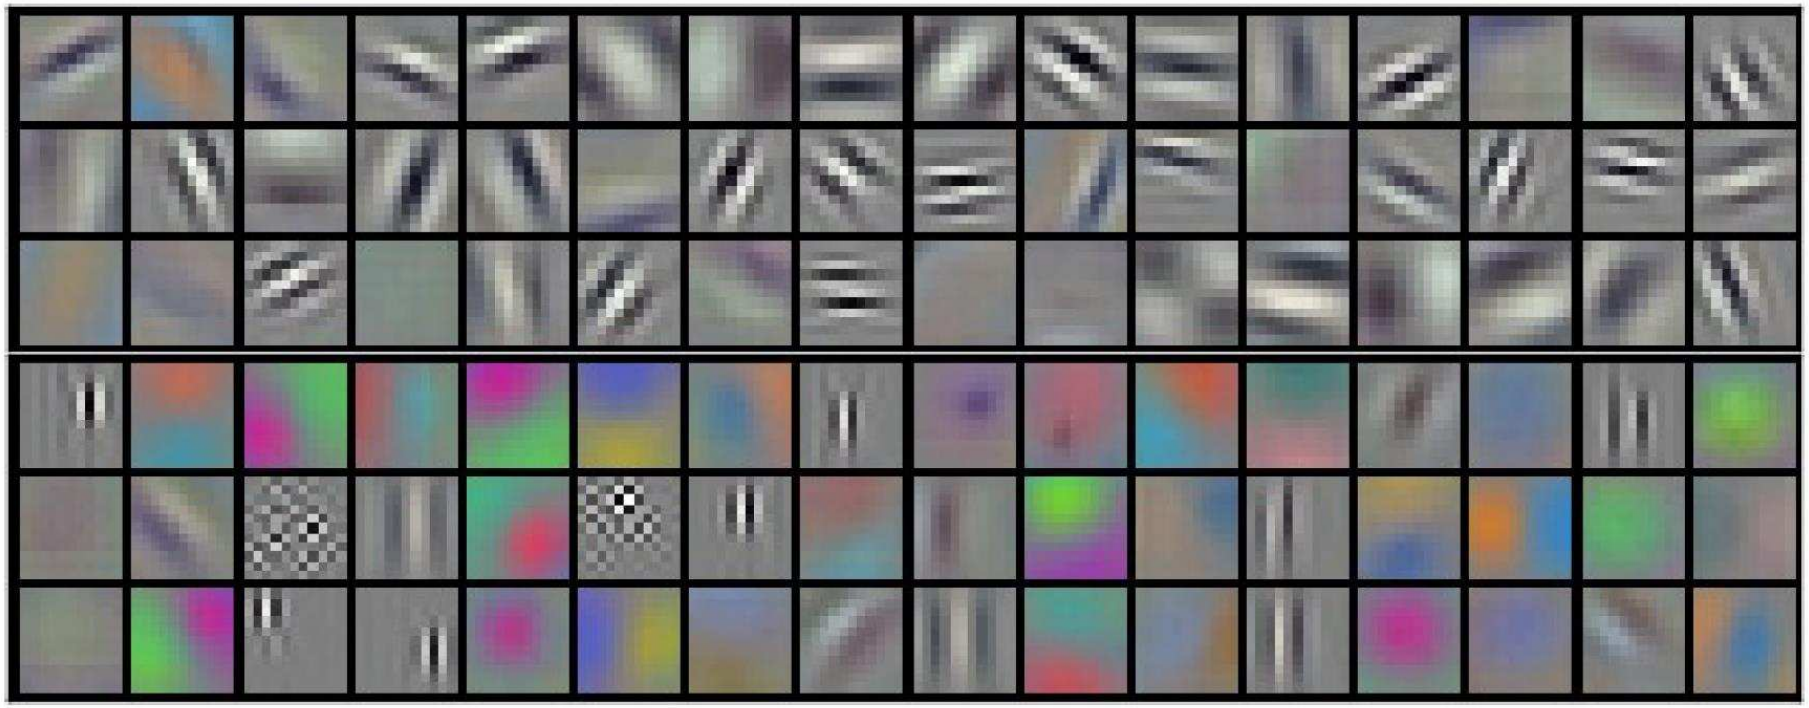

VISUALIZING CONVOLUTIONAL NETWORKS

- 9 feature maps and top 9 activations for each of them across validation data

Layer 1

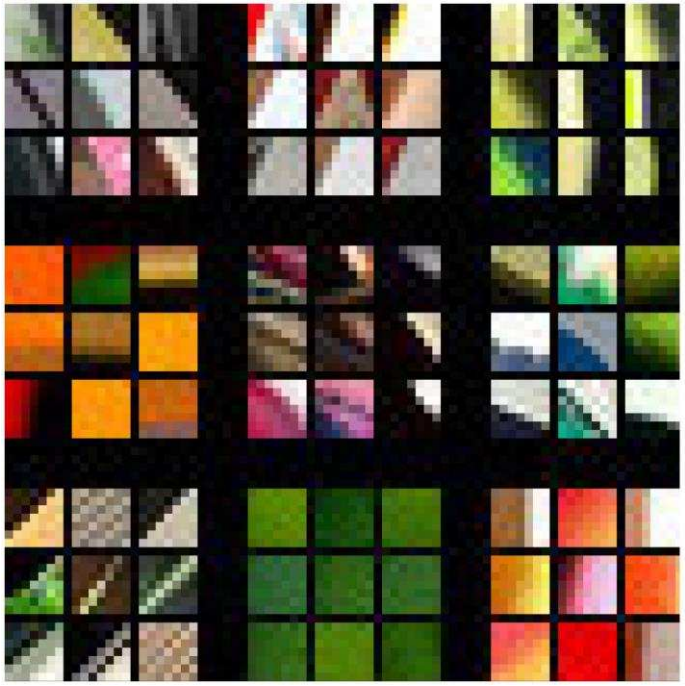

Visualizing Convolutional Networks
Zeiler, Matthew D., and Rob Fergus.
Computer vision – ECCV 2014 (cited by **639**)

VISUALIZING CONVOLUTIONAL NETWORKS

- Top 9 activations for some random feature maps across the validation data

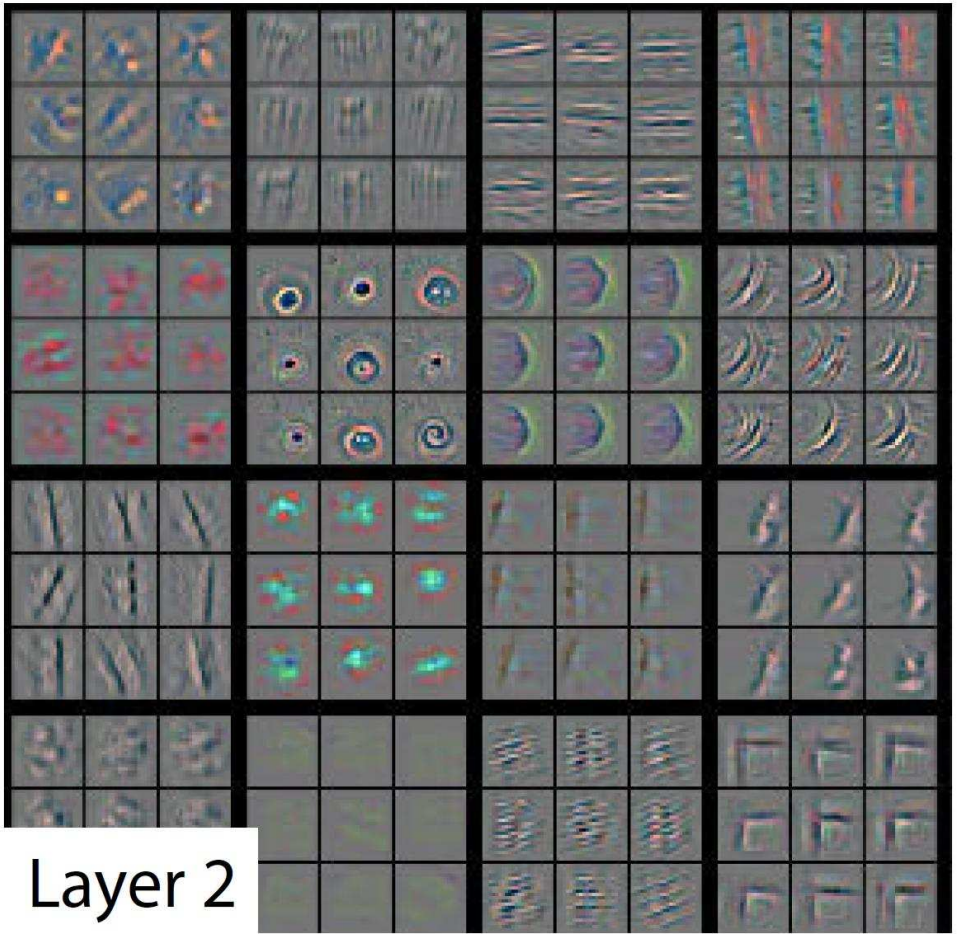

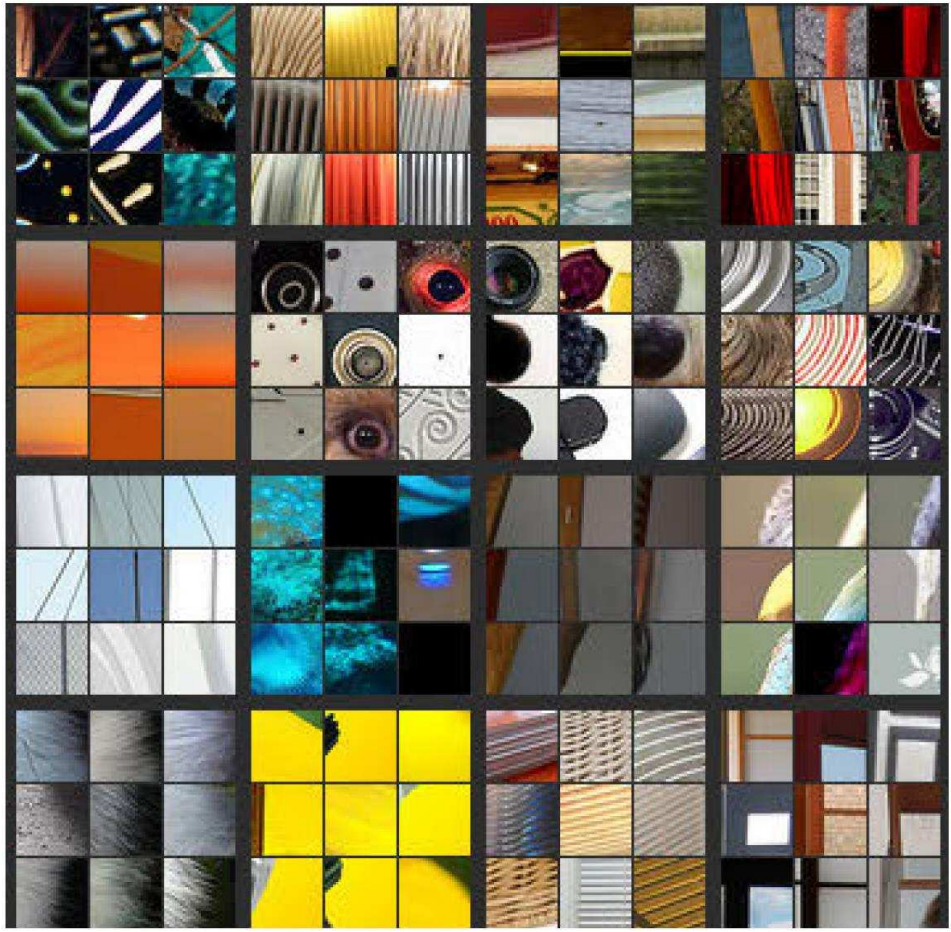

Visualizing Convolutional Networks
Zeiler, Matthew D., and Rob Fergus.
Computer vision – ECCV 2014 (cited by **639**)

VISUALIZING CONVOLUTIONAL NETWORKS

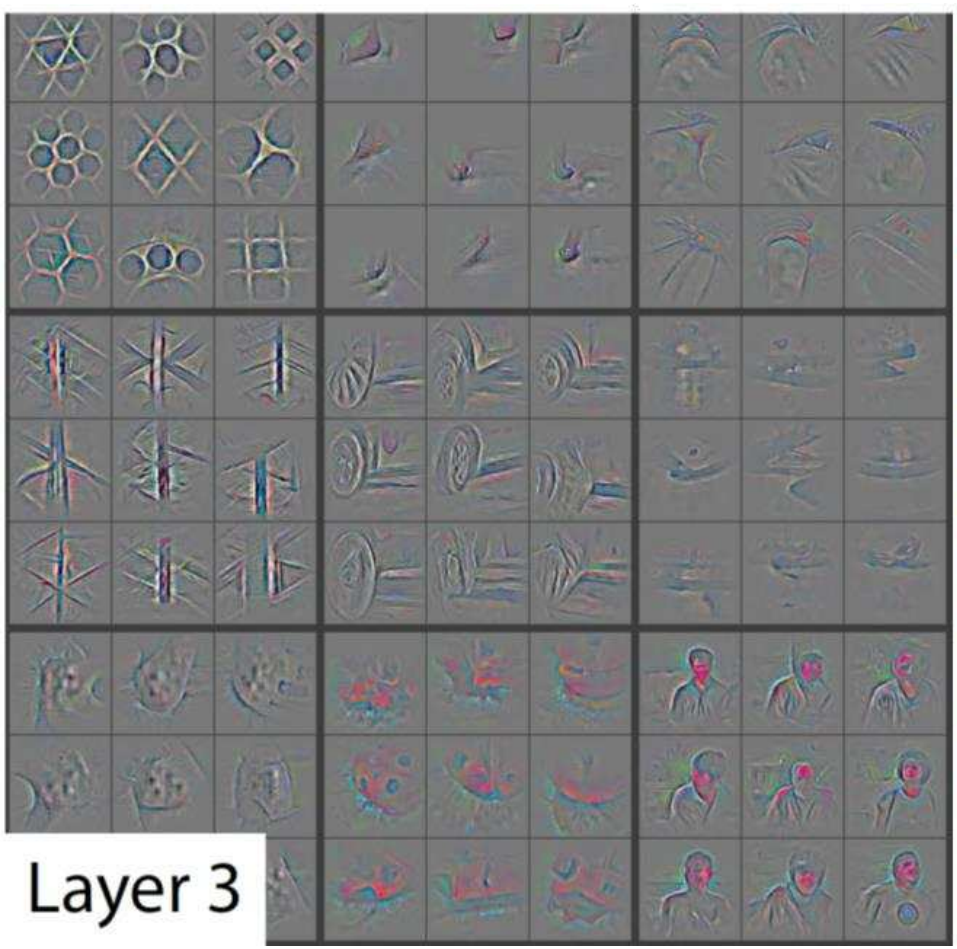

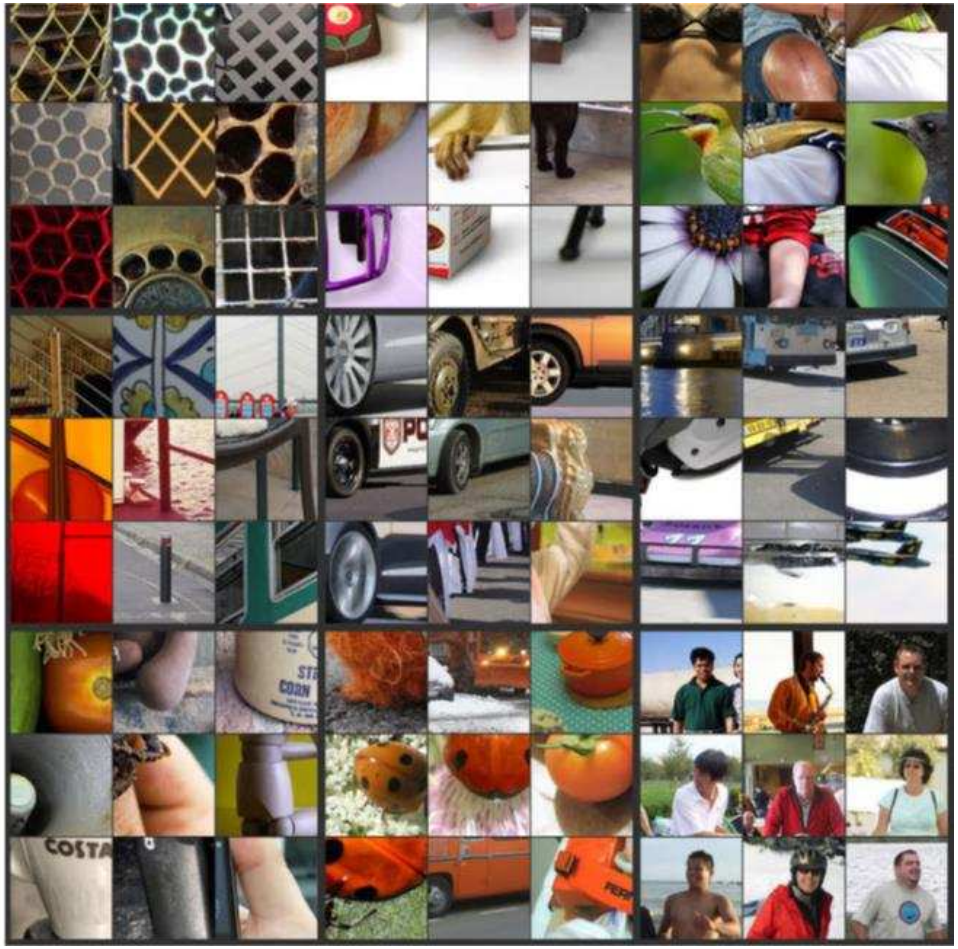

Visualizing Convolutional Networks
Zeiler, Matthew D., and Rob Fergus.
Computer vision – ECCV 2014 (cited by **639**)

VISUALIZING CONVOLUTIONAL NETWORKS

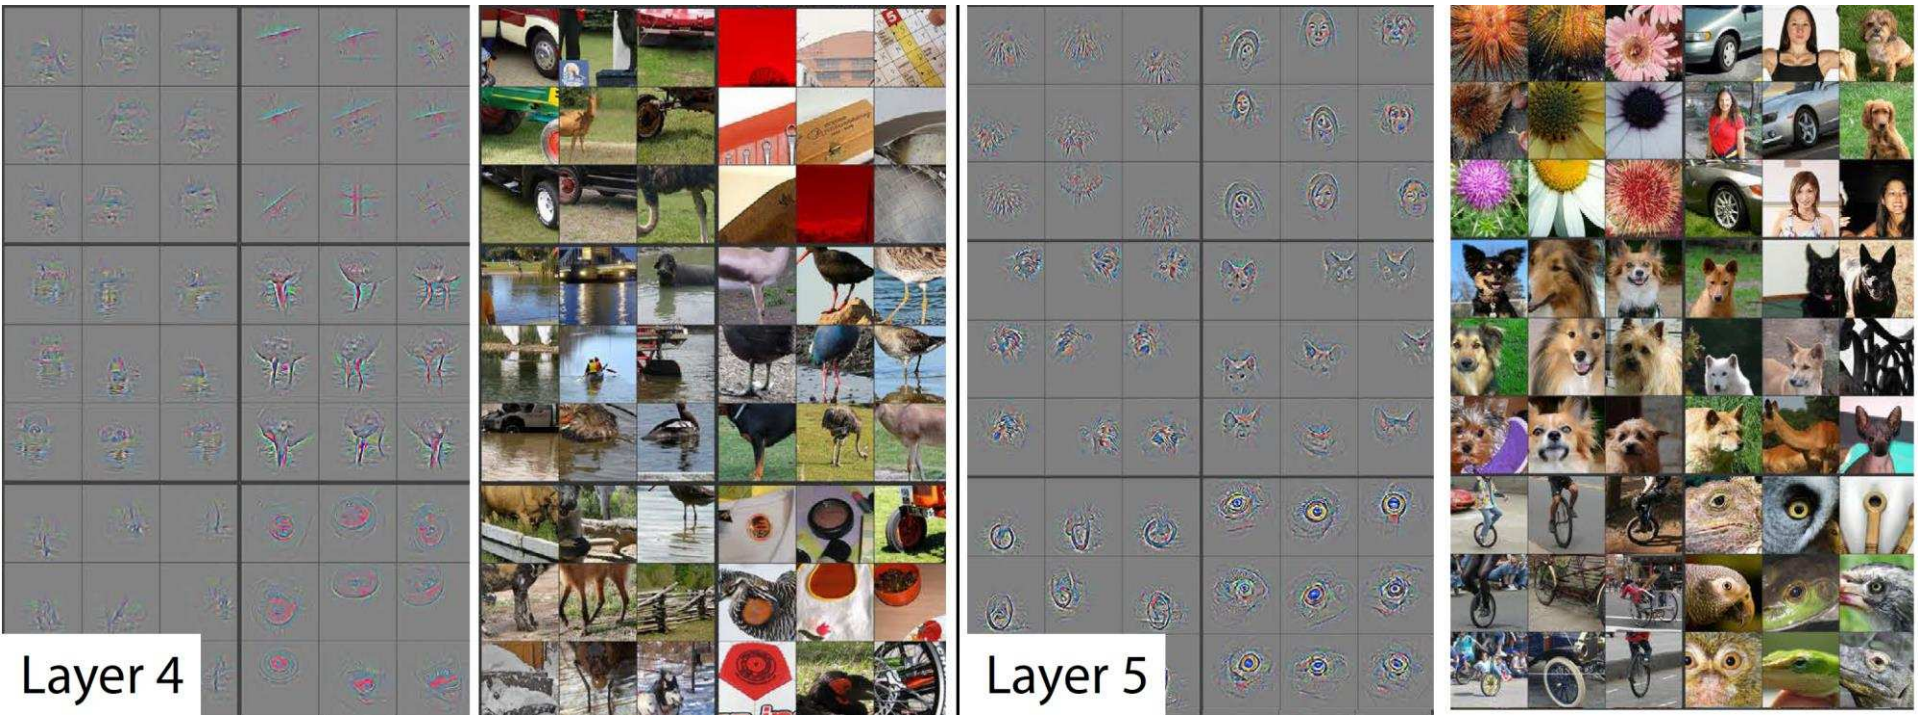

Visualizing Convolutional Networks
Zeiler, Matthew D., and Rob Fergus.
Computer vision – ECCV 2014 (cited by **639**)

ILLUSIONI

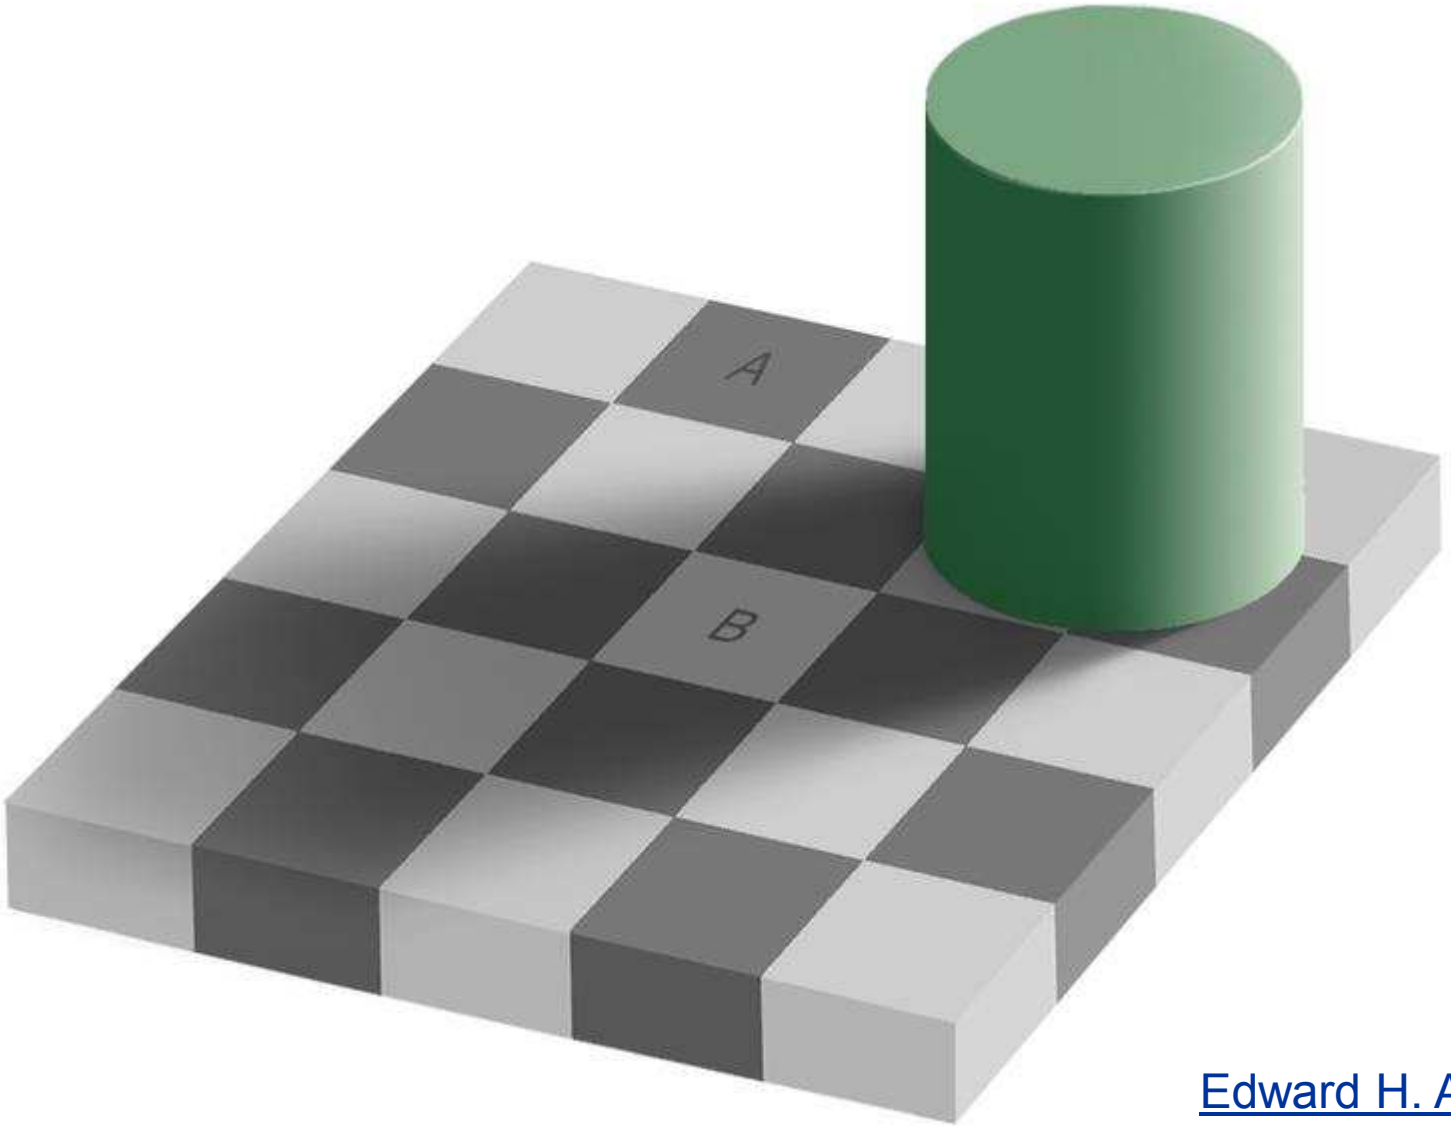

Edward H. Adelson

ILLUSIONI

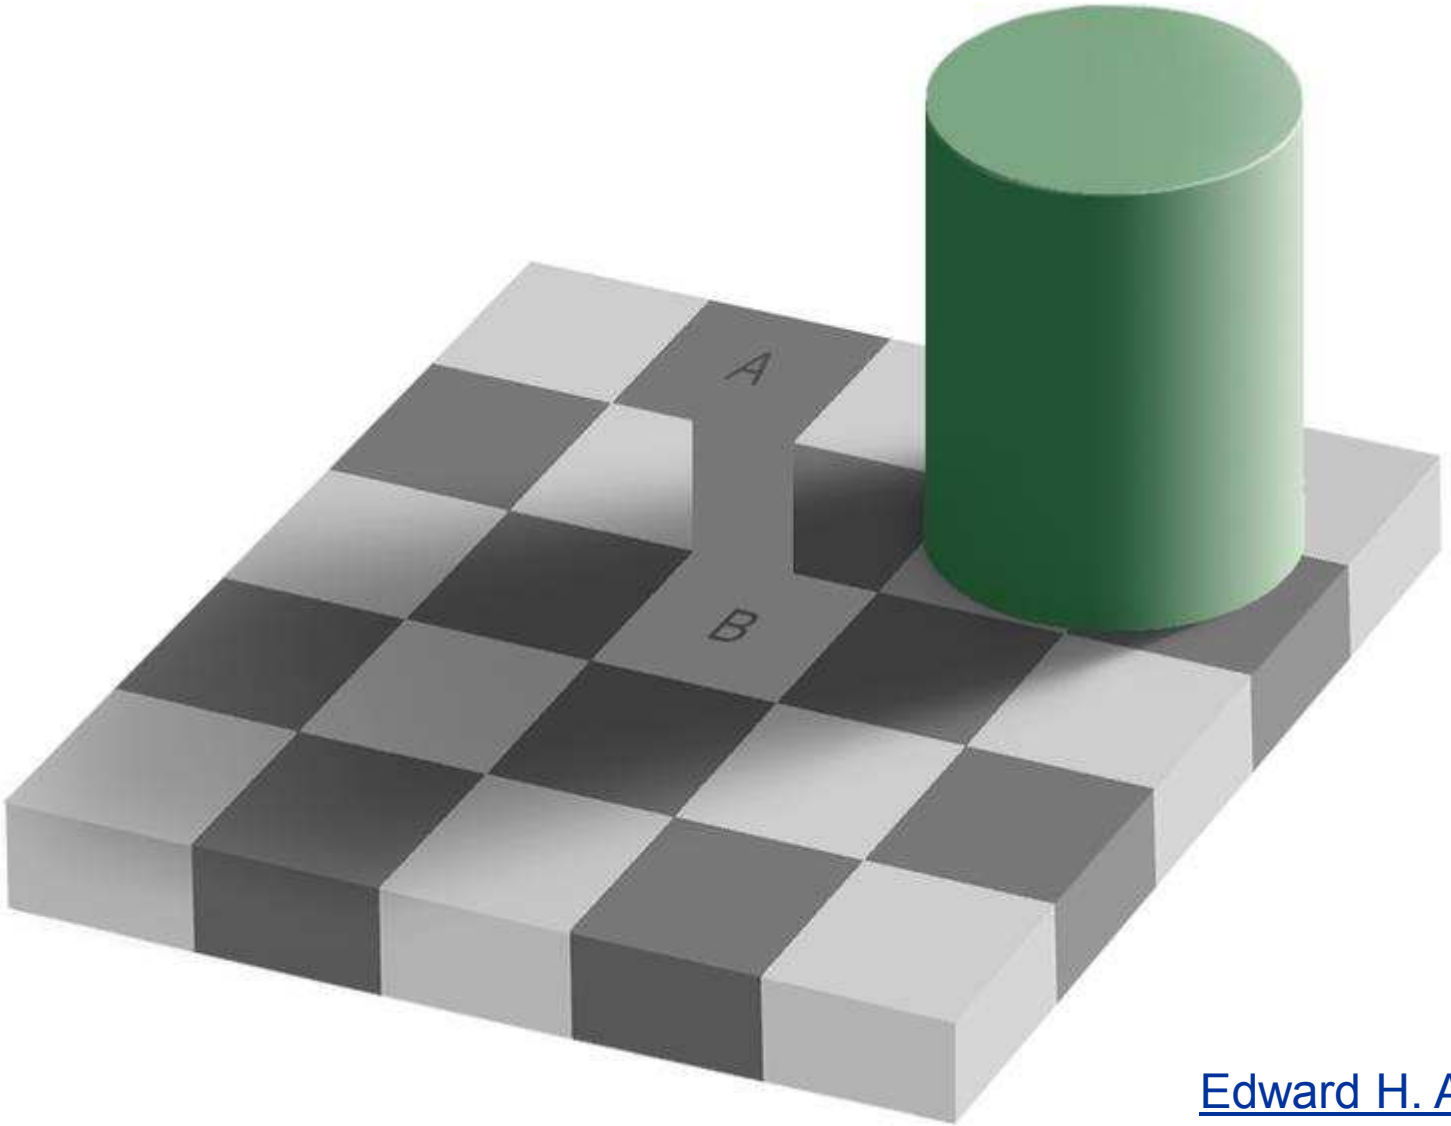

Edward H. Adelson

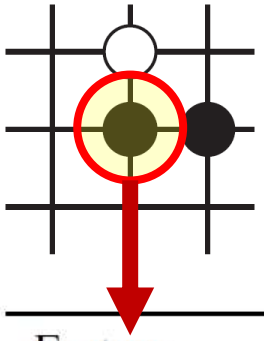

Cosa “vede” AlphaGo Deep CNN?

Feature	# of planes	Description
Stone colour	3	Player stone / opponent stone / empty
Ones	1	A constant plane filled with 1
Turns since	8	How many turns since a move was played
Liberties	8	Number of liberties (empty adjacent points)
Capture size	8	How many opponent stones would be captured
Self-atari size	8	How many of own stones would be captured
Liberties after move	8	Number of liberties after this move is played
Ladder capture	1	Whether a move at this point is a successful ladder capture
Ladder escape	1	Whether a move at this point is a successful ladder escape
Sensibleness	1	Whether a move is legal and does not fill its own eyes
Zeros	1	A constant plane filled with 0
Player color	1	Whether current player is black

Feature planes used by the policy network (all but last feature) and value network (all features).

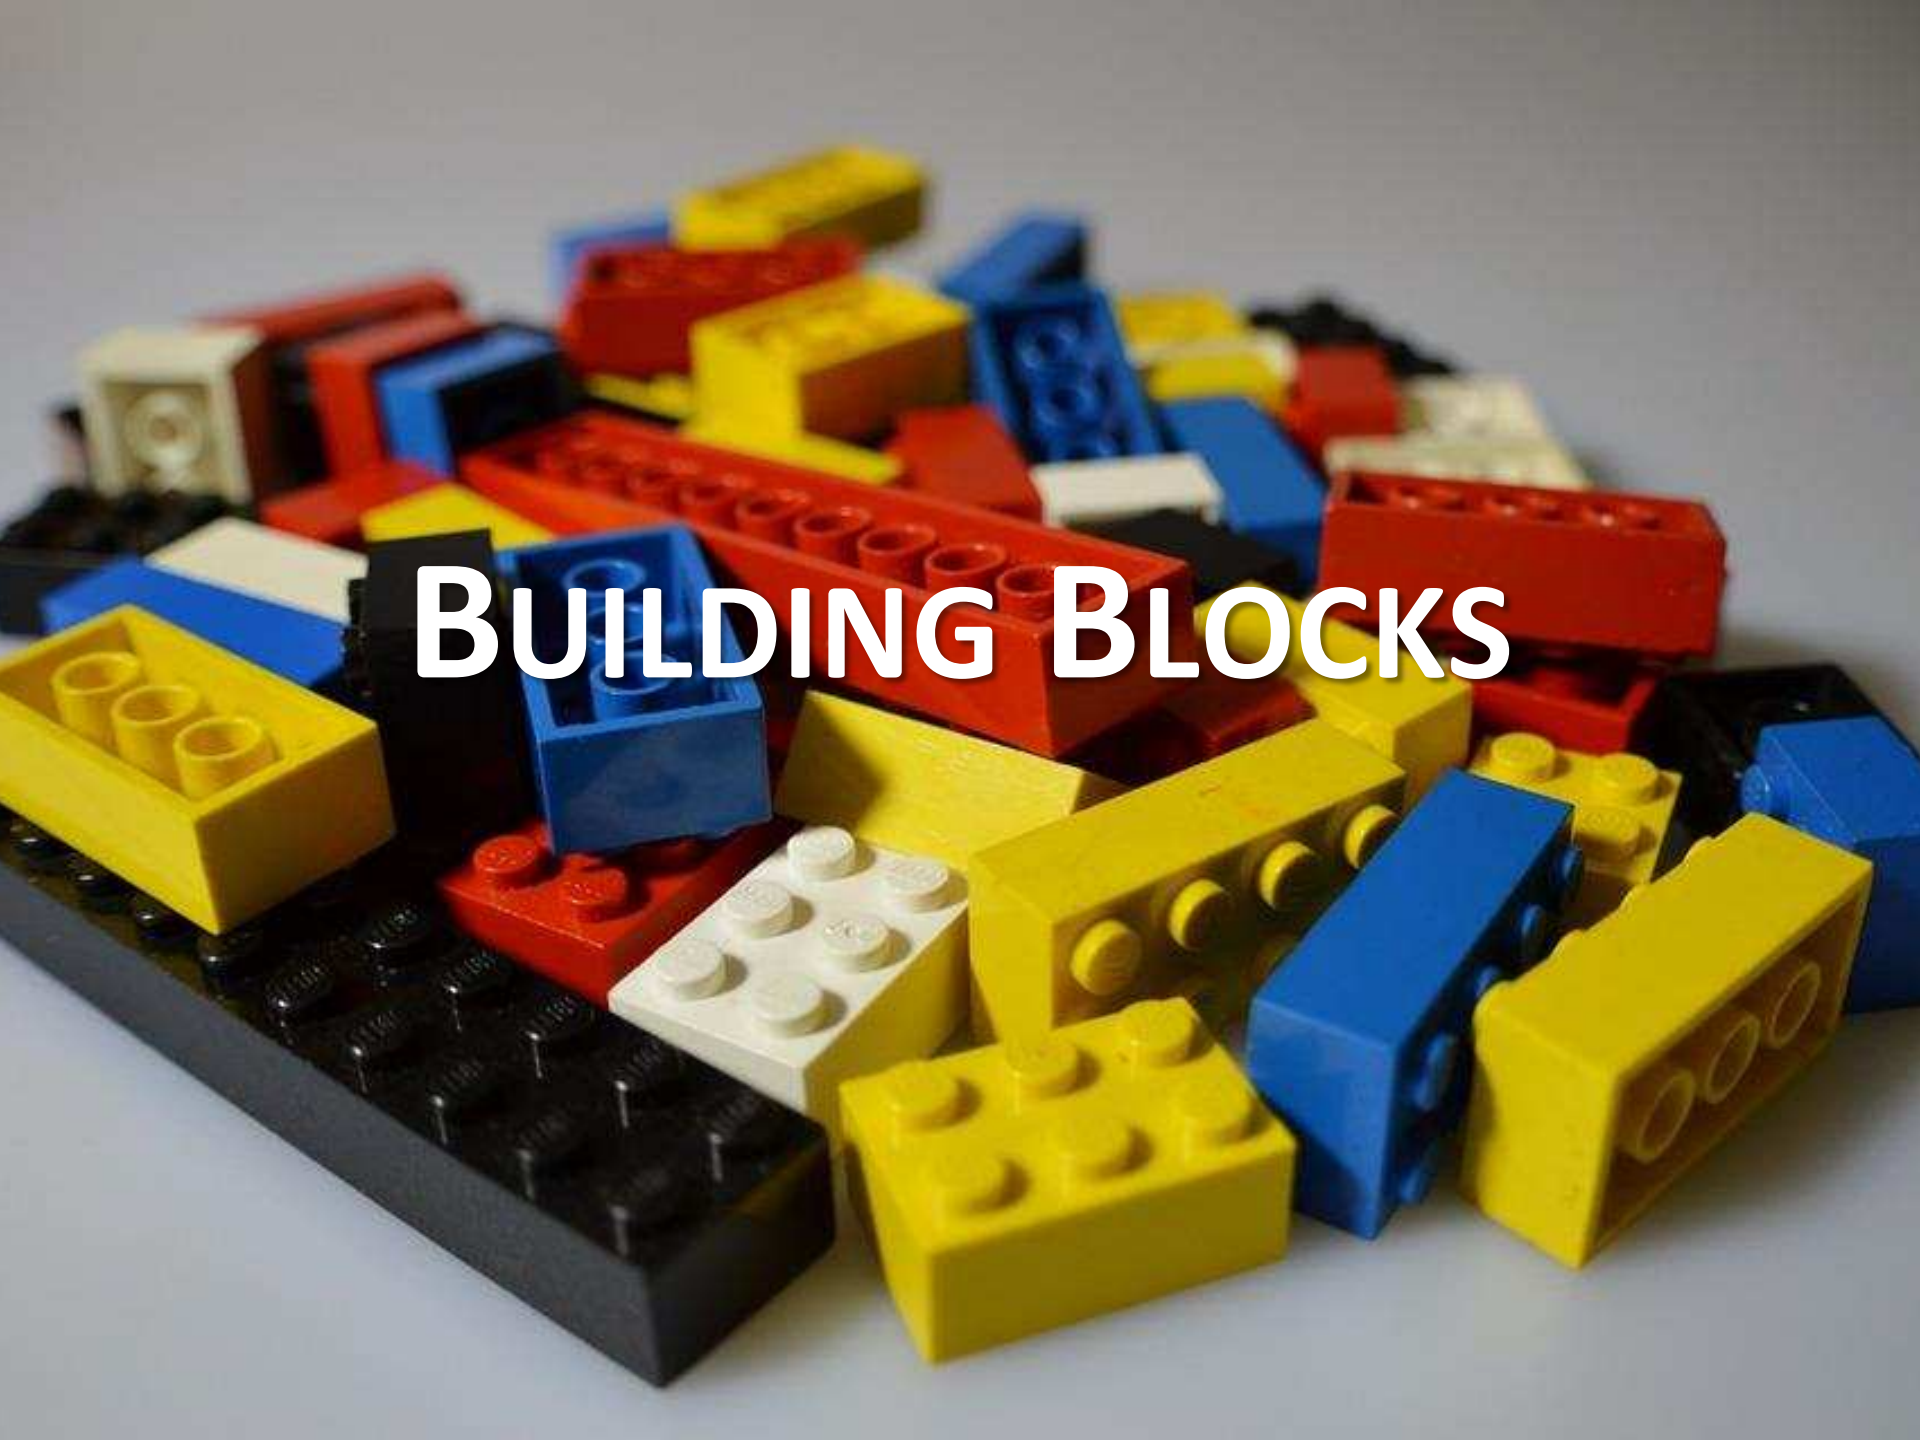
BUILDING BLOCKS

NEURONE

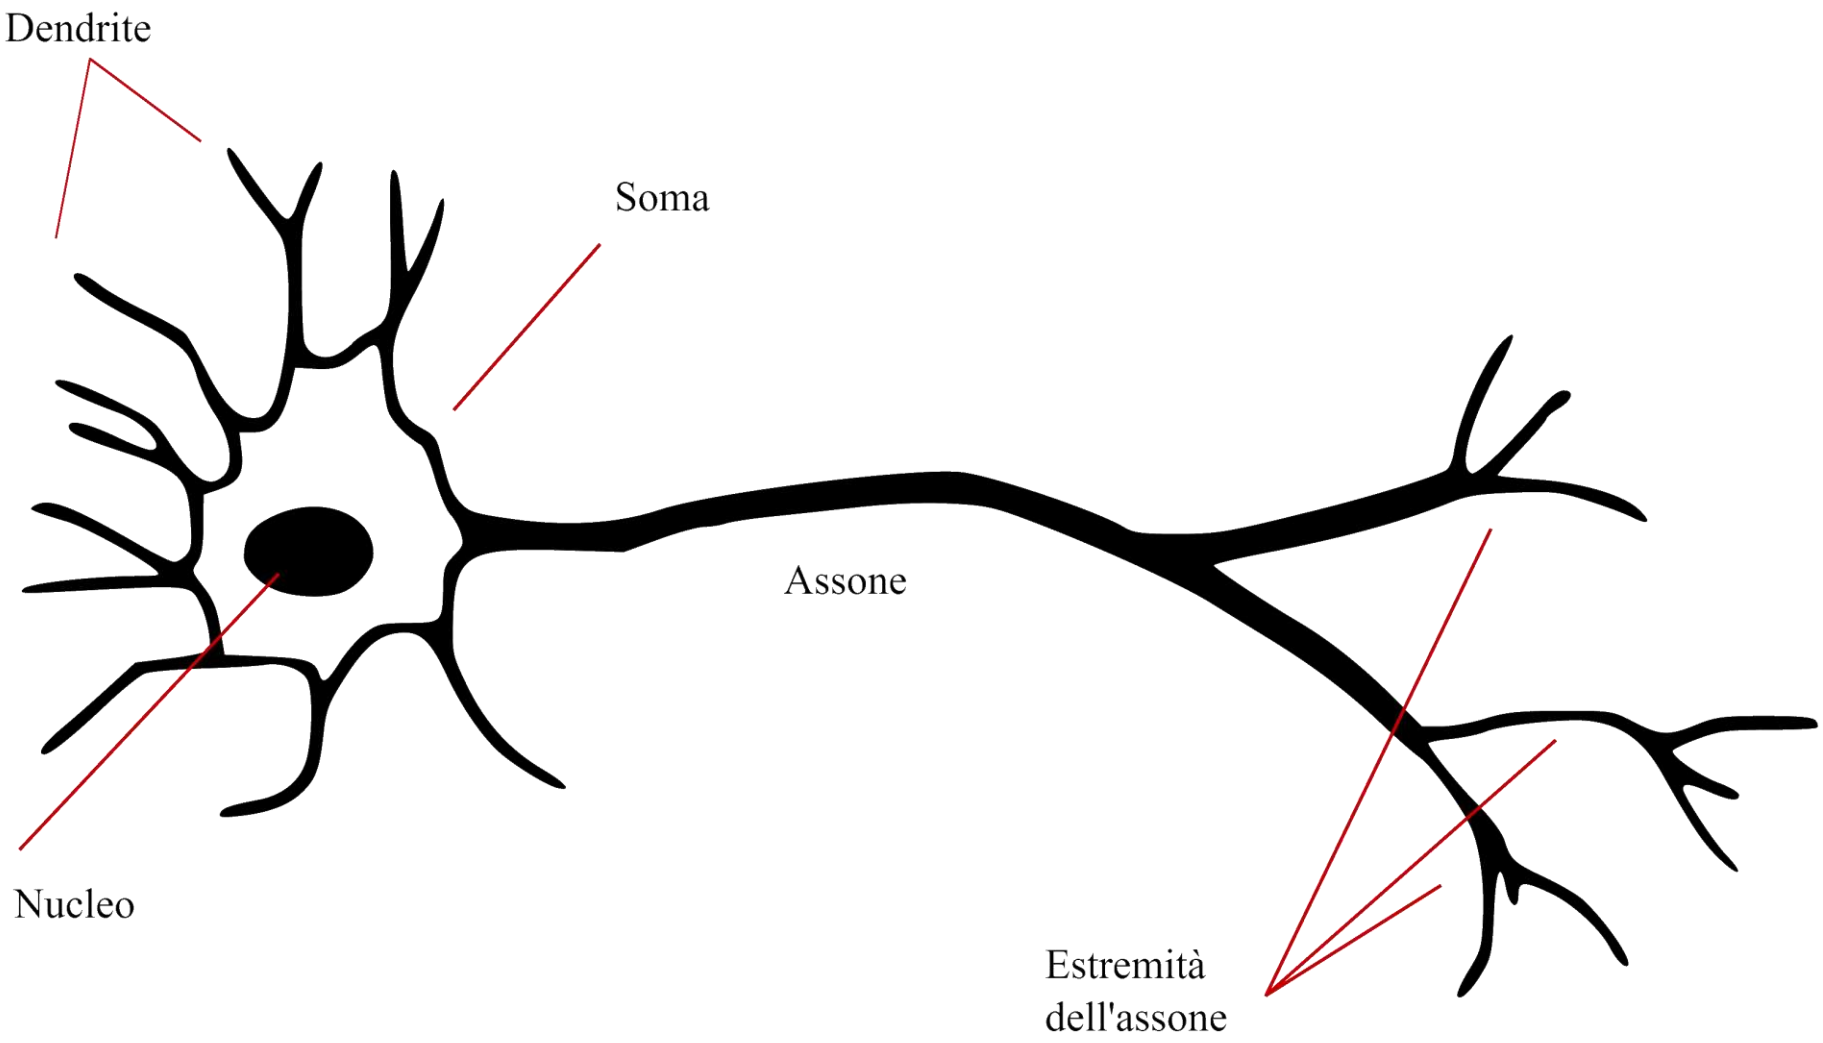

NEURONE

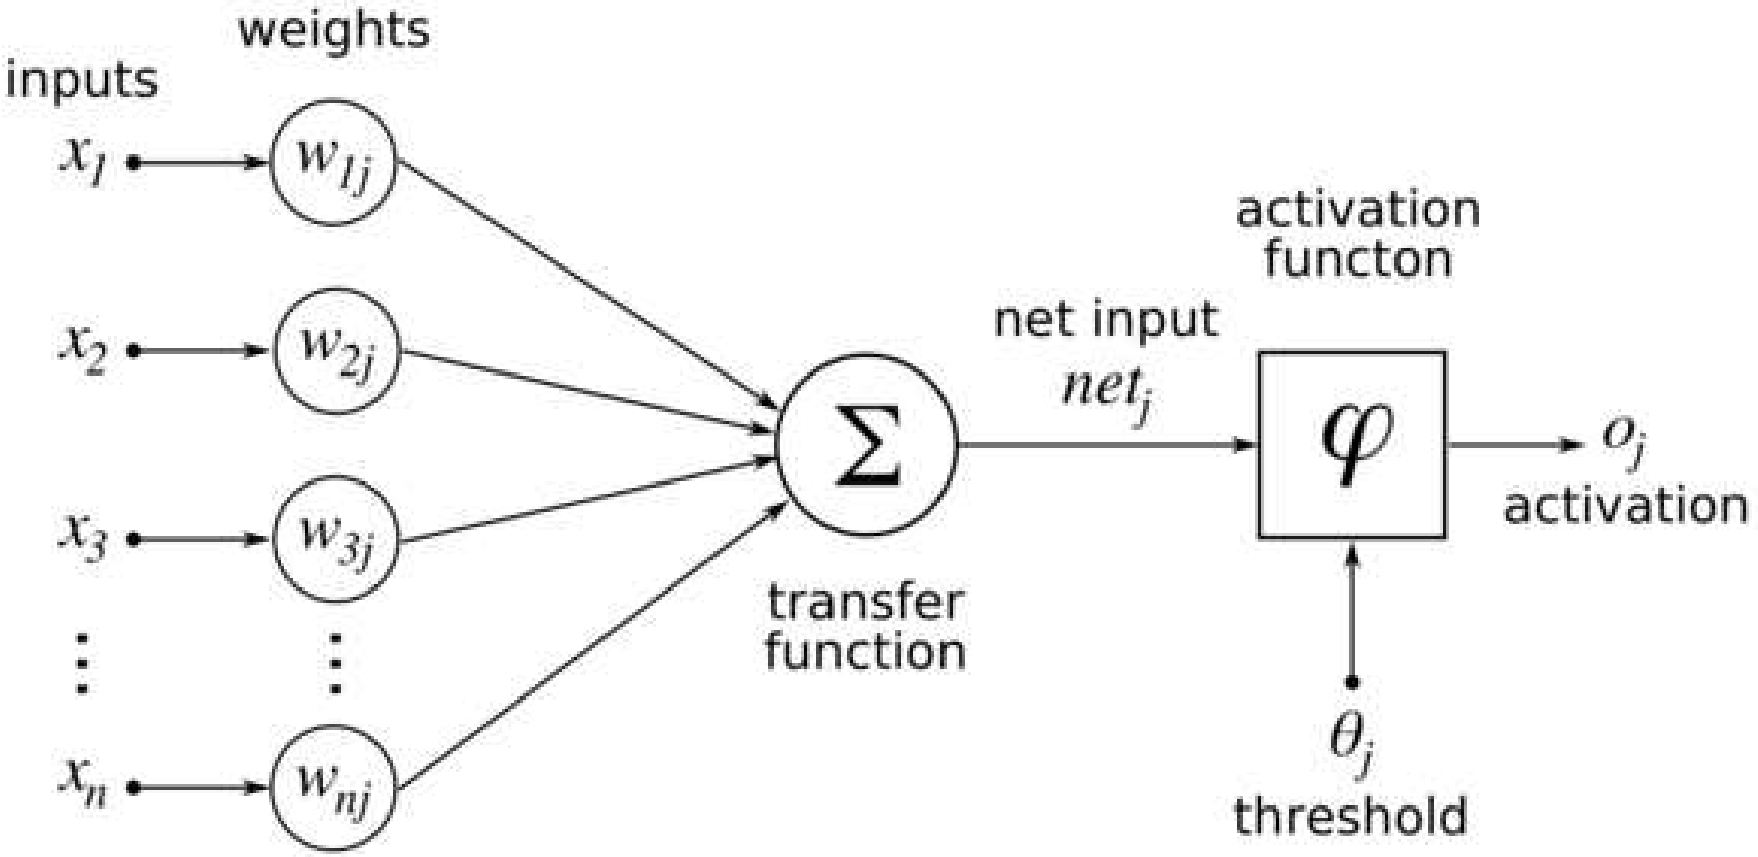
weights
inputs
x_1
x_2
x_3
x_n
w_{1j}
w_{2j}
w_{3j}
w_{nj}
Σ
transfer function
net input
net_j
activation functon
φ
o_j
activation
θ_j
threshold

Rete Neuronale (Classica)

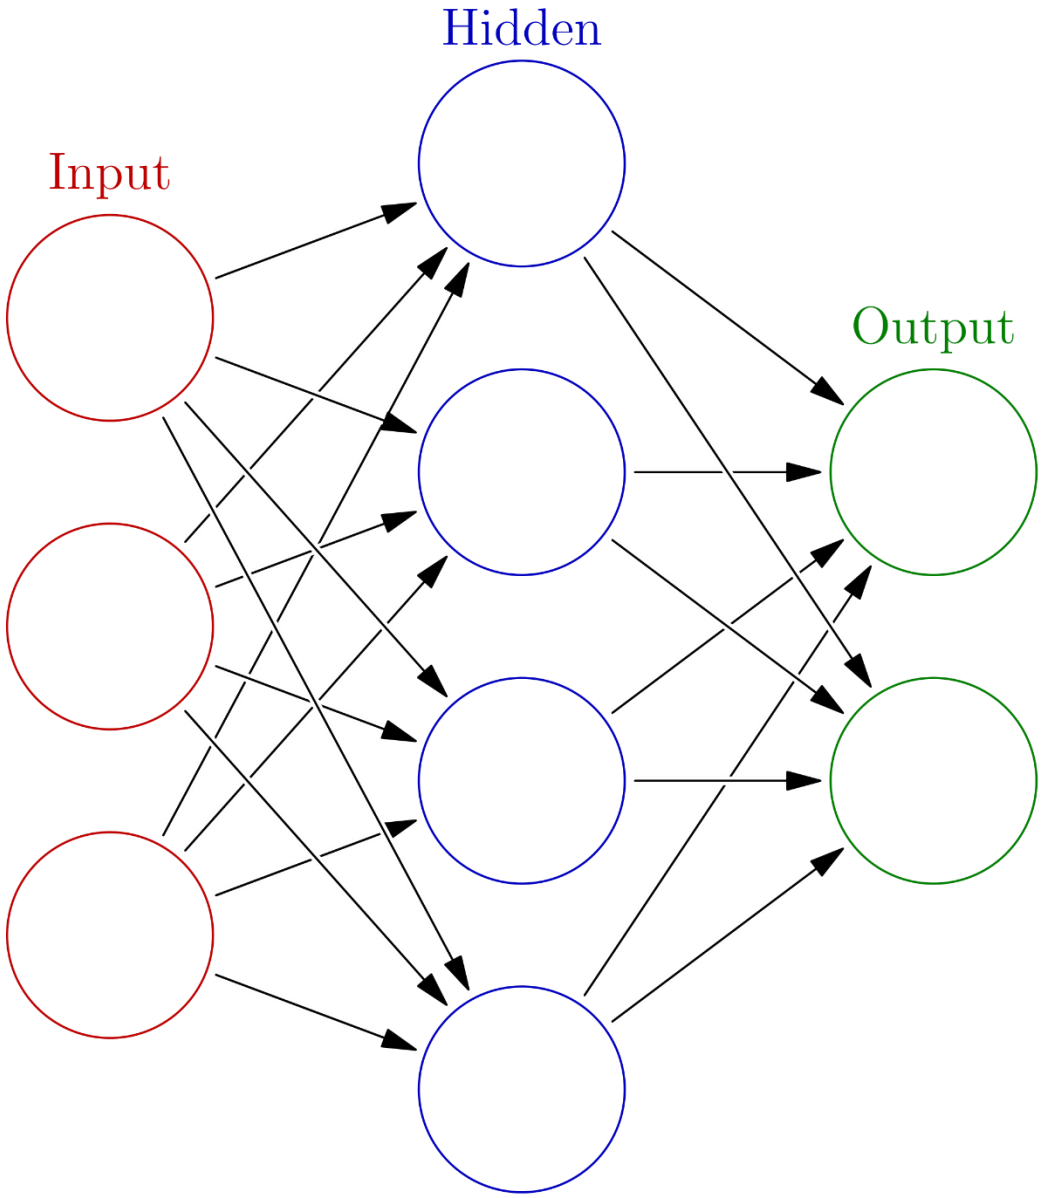

CONVOLUTIONAL NEURAL NETWORK

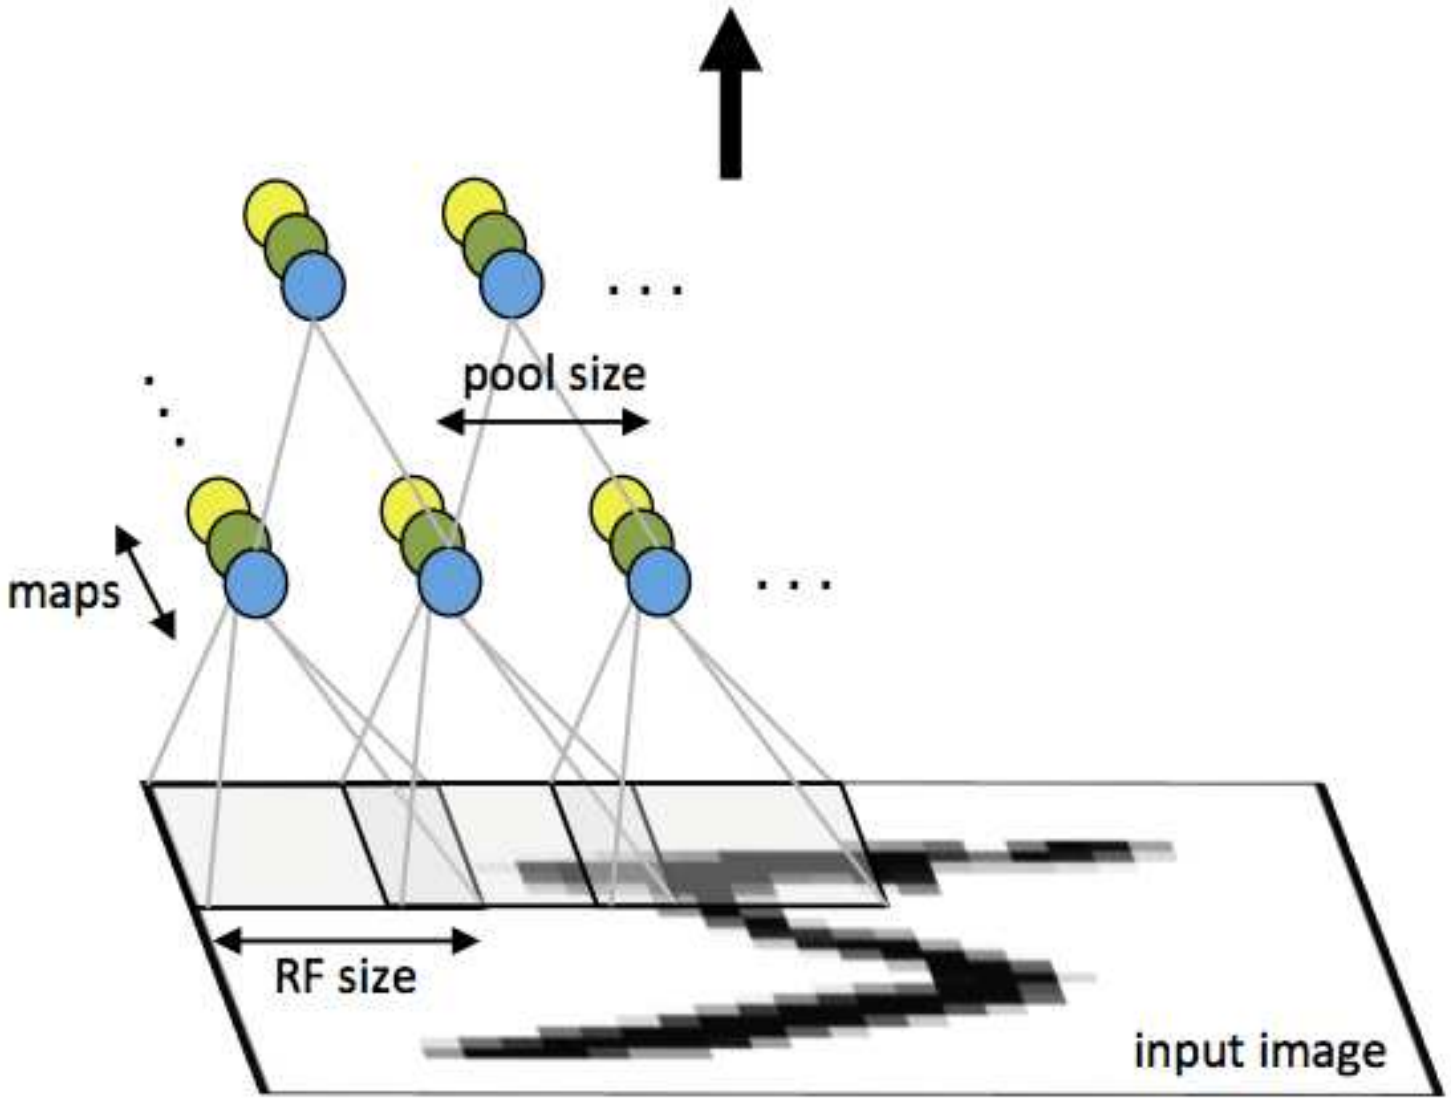

http://ufldl.stanford.edu/tutorial

From Tutorial by: Andrew Ng, Jiquan Ngiam, Chuan Yu Foo, Yifan Mai, Caroline Suen, Adam Coates, Andrew Maas, Awni Hannun, Brody Huval, Tao Wang, Sameep Tandon

AlexNet 2012

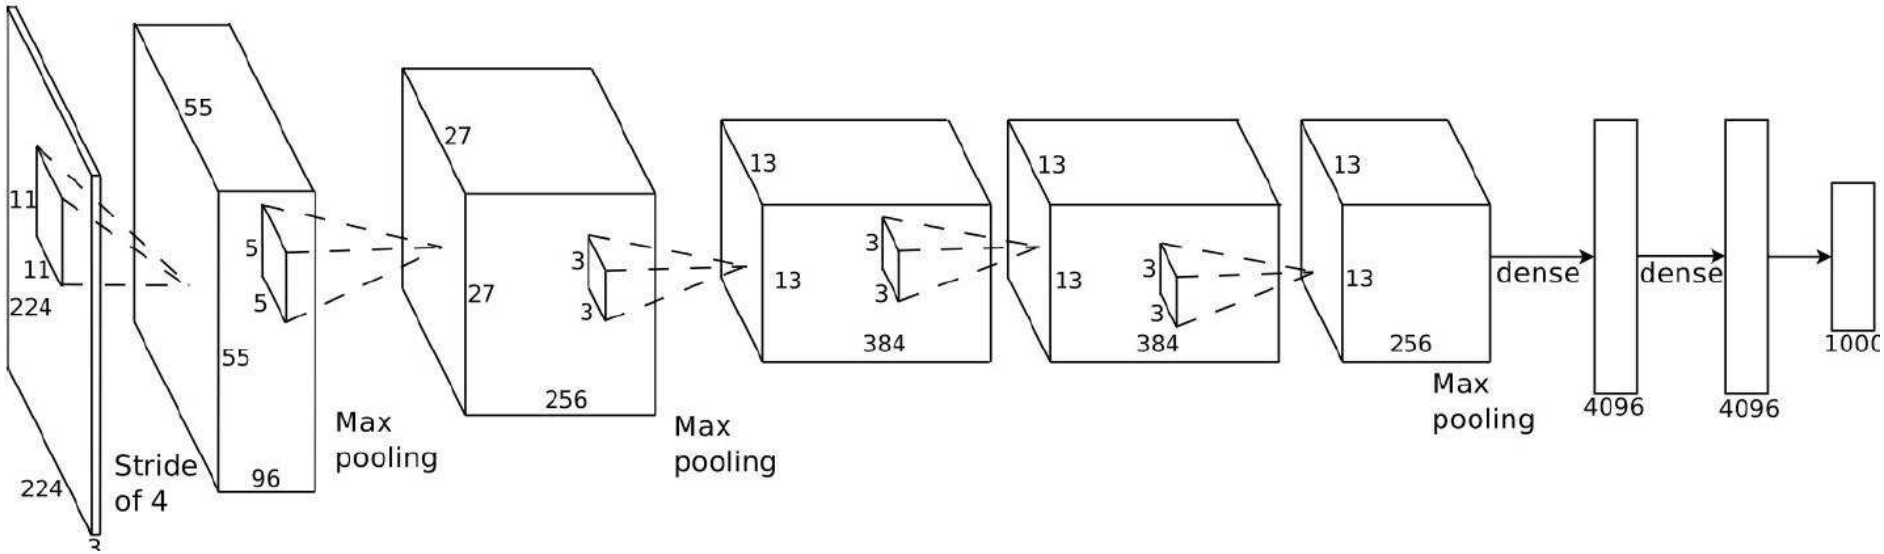

- Livelli: 6
- Filtri: 96, 256, 384, 384, 256
- Neuroni: 650,000
- Valori da imparare: 60,000,000

- AlphaGO: 13 livelli, 196 filtri per layer

COSA FACCIAMO NOI

Istituto di Scienza e Tecnologie dell'Informazione "A. Faedo"

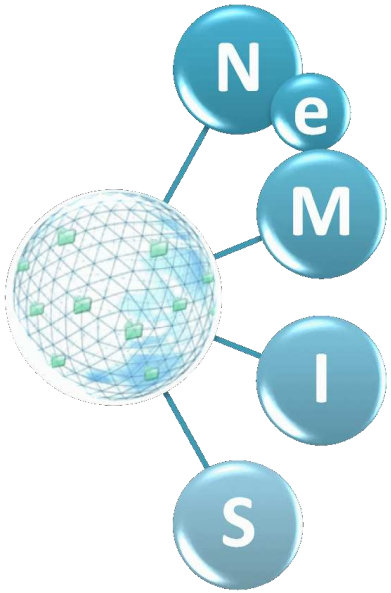

Dir: Fausto Rabitti

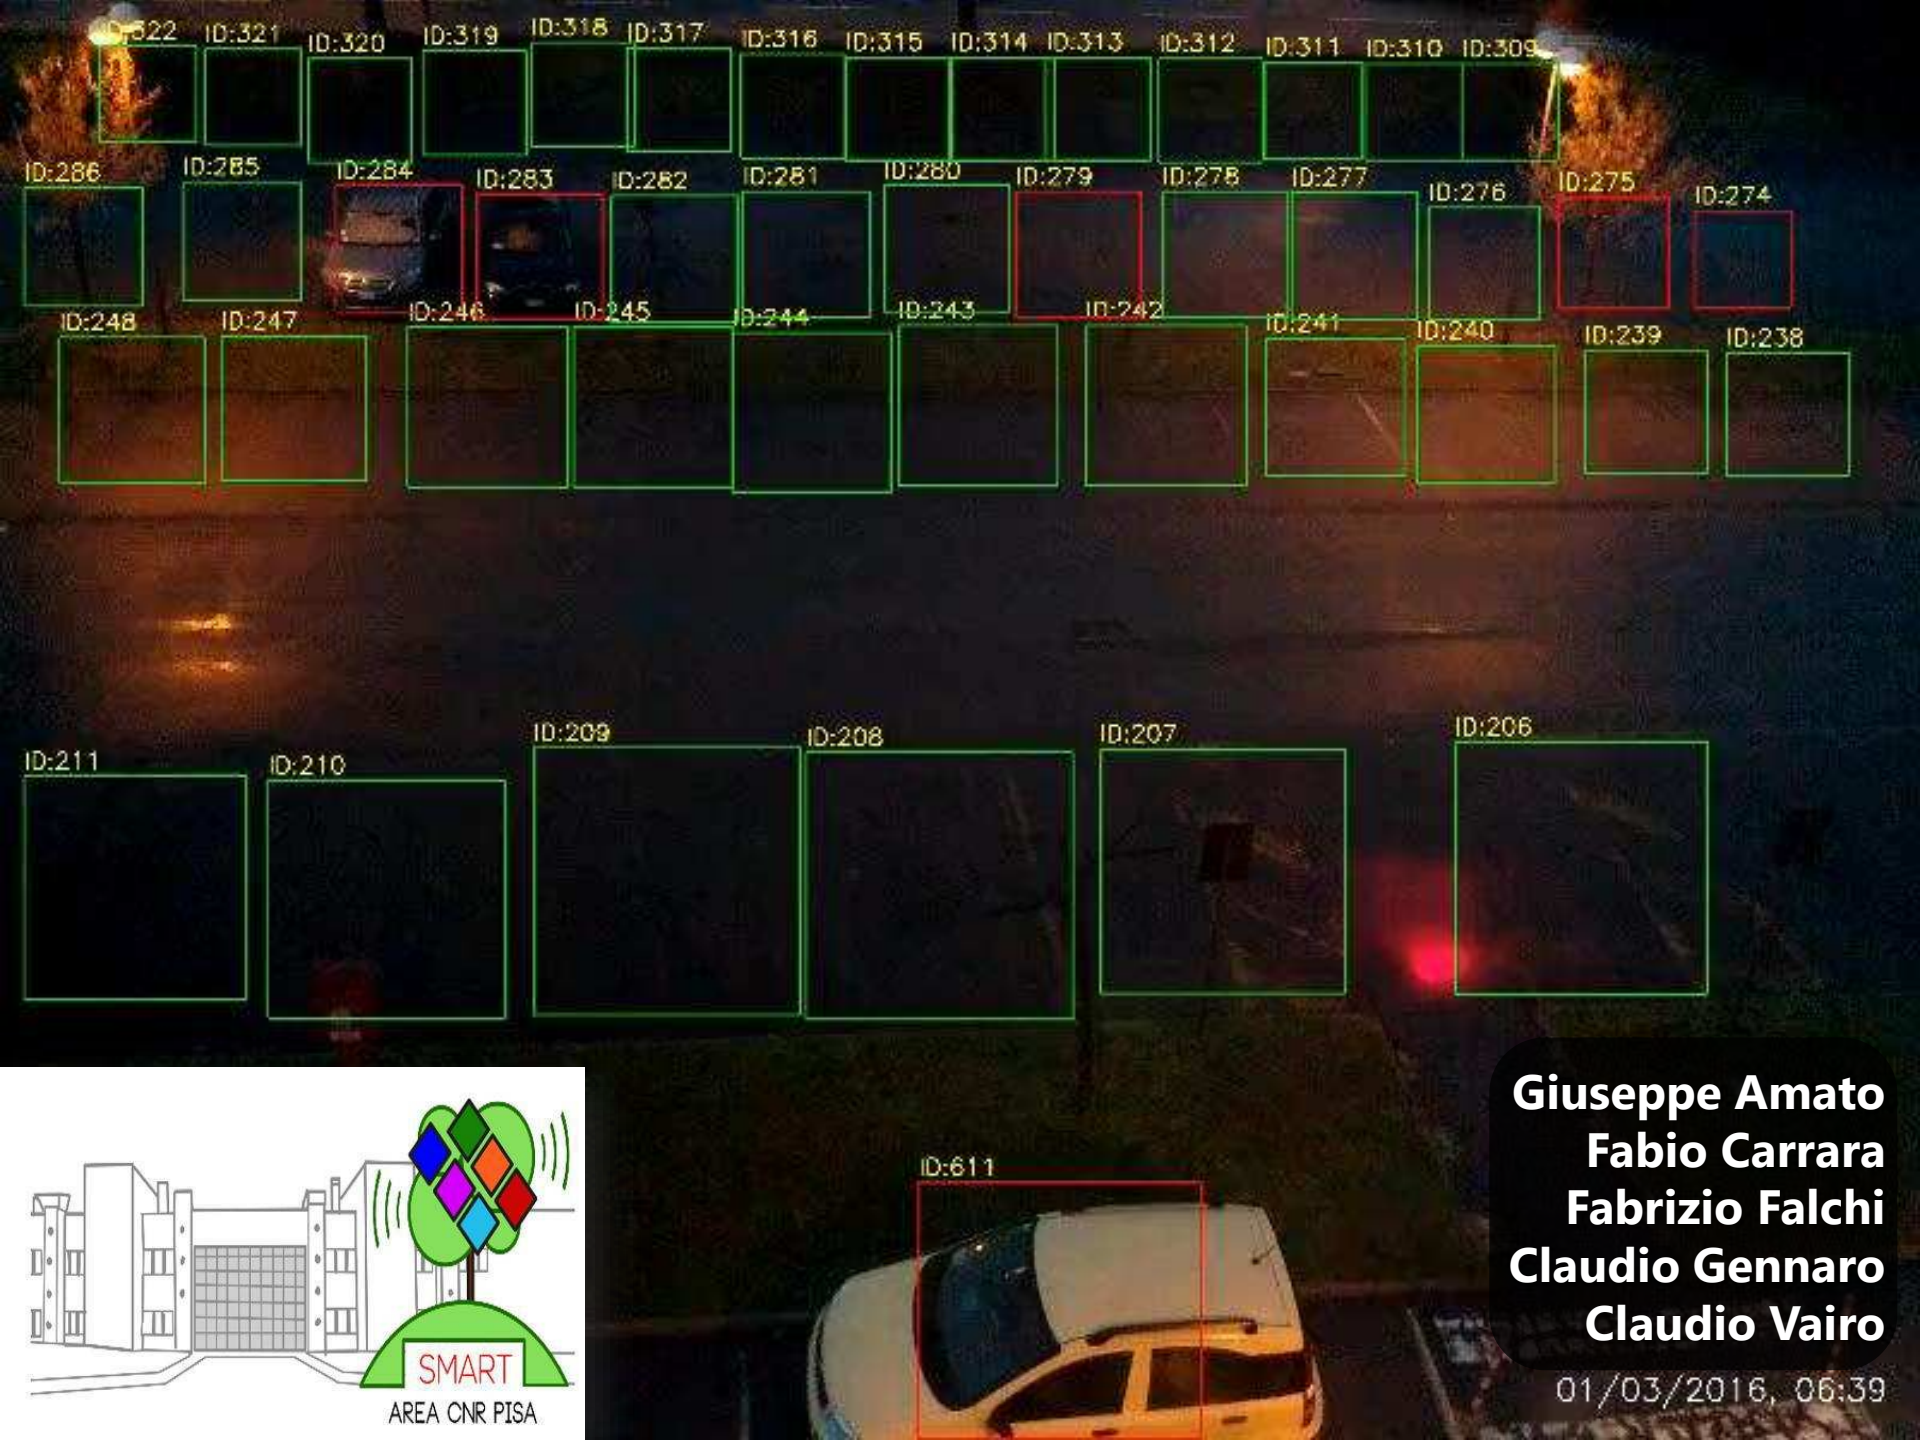
ID:322
ID:321
ID:320
ID:319
ID:318
ID:317
ID:316
ID:315
ID:314
ID:313
ID:312
ID:311
ID:310
ID:309
ID:286
ID:285
ID:284
ID:283
ID:282
ID:281
ID:280
ID:279
ID:278
ID:277
ID:276
ID:275
ID:274
ID:248
ID:247
ID:246
ID:245
ID:244
ID:243
ID:242
ID:241
ID:240
ID:239
ID:238
ID:211
ID:210
ID:209
ID:208
ID:207
ID:206
ID:611
SMART
AREA CNR PISA
Giuseppe Amato
Fabio Carrara
Fabrizio Falchi
Claudio Gennaro
Claudio Vairo
01/03/2016, 06:39

europeana
eagle project

Giuseppe Amato
Fabrizio Falchi
Lucia Vadicamo

pool5, FV256

(pool5→ 1,024), (FV256→4,096)

pool5→2,048

FV256→4,096

VLAD256

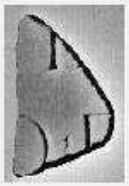

SEARCHING IN 100M PHOTOS

SEARCHING IN 100M PHOTOS

SEARCHING IN 100M PHOTOS

SEARCHING IN 100M PHOTOS

0.000000 0.168457 0.170166 0.185547 0.186279 0.187012 0.190186 0.192627 0.192871 0.192871

0.193359 0.194580 0.194824 0.195801 0.196045 0.196289 0.196533 0.197266 0.197510 0.197510

0.197754 0.198242 0.198486 0.198730 0.198730 0.198730 0.199219 0.199219 0.199463 0.199707

0.200195 0.200195 0.200439 0.200439 0.200684 0.201172 0.201416 0.201416 0.201416 0.201660

0.201660 0.201660 0.201904 0.201904 0.202148 0.202393 0.202393 0.202393 0.202637 0.202637

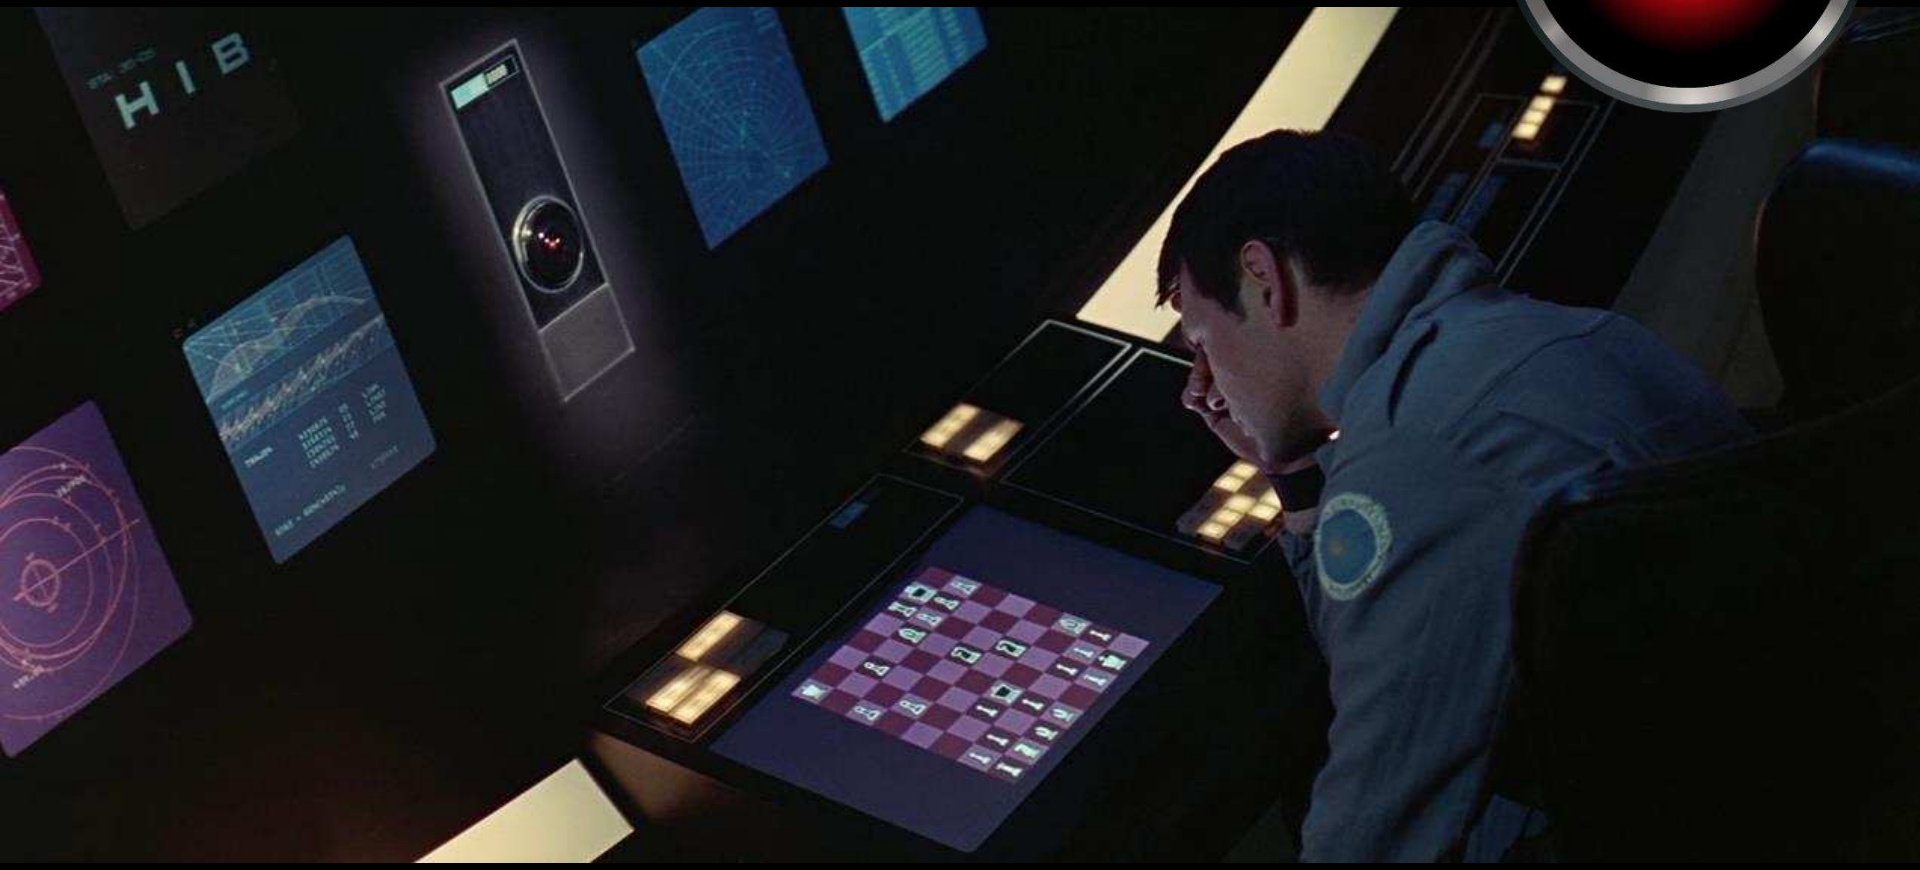
HIB

GRAZIE!

Domande?

Fabrizio Falchi
fabrizio.falchi@cnr.it

DIGITAL INTUITION

nature International weekly journal of science

Archive › Volume 529 › Issue 7587 › Editorial › Article

As shown by its results, the moves that AlphaGo selects are invariably correct. But the interplay of its neural networks means that a ***human can hardly check its working****, or verify its decisions before they are followed through. As the use of deep neural network systems spreads into everyday life — they are already used to analyze and recommend financial transactions — it raises an interesting concept for humans and their relationships with machines. The machine becomes an* ***oracle****; its pronouncements have to be believed.*

Præludium 2

Das Wohltemperierte Klavier I, BWV 847

Johann Sebastian BACH
(1685-1750)

BACH - PRELUDIO

Præludium 2

Das Wohltemperierte Klavier I, BWV 847

Johann Sebastian BACH
(1685-1750)

BACH - PRELUDIO

Staatsbibliothek zu Berlin, D-B Mus. ms. Bach P 202
middle of the 18th century (ca. 1740–1759)

BACH - PRELUDIO

Staatsbibliothek zu Berlin, D-B Mus. ms. P 415

Autograph, 1720-22

BACH - PRELUDIO

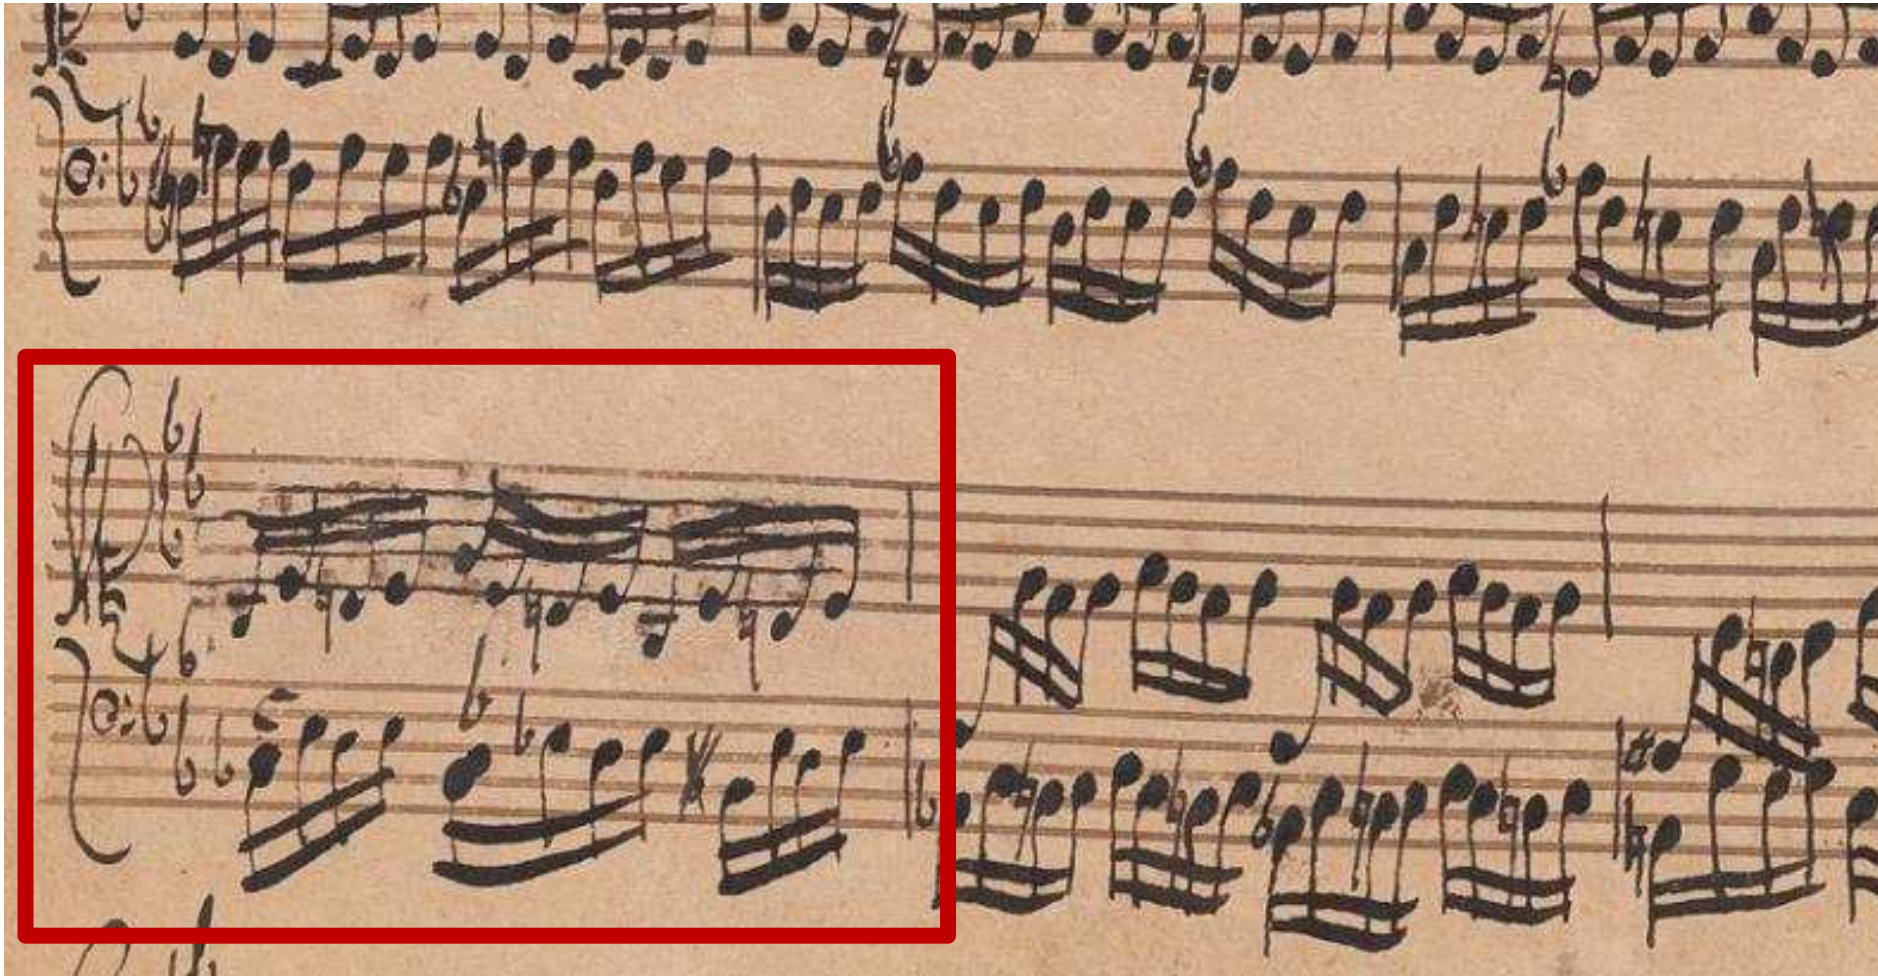

Yale Music Music Deposit 31
W. F. Bach's Klavierbüchlein (1720)


SOFTWARE
&
RESOURCES

GOOGLE TENSORFLOW

- November 9, 2015
- White paper: http://www.tensorflow.org/whitepaper2015.pdf
- API
  - Python (preferred) and C++
  - (Go, Java, JavaScript, Lua R) community is invited to work on this using SWIG
- “Google’s goal with TF seems to be recruiting”
- is about more than deep learning

GOOGLE DEEP LEARNING COURSE

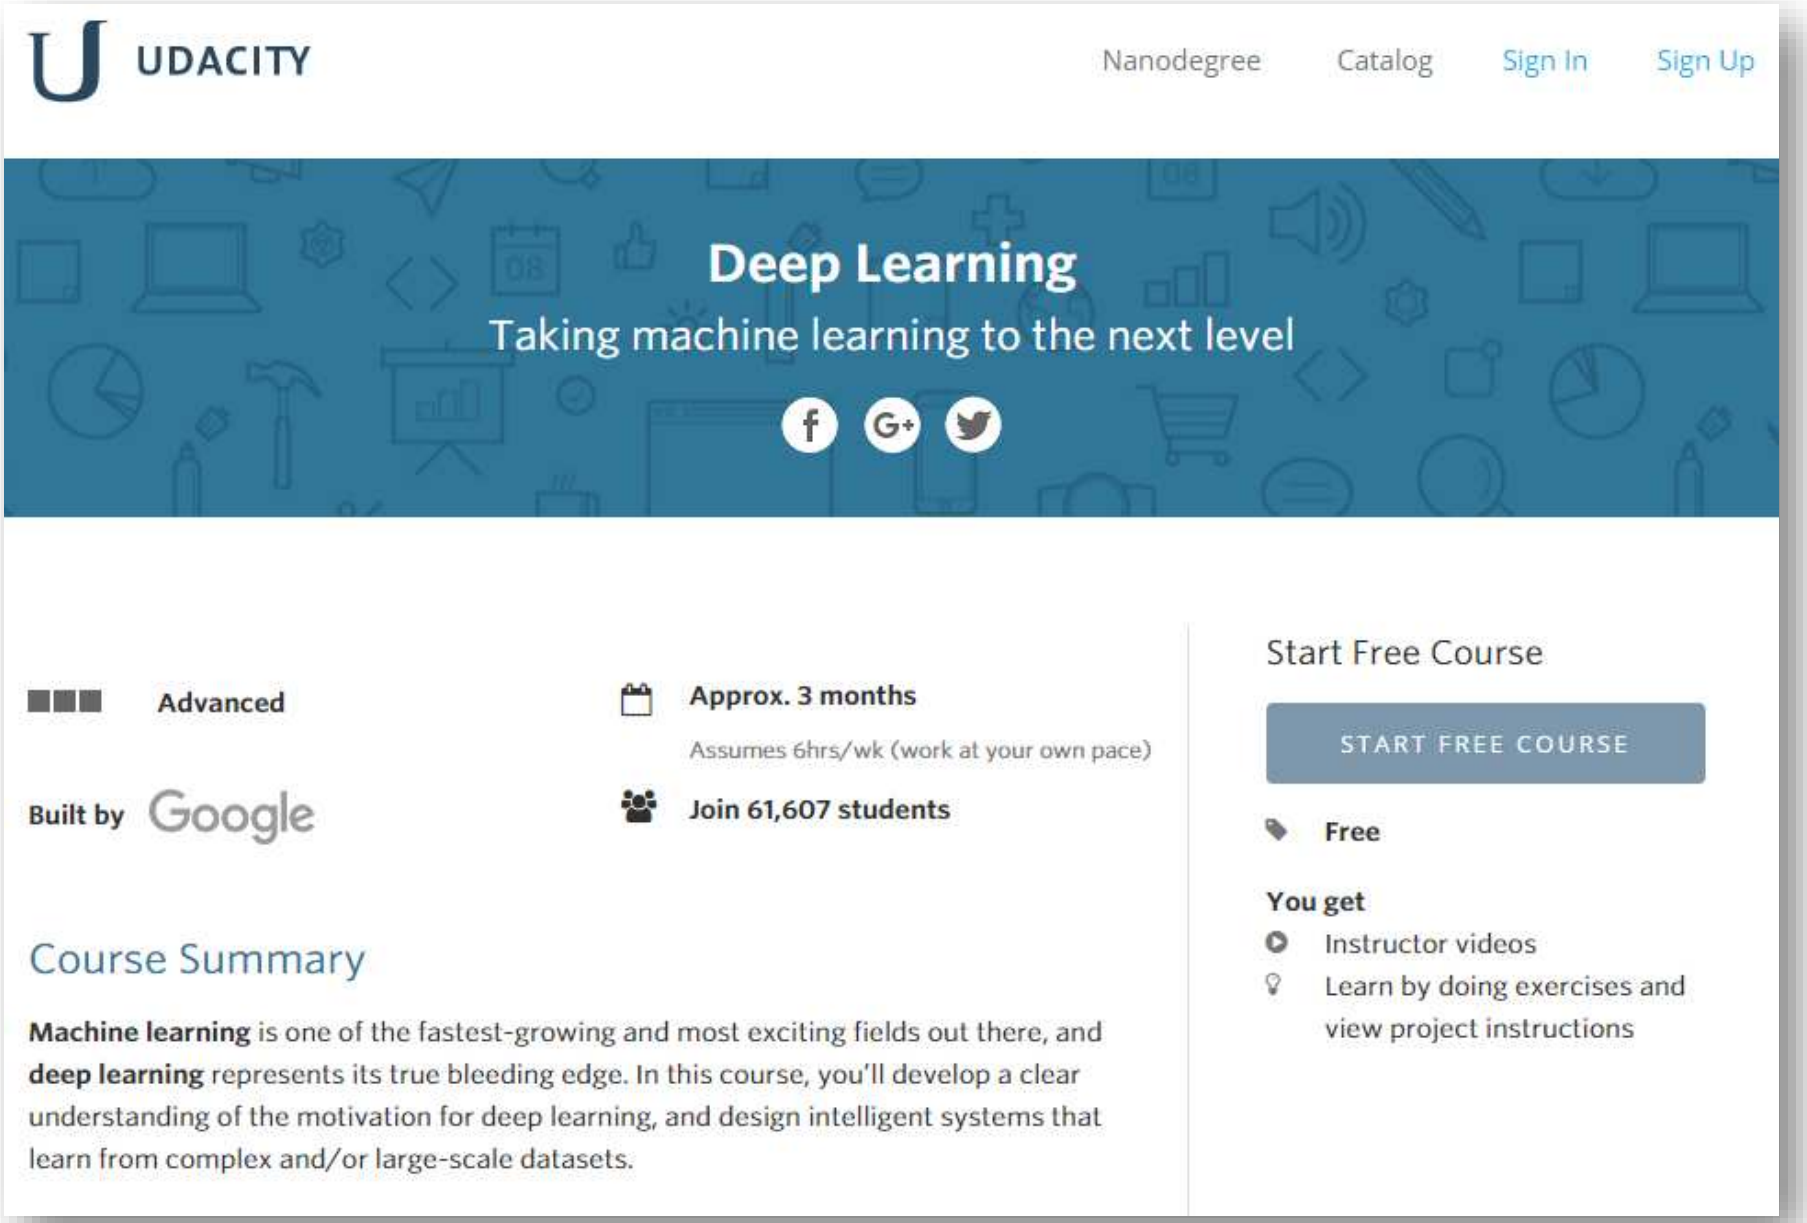

GOOGLE CLOUD VISION API

CLOUD VISION API FEATURES

Derive insight from images with our powerful Cloud Vision API

Label Detection

Detect broad sets of categories within an image, ranging from modes of transportation to animals.

Explicit Content Detection

Detect explicit content like adult content or violent content within an image.

Logo Detection

Detect popular product logos within an image.

Landmark Detection

Detect popular natural and manmade structures within an image.

Image Attributes

Detect general attributes of the image, such as dominant color.

Optical Character Recognition

Detect and extract text within an image, with support for a broad range of languages, along with support for automatic language identification.

Face Detection

Detect multiple faces within an image, along with the associated key facial attributes like emotional state or wearing headwear.

Facial Recognition is not supported.

Integrated REST API

Access via REST API to request one or more annotation types per image. Images can be uploaded in the request or integrated with Google Cloud Storage.

CAFFE, BERKELEY

Caffe

build passing | license BSD

Caffe is a deep learning framework made with expression, speed, and modularity in mind. It is developed by the Berkeley Vision and Learning Center (BVLC) and community contributors.

Check out the project site for all the details like

- DIY Deep Learning for Vision with Caffe
- Tutorial Documentation
- BVLC reference models and the community model zoo
- Installation instructions

and step-by-step examples.

gitter join chat

Please join the caffe-users group or gitter chat to ask questions and talk about methods and models. Framework development discussions and thorough bug reports are collected on Issues.

Happy brewing!

License and Citation

Caffe is released under the BSD 2-Clause license. The BVLC reference models are released for unrestricted use.

NVIDIA® DIGITS™ – Interactive Deep Learning GPU Training System

The NVIDIA Deep Learning GPU Training System (DIGITS) puts the power of deep learning in the hands of data scientists and researchers. Quickly design the best deep neural network (DNN) for your data using real-time network behavior visualization. Best of all, DIGITS is a complete system so you don't have to write any code. Get started with DIGITS in under an hour.

> *DIGITS makes it way easier to design the best network for the job. The DIGITS interface makes it super easy to track key diagnostics during training. The field will definitely benefit from having tools like this for configuration and introspection.*
>
> — Simon Osindero, AI Architect at Flickr

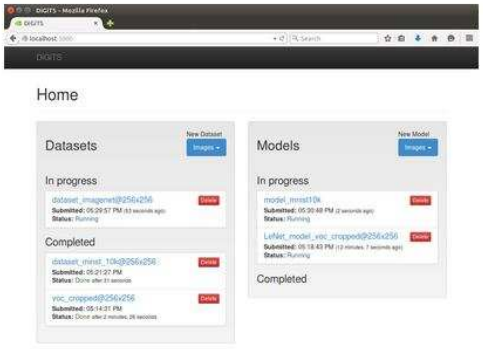

DIGITS training multiple DNNs and building a new dataset in parallel.

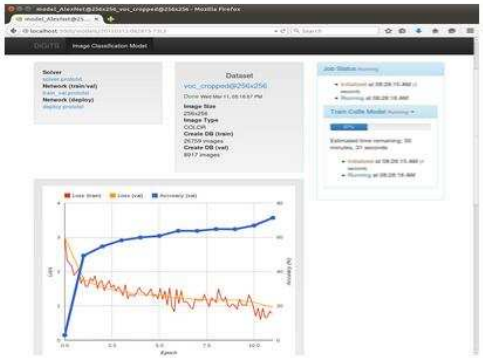

DIGITS real-time monitoring of DNN model training

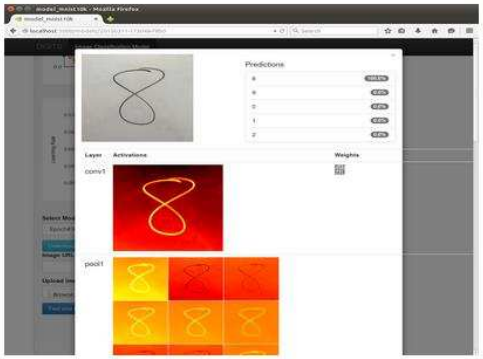

DIGITS visualizing a test image

MICROSOFT CNTK

- January 25, 2016 (academic licens since April 2015)
- one of the advantages of CNTK is its ability to run on a single core, as well as on a large cluster of GPU-based machines
- C++

QUALCOMM SNAPDRAGON 820 AND 80A

- January 2016
- Snapdragon 820 is a phone
- Snapdragon 820A is a chip for cars and drones

MIT EYERISS GPU

- 10 times as efficient as a mobile GPU could allow mobile devices to run artificial intelligence algorithms locally, without a major drain to the phone's battery
- “This work is very important, showing how embedded processors for deep learning can provide power and performance optimizations that will bring these complex computations ***from the cloud to mobile devices***.”, Mike Polley
- Funded by Defense Advanced Research Projects Agency (**DARPA**)

AlphaGo

nature
International weekly journal of science

NATURE | NEWS

Google AI algorithm masters ancient game of Go

Deep-learning software defeats human professional for first time.

Elizabeth Gibney

27 January 2016

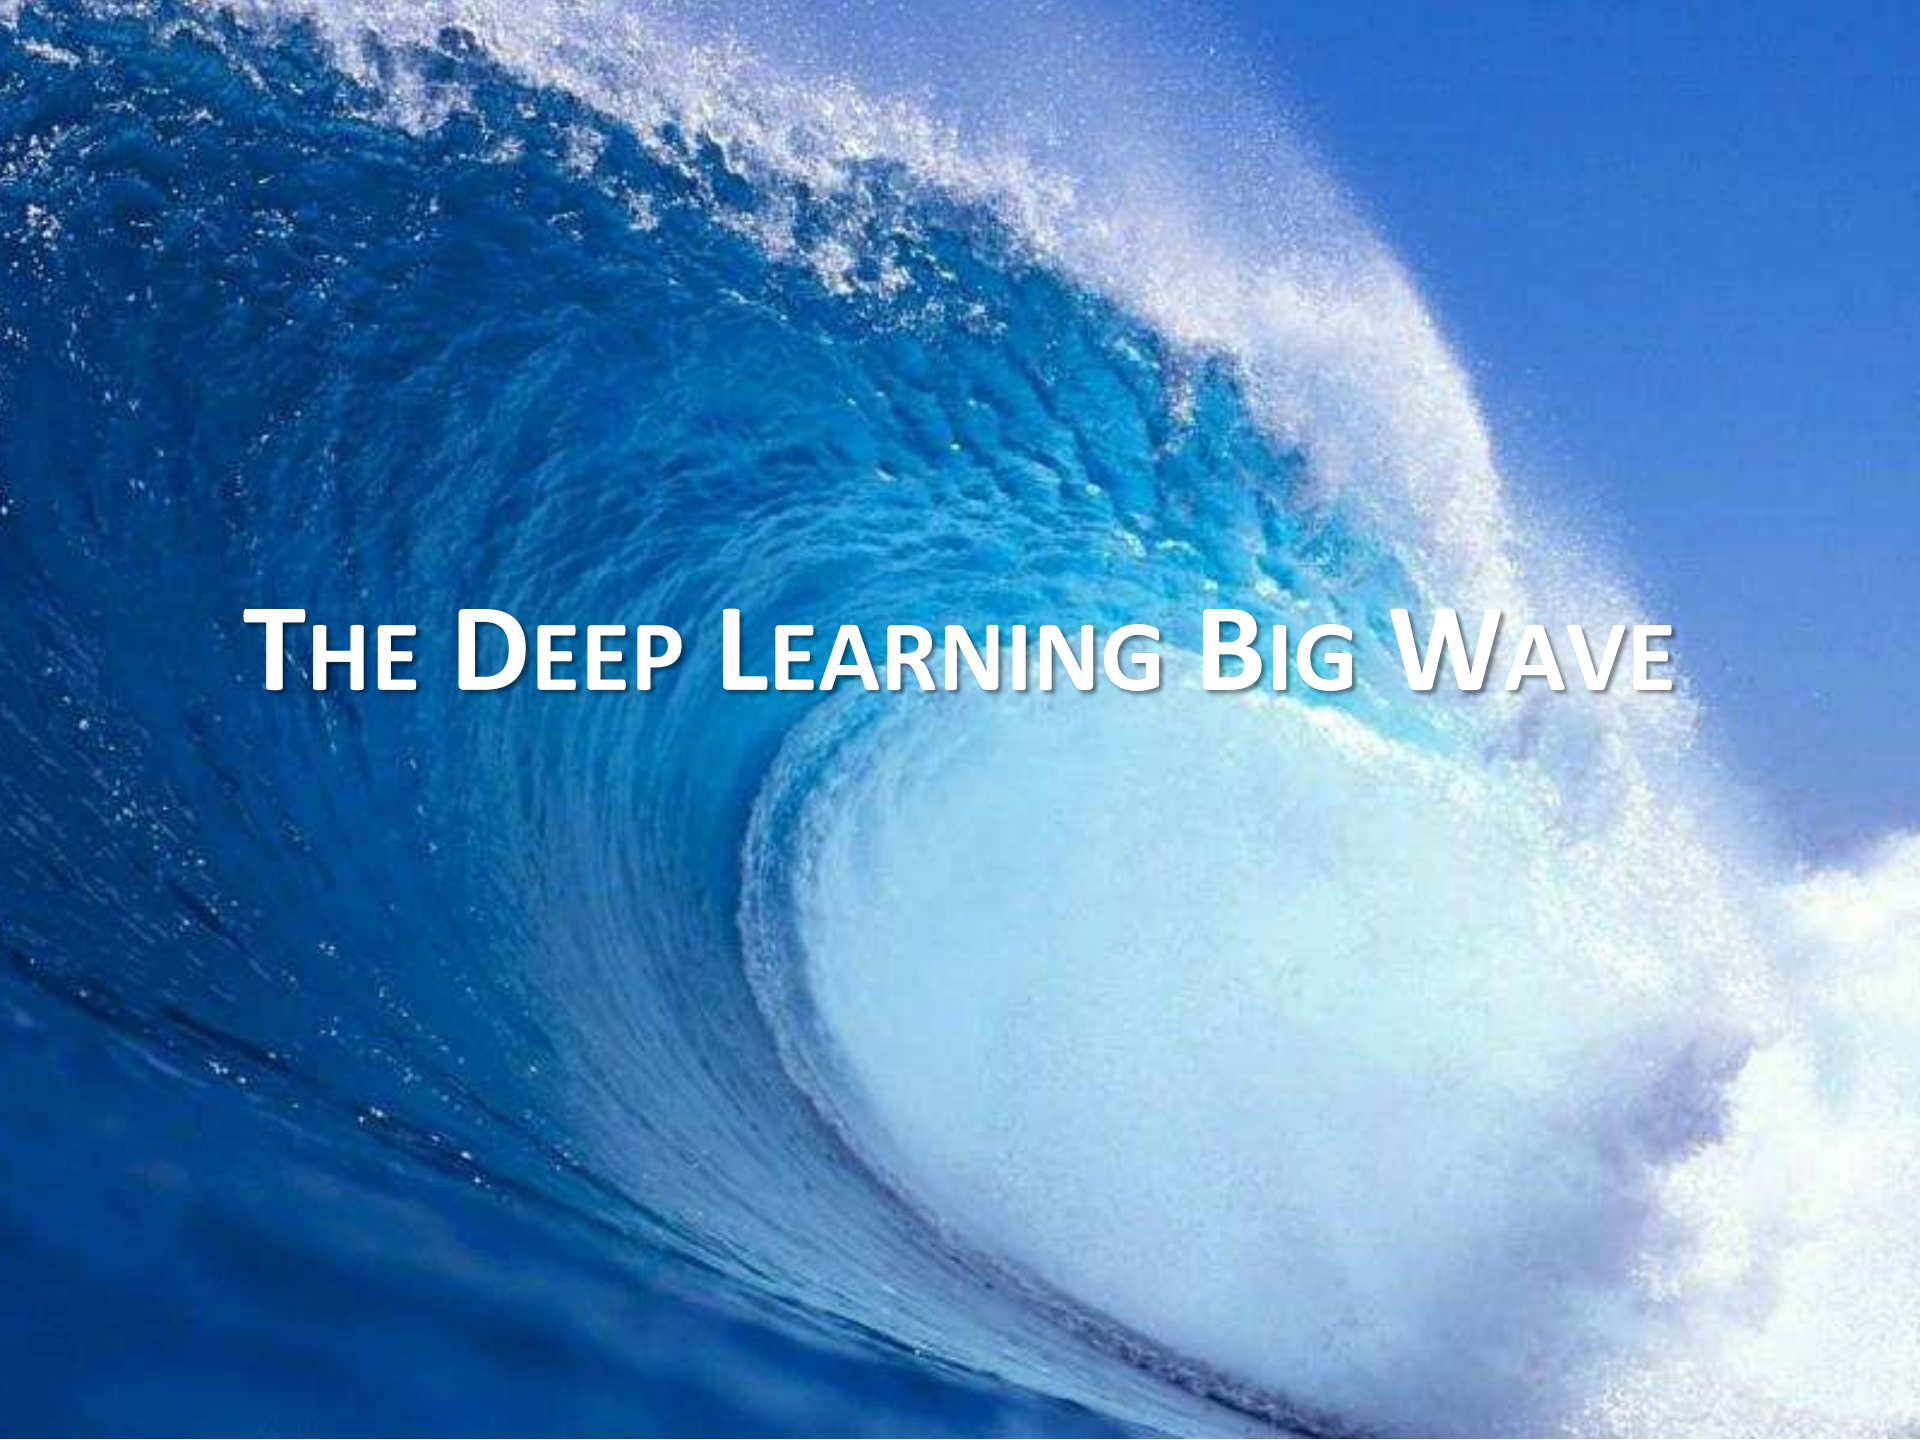
The Deep Learning Big Wave

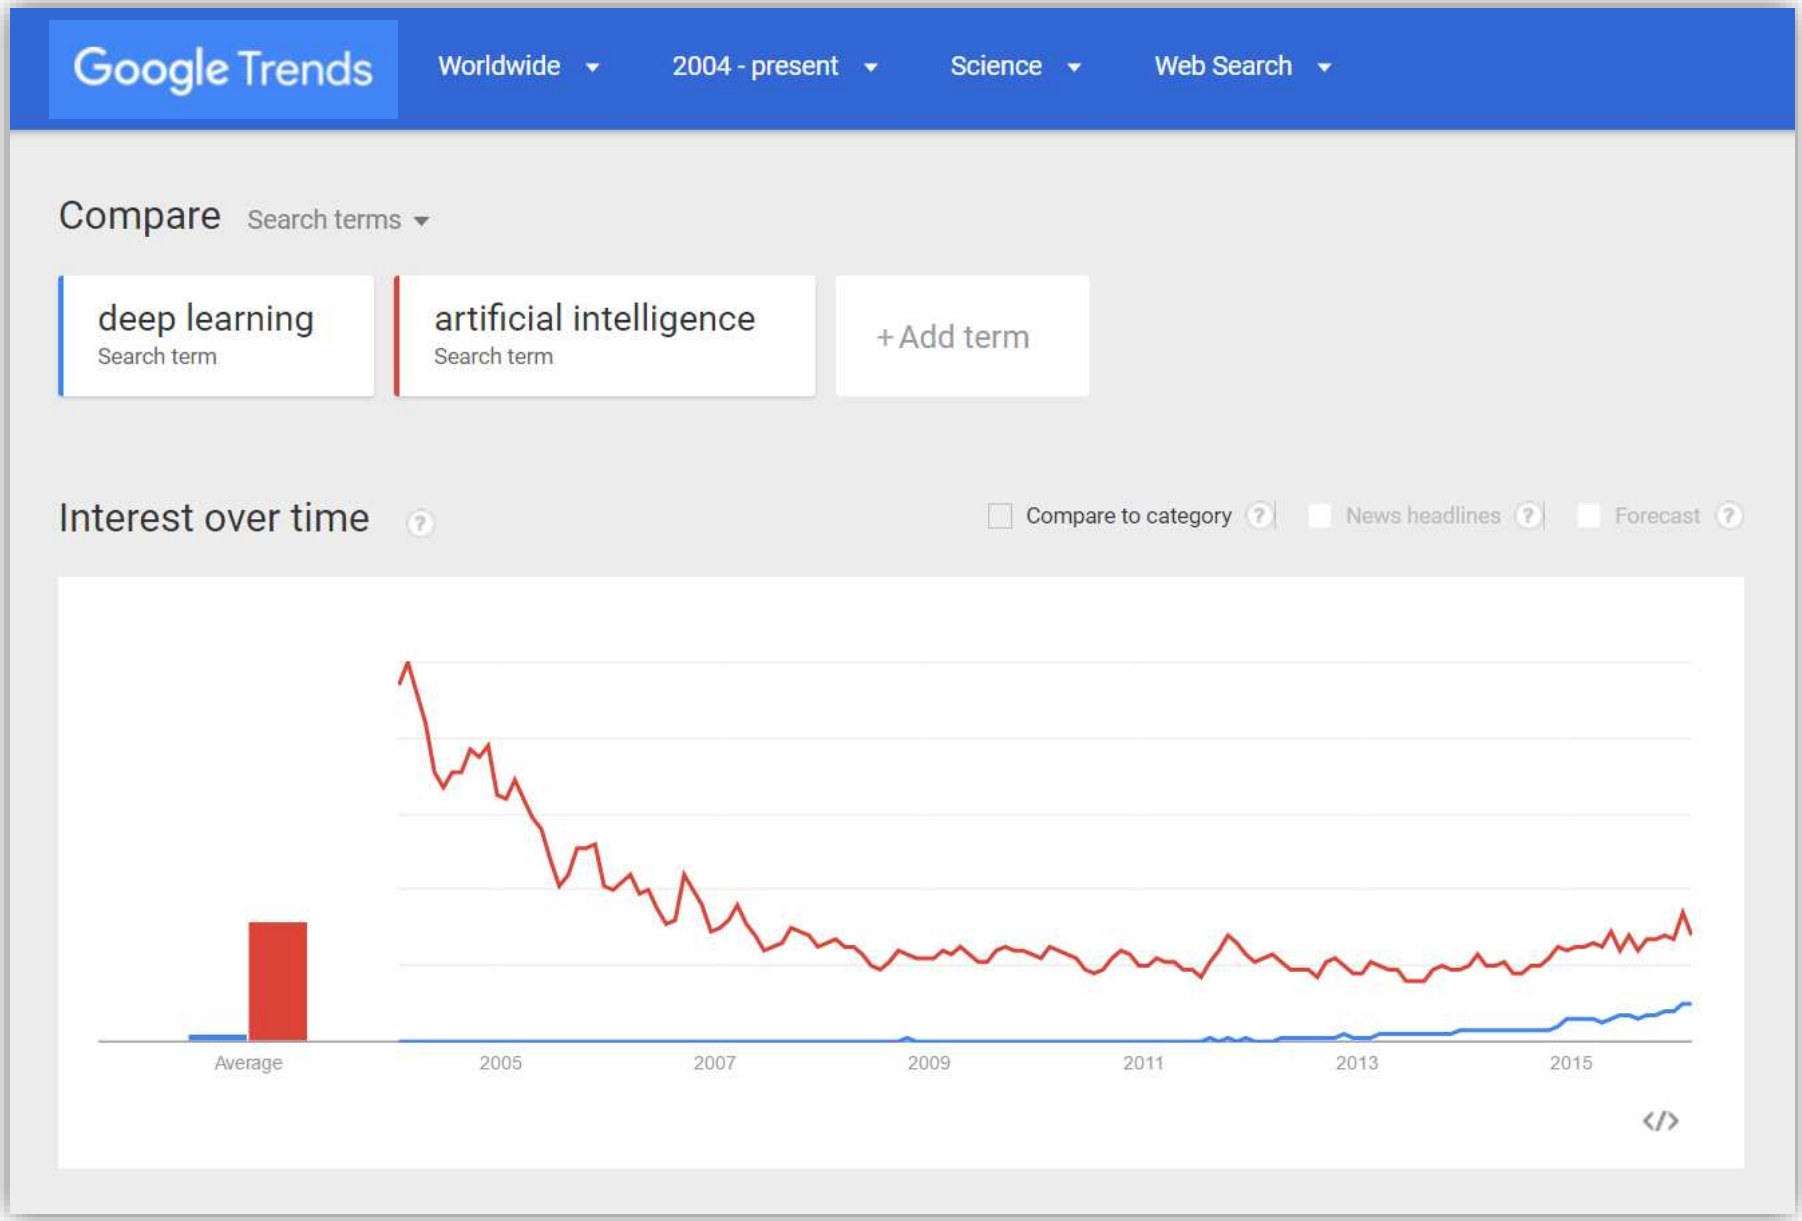
Google Trends
Worldwide
2004 - present
Science
Web Search
Compare
Search terms
deep learning
Search term
artificial intelligence
Search term
+ Add term
Interest over time
Compare to category
News headlines
Forecast
Average
2005
2007
2009
2011
2013
2015

Google Trends in Science - Web Search

Worldwide ▾ Jan. 2012 - Feb. 2016 ▾ Science ▾ Web Search ▾

Compare Search terms ▾

Deep Lear...	Big Data	Computer ...	Multimedia	+ Add term
Search term	Search term	Search term	Search term	

Interest over time

Compare to category | News headlines | Forecast

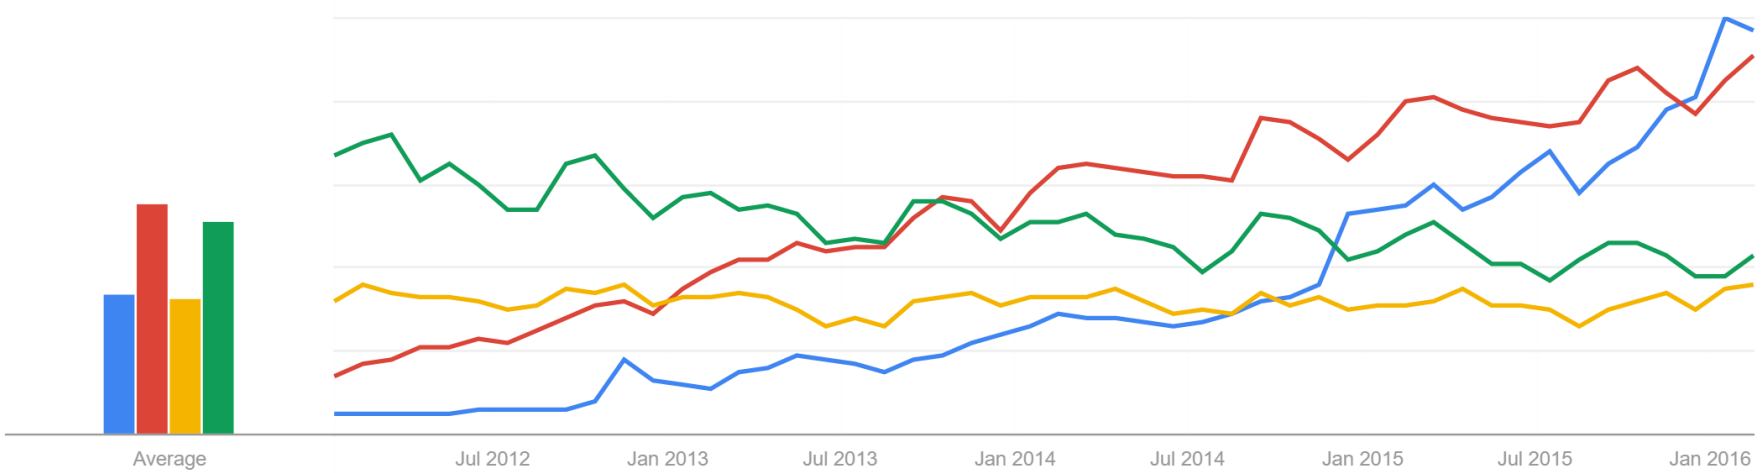

HORIZON 2020

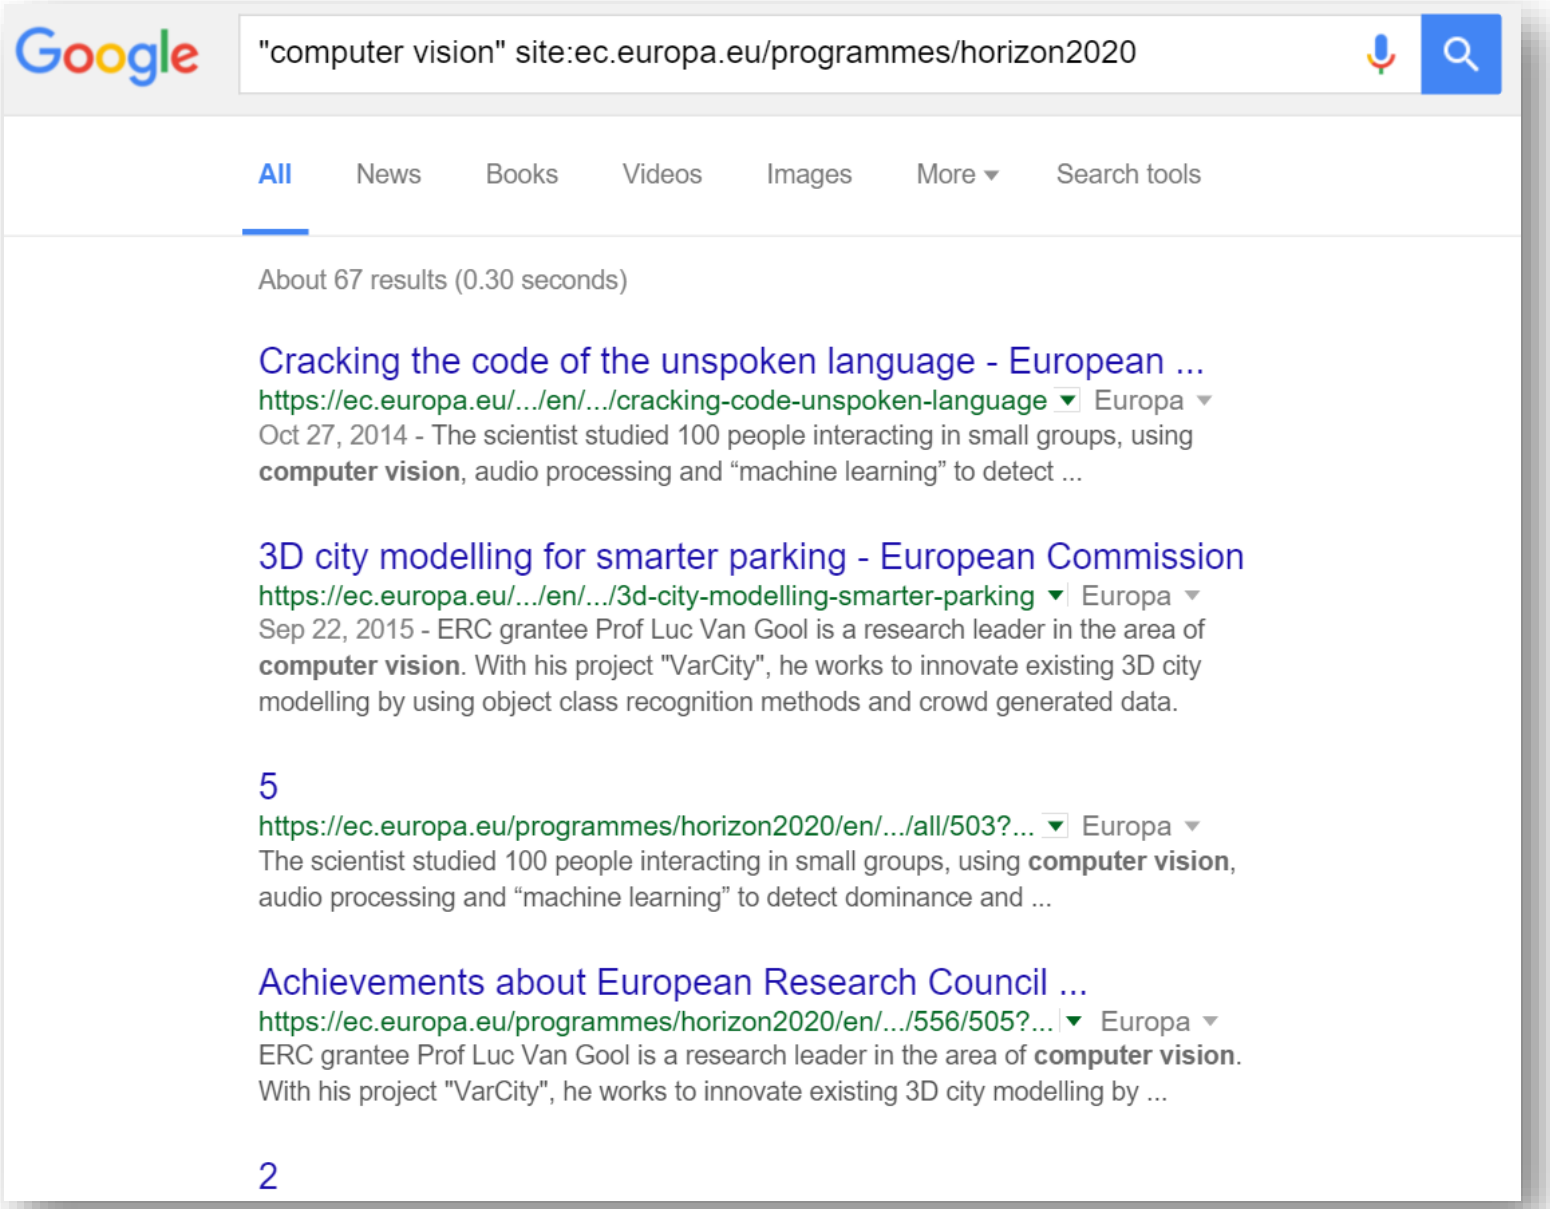

HORIZON 2020

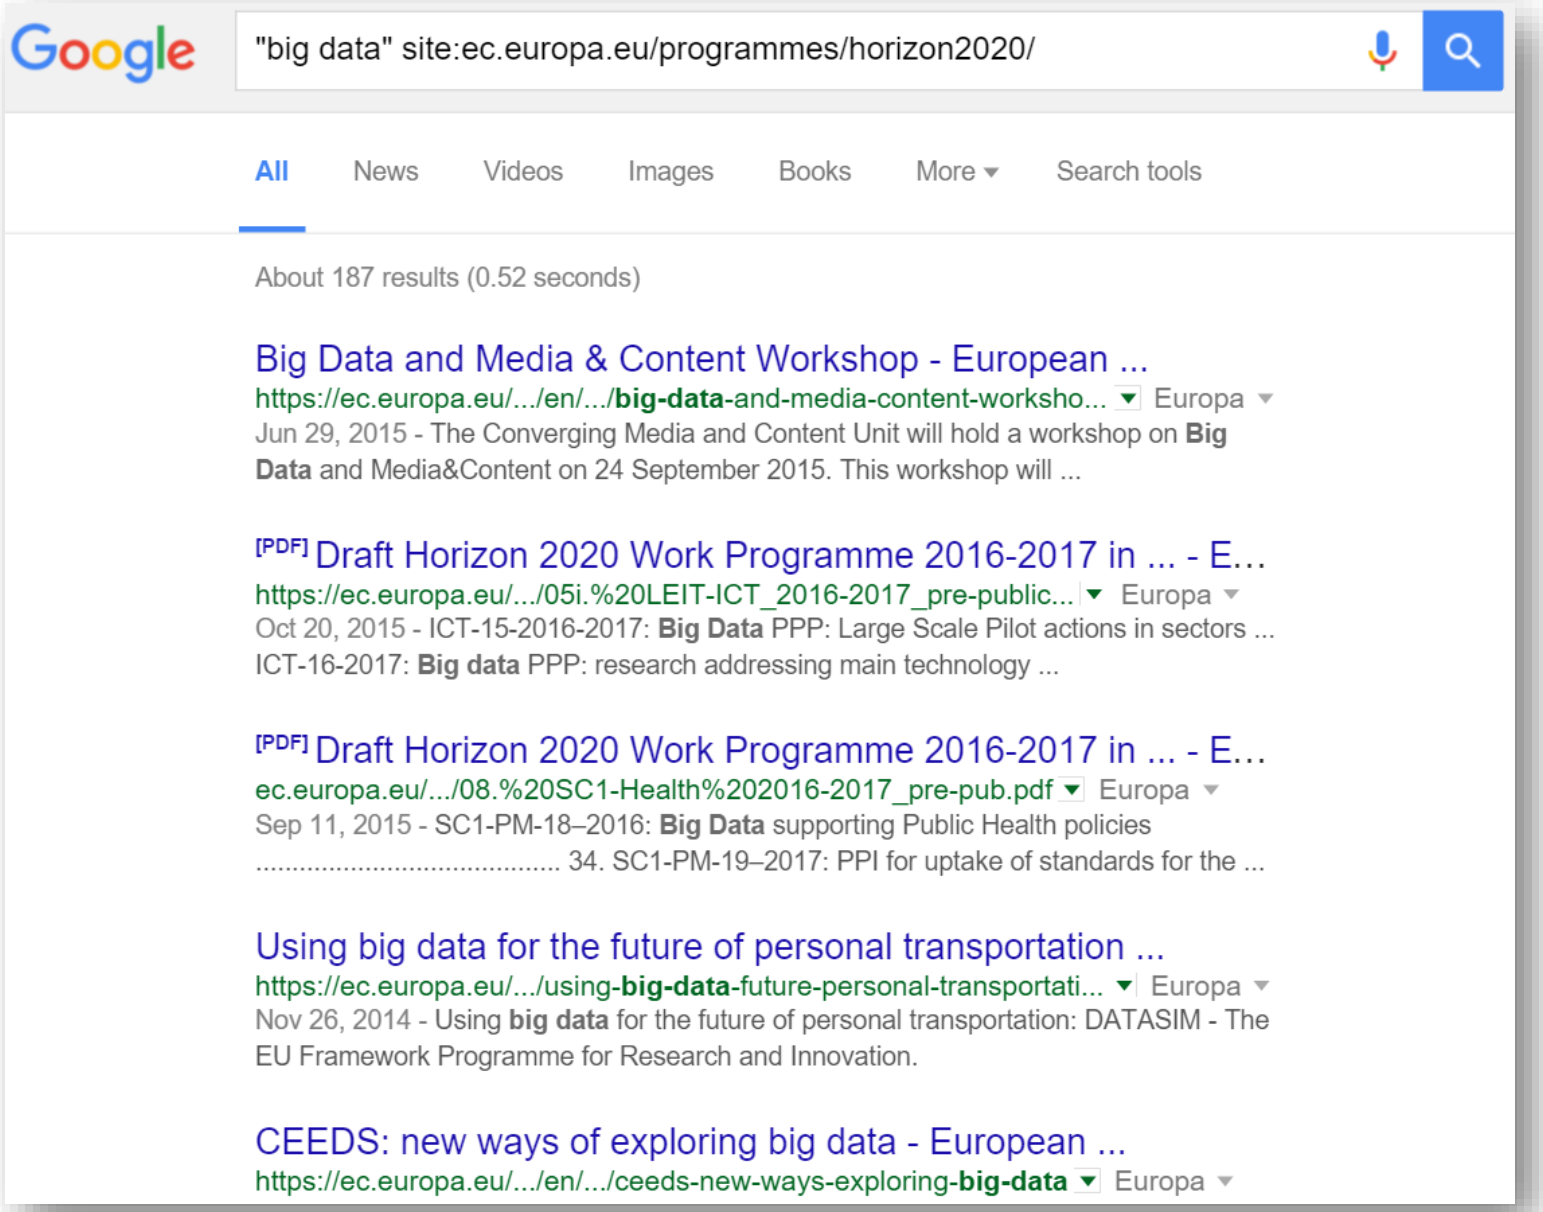

HORIZON 2020

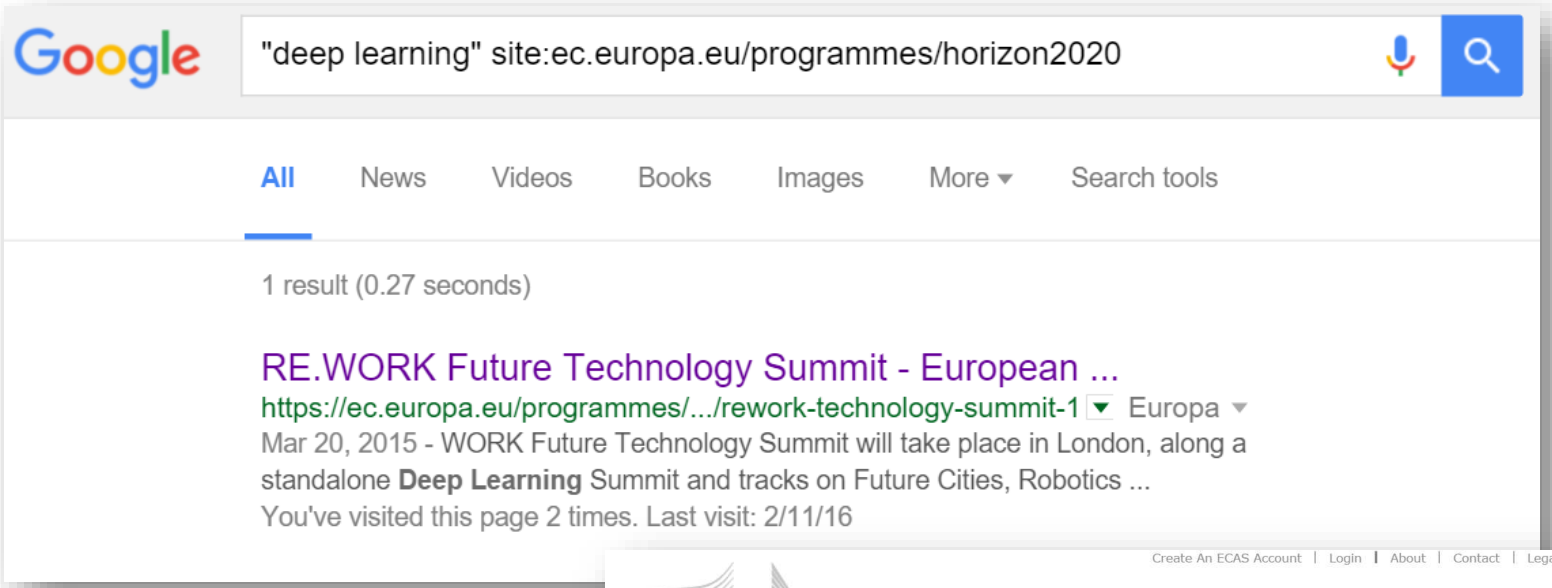

DEEP DREAM HTTP://DEEPDREAMGENERATOR.COM/

- Instead of exactly prescribing which feature we want the network to amplify, we can also let the network make that decision.
- In this case we simply feed the network an arbitrary image or photo and let the network analyze the picture.
- **We then pick a layer and ask the network to enhance whatever it detected.**
- Each layer of the network deals with features at a different level of abstraction, so the complexity of features we generate depends on which layer we choose to enhance.
- For example, lower layers tend to produce strokes or simple ornament-like patterns, because those layers are sensitive to basic features such as edges and their orientations.

Inceptionism: Going Deeper into Neural Networks
A. Mordvintsev, C. Olah, M. Tyka, Wednesday, June 17, 2015

Google Research Blog

BENGIO ON "WHY DEEP LEARNING WORKS"

Quora

- We understand that distributed representations, depth, and elements of the convolutional architecture and recurrent architectures correspond to
  - preferences in the space of functions (or informally, **priors**)
- and we have theory explaining why some of these preferences can buy an important (sometimes exponential) **statistical advantage** (in the sense of needing less data to achieve some level of accuracy);
- more details in my book, pointing to some of the recent papers, my favourite being
  - "On the Number of Linear Regions of Deep Neural Networks" (NIPS'2014).

BENGIO ON "WHY DEEP LEARNING WORKS"

Quora

- We understand better why the **optimization** problem involved in training deep networks is probably not as intractable as was previously believed, in the sense that the **vast majority of local minima** would actually correspond to very **good solutions**. Two good papers on the subject:
  - Identifying and attacking the saddle point problem in high-dimensional non-convex optimization (NIPS'2014)
  - The loss surface of multilayer networks (AISTATS'2015).

Neural network architecture. The input to the policy network is a $19 \times 19 \times 48$ image stack consisting of 48 feature planes. The first hidden layer zero pads the input into a 23×23 image, then convolves k filters of kernel size 5×5 with stride 1 with the input image and applies a rectifier nonlinearity. Each of the subsequent hidden layers 2 to 12 zero pads the respective previous hidden layer into a 21×21 image, then convolves k filters of kernel size 3×3 with stride 1, again followed by a rectifier nonlinearity. The final layer convolves 1 filter of kernel size 1×1 with stride 1, with a different bias for each position, and applies a softmax function. The match version of AlphaGo used $k = 192$ filters; Fig. 2b and Extended Data Table 3 additionally show the results of training with $k = 128$, 256 and 384 filters.

Deep learning (from Nature)

An **image**, for example, comes in the form of an array of pixel values, and

1. the learned features in the first layer of representation typically represent the presence or absence of **edges** at particular orientations and locations in the image.
2. The second layer typically detects **motifs** by spotting particular arrangements of edges, regardless of small variations in the edge positions.
3. The third layer may assemble motifs into larger combinations that correspond to **parts** of familiar objects, and subsequent layers would detect objects as combinations of these parts.
4. The key aspect of deep learning is that these layers of features are not designed by human engineers: they are learned from data using a **general-purpose learning** procedure.

DIGITAL INTUITION

nature International weekly journal of science

NATURE | EDITORIAL

Digital intuition

A computer program that can outplay humans in the abstract game of Go will redefine our relationship with machines.

27 January 2016