



Introduction to Deep Learning

Princeton University COS 495

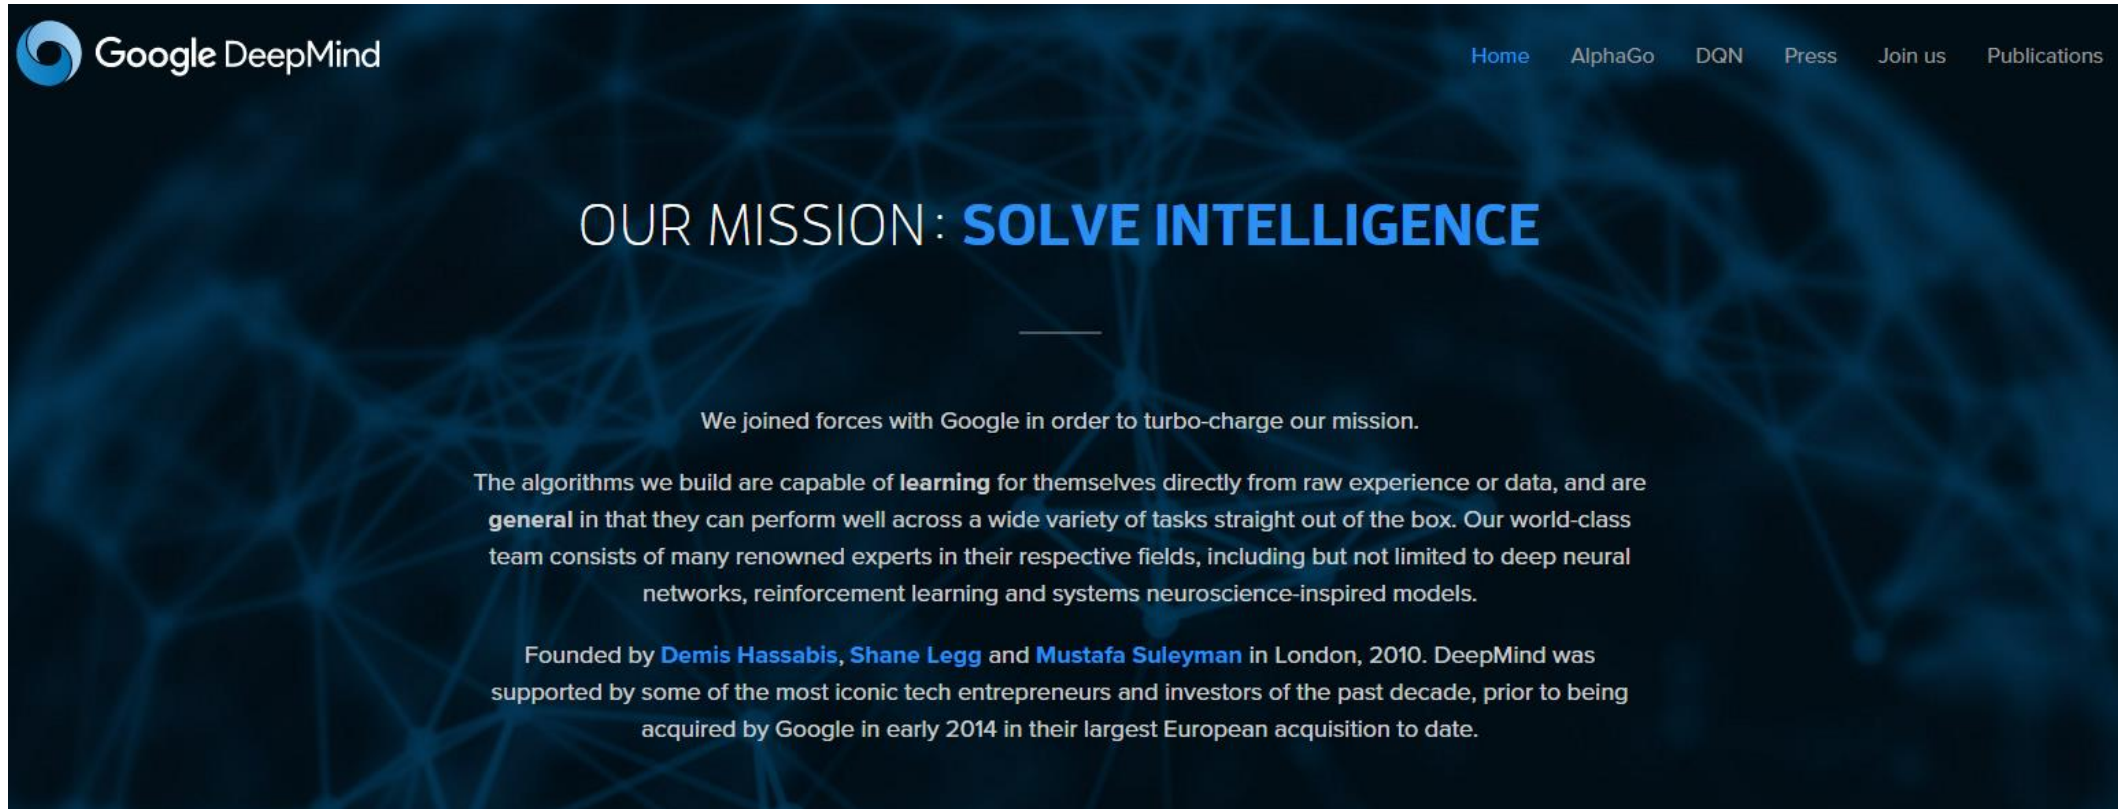
Instructor: Yingyu Liang

What is deep learning?

- Short answer: recent buzz word

Industry

- Google



The image shows a screenshot of the Google DeepMind website homepage. The background is dark blue with a faint, glowing neural network pattern. In the top left corner, the Google DeepMind logo is displayed. In the top right corner, there is a navigation menu with links for Home, AlphaGo, DQN, Press, Join us, and Publications. The main heading in the center reads "OUR MISSION: SOLVE INTELLIGENCE". Below this, a paragraph states: "We joined forces with Google in order to turbo-charge our mission." This is followed by a paragraph describing the algorithms: "The algorithms we build are capable of learning for themselves directly from raw experience or data, and are general in that they can perform well across a wide variety of tasks straight out of the box. Our world-class team consists of many renowned experts in their respective fields, including but not limited to deep neural networks, reinforcement learning and systems neuroscience-inspired models." The final paragraph mentions the founders: "Founded by Demis Hassabis, Shane Legg and Mustafa Suleyman in London, 2010. DeepMind was supported by some of the most iconic tech entrepreneurs and investors of the past decade, prior to being acquired by Google in early 2014 in their largest European acquisition to date."

Google DeepMind

Home AlphaGo DQN Press Join us Publications

OUR MISSION: **SOLVE INTELLIGENCE**

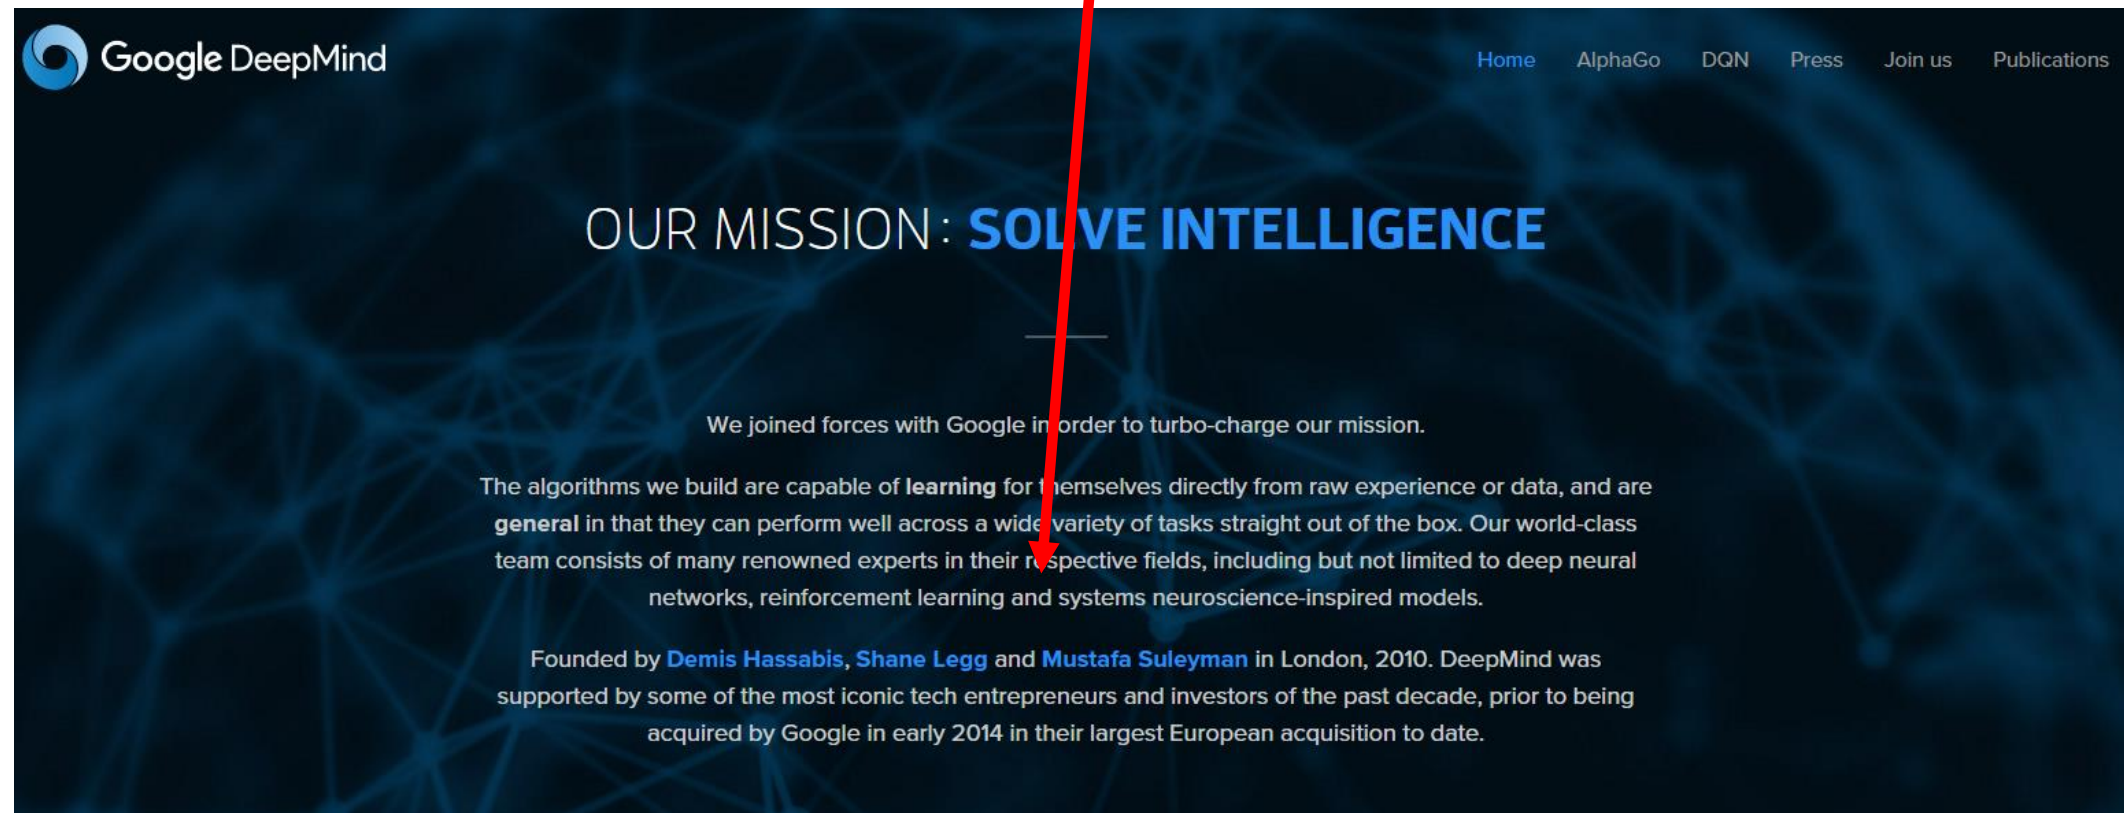
We joined forces with Google in order to turbo-charge our mission.

The algorithms we build are capable of **learning** for themselves directly from raw experience or data, and are **general** in that they can perform well across a wide variety of tasks straight out of the box. Our world-class team consists of many renowned experts in their respective fields, including but not limited to deep neural networks, reinforcement learning and systems neuroscience-inspired models.

Founded by **Demis Hassabis**, **Shane Legg** and **Mustafa Suleyman** in London, 2010. DeepMind was supported by some of the most iconic tech entrepreneurs and investors of the past decade, prior to being acquired by Google in early 2014 in their largest European acquisition to date.

team consists of many renowned experts in their respective fields, including but not limited to deep neural networks, reinforcement learning and systems neuroscience-inspired models.

- Google



Google DeepMind

Home AlphaGo DQN Press Join us Publications

OUR MISSION: SOLVE INTELLIGENCE

We joined forces with Google in order to turbo-charge our mission.

The algorithms we build are capable of **learning** for themselves directly from raw experience or data, and are **general** in that they can perform well across a wide variety of tasks straight out of the box. Our world-class team consists of many renowned experts in their respective fields, including but not limited to deep neural networks, reinforcement learning and systems neuroscience-inspired models.

Founded by **Demis Hassabis**, **Shane Legg** and **Mustafa Suleyman** in London, 2010. DeepMind was supported by some of the most iconic tech entrepreneurs and investors of the past decade, prior to being acquired by Google in early 2014 in their largest European acquisition to date.

Industry

- Facebook

Facebook AI Research (FAIR)

Home Publications People Research Downloads Blog

Teaching Machines To See and Understand
by Ari Entin about 2 months ago
Facebook AI Research (FAIR)

Highlights

Teaching Machines To See and Understand
by Ari Entin about 2 months ago
Blog post

Simple bag-of-words baseline for visual question answering
by Bolei Zhou, Yuandong Tian, Sainbayar Sukhbaatar, Arthur Szlam, Rob Fergus about 2 months ago
Publication

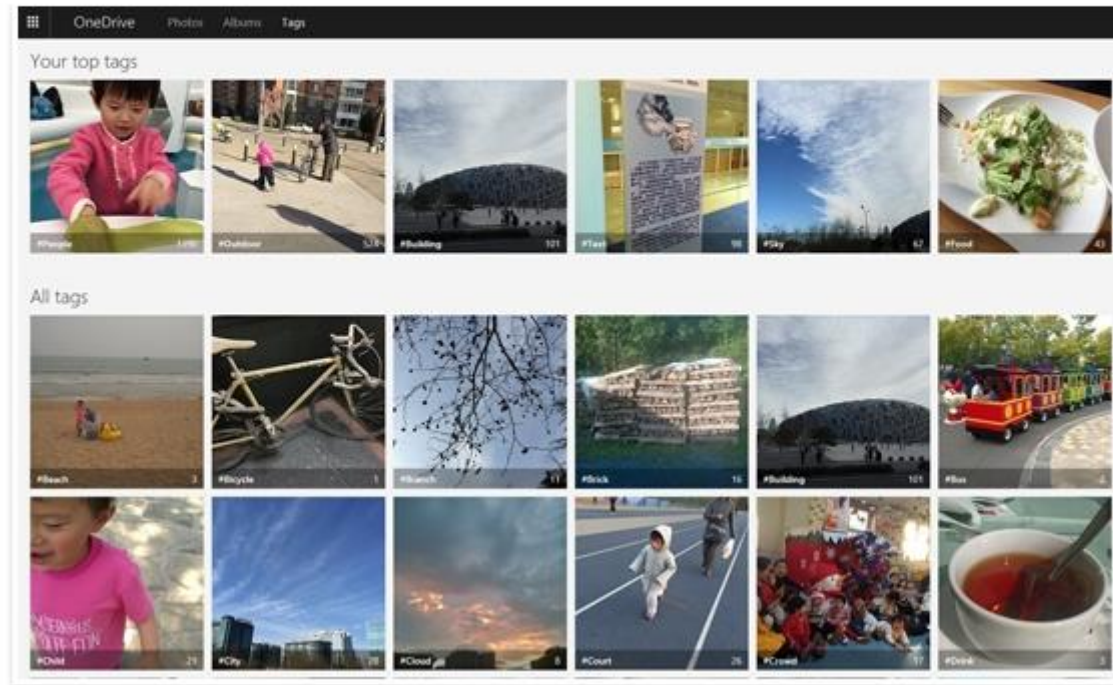
A Roadmap towards Machine Intelligence
by Tomas Mikolov, Armand Joulin, Marco Baroni about 2 months ago
Publication

MazeBase: A Sandbox for Learning from Games
by Sainbayar Sukhbaatar, Arthur Szlam, Gabriel Synnaeve, Soumith Chintala, Rob Fergus about 2

Industry

- Microsoft

Microsoft Researchers' Algorithm Sets ImageNet Challenge Milestone



Industry

- Elon Musk



Elon Musk @elonmusk · 3m
Announcing formation of @open_ai ...
openai.com/blog/introduci...

OpenAI

OpenAI

OpenAI is a non-profit artificial intelligence group.
openai.com

Forbes / Tech

Top 20 Stocks for 2016

DEC 11, 2015 @ 05:04 PM 4,715 VIEWS

Elon Musk And Peter Thiel Launch OpenAI, A Non-Profit Artificial Intelligence Research Company

Industry

- Toyota

SEARCH

The New York Times

TECHNOLOGY

Toyota Invests \$1 Billion in Artificial Intelligence


By JOHN MARKOFF NOV. 6, 2015




Gill Pratt, a roboticist who will oversee Toyota's new research laboratory in the United States, at a news conference Friday in Tokyo. Yuya Shino/Reuters

Academy

- NIPS 2015: ~4000 attendees, double the number of NIPS 2014



Tutorial: Deep Learning



Academy

- Science special issue
- Nature invited review

REVIEW

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{4,5}



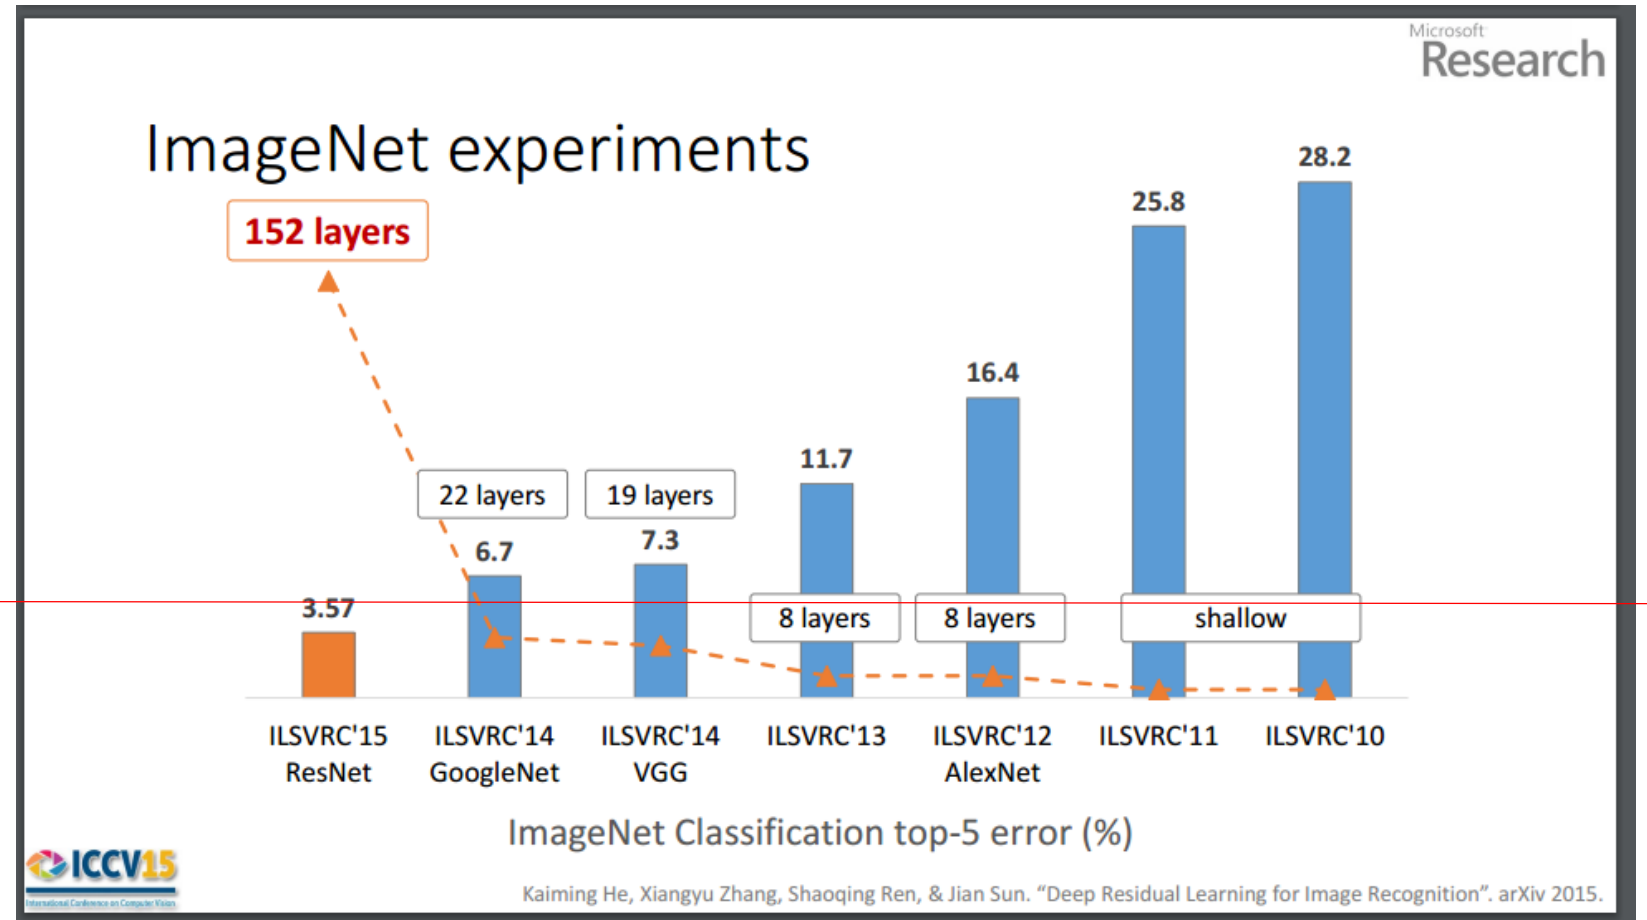
What is deep learning?

- Longer answer: machine learning framework that shows impressive performance on many Artificial Intelligence tasks

Image

- Image classification
 - 1000 classes

Human performance: ~5%

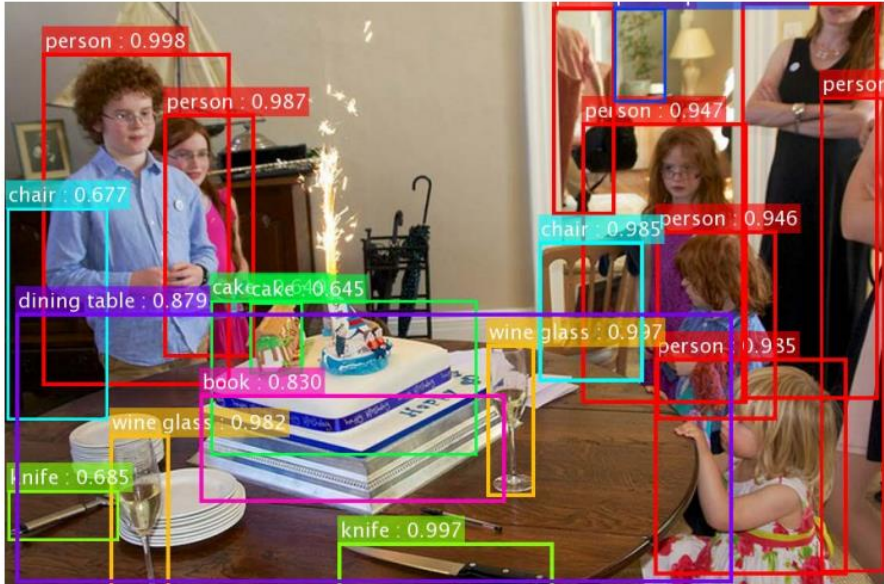


Slides from Kaimin He, MSRA

Image

- Object location

Microsoft Research



person : 0.998
person : 0.987
person : 0.947
person
chair : 0.677
chair : 0.985
person : 0.946
dining table : 0.879
caketake : 0.645
wine glass : 0.997
person : 0.935
book : 0.830
wine glass : 0.982
knife : 0.685
knife : 0.997

Our results on COCO – too many objects, let's check carefully!

*the original image is from the COCO dataset

ICCV15 International Conference on Computer Vision

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. "Deep Residual Learning for Image Recognition". arXiv 2015.
Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks". NIPS 2015.

Slides from Kaimin He, MSRA

Image

- Image captioning

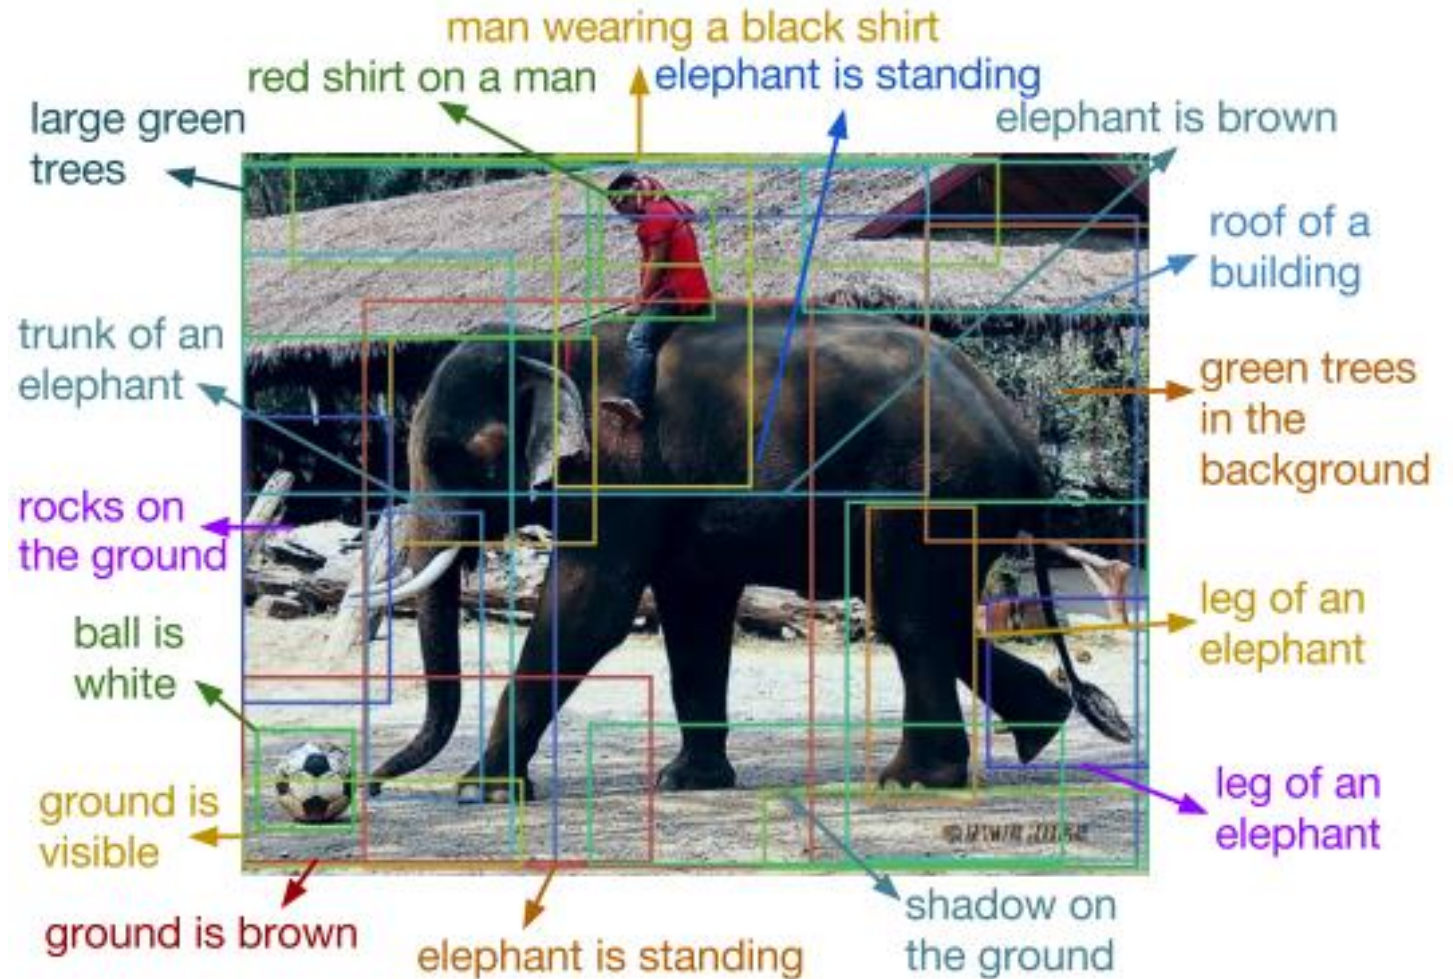


Figure from the paper “DenseCap: Fully Convolutional Localization Networks for Dense Captioning”, by Justin Johnson, Andrej Karpathy, Li Fei-Fei

Text


- Question & Answer

I: Jane went to the hallway.
I: Mary walked to the bathroom.
I: Sandra went to the garden.
I: Daniel went back to the garden.
I: Sandra took the milk there.
Q: Where is the milk?
A: garden

I: The answer is far from obvious.
Q: In French?
A: La réponse est loin d'être évidente.

Figures from the paper "Ask Me Anything: Dynamic Memory Networks for Natural Language Processing", by Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Richard Socher

Game



[Google DeepMind's Deep Q-learning playing Atari Breakout](#)

From the paper “Playing Atari with Deep Reinforcement Learning”,
by Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou,
Daan Wierstra, Martin Riedmiller

Game



CONSERVATION

**SONGBIRDS
À LA CARTE**
*Illegal harvest of millions
of Mediterranean birds*
PAGE 452

RESEARCH ETHICS

**SAFEGUARD
TRANSPARENCY**
*Don't let openness backfire
on individuals*
PAGE 459

POPULAR SCIENCE

**WHEN GENES
GOT 'SELFISH'**
*Dawkins's calling
card forty years on*
PAGE 462

NATURE.COM/NATURE

28 January 2016 £10

Vol. 529, No. 7587



The impact

- Revival of Artificial Intelligence
- Next technology revolution?

- A big thing ongoing, should not miss

Questions behind the scene

- Return of artificial neural network
 - What's different
 - Why get great performance
- Future development
 - The road to general-purpose AI?

Goal of the course

- Introduction
- Key concepts
- Ticket to the party

Syllabus

- Part I: machine learning basics
 - Linear model, Perceptron, SVM
 - Multi-class
 - Training by gradient descent
 - overfitting
- Part II: supervised deep learning (feedforward network)
- Part III: unsupervised learning
- Part IV: deep learning in the wild

Syllabus

- Part I: machine learning basics
- Part II: supervised deep learning (feedforward network)
 - Multiple-layer and Backpropogation
 - Regularization
 - Convolution
- Part III: unsupervised deep learning
- Part IV: deep learning in the wild

Syllabus

- Part I: machine learning basics
- Part II: supervised deep learning (feedforward network)
- Part III: unsupervised deep learning
 - PCA
 - Boltzmann machine, Deep Boltzmann machine
 - autoencoder
- Part IV: deep learning in the wild

Syllabus

- Part I: machine learning basics
- Part II: supervised deep learning (feedforward network)
- Part III: unsupervised deep learning
- Part IV: deep learning in the wild
 - Read papers on advanced topics
 - Play with the code
 - Presentation

Textbook and materials

- Deep Learning:

<http://www.deeplearningbook.org/>

- Suggested software framework: Tensorflow
 - in Python
 - Easy to install/use
 - Can try it on your laptop
- Other software frameworks: Theano, Caffe, Torch, Marvin, ...

Grading

- Problem Sets (5 sets): 70%
- Design Projects: 25%
- Oral Presentation: 5%