

Introduction

Nando de Freitas



UNIVERSITY OF
OXFORD

O X


Application: Invariant recognition in natural images




2



[Thomas Serre 2012]






Millions of labeled examples are used to build real-world applications, such as pedestrian detection



[Tomas Serre]



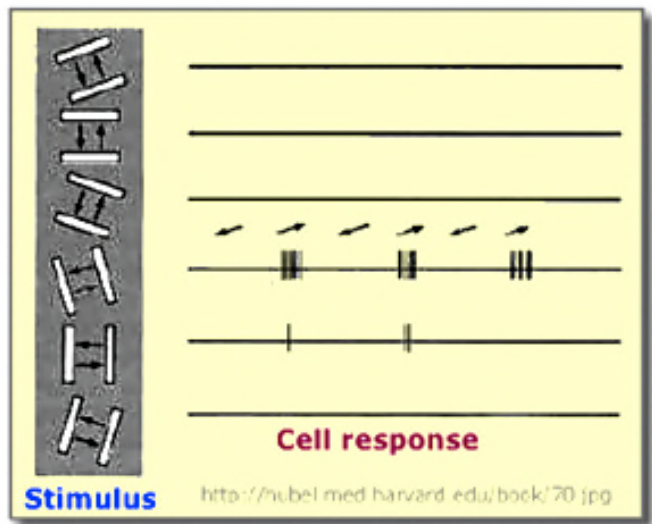

More applications of the same idea







[Okuma, Taleghani, dF, Little, Lowe, 2004]
Best Cognitive Vision Paper- ECCV

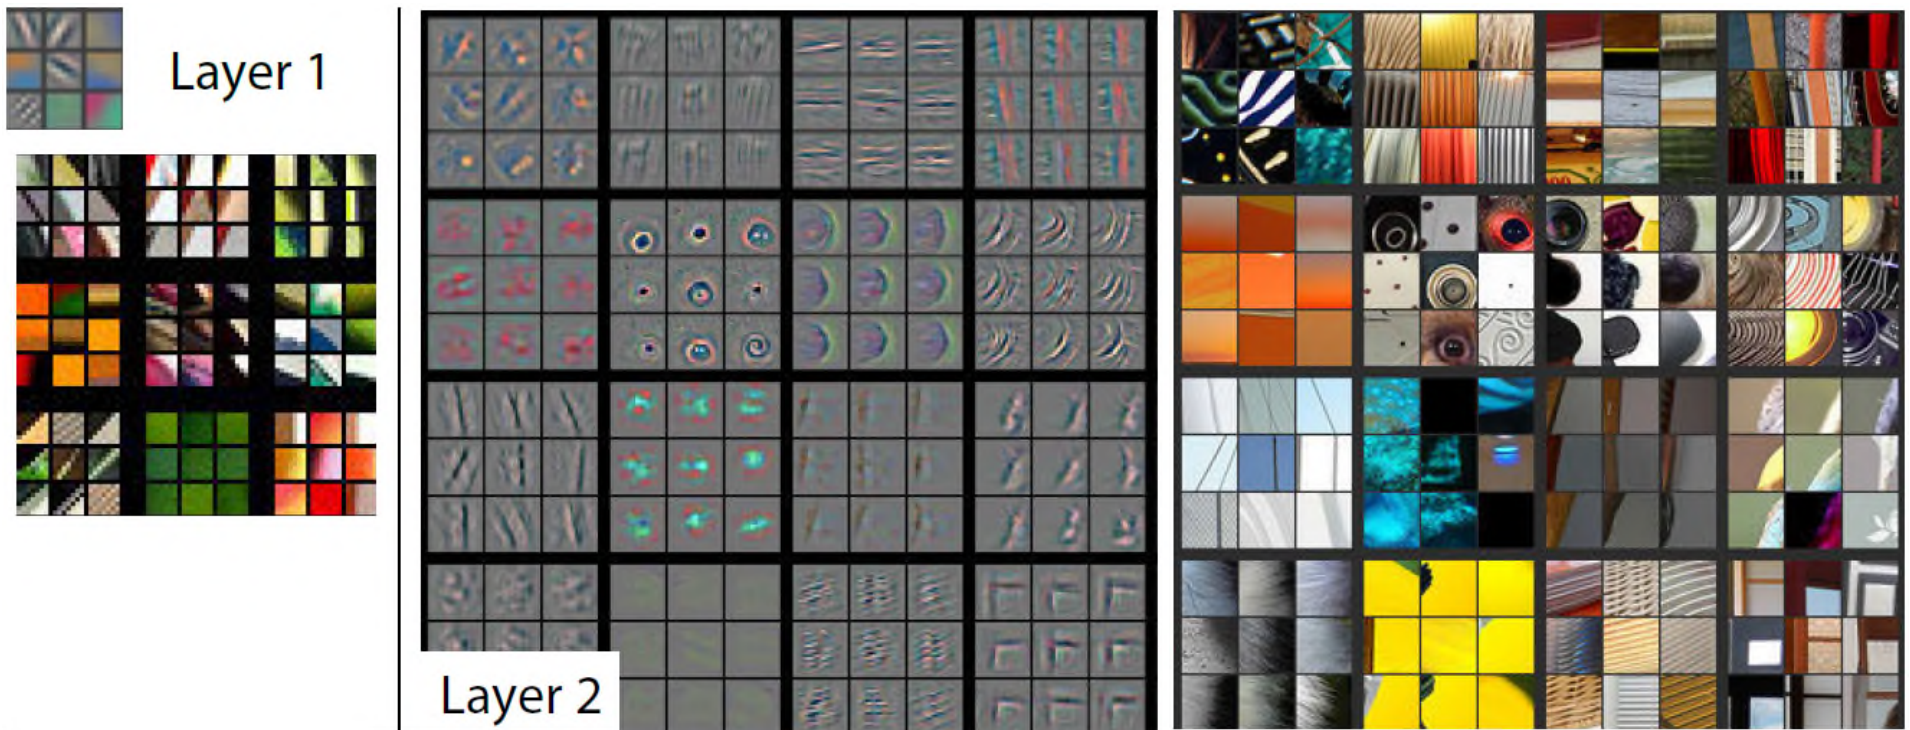
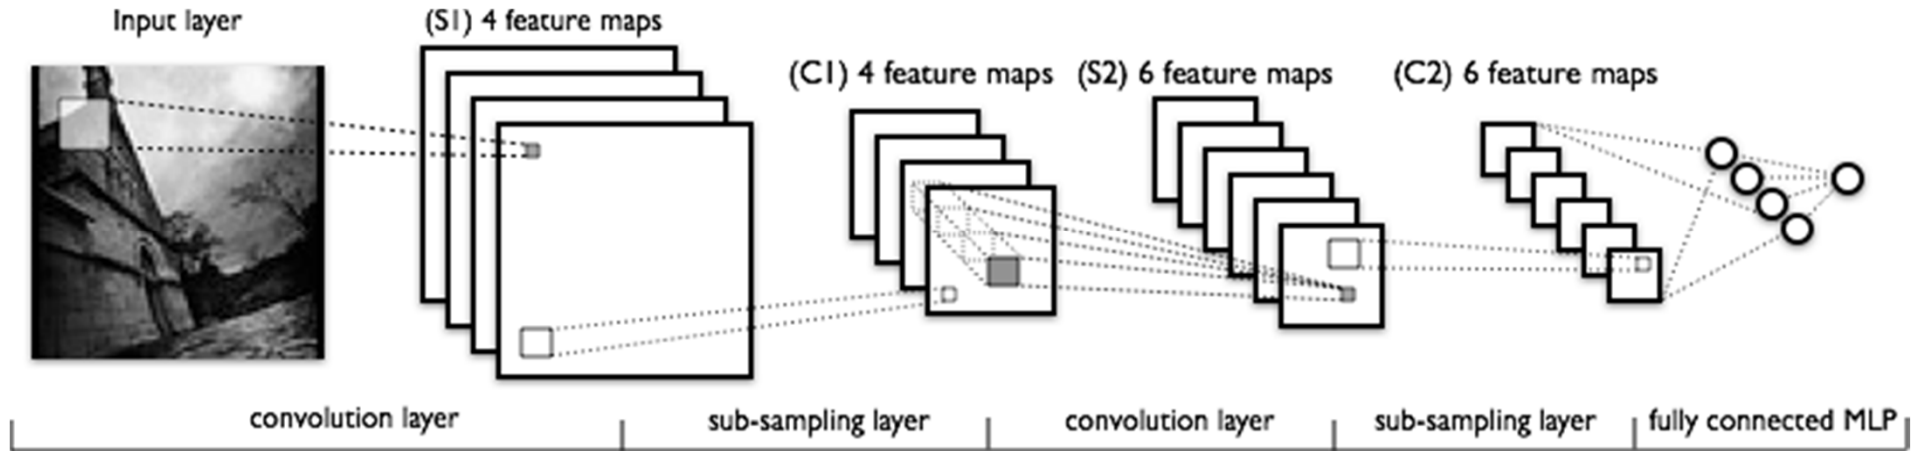
Place cells in the hippocampus and grid cells in the entorhinal cortex
[John O'Keefe, May-Britt Moser and Edvard Mosel, NOBEL 2014]



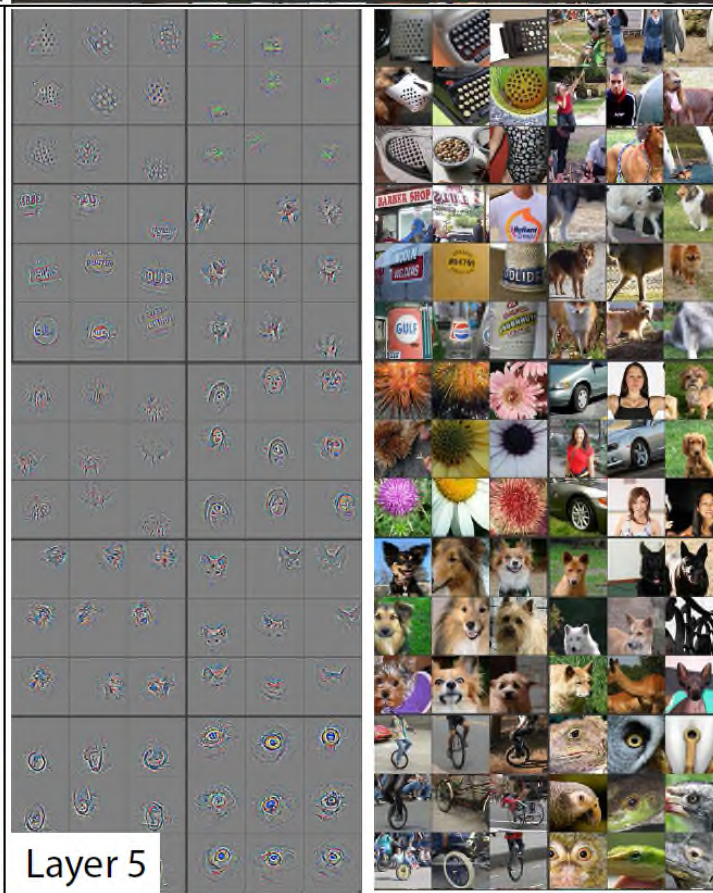


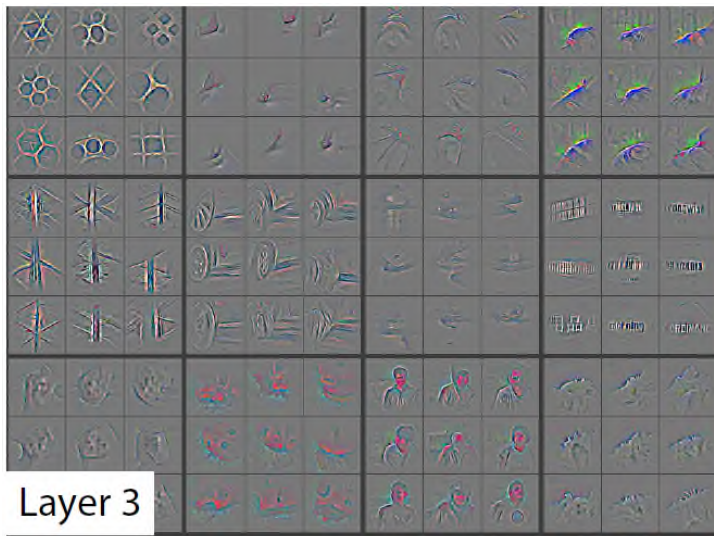






Convolutional networks



[Matthew Zeiler & Rob Fergus]




Google's neural net learns just by watching youtube videos



Machines that learn to recognise what they **see** and **hear** are at the heart of Apple, Google, Amazon, Facebook, Netflix, Microsoft, etc.

The image shows a screenshot of the Baidu Image Search interface. At the top, there is a search bar with the text "粘帖图片网址" and "从本地上传". Below the search bar, there are tabs for "全部", "相似图片", and "人脸搜索". The search results are displayed in a grid of images, primarily showing a woman in a white dress. In the bottom left corner, there is a social media post from a user named "Dan" showing a photo of a child in a striped shirt sitting at a table. The post includes a comment: "Arhmy..... I dont think they are me" and a timestamp of "2 hours ago".



Review sentiment and summarization



my reading was similar to everyones. she told me she was going to take her time and not rush me out of there. i was there not even 8 minutes she told me i was pregnant then she changed her mind and said i had a miscarriage. im 17 years old i told her she was wrong she then went on and said "i see you and your brother fight alot just know he loves you" i dont even have a brother.

she then told my friend she was going to get stabbed

Was this review helpful?
Ask taydube about Fatima's Psychic Studio

[Problem with this review?](#)

Paul Bettany did a great role as the tortured father whose favorite little girl dies tragically of disease. For that, he deserves all the credit.


However, the movie was mostly about exactly that, keeping the adventures of Darwin as he gathered data for his theories as incomplete stories told to children and skipping completely the disputes regarding his ideas. Two things bothered me terribly: the soundtrack, with its whiny sound, practically shoving sadness down the throat of the viewer, and the movie trailer, showing some beautiful sceneries, the theological musings of him and his wife and the enthusiasm of his best friends as they prepare for a battle against blind faith, thus misrepresenting the movie completely.

To put it bluntly, if one were to remove the scenes of the movie trailer from the movie, the result would be a non descript family drama about a little child dying and the hardships of her parents as a result.


Clearly, not what I expected from a movie about Darwin, albeit the movie was beautifully interpreted.

[Kotzias, Denil, Blunsom & NdF, 2014]

Structured queries and outputs

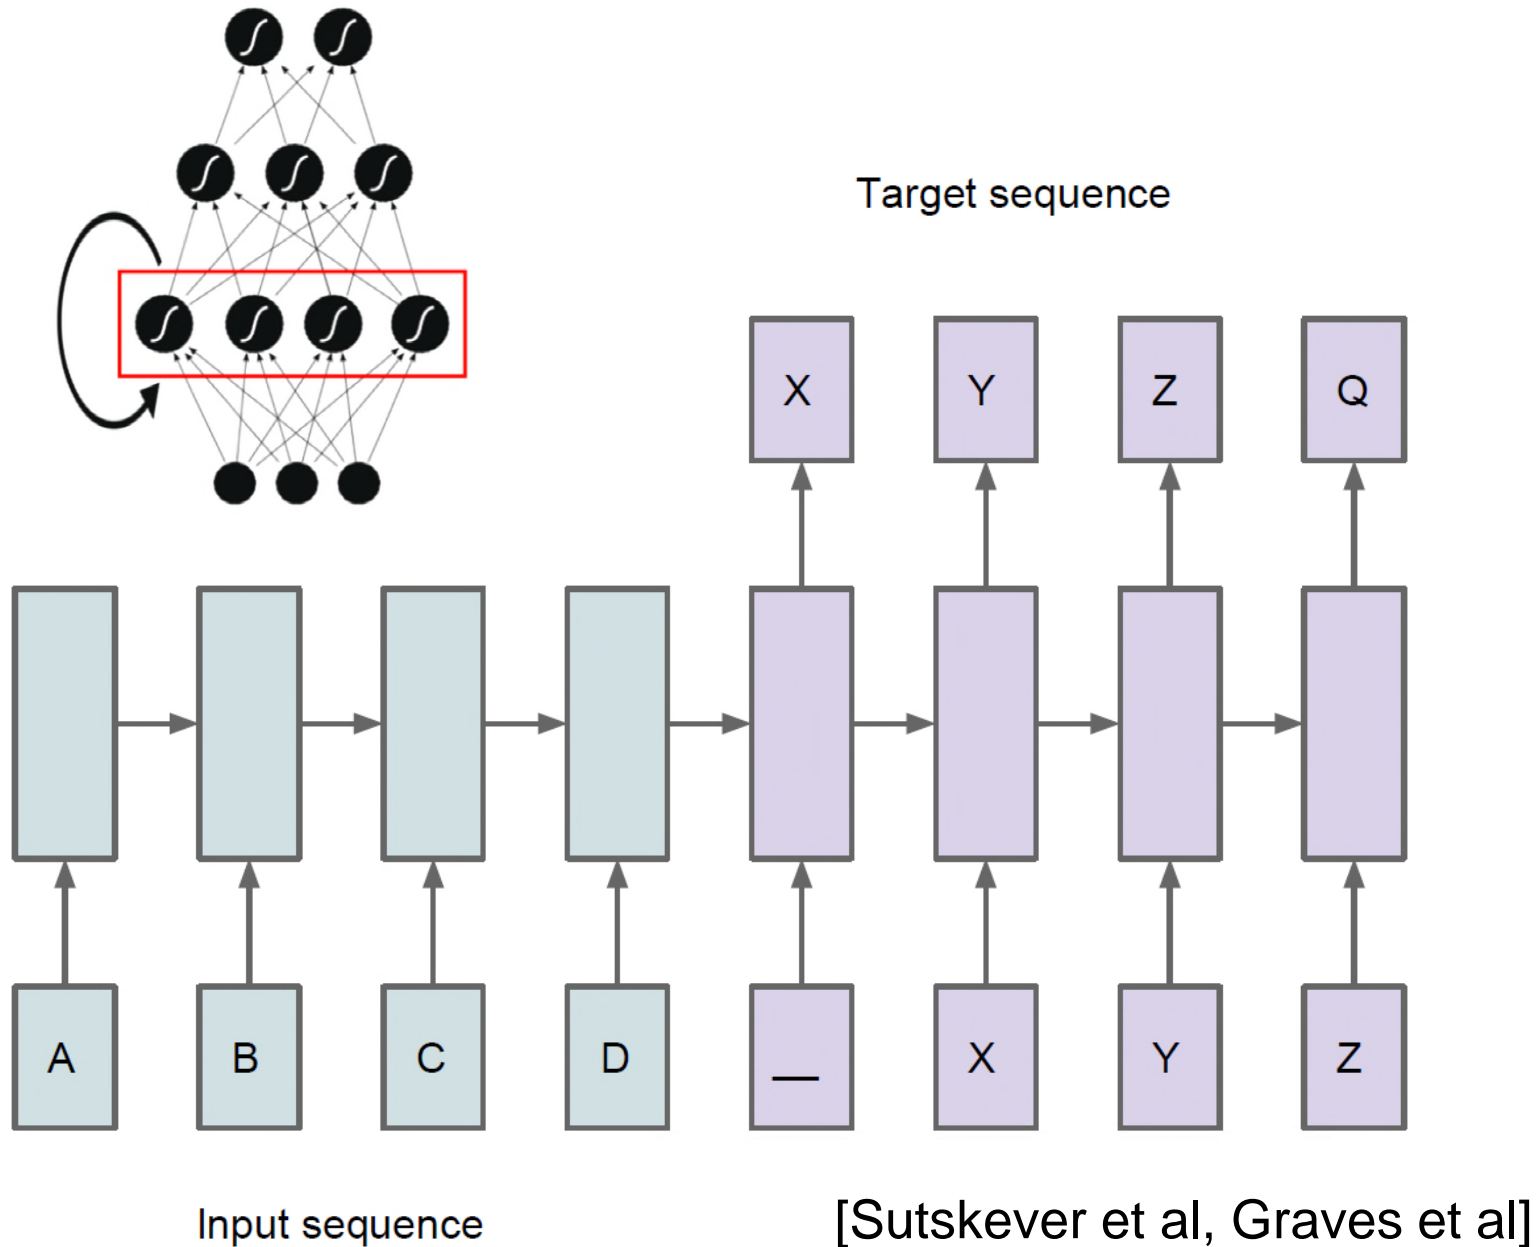


Who's likely
to want
to watch this
movie with
me on
Friday?




Phil is available and he likes movies with Downey JR

Sequence learning and recurrent nets



Sequence learning and recurrent nets

Real



Generated



Karol Gregor, Ivo Danihelka, Andriy Mnih, Daan Wierstra...

Google DeepMind 

Sequence learning and recurrent nets

Which is Real?

from his travels it might have been

from his travels it might have been

from his travels it might have been

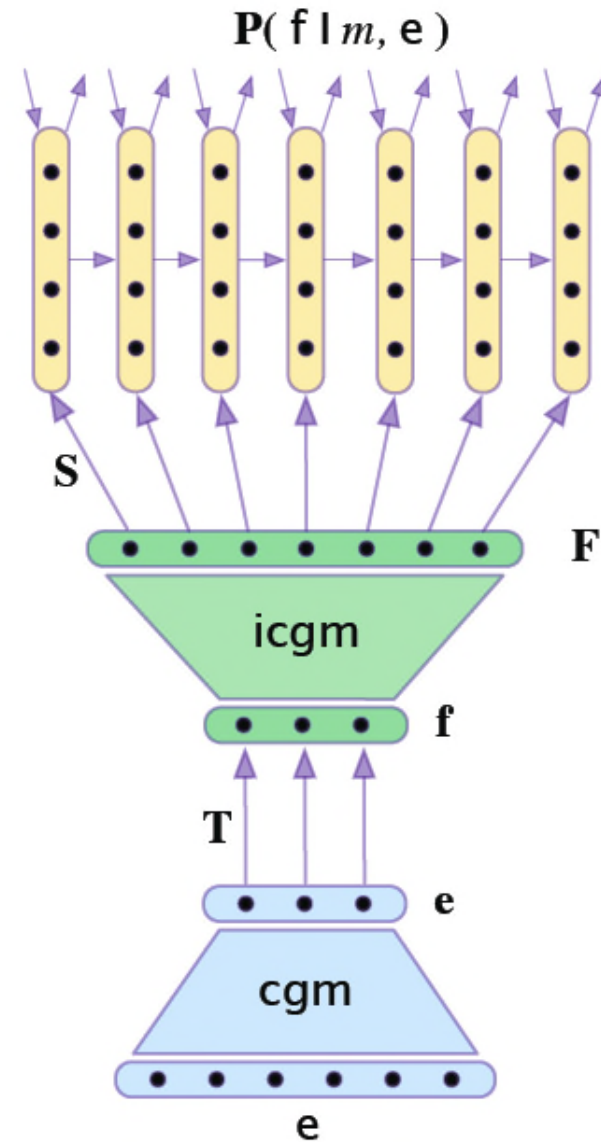
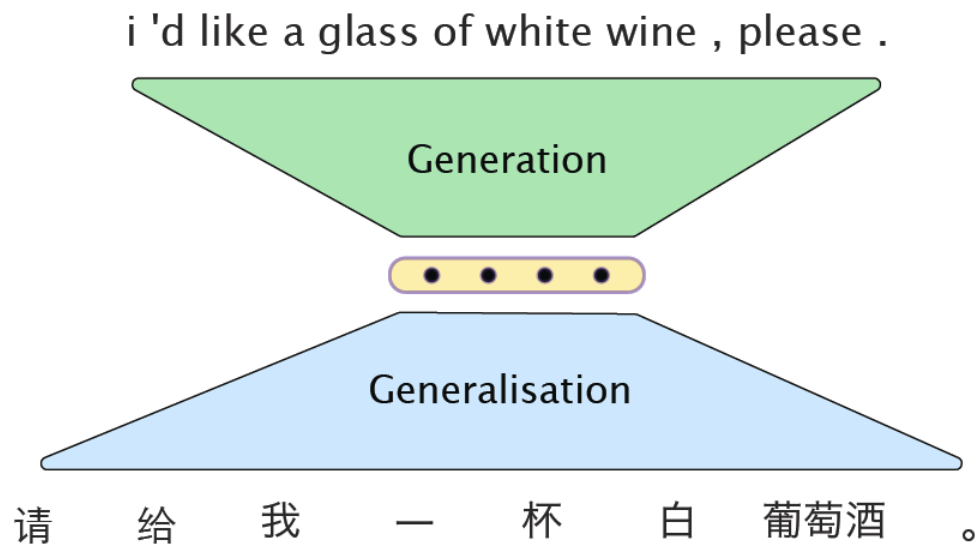
from his travels it might have been

from his travels it might have been

from his travels - it might have been


[Alex Graves]

Siamese nets, machine translation




[Le Cun et al, Blunsom et al]

Imitation learning & mirror neurons



Imitation learning for Atari





[Dejan Markovikj, Miroslav Bogdanovic, Misha Denil, NdF 2014]

Machine learning


Machine learning deals with the problem of extracting *features* from data so as to solve many different *predictive* tasks:

- ❑ Forecasting (e.g. *Energy demand prediction, finance*)
- ❑ Imputing missing data (e.g. *Netflix recommendations*)
- ❑ Detecting anomalies (e.g. *Security, fraud, virus mutations*)
- ❑ Classifying (e.g. *Credit risk assessment, cancer diagnosis*)
- ❑ Ranking (e.g. *Google search, personalization*)
- ❑ Summarizing (e.g. *News zeitgeist, social media sentiment*)
- ❑ Decision making (e.g. *AI, robotics, compiler tuning, trading*)




When to apply machine learning

- ❑ Human expertise is absent (e.g. *Navigating on Mars*)
- ❑ Humans are unable to explain their expertise (e.g. *Speech recognition, vision, language*)
- ❑ Solution changes with time (e.g. *Tracking, temperature control, preferences*)
- ❑ Solution needs to be adapted to particular cases (e.g. *Biometrics, personalization*)
- ❑ The problem size is too vast for our limited reasoning capabilities (e.g. *Calculating webpage ranks, matching ads to facebook pages*)



Challenge: One-shot learning

“tufa”



Can you pick out the tufas?

Other challenges

- Multi-task & transfer learning (generalization)
- Scaling and energy efficiency
- Ability to generate data (e.g., vision as inverse graphics)
- Architectures for AI

Next lecture

In the next lecture, we look at the simplest model in machine learning: Linear regression.