Machine Learning and AI

Qiang Yang, HKUST
Chair Professor
http://www.cse.ust.hk/~qyang

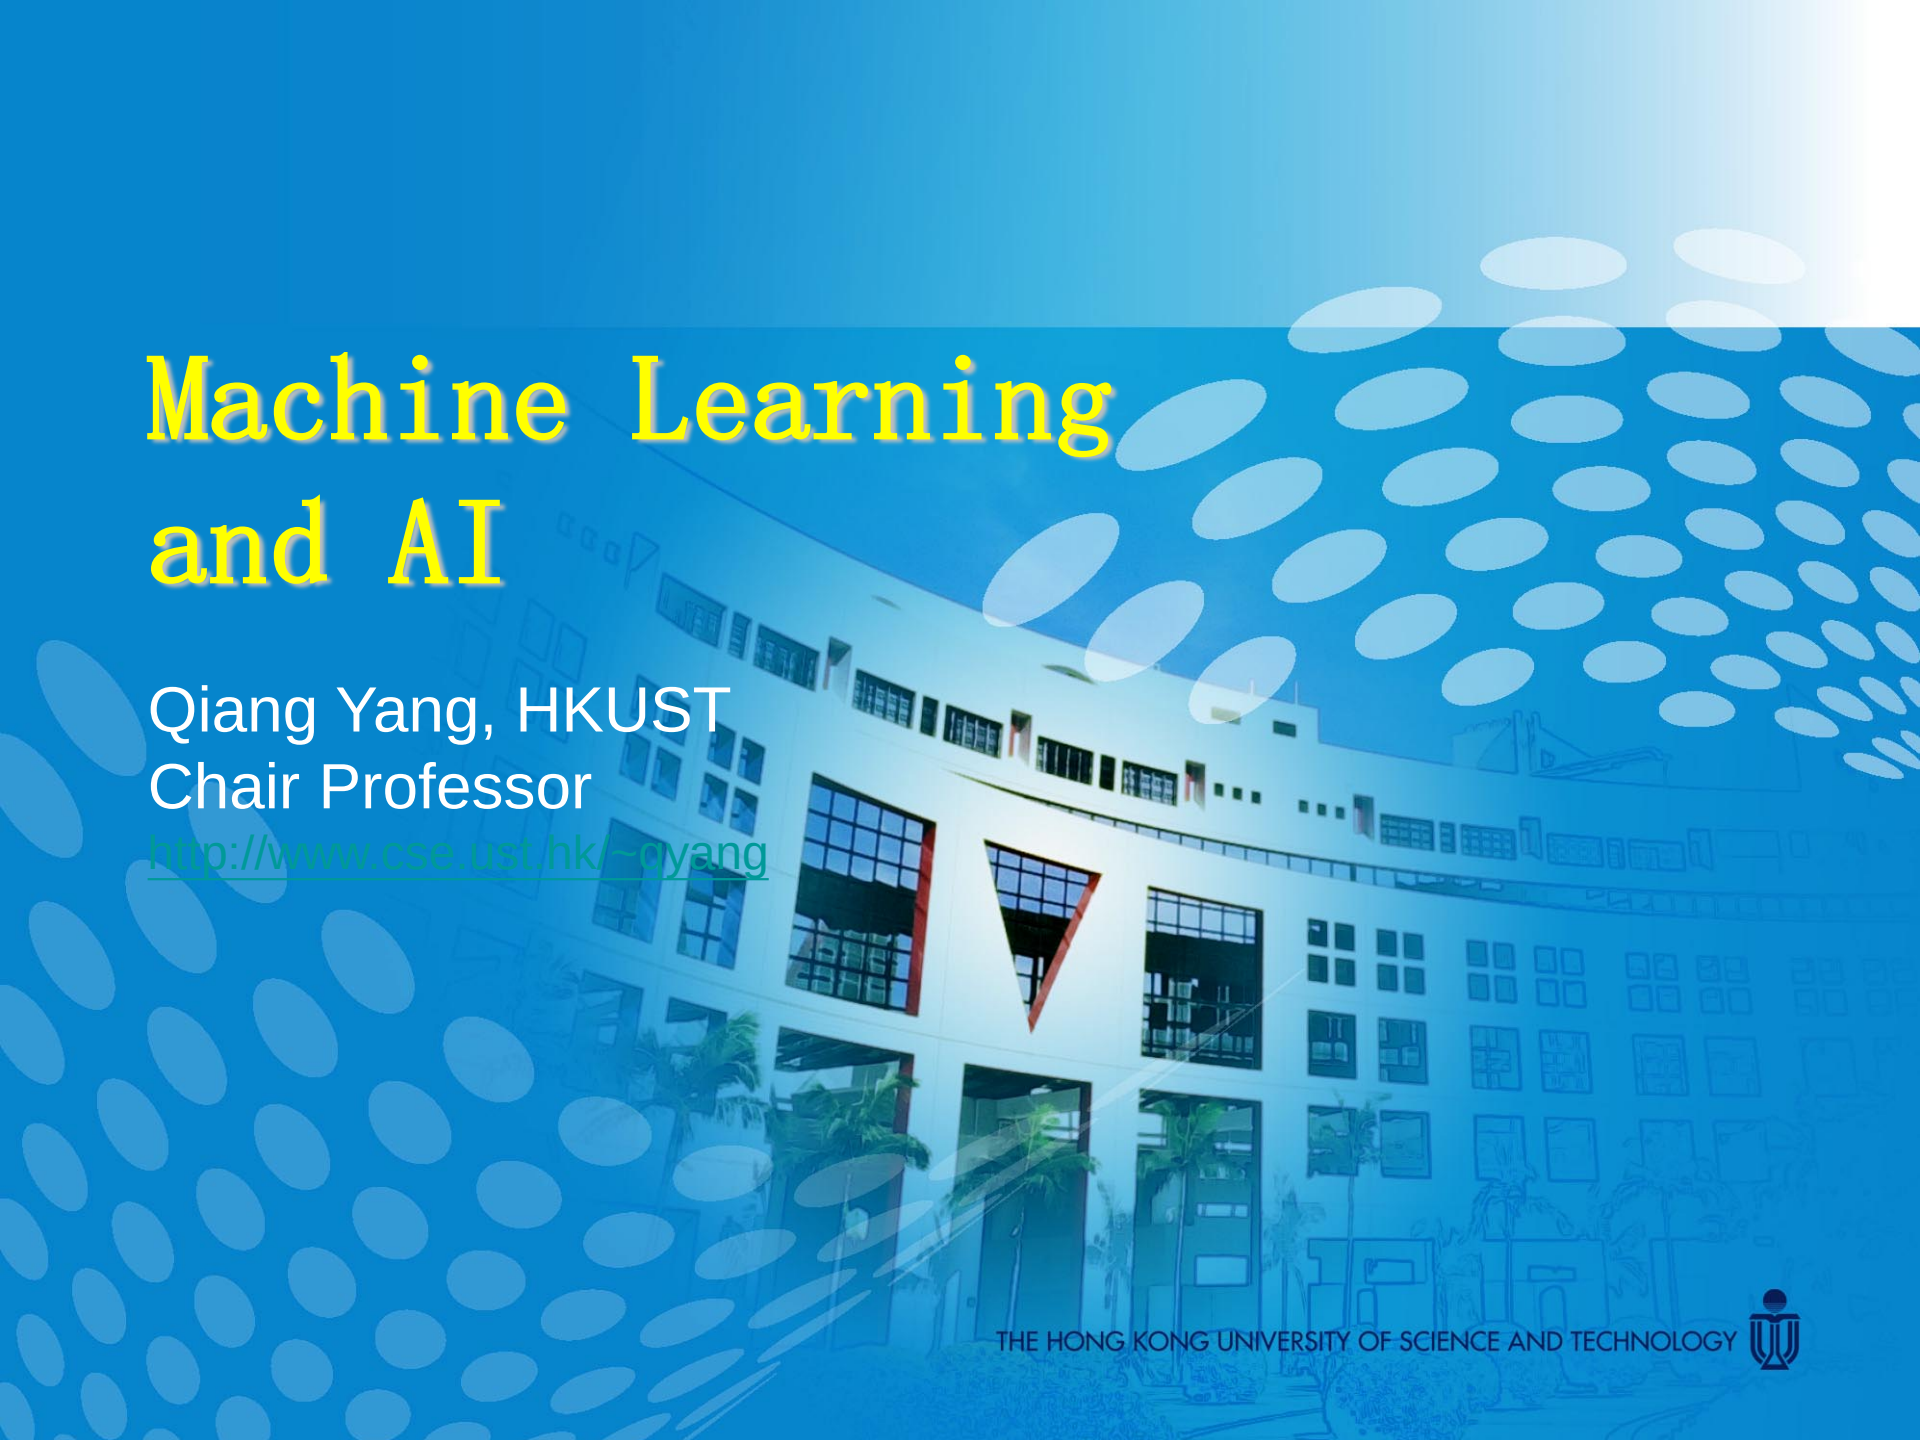

THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

60s： Intelligence from Logic

- Herbert Simon and Allen Newell:
  - Logic Theorist
  - General Problem Solver: Heuristic Search
  - Physical Symbol System Hypothesis

70s: Neural Networks

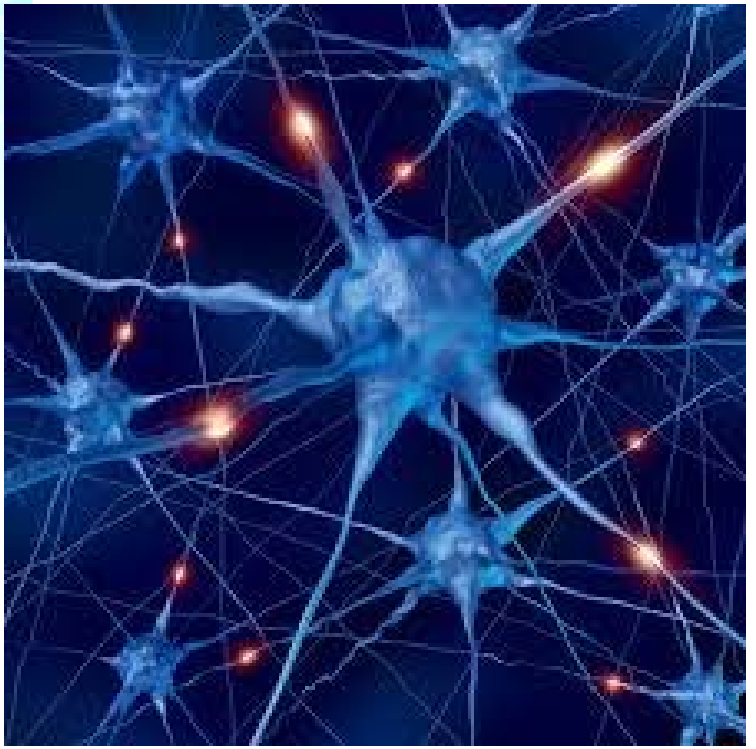

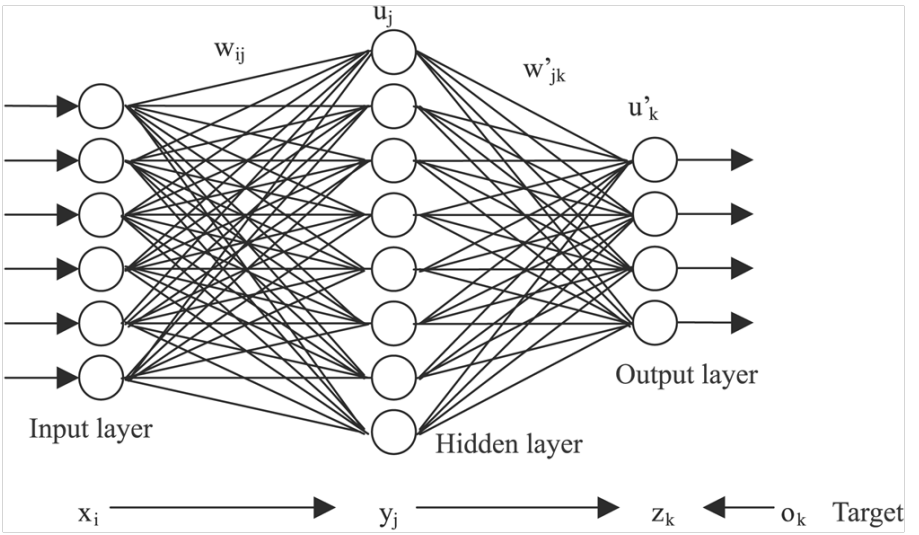

Search and Massive Parallelism

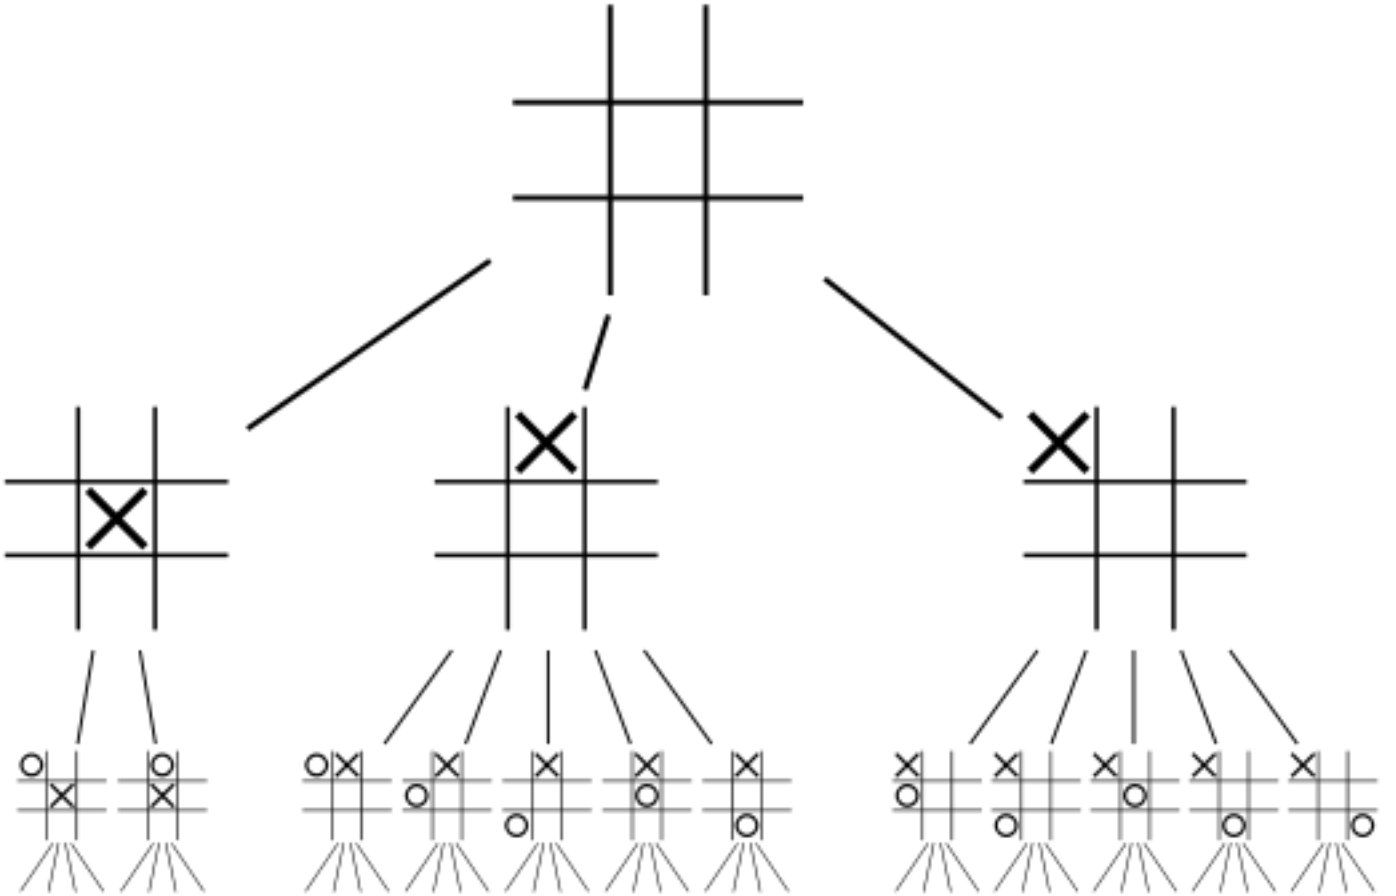

1997: DeepBlue

IBM

Kasparov

2011: IBM Watson/ Knowledge is Power!

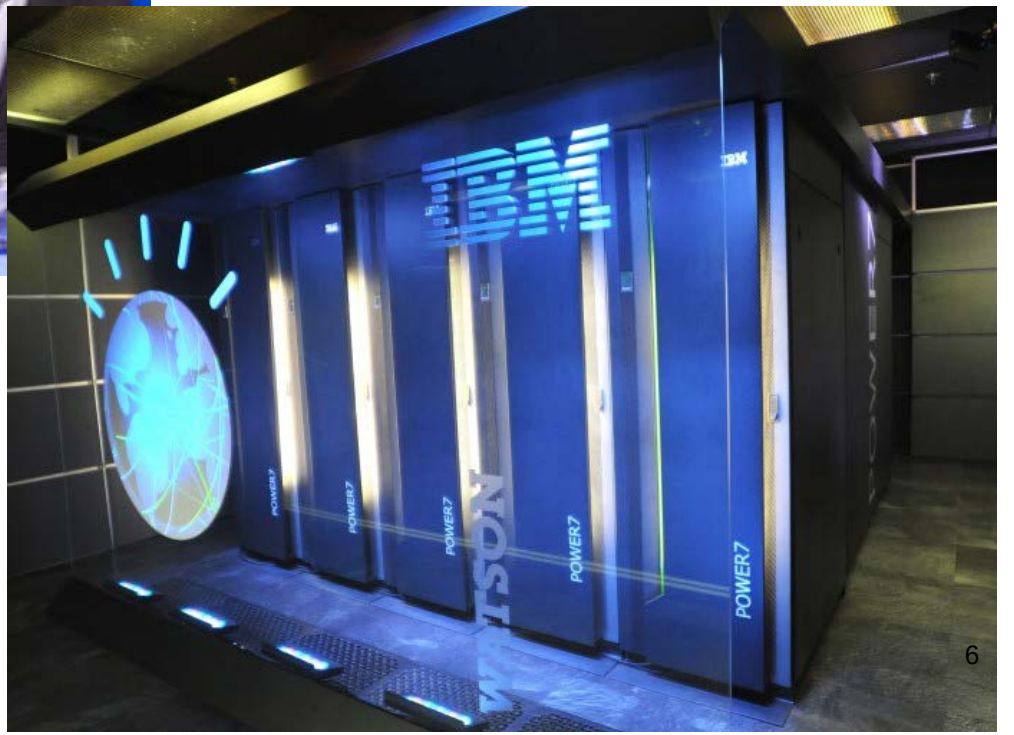

2010： CNN and Deep Learning

Convolutional Neural Nets CNN

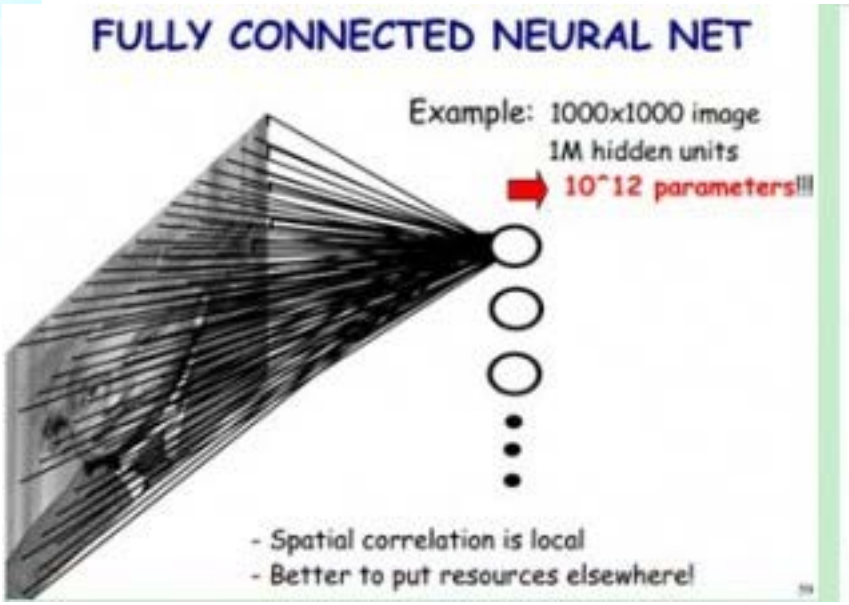

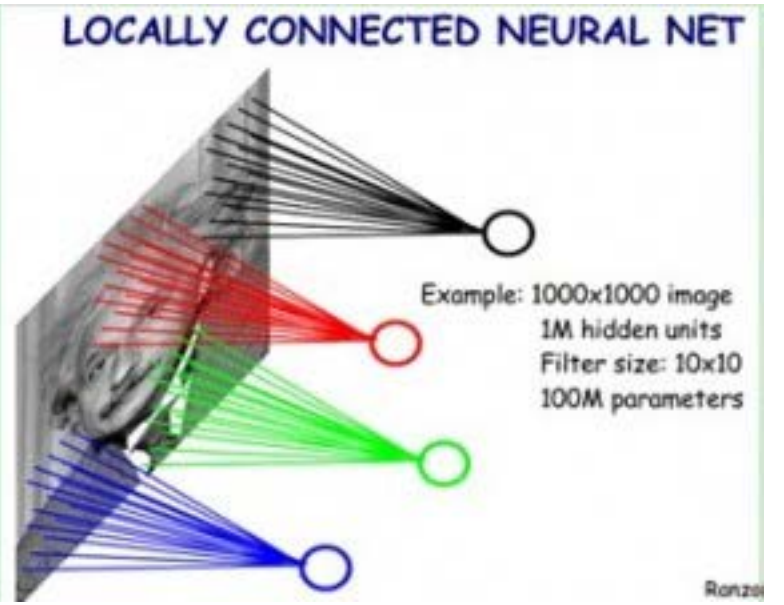

Features at Different Levels

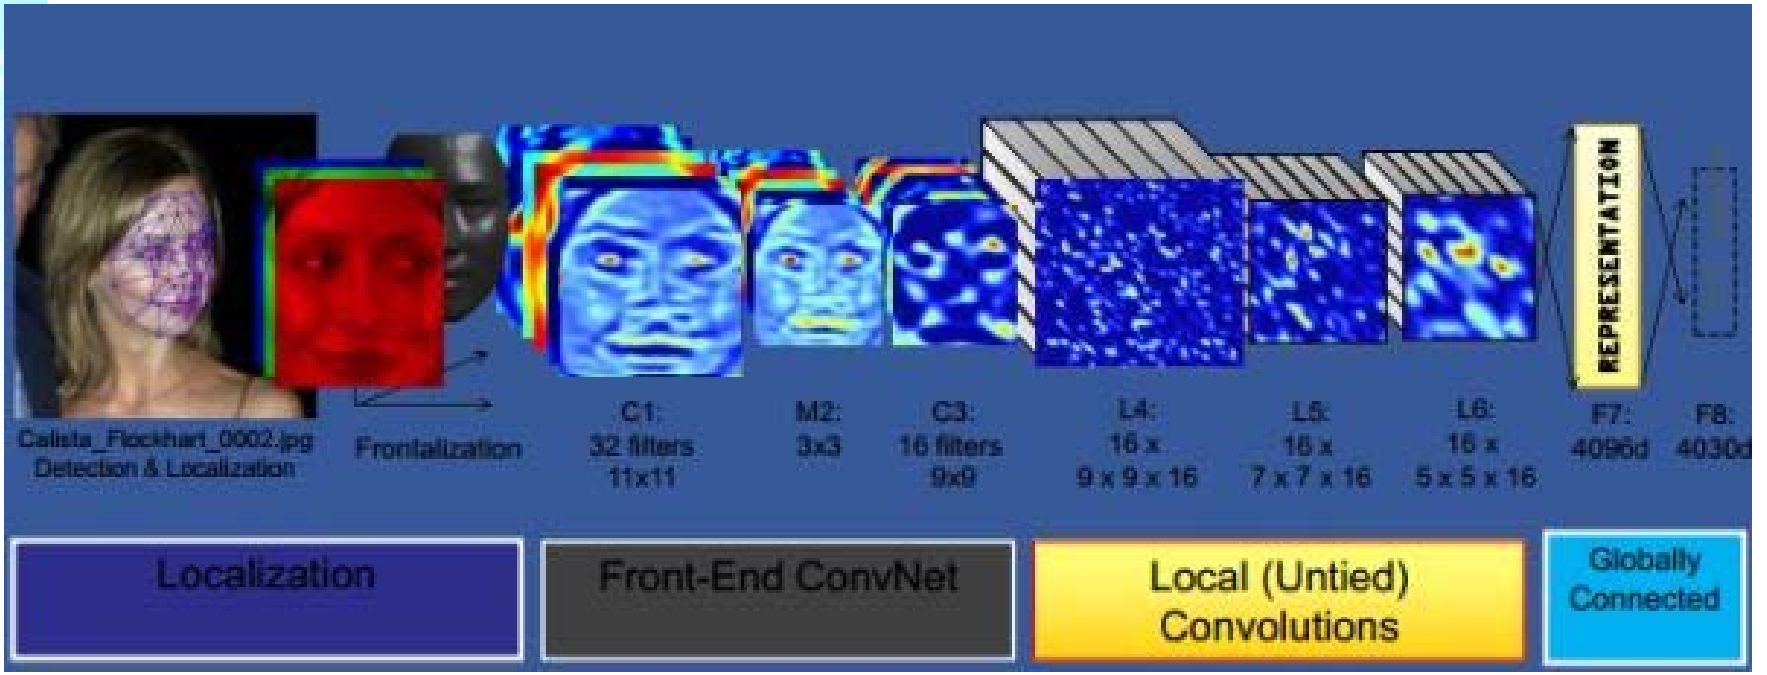

- Vision
- Speech
- Understanding

Next: Reinforcement Learning

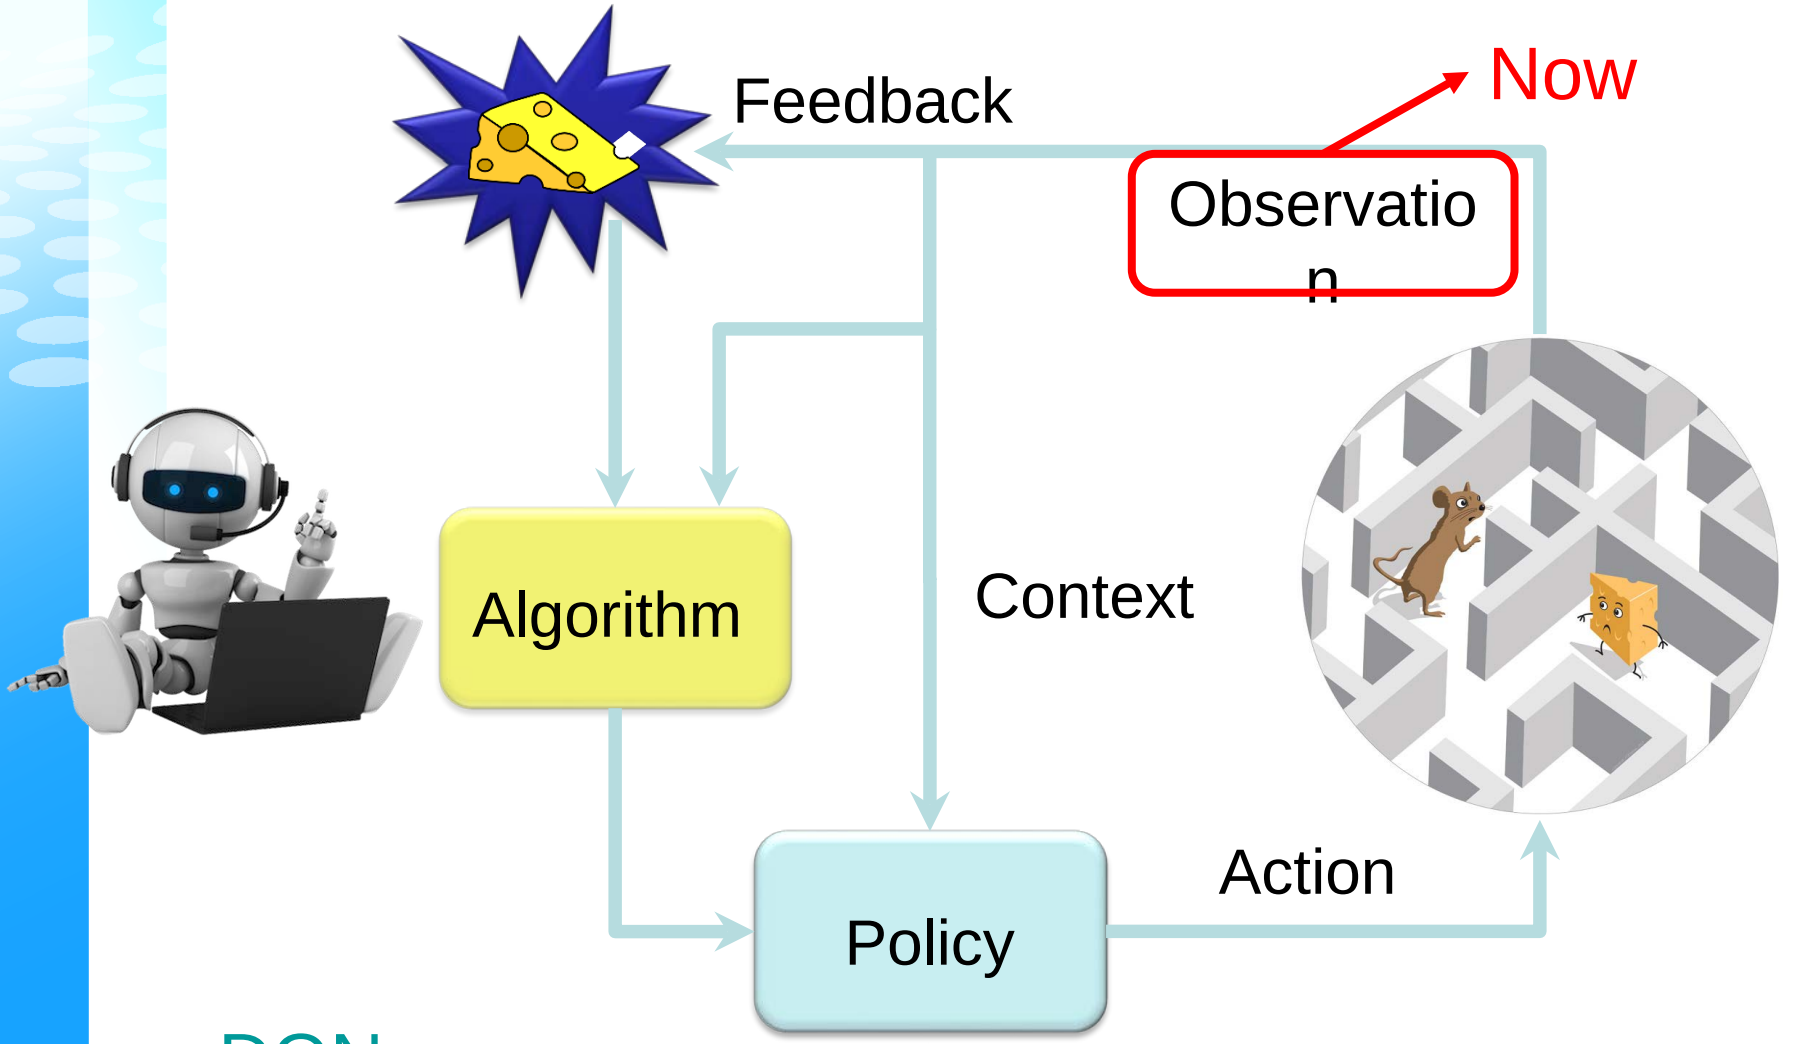

DQN

DeepMind: End to End

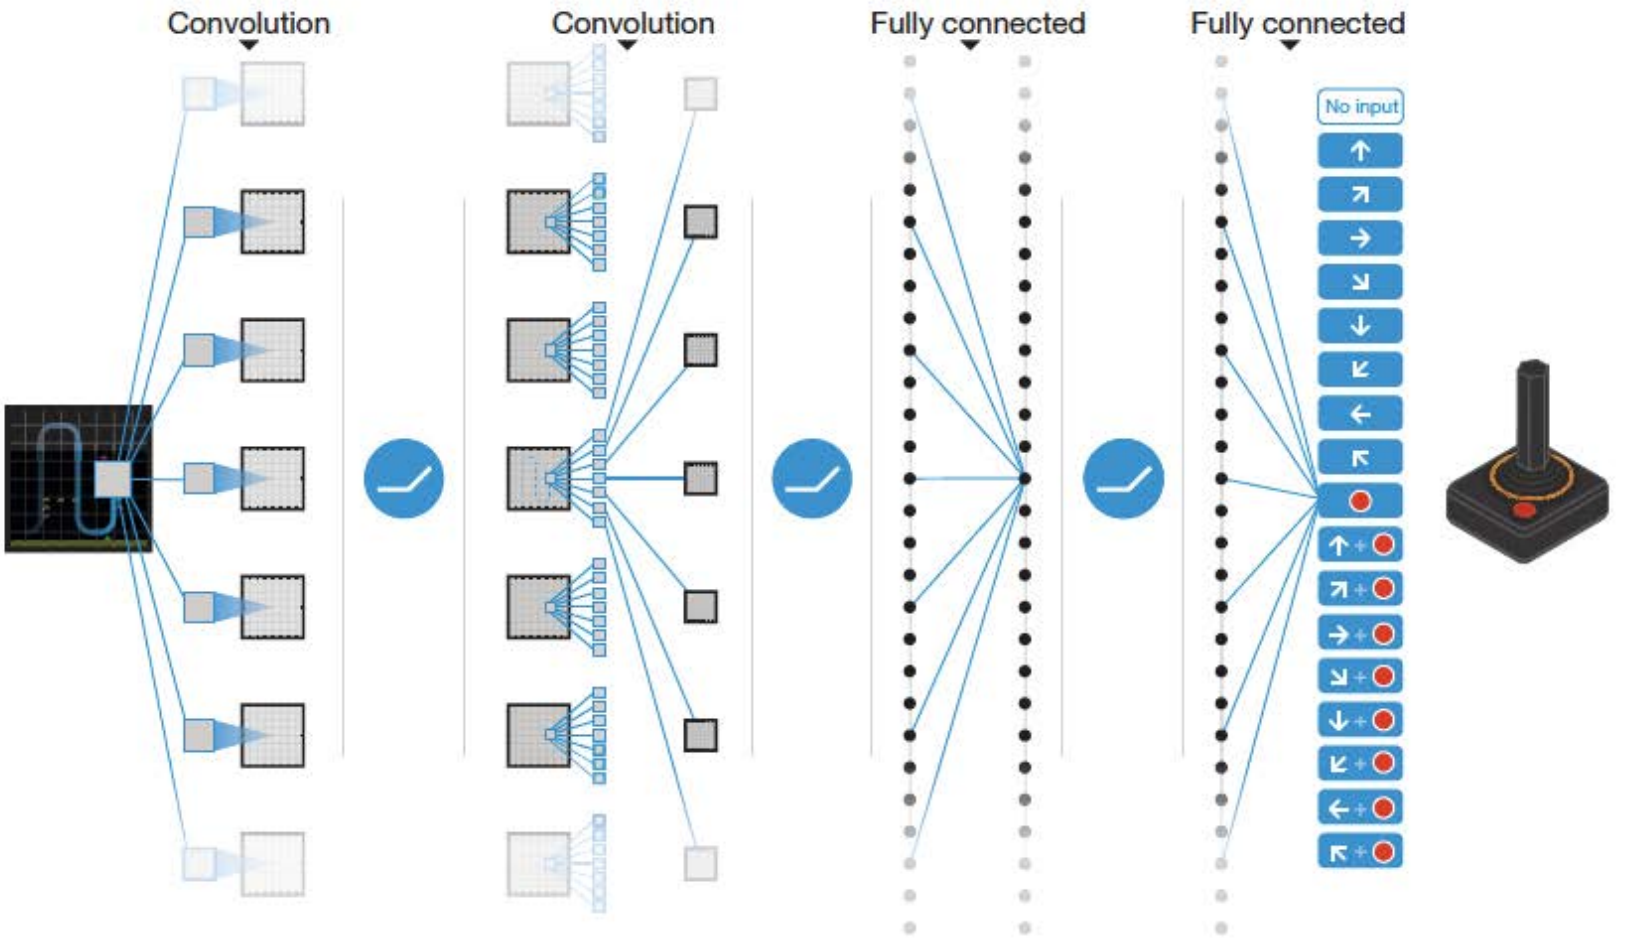

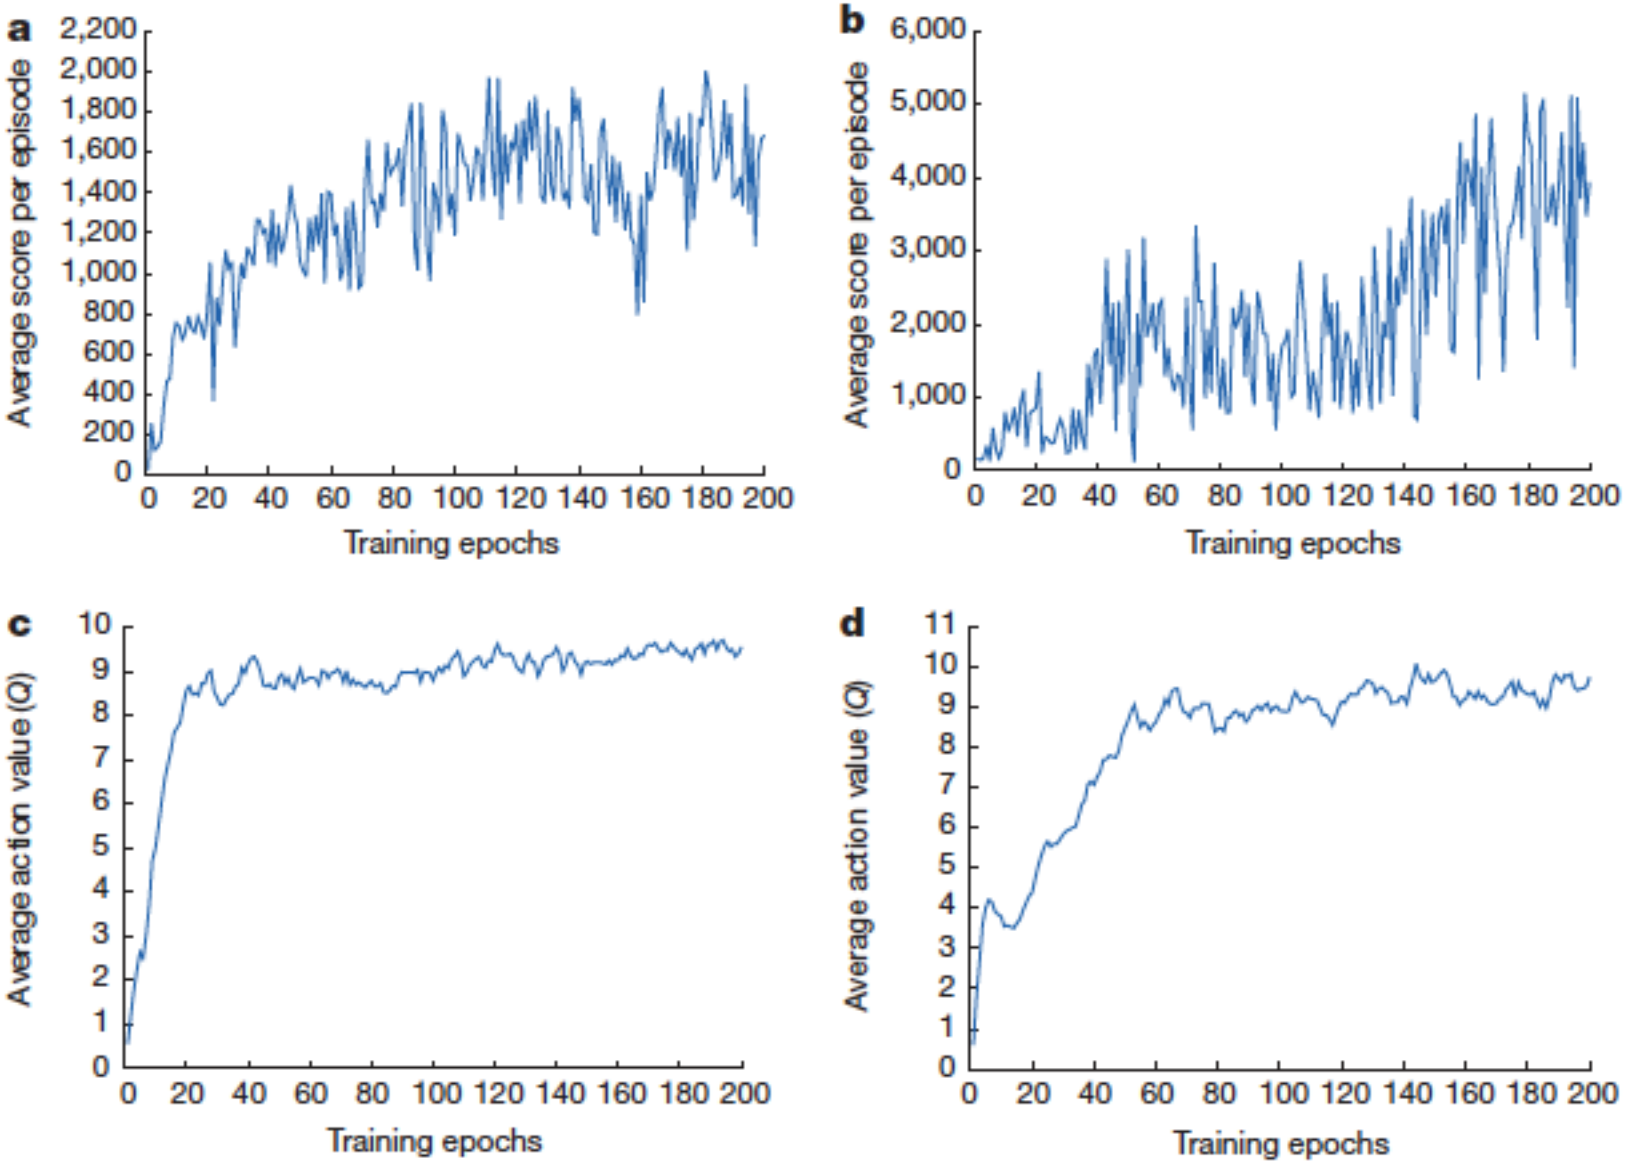
a
Average score per episode
2,200
2,000
1,800
1,600
1,400
1,200
1,000
800
600
400
200
0
0 20 40 60 80 100 120 140 160 180 200
Training epochs
b
Average score per episode
6,000
5,000
4,000
3,000
2,000
1,000
0
0 20 40 60 80 100 120 140 160 180 200
Training epochs
c
Average action value (Q)
10
9
8
7
6
5
4
3
2
1
0
0 20 40 60 80 100 120 140 160 180 200
Training epochs
d
Average action value (Q)
11
10
9
8
7
6
5
4
3
2
1
0
0 20 40 60 80 100 120 140 160 180 200
Training epochs

AlphaGo：Generality in AI

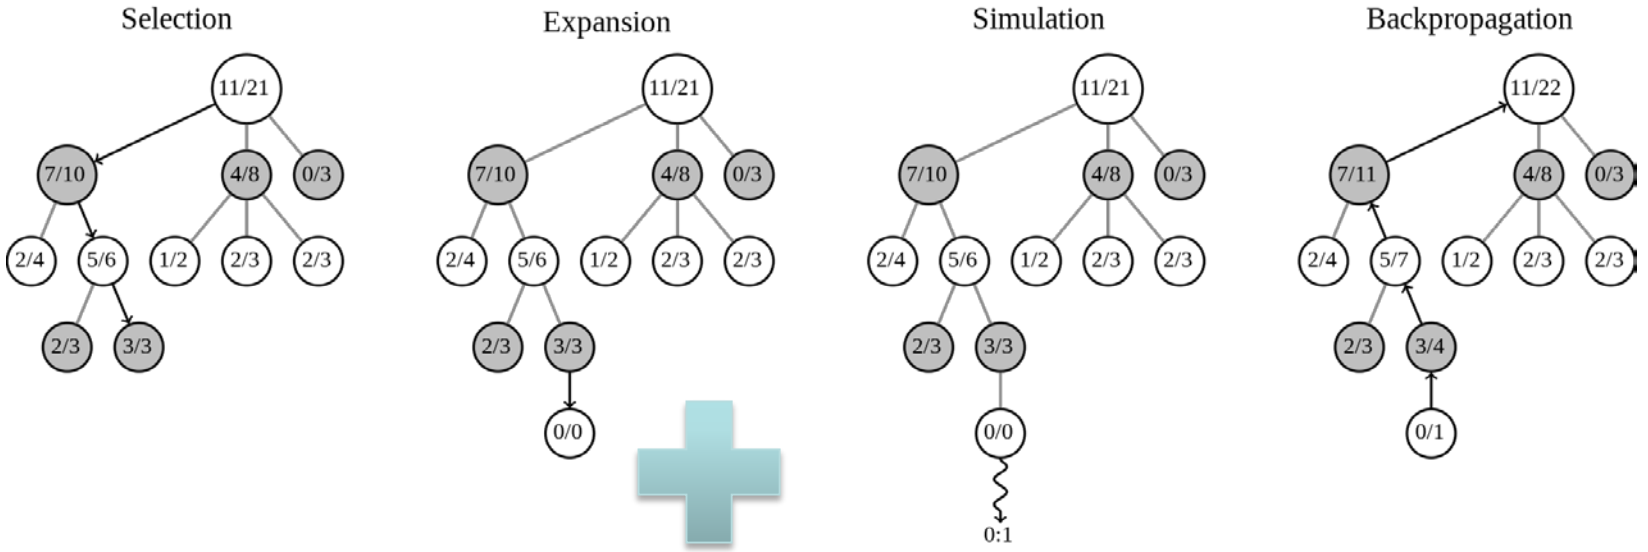

CMU Never Ending Learning Machine

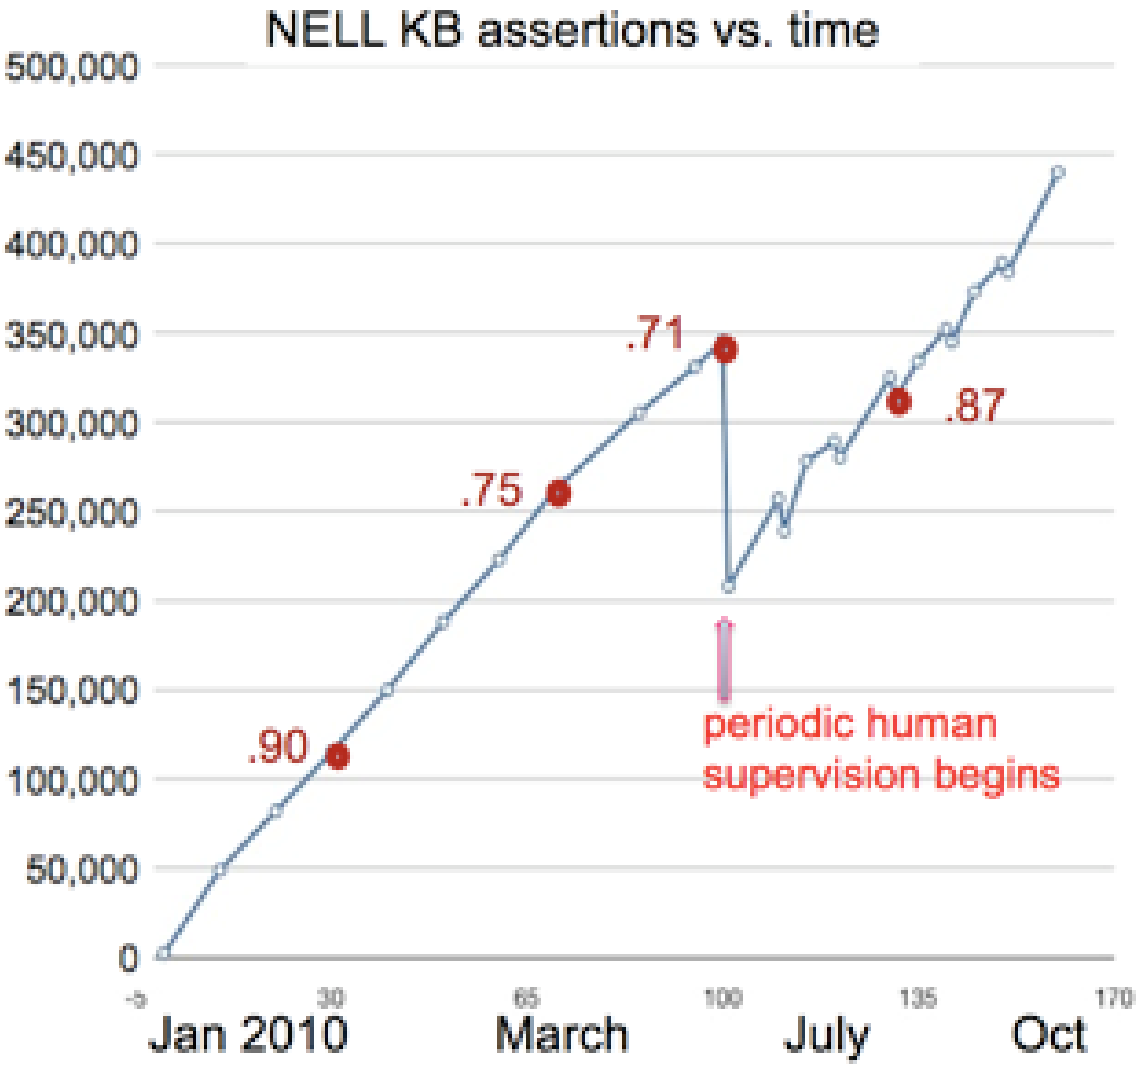

Biased Data → Big Data

Transfer Learning

Oquab, Bottou, Laptev, Sivic: Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks. CVPR 2014.

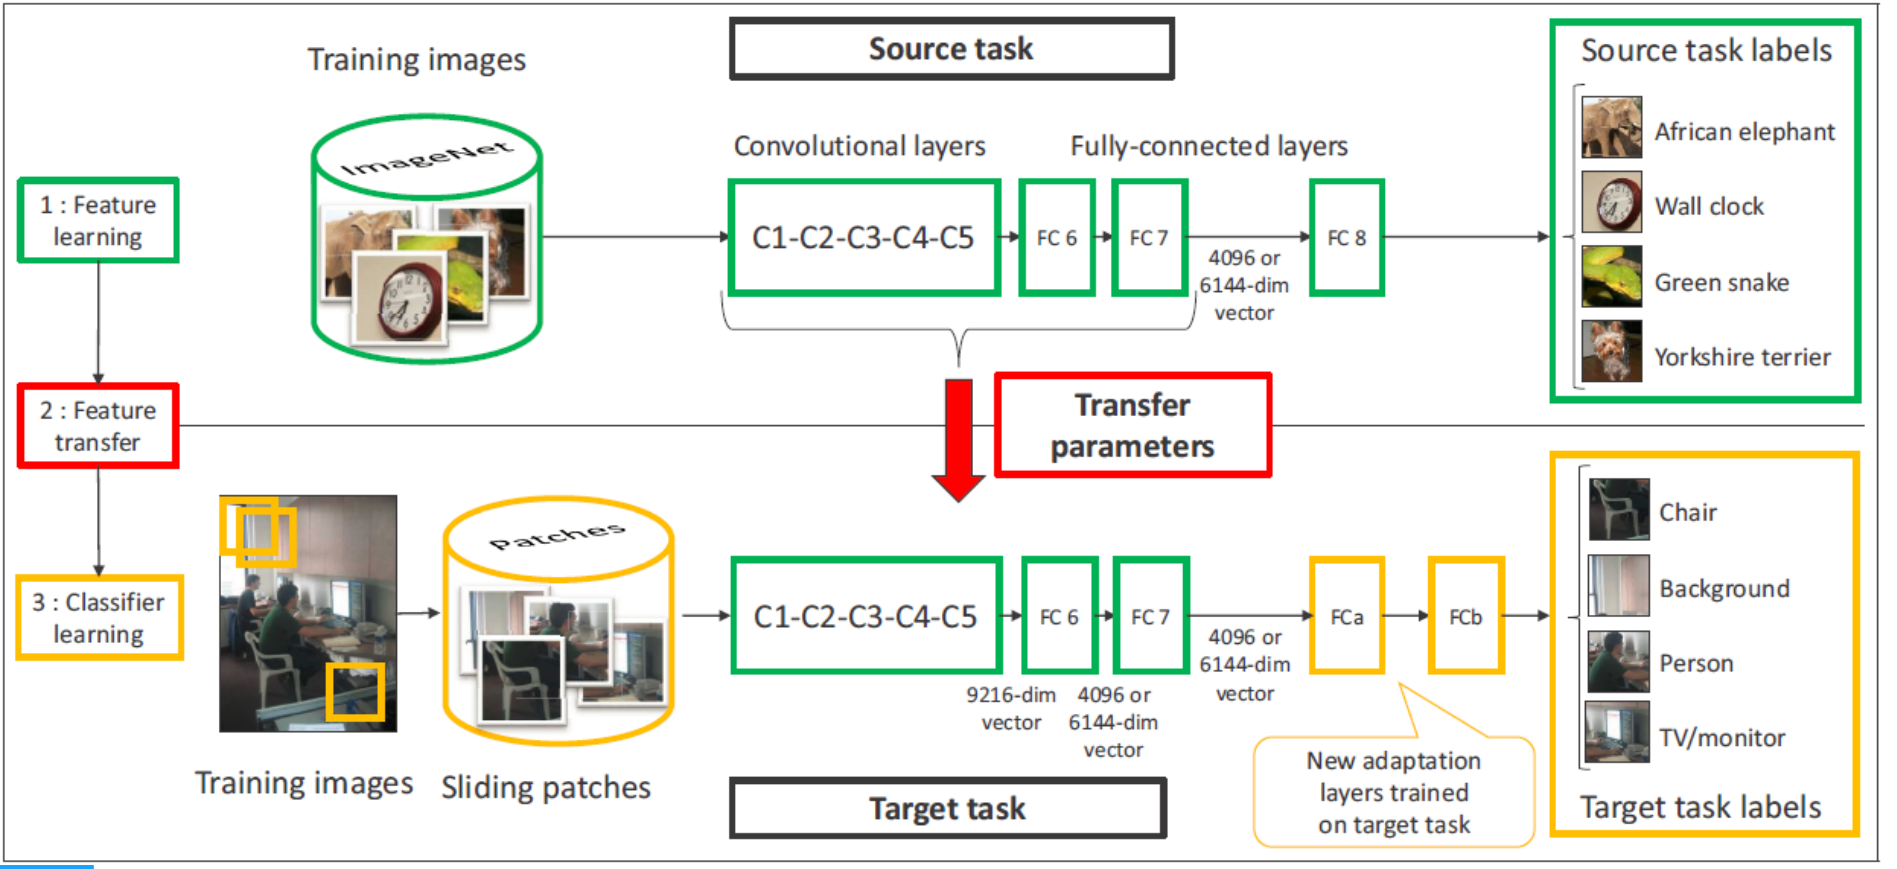

Transfer Learning: text to images

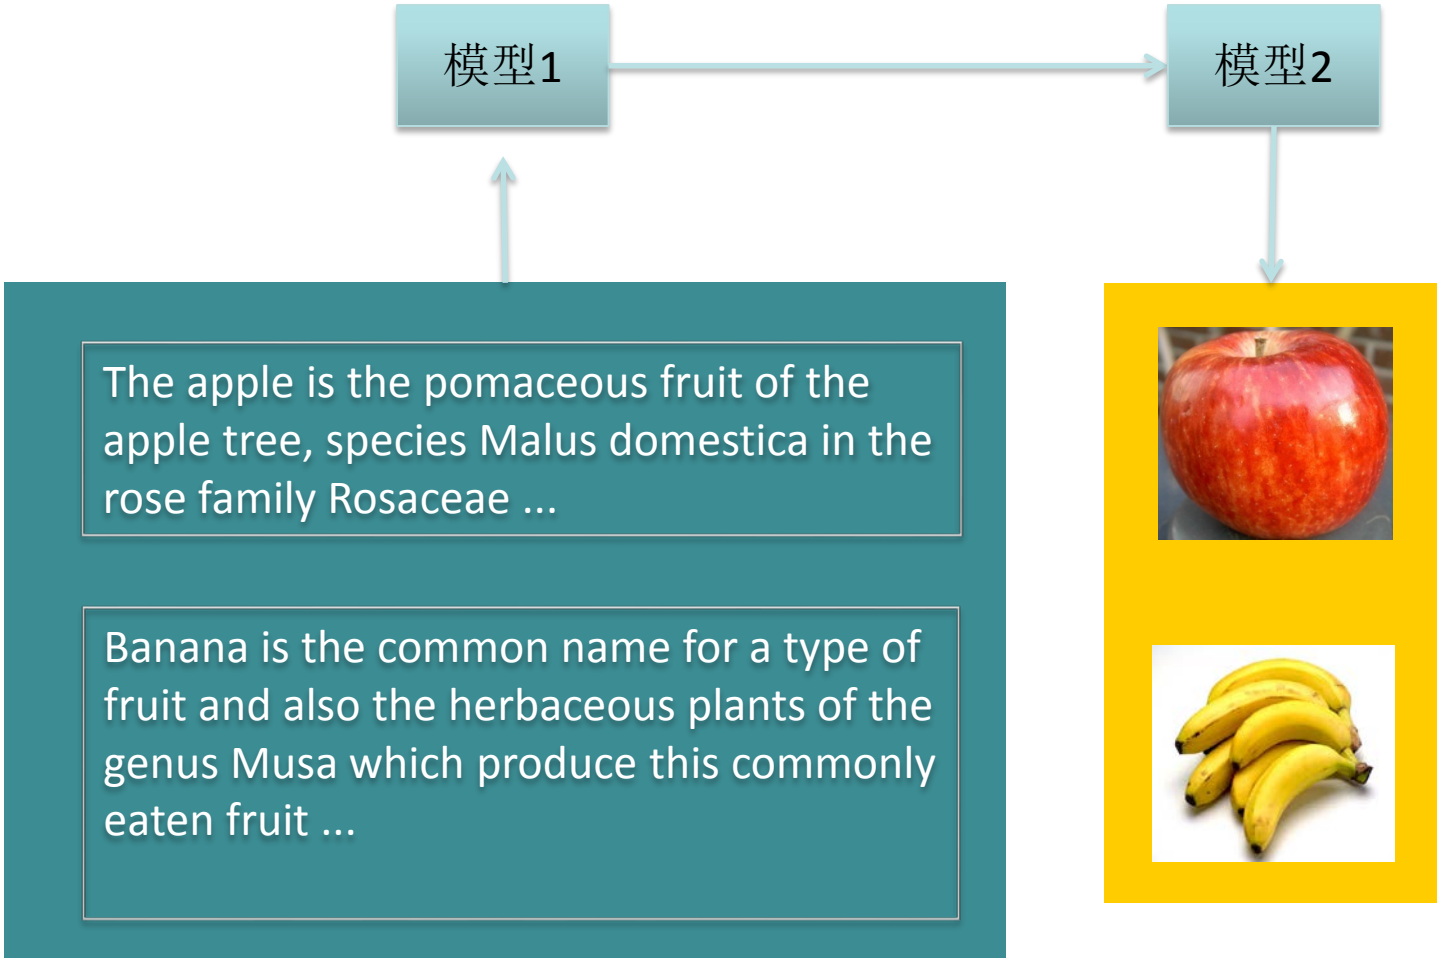

苹果

香蕉

Large-Scale Recommendation Engines

Over 100 Billion Features

$$x = (1,0,0,0,1,\ldots,1,0,\ldots,0,1,0,0)$$

User Demographic

User Relevance to Target

Recent User Behavior

Human Managed
Vs.
Automatically Generated and Maintained

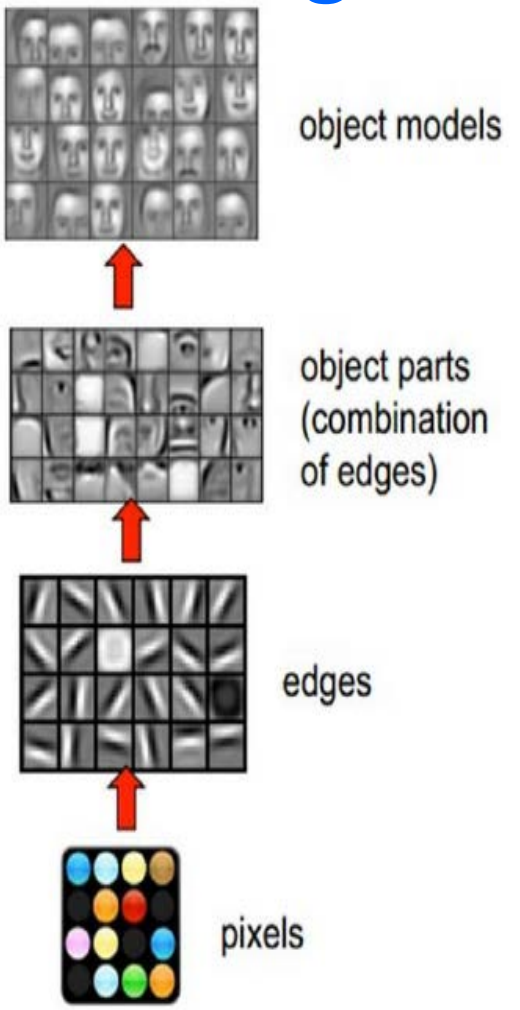

Discriminative → Generative Models

VQA: Visual Question and Answers

What color are her eyes?
What is the mustache made of?

How many slices of pizza are there?
Is this a vegetarian pizza?

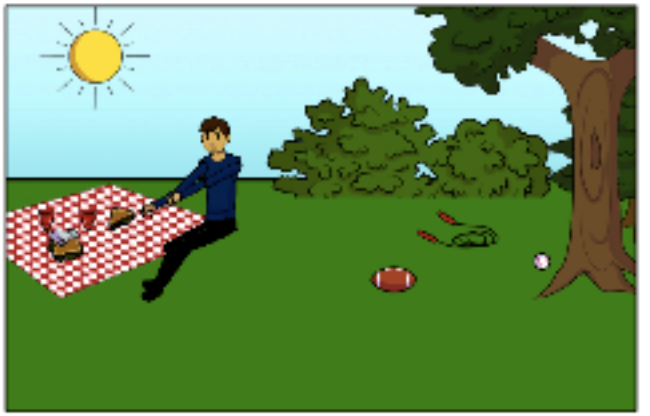

Is this person expecting company?
What is just under the tree?

Does it appear to be rainy?
Does this person have 20/20 vision?

Figure 1: Examples of free-form, open-ended questions collected for images via Amazon Mechanical Turk. Note that common-sense knowledge is needed along with a visual understanding of

Single Sample Learning: Bayesian Program Learning

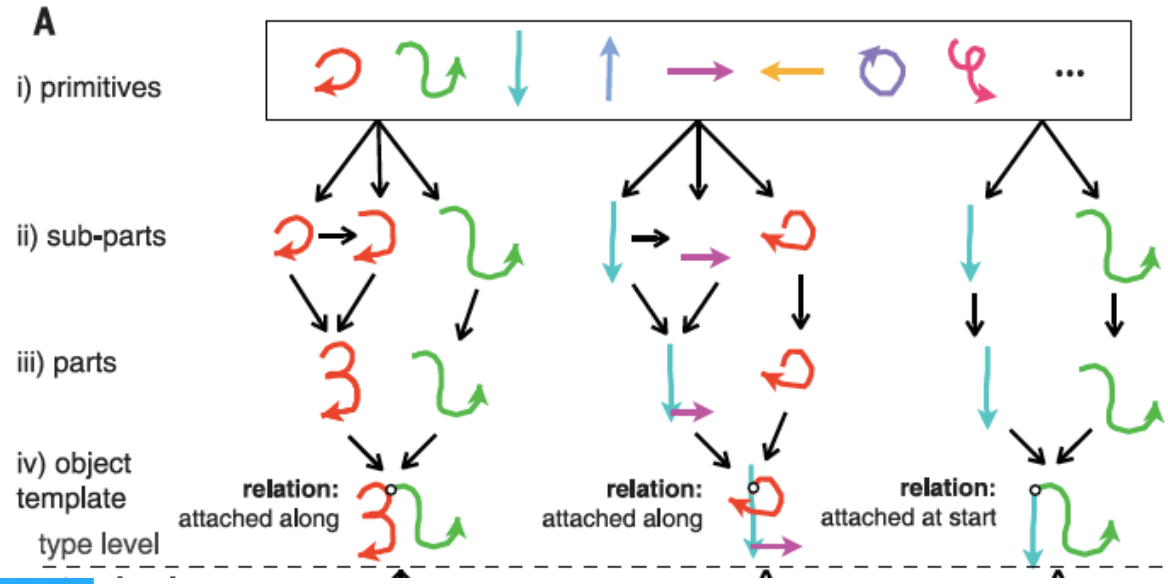

B

procedure GENERATETYPE
- $\kappa \leftarrow P(\kappa)$ ▷ Sample number of parts
- **for** $i = 1 \ldots \kappa$ **do**
  - $n_i \leftarrow P(n_i|\kappa)$ ▷ Sample number of sub-parts
  - **for** $j = 1 \ldots n_i$ **do**
    - $s_{ij} \leftarrow P(s_{ij}|s_{i(j-1)})$ ▷ Sample sub-part sequence
  - **end for**
  - $R_i \leftarrow P(R_i|S_1, \ldots, S_{i-1})$ ▷ Sample relation
- **end for**
- $\psi \leftarrow \{\kappa, R, S\}$
- **return** @GENERATETOKEN(ψ) ▷ Return program

Science December 2015

Robots: What Amazon Teaches US

Finance + AI

Millions of Users, Credit ratings **+61.7%**

Partially Observable Markov Decision Models (training 6 years data; test: 2 years data)

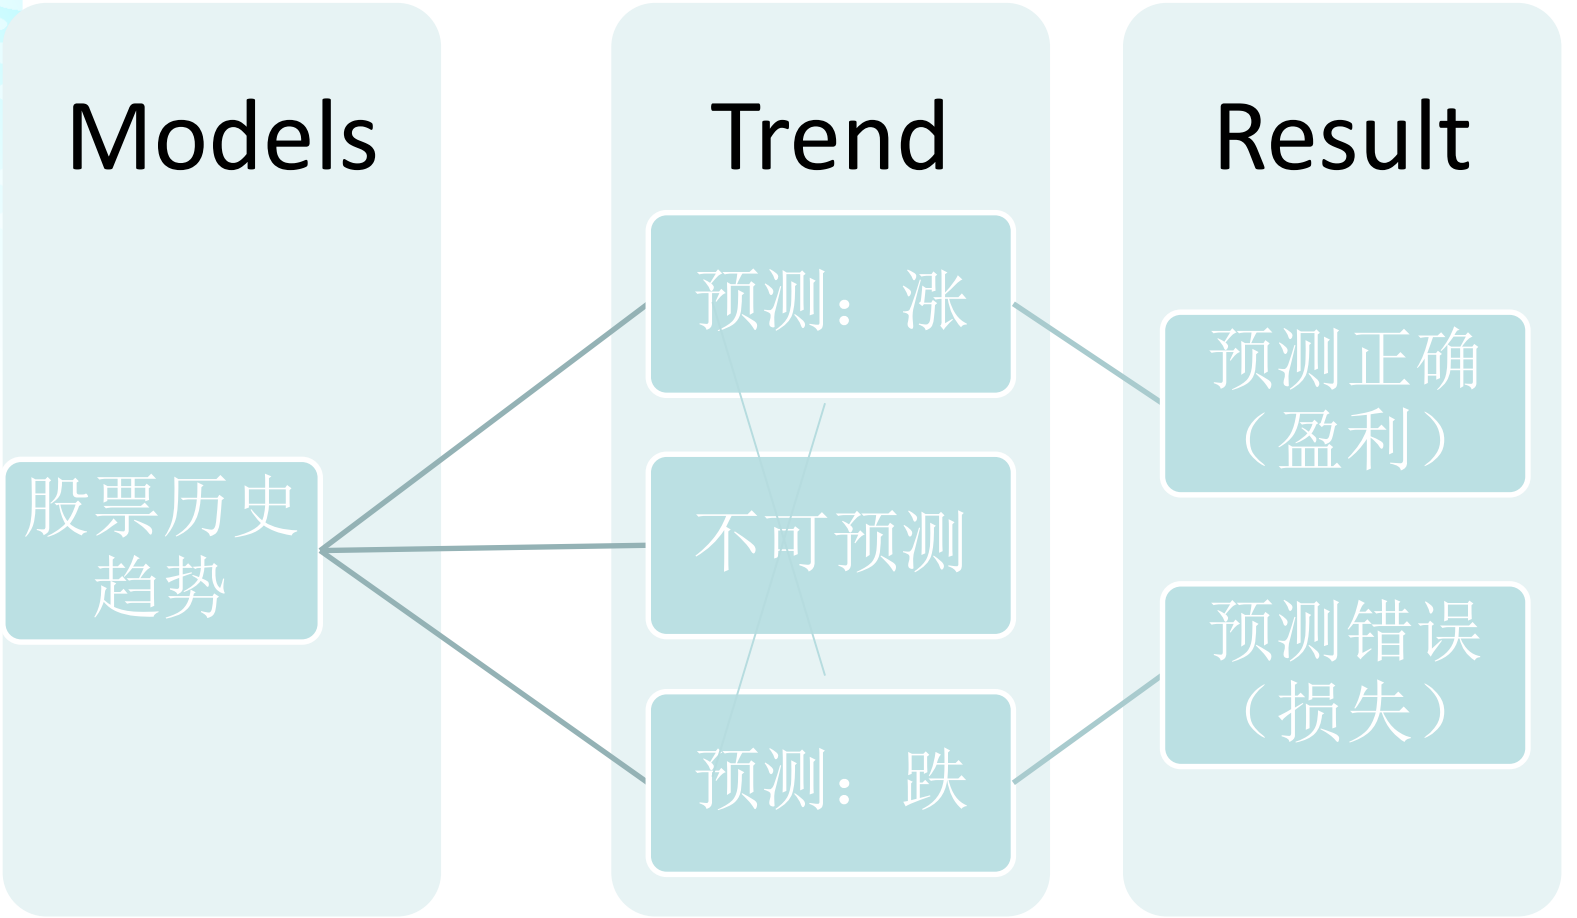

Case Study: Shenzhen/Shanghai Market, POMDP Model

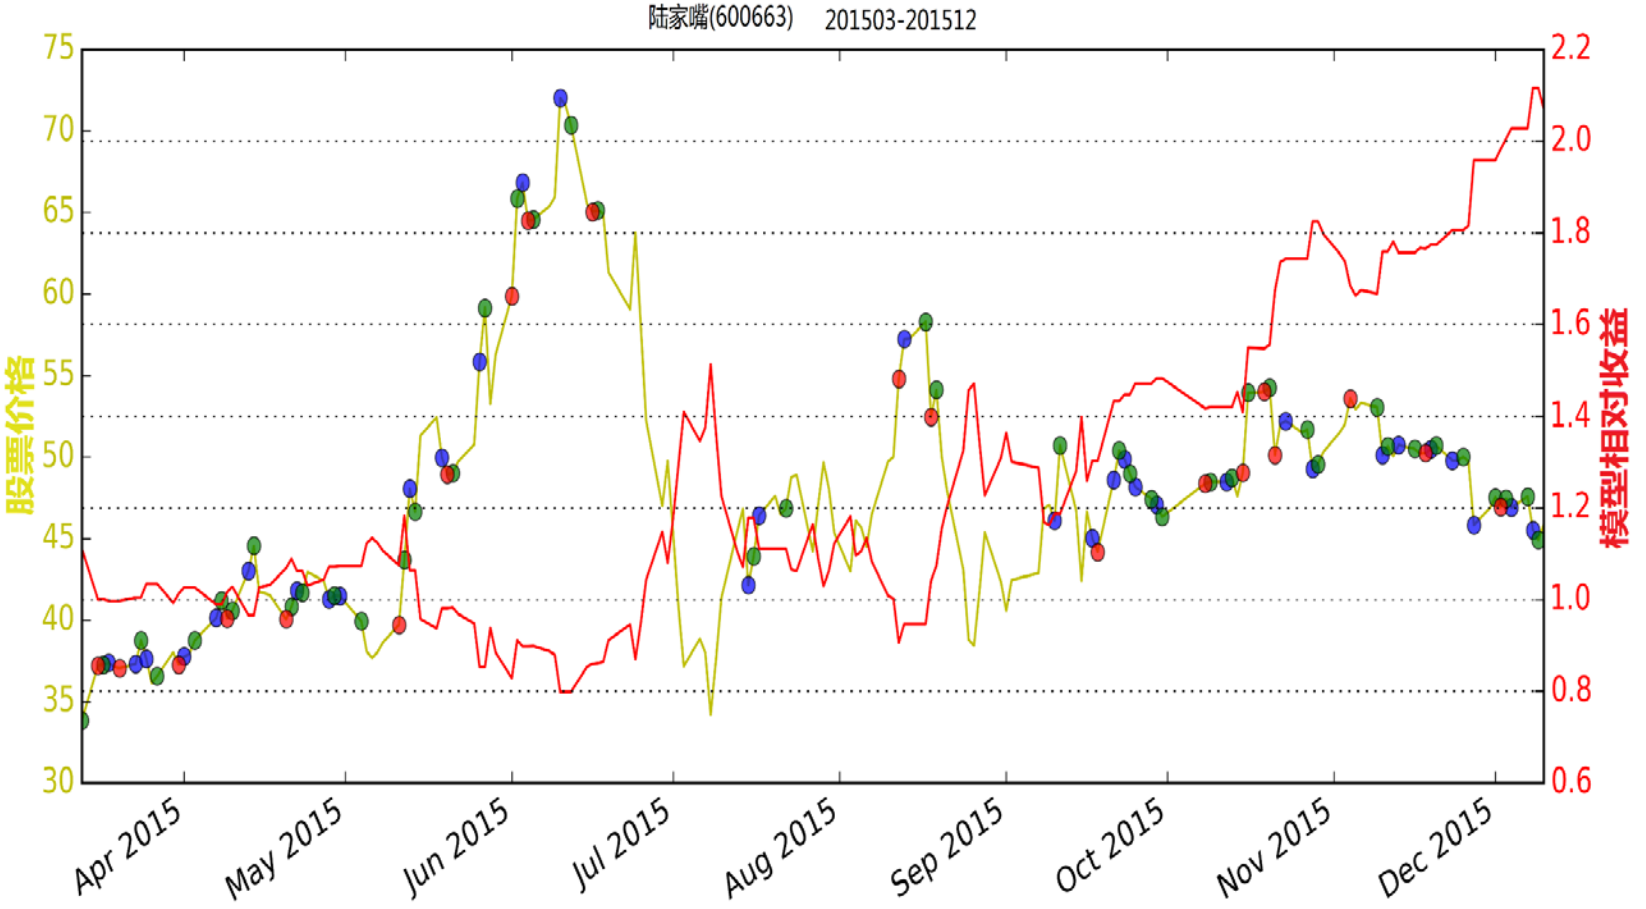

- 模型预测单只股票下一交易日的涨跌 (红线-模型相对收益，黄线-股票实际价格)
  - 红点 – （平空仓，并且）开多仓
  - 绿点 – （平多仓，并且）开空仓
  - 蓝点 – （平仓）空仓

NLP Dialog Agent (based on Data >> 10k dialogs)

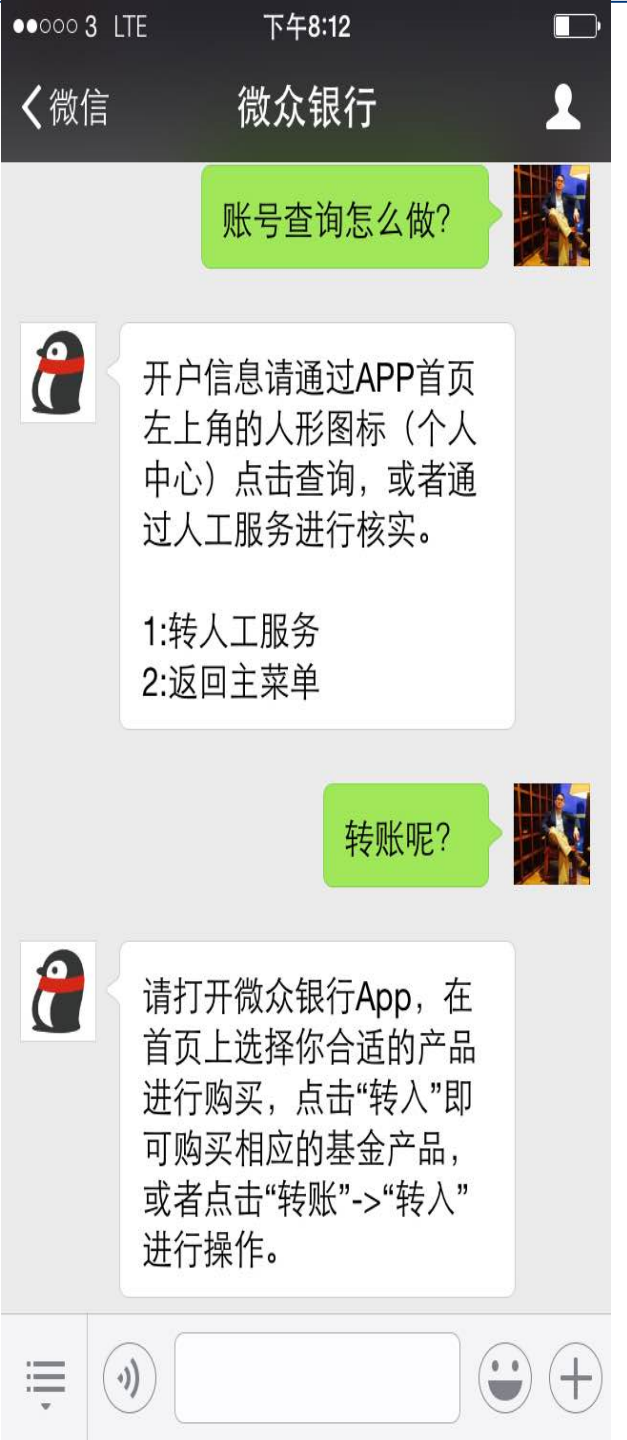

Machine Reading

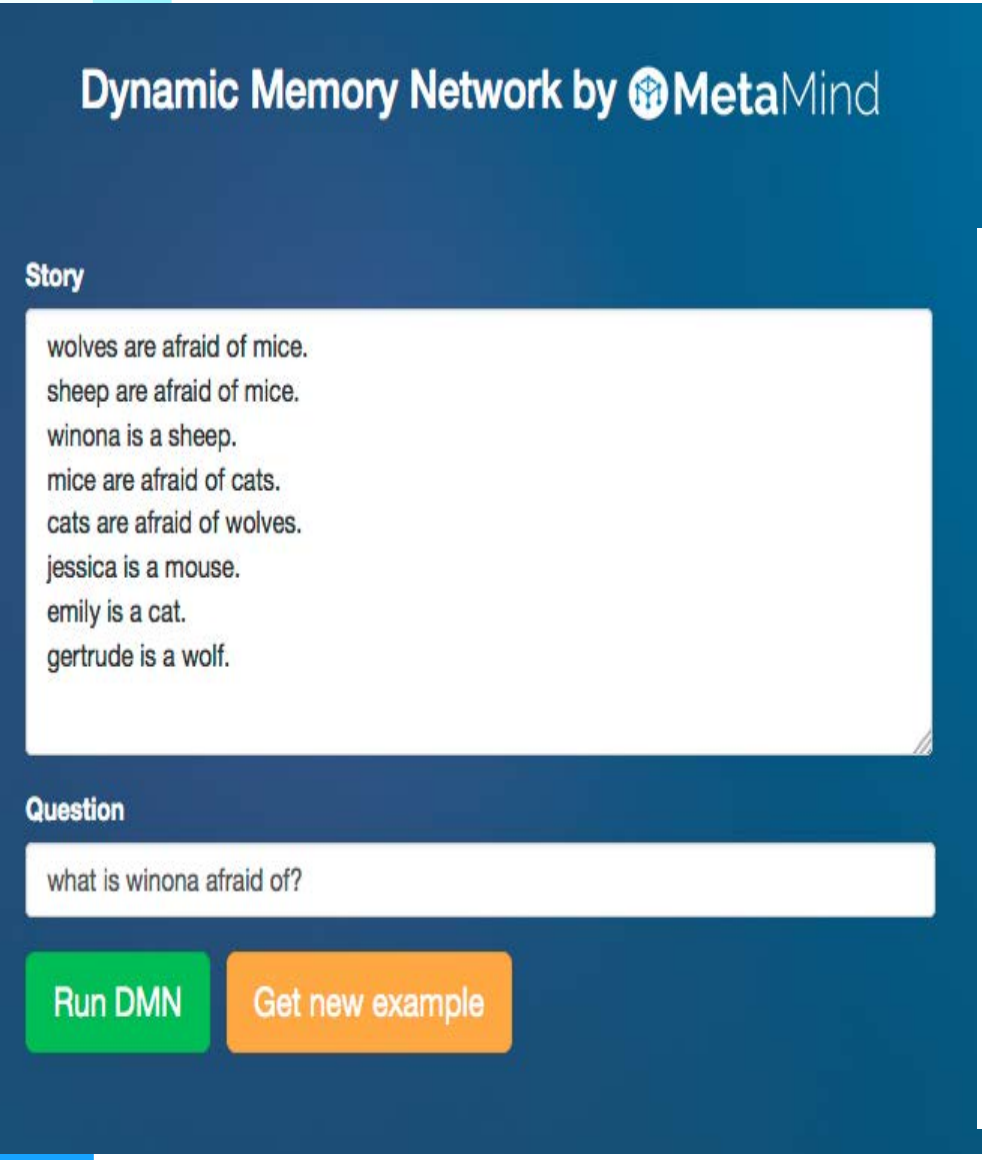

Answer: mouse

Episode 1	Episode 2
0.00: wolves are afraid of mice	**0.00:** wolves are afraid of mice
0.00: sheep are afraid of mice	**1.00:** sheep are afraid of mice
0.99: winona is a sheep	**0.00:** winona is a sheep
0.00: mice are afraid of cats	**0.00:** mice are afraid of cats
0.00: cats are afraid of wolves	**0.00:** cats are afraid of wolves
0.00: jessica is a mouse	**0.00:** jessica is a mouse
0.00: emily is a cat	**0.00:** emily is a cat
0.01: gertrude is a wolf	**0.00:** gertrude is a wolf

Genotype and Phenotype

Input: Very high dimension and low sample size labeled data $(N \approx 2000, D \approx 240K)$

Machine Learning

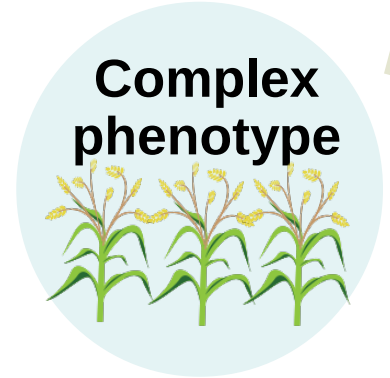

Accelerate Hybridization Breeding

Task: Train an accurate phenotype predictor using genetic data.

E.g. Facilitate understand biology process

markers underlying specific phenotype.

Conclusions

- Current AI technology:
  - Deep Learning: Needs BIG DATA
    - Samples must be sufficient to ensure convergence
  - Need to find complimentary points of Man and Machine
    - AMAZON example
- Future
  - Transfer Learning, One-Example Learning
  - Reinforcement Learning (complete feedback loop)