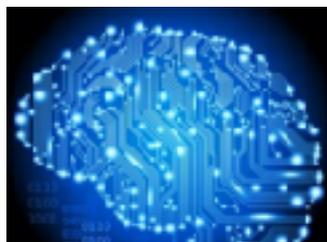


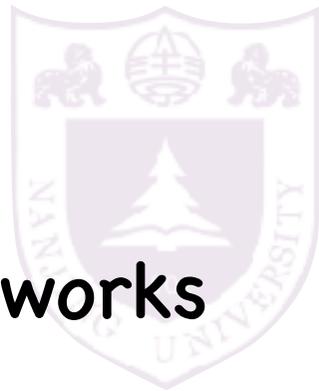
# Lecture 16: Final

# On Artificial Intelligence

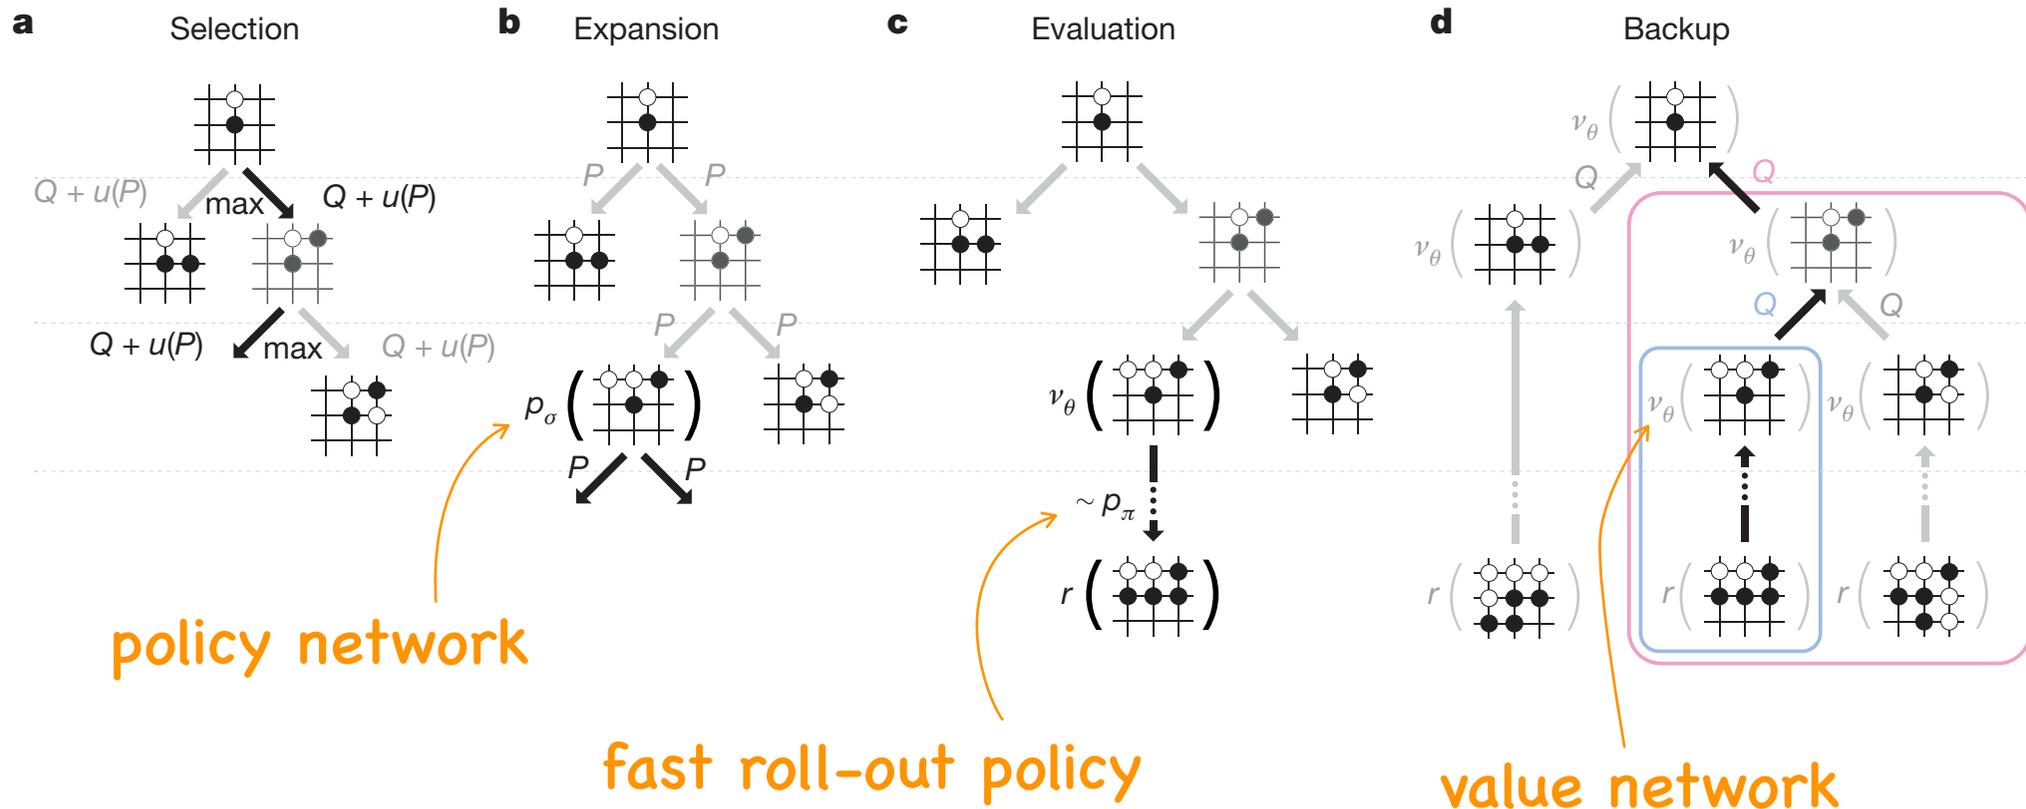
[http://cs.nju.edu.cn/yuy/course\\_ai18.ashx](http://cs.nju.edu.cn/yuy/course_ai18.ashx)



# AlphaGo



A combination of tree search, deep neural networks and reinforcement learning





fast roll-out policy:

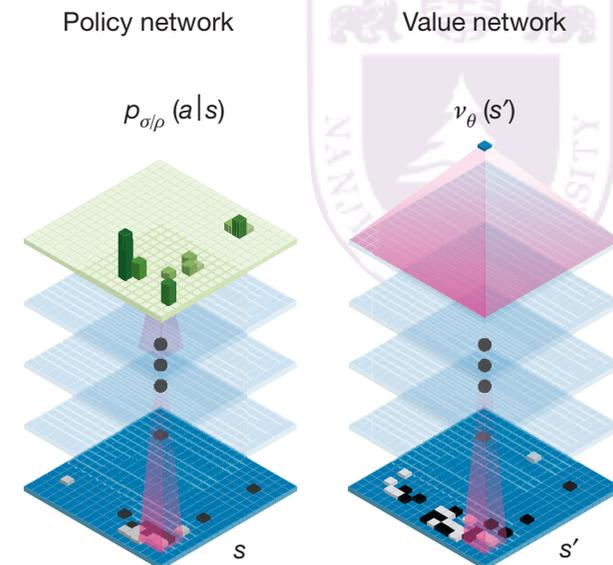
supervised learning from human v.s. human data

Feature	# of patterns	Description
Response	1	Whether move matches one or more response pattern features
Save atari	1	Move saves stone(s) from capture
Neighbour	8	Move is 8-connected to previous move
Nakade	8192	Move matches a <i>nakade</i> pattern at captured stone
Response pattern	32207	Move matches 12-point diamond pattern near previous move
Non-response pattern	69338	Move matches $3 \times 3$ pattern around move
Self-atari	1	Move allows stones to be captured
Last move distance	34	Manhattan distance to previous two moves
Non-response pattern	32207	Move matches 12-point diamond pattern centred around move

# AlphaGo

policy network: a CNN output  $\pi(s,a)$

value network: a CNN output  $V(s)$



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

# AlphaGo



## policy network: initialization

supervised learning from human v.s. human data

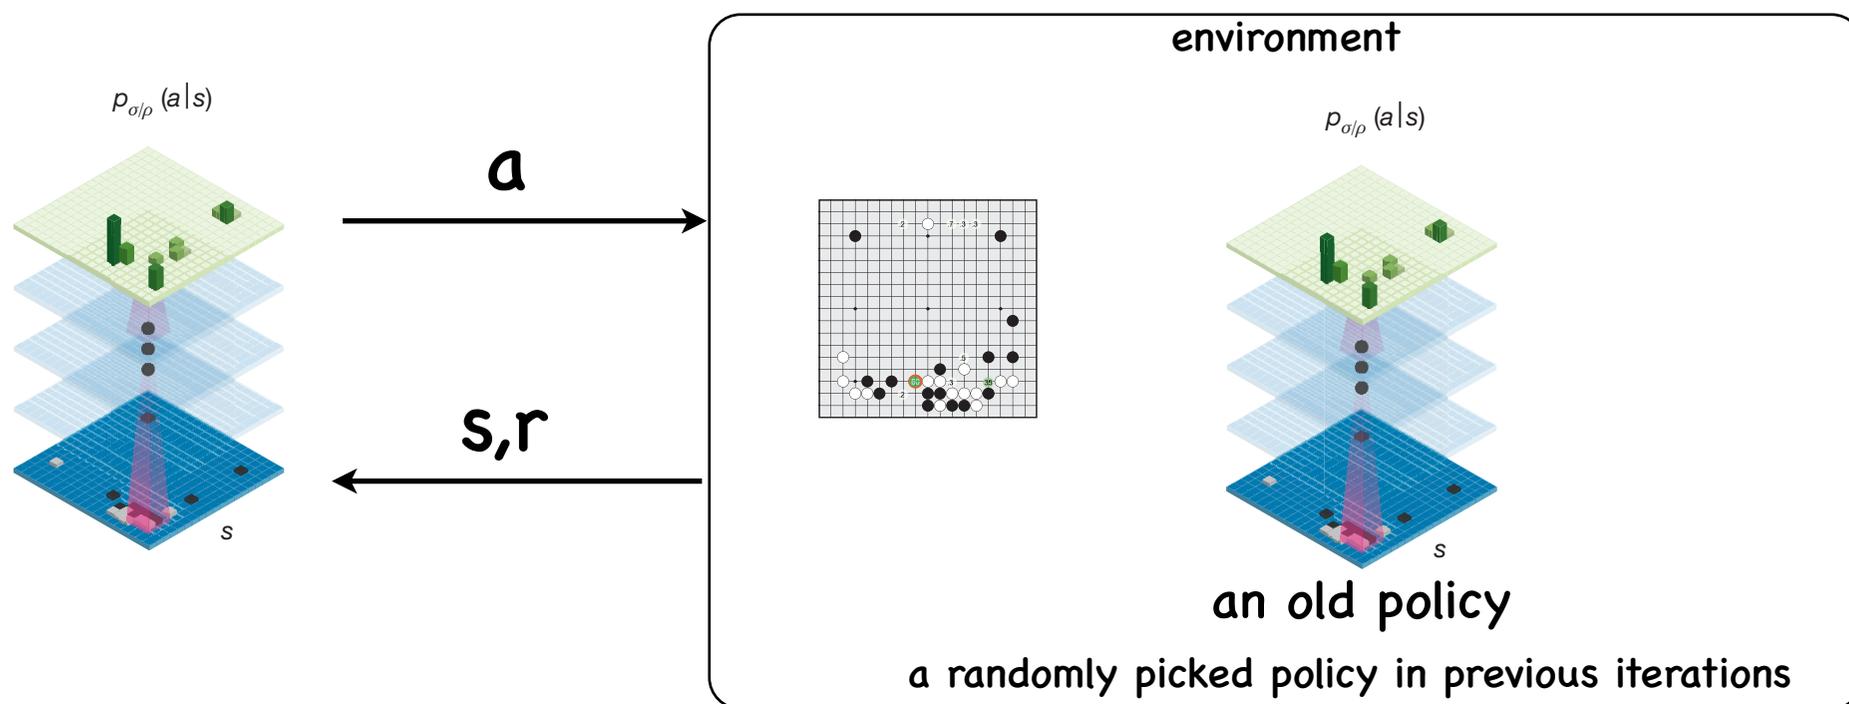
Architecture			Evaluation				
Filters	Symmetries	Features	Test accuracy %	Train accuracy %	Raw net wins %	<i>AlphaGo</i> wins %	Forward time (ms)
128	1	48	54.6	57.0	36	53	2.8
192	1	48	55.4	58.0	50	50	4.8
256	1	48	55.9	59.1	67	55	7.1
256	2	48	56.5	59.8	67	38	13.9
256	4	48	56.9	60.2	69	14	27.6
256	8	48	57.0	60.4	69	5	55.3
192	1	4	47.6	51.4	25	15	4.8
192	1	12	54.7	57.1	30	34	4.8
192	1	20	54.7	57.2	38	40	4.8
192	8	4	49.2	53.2	24	2	36.8
192	8	12	55.7	58.3	32	3	36.8
192	8	20	55.8	58.4	42	3	36.8

# AlphaGo



## policy network: further improvement

### reinforcement learning



a.k.a. self-play

reward:

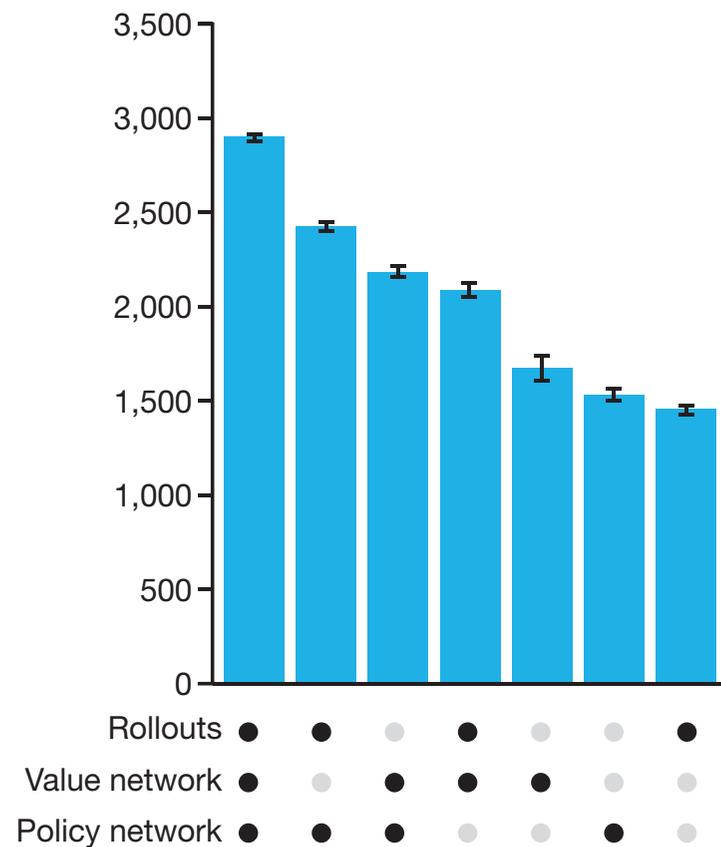
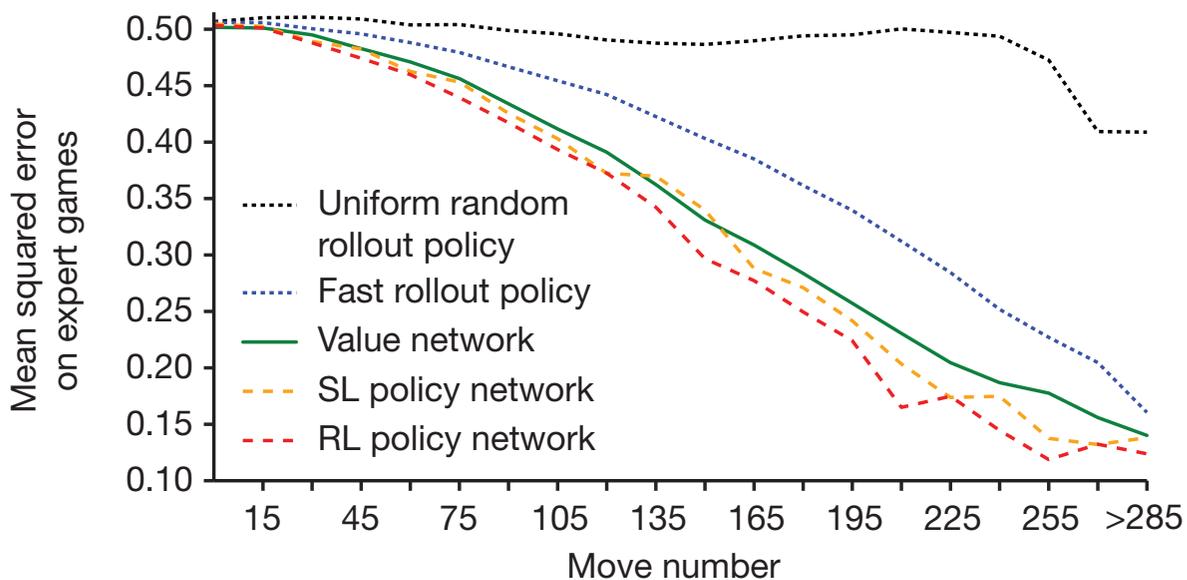
+1 -- win at terminate state

-1 -- loss at terminate state

# AlphaGo

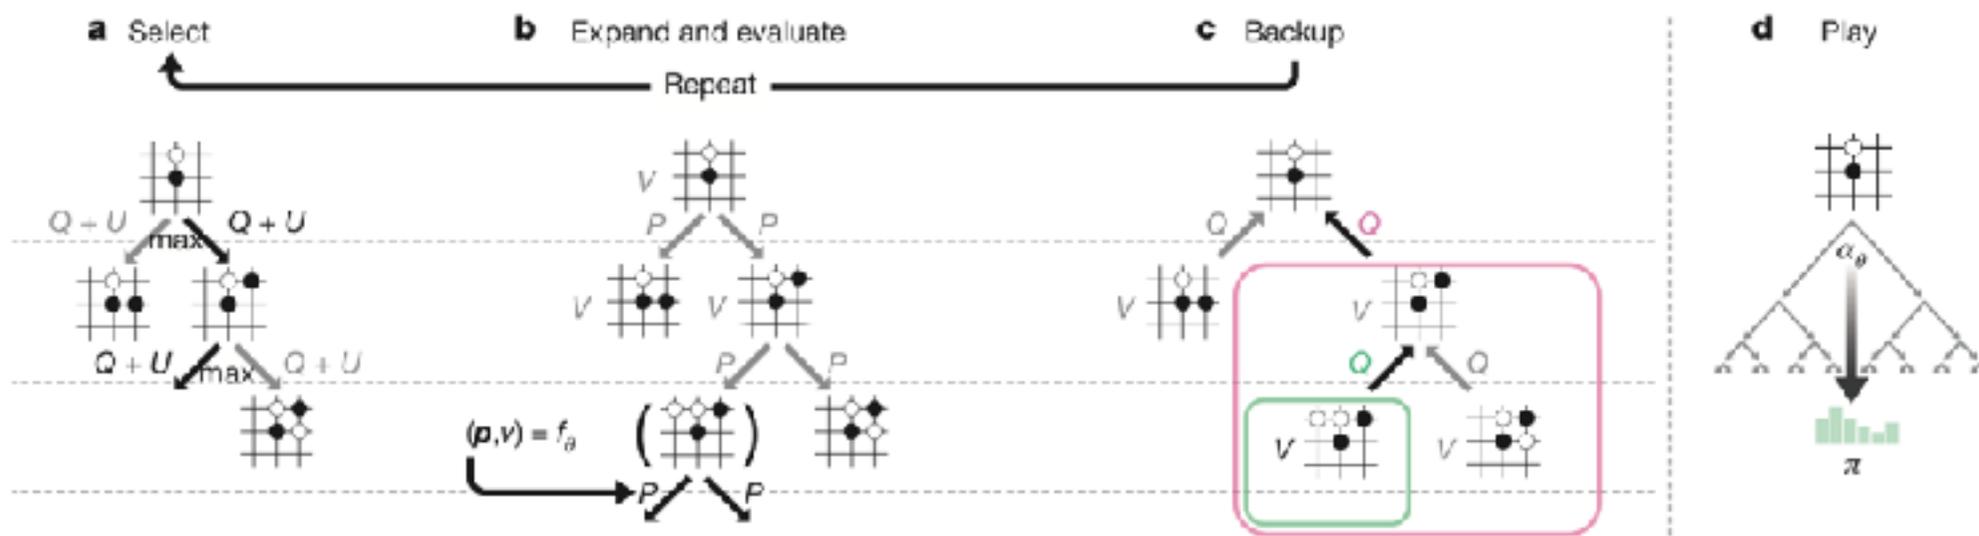


value network: supervised learning from RL data



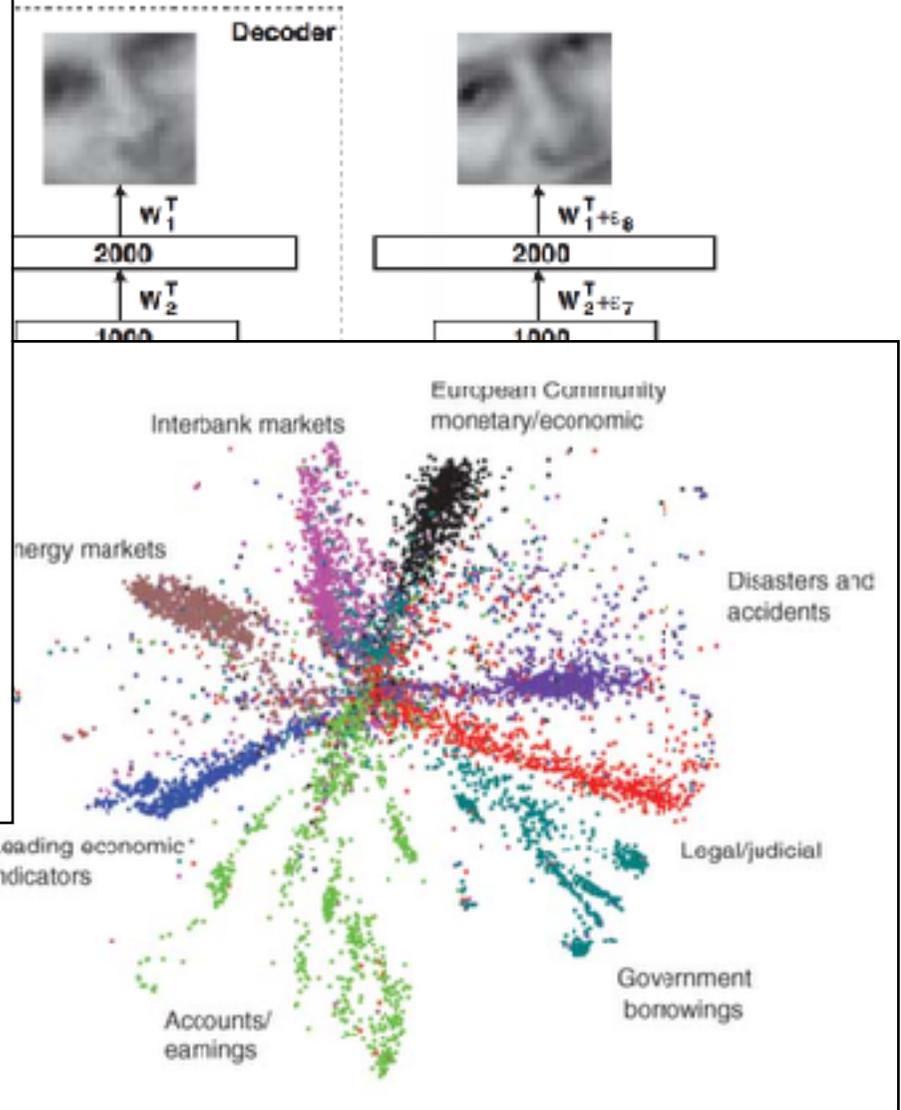
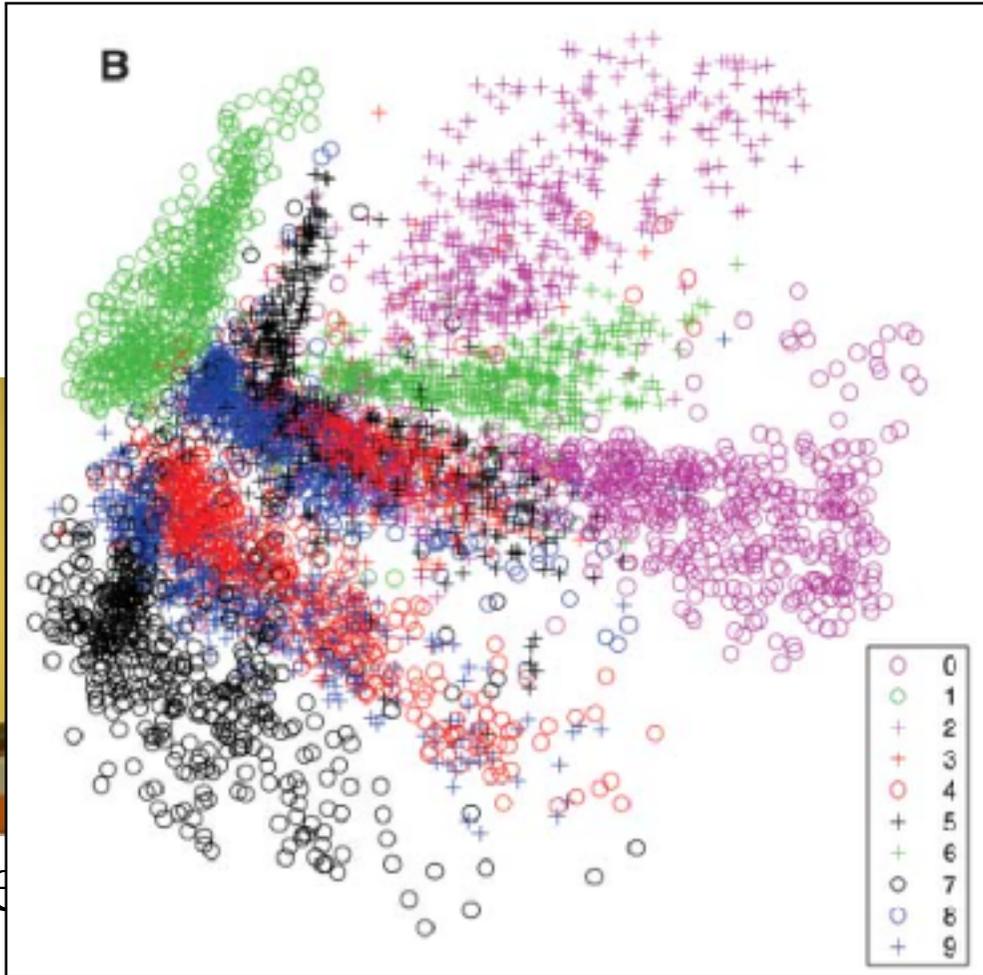
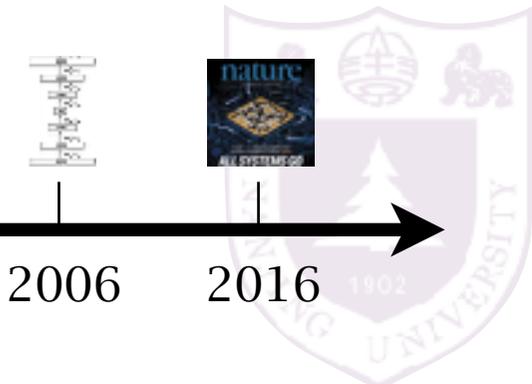
# AlphaGo Zero

No pre-training, No roll-out



Reinforcement learning:  
collect data by MCTS,  
remember by neural networks

# History



Pretraining

RBM End

Ge



**Describes without errors**

**Describes with minor errors**

**Somewhat related to the image**

**Unrelated to the image**



**A person riding a motorcycle on a dirt road.**



**Two dogs play in the grass.**



**A skateboarder does a trick on a ramp.**



**A dog is jumping to catch a frisbee.**



**A group of young people playing a game of frisbee.**



**Two hockey players are fighting over the puck.**



**A little girl in a pink hat is blowing bubbles.**



**A refrigerator filled with lots of food and drinks.**



**A herd of elephants walking across a dry grass field.**



**A close up of a cat laying on a couch.**



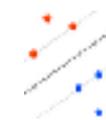
**A red motorcycle parked on the side of the road.**



**A yellow school bus parked in a parking lot.**

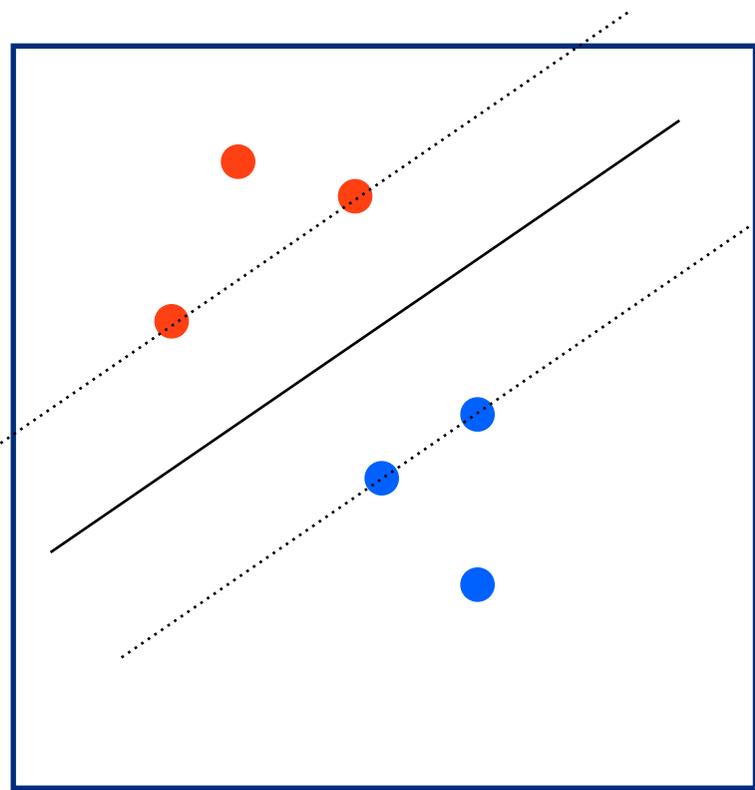
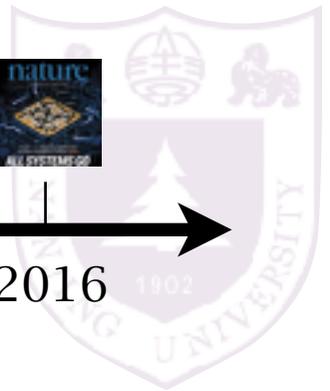


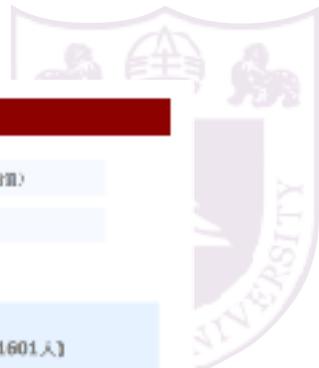
90年代  
中期



1998 2006

2016





### 店铺2级服务概述

好评退款率: 0.0022% 小于 0.0064% (行业均值)      退款自主完成率: 99.90% 大于 95.47% (行业均值)

退款纠纷时长: 0.64天 小于 1.66天 (行业均值)

店铺综合评分: (所属行业: 书籍音像)

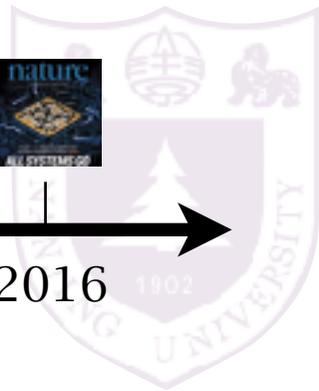
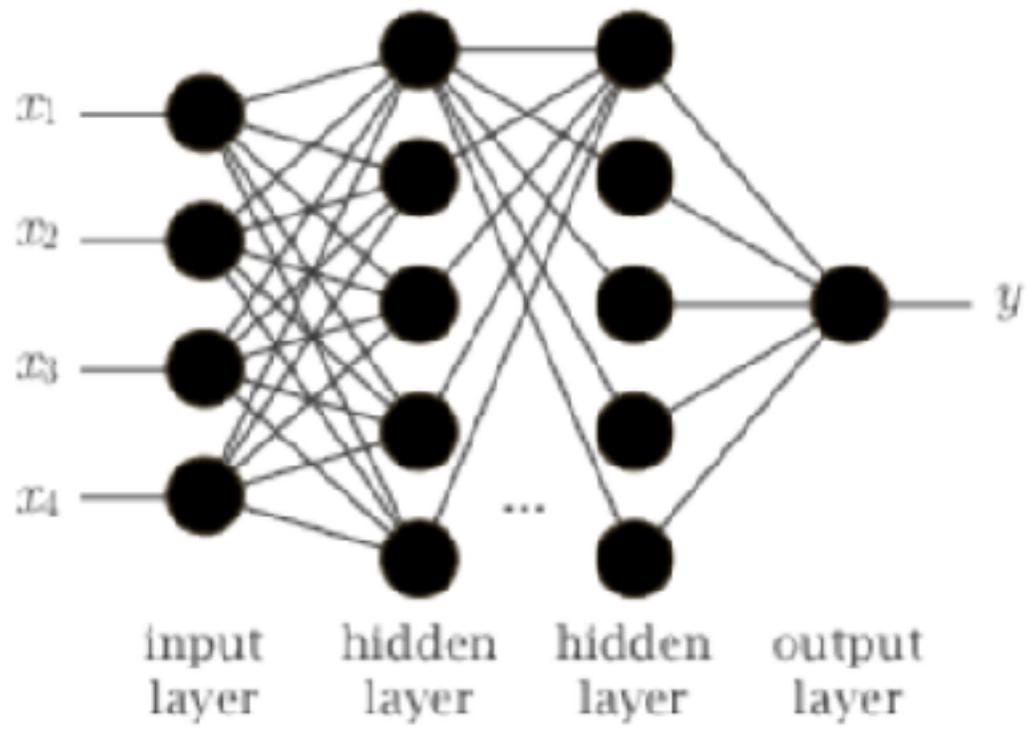
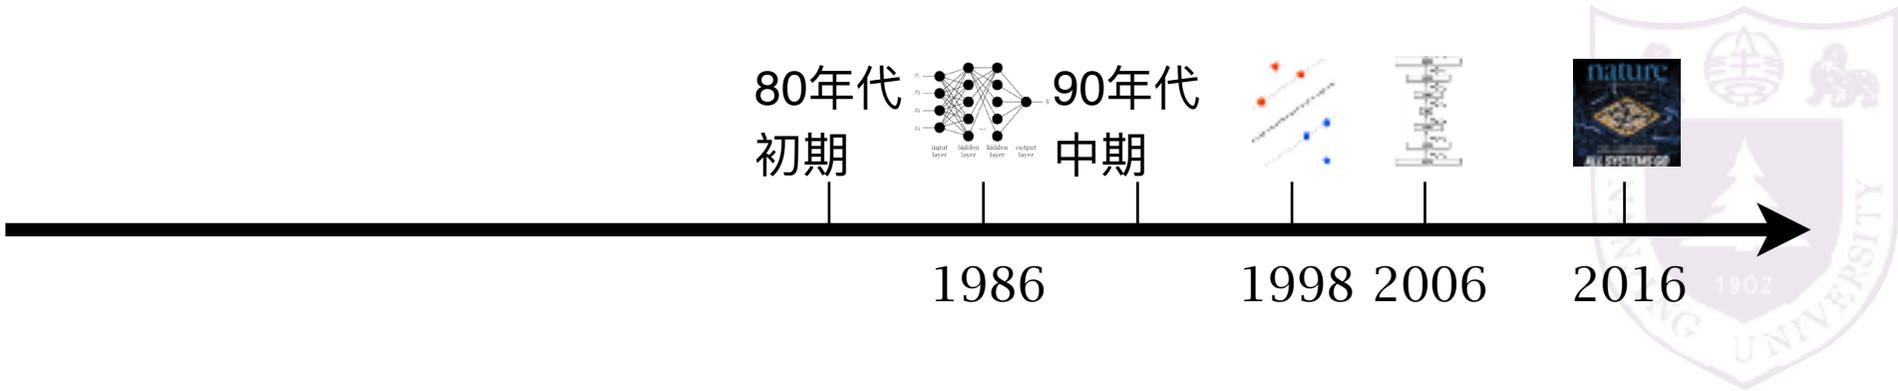
商品与服务指标: 4.8分	比同行平均水平 低0.86%	5星评价: ★★★★★ 共3476人
商家响应速度: 4.7分	比同行平均水平 低0.82%	5分: 3.99% (1362人)
商家发货速度: 4.7分	比同行平均水平 低1.89%	4分: 1.61% (551人)
		3分: 0.58% (198人)
		2分: 1.36% (464人)
		1分: 1.36% (464人)

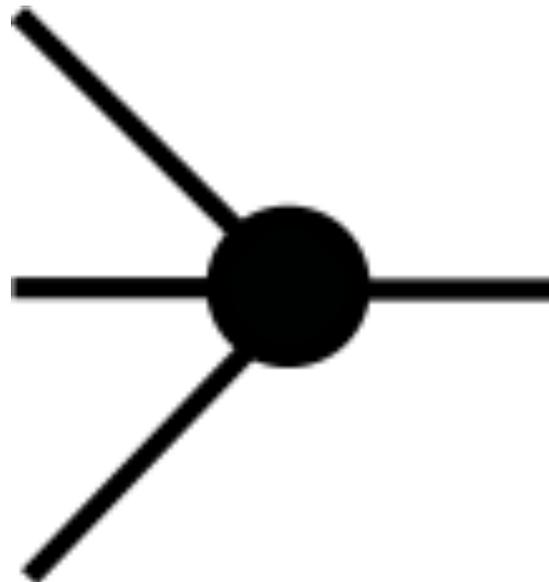
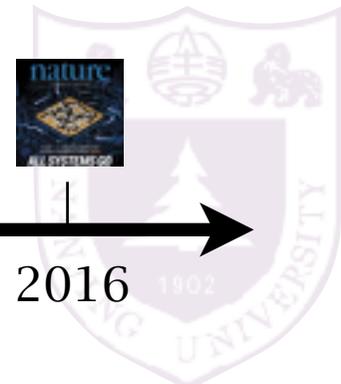
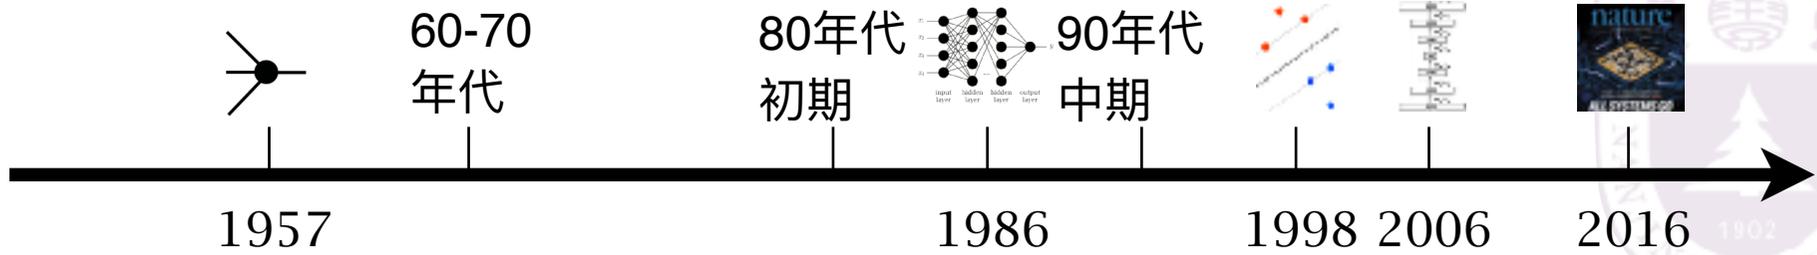
商家承诺: 凡使用支付宝交易付款购买商品, 若存在该商家发货延迟不期, 本店支持退货退款服务并承诺包邮费!

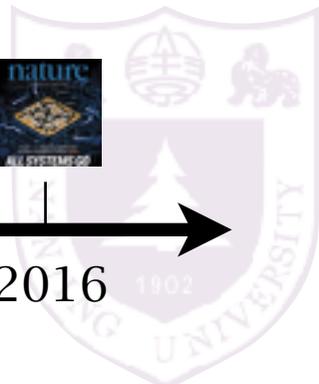
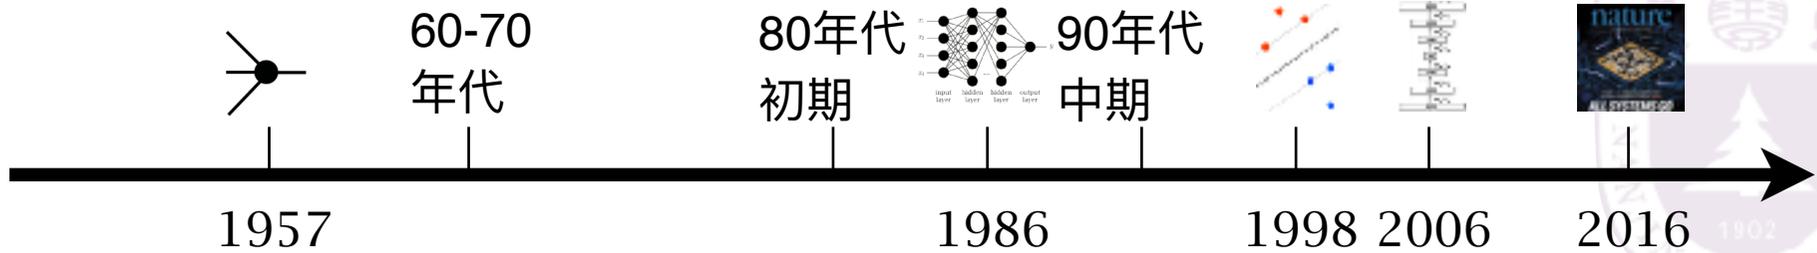
### 畅销榜单

<p>Head First C#</p> <p>¥38.8</p>	<p>人脸识别原理及算法</p> <p>¥65.9</p>	<p>数学之美</p> <p>¥35</p>
<p>智能车辆导航技术</p> <p>¥51.8</p>	<p>视觉机器学习</p> <p>¥42</p>	<p>Offer</p> <p>¥41</p>
<p>统计学习方法</p> <p>¥28.5</p>	<p>程序员的数学</p> <p>¥153.2</p>	<p>深度学习 方法及应用</p> <p>¥35.9</p>





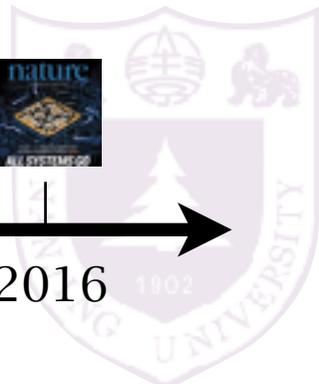
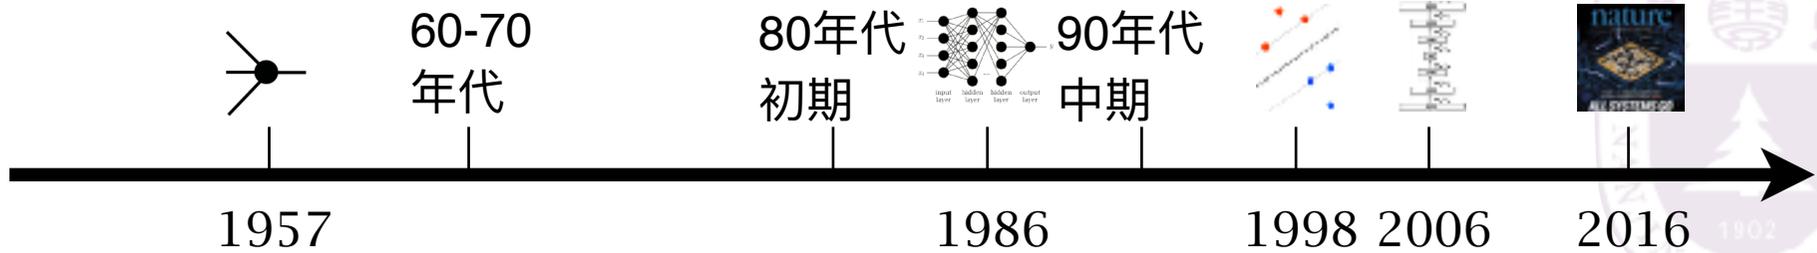




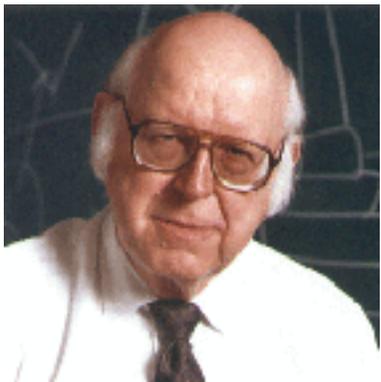
Edward Albert Feigenbaum

## 专家系统





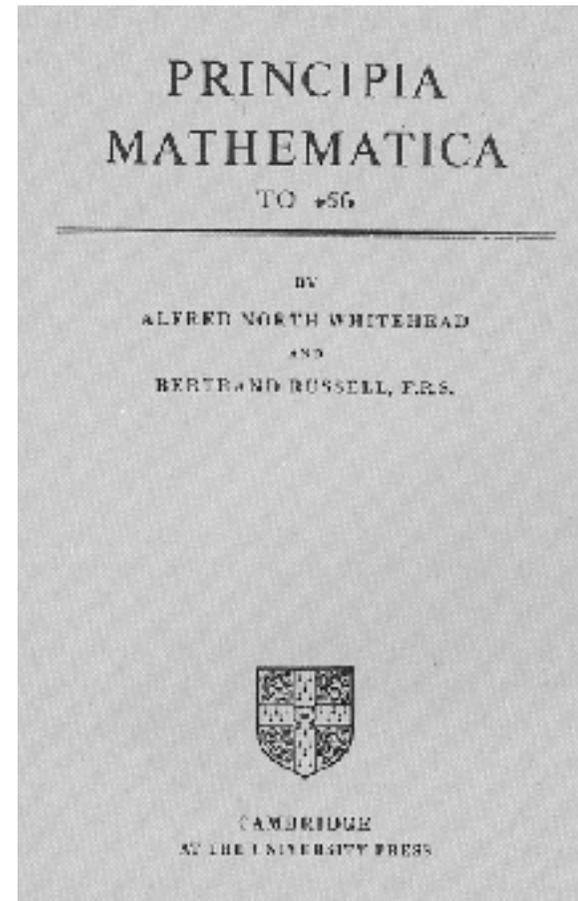
## 逻辑学家



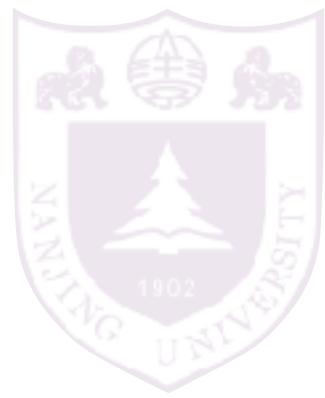
Allen Newell



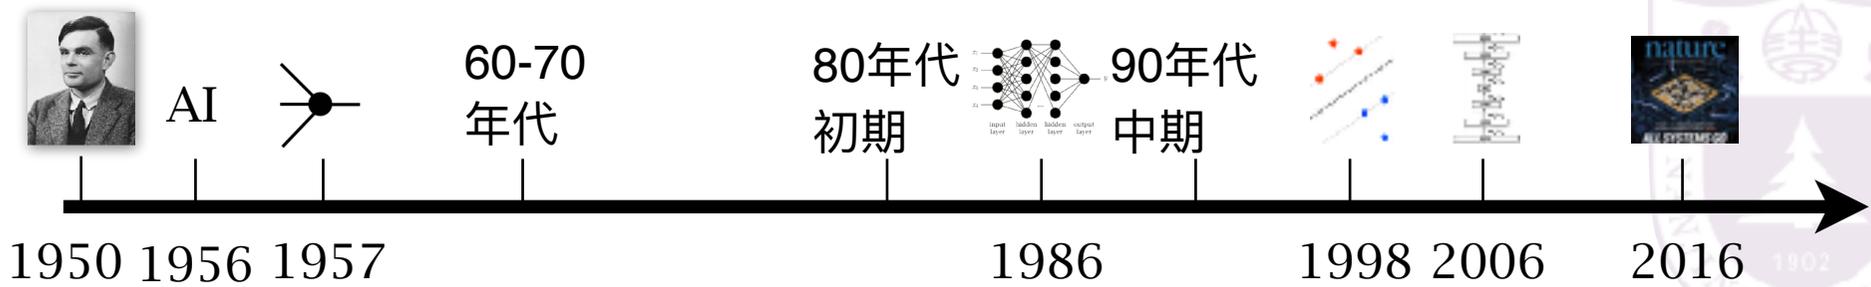
Herbert Simon



# 1956 Dartmouth meeting: “Artificial Intelligence” adopted



- |                      |  |
|----------------------|--|
| <b>John McCarthy</b> | Turing Award (1971)                                  |
| Marvin Minsky        | Turing Award (1969)                                  |
| Claude Shannon       | the father of information theory                     |
| Oliver Selfridge     | father of machine perception                         |
| Herbert A. Simon     | Turing Award (1975), Nobel Prize in Economics (1978) |
| Allen Newell         | Turing Award (1975)                                  |

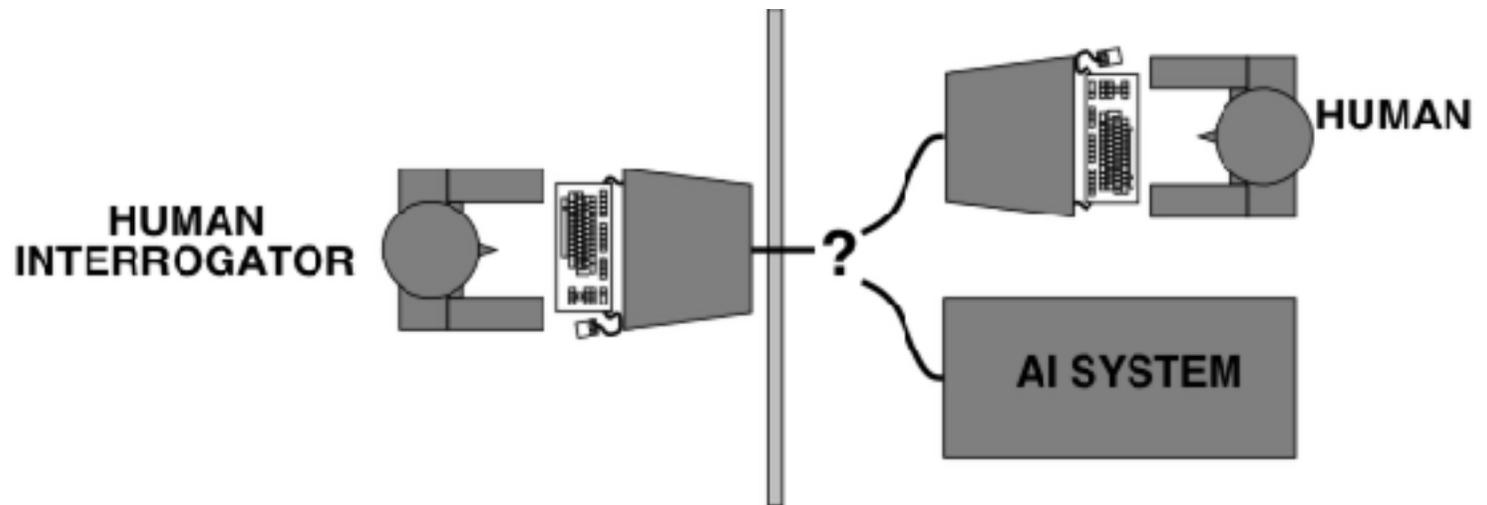


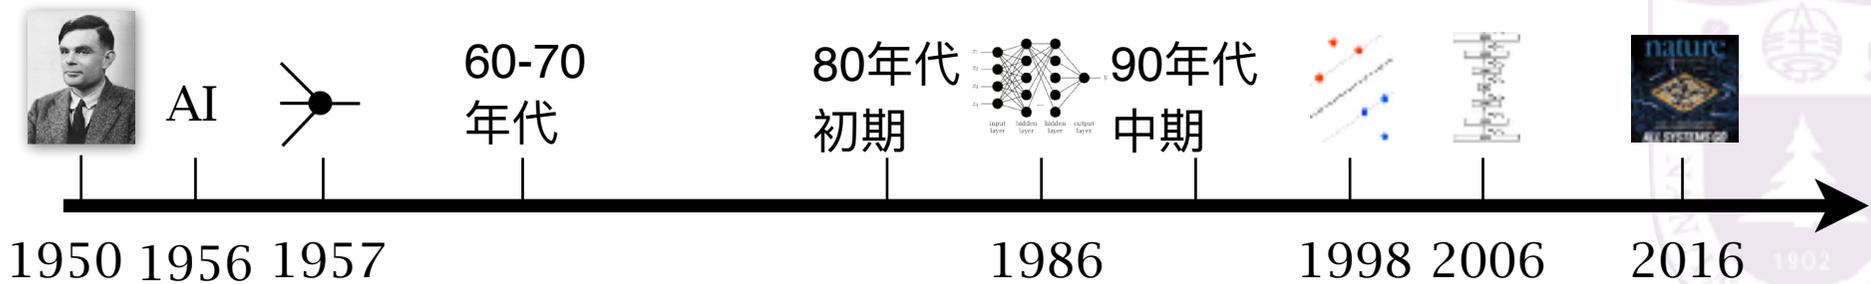
[*Computing machinery and intelligence. Mind* 49: 433-460, 1950.]



Alan Turing  
1912-1954

## Section 1: Imitation game





[*Computing machinery and intelligence. Mind* 49: 433-460, 1950.]



Alan Turing  
1912-1954

## Section 7: Learning machines

“In the process of trying to imitate an adult human mind we are bound to think a good deal about the process which has brought it to the state that it is in. We may notice three components.

- (a) The initial state of the mind, say at birth,
- (b) The education to which it has been subjected,
- (c) Other experience, not to be described as education, to which it has been subjected.”

Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's?

# History

学习期



知识期

推理期

AI之冬

AI之冬

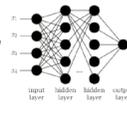


AI

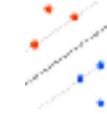


60-70年代

80年代初期



90年代中期



1950 1956 1957

1986

1998 2006

2016

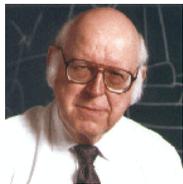
“解决了神秘的心/身问题，解释了物质构成的系统如何获得心灵的性质。”

“十年之内，数字计算机将成为国际象棋世界冠军。”

“二十年内，机器将能完成人能做到的一切工作。”

“一代之内……创造‘人工智能’的问题将获得实质上的解决。”

“在三到八年的时间里我们将得到一台具有人类平均智能的机器。”



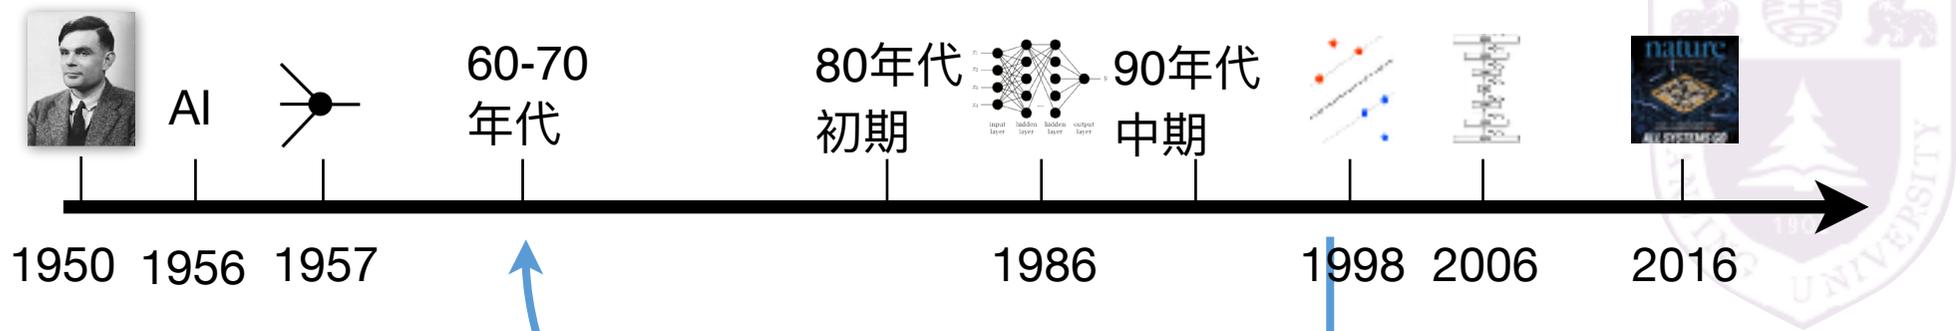
Allen Newell



Herbert Simon



Marvin Minsky



1997 深蓝

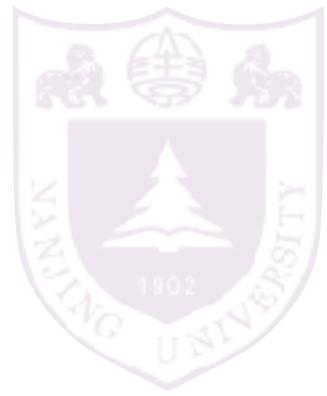


# Potted history of AI

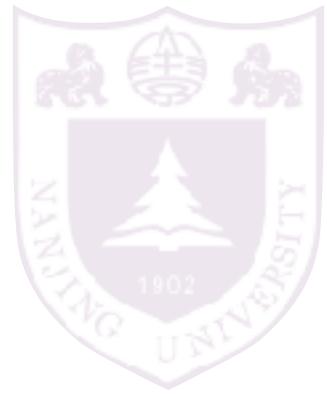


- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1952–69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 **Dartmouth meeting: "Artificial Intelligence" adopted**
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity  
Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1980–88 Expert systems industry booms
- 1988–93 Expert systems industry busts: "AI Winter"
- 1985–95 Neural networks return to popularity
- 1988– Resurgence of probability; general increase in technical depth  
"Nouvelle AI": ALife, GAs, soft computing
- 1995– Agents, agents, everywhere . . .
- 1990– Machine learning quickly develops
- 2003– Human-level AI back on the agenda
- 2006– Machine learning industry booms

# What is intelligence?



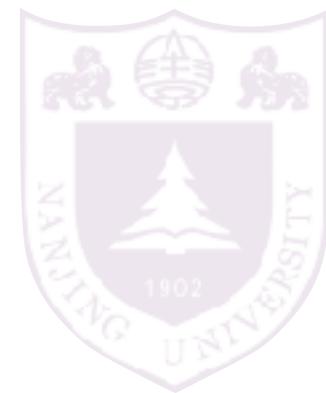
# What is intelligence?



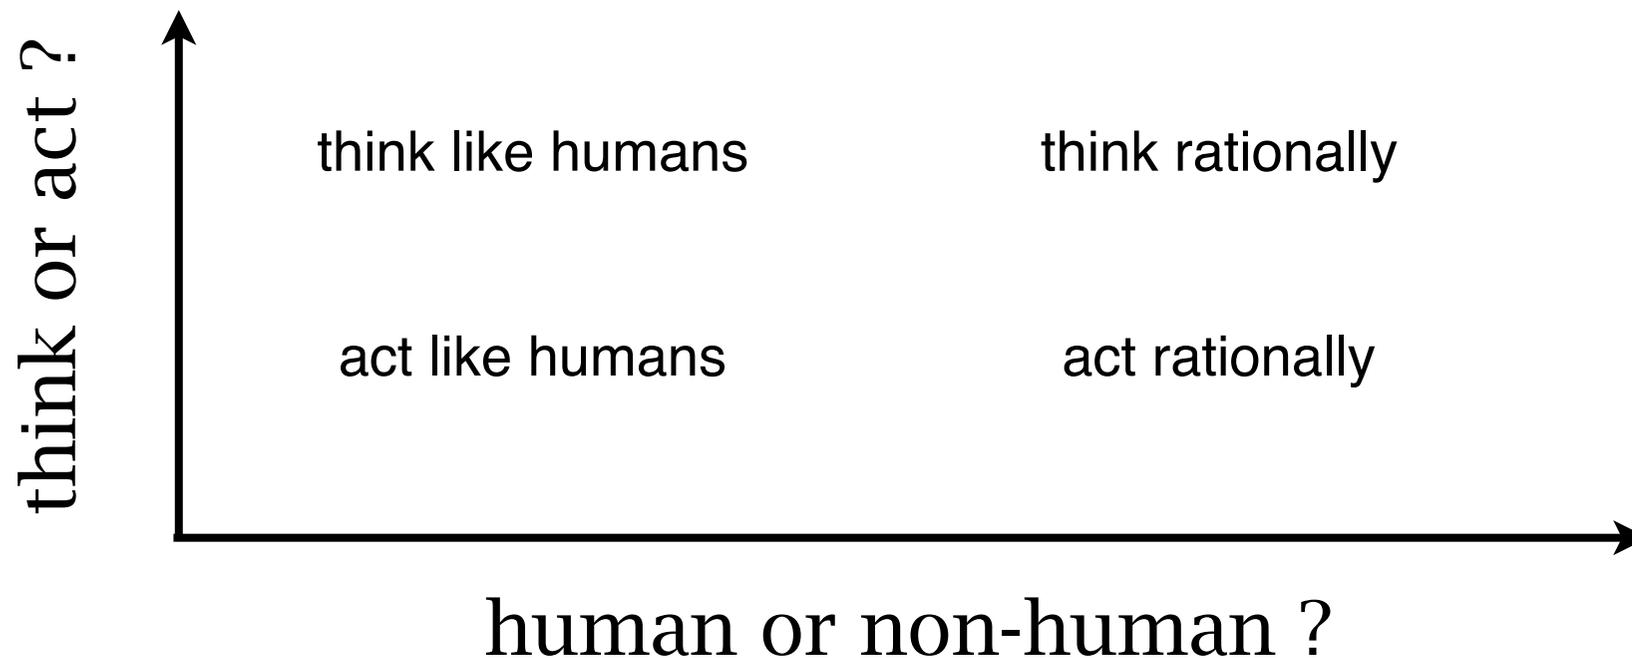
The uncertain about intelligence is a fundamental problem of AI



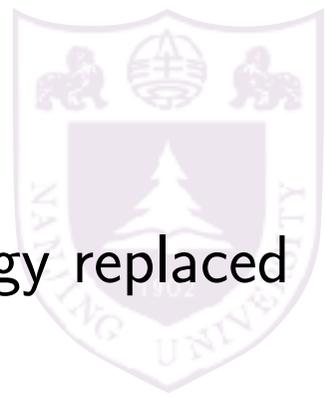
# What is AI?



AI is a system that



# Thinking humanly: Cognitive Science



1960s “cognitive revolution”: information-processing psychology replaced prevailing orthodoxy of behaviorism

Requires scientific theories of internal activities of the brain

- What level of abstraction? “Knowledge” or “circuits”?
- How to validate? Requires
  - 1) Predicting and testing behavior of human subjects (top-down)
  - or 2) Direct identification from neurological data (bottom-up)

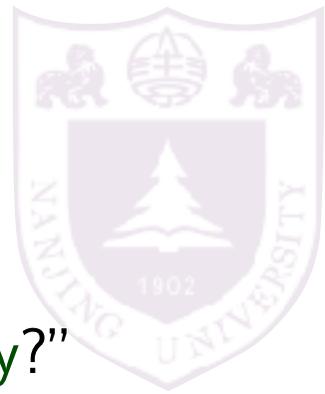
Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Both share with AI the following characteristic:

**the available theories do not explain (or engender) anything resembling human-level general intelligence**

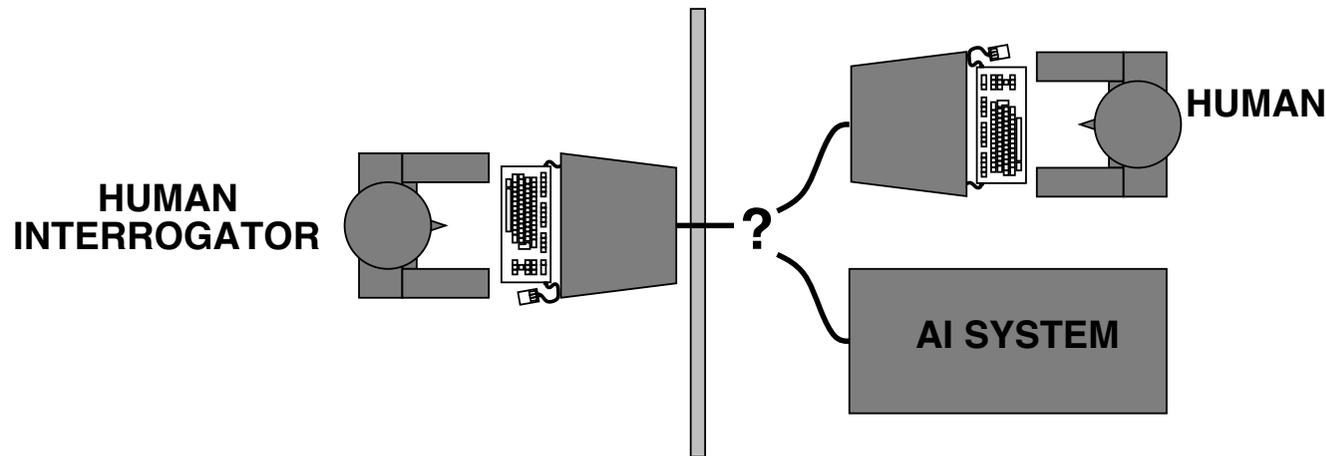
Hence, all three fields share one principal direction!

# Acting humanly: The Turing test



Turing (1950) “Computing machinery and intelligence”:

- ◇ “Can machines think?” → “Can machines behave intelligently?”
- ◇ Operational test for intelligent behavior: the **Imitation Game**



- ◇ Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- ◇ Anticipated all major arguments against AI in following 50 years
- ◇ Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

# Thinking rationally: Laws of Thought



Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:

**notation** and **rules of derivation** for thoughts;  
may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:

- 1) Not all intelligent behavior is mediated by logical deliberation
- 2) **What is the purpose of thinking?** What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

# Acting rationally



Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

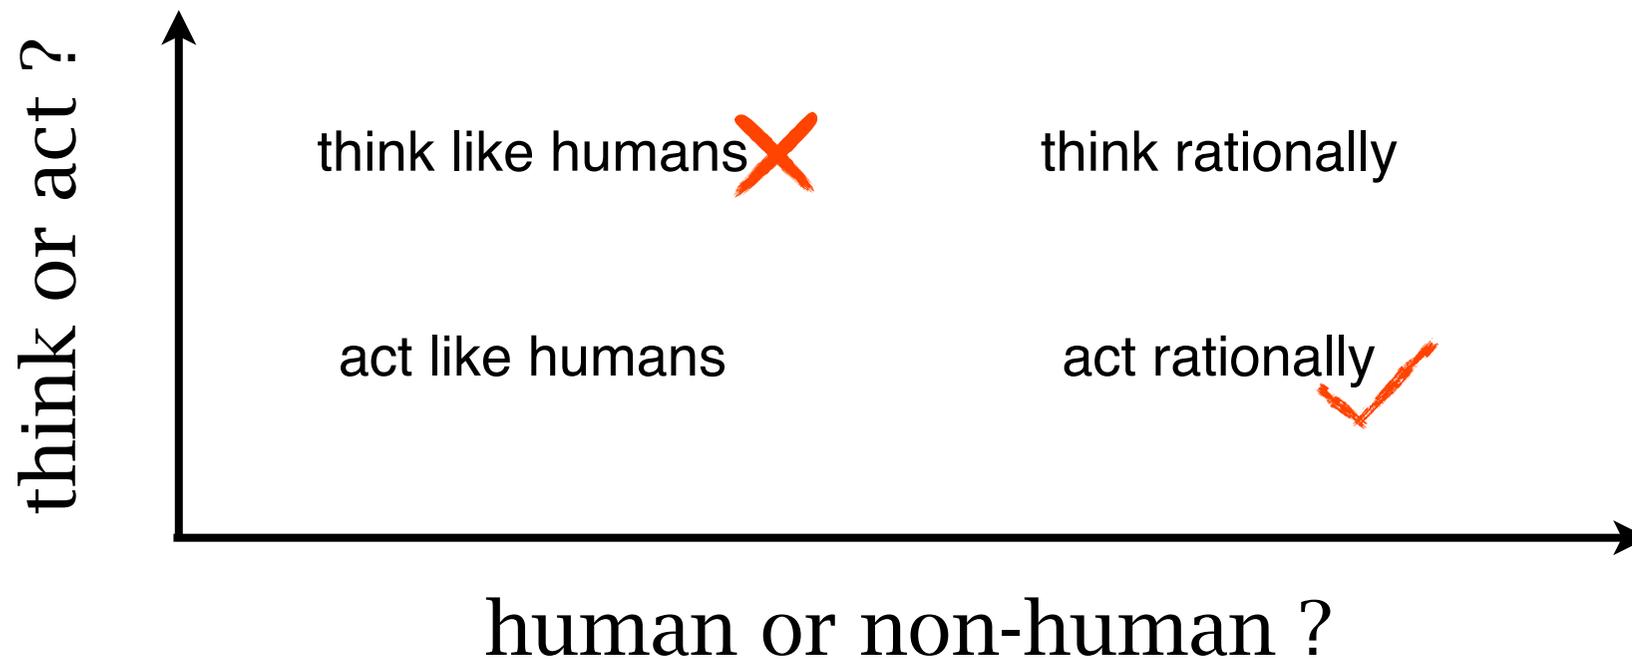
Aristotle (Nicomachean Ethics):

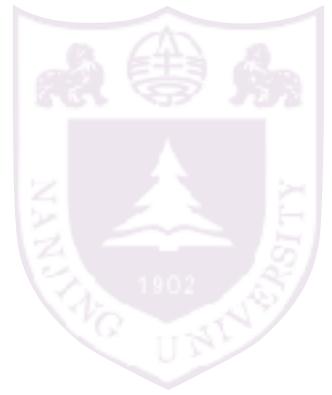
**Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good**

# What is AI?



AI is a system that





AI IS BLOOMING  
HOPE YOU ENJOY  
THANK YOU ALL!