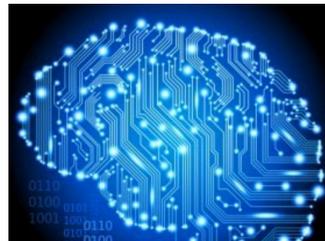




Lecture 19: Final On Artificial Intelligence

http://cs.nju.edu.cn/yuy/course_ai16.ashx



Review on AlphaGo

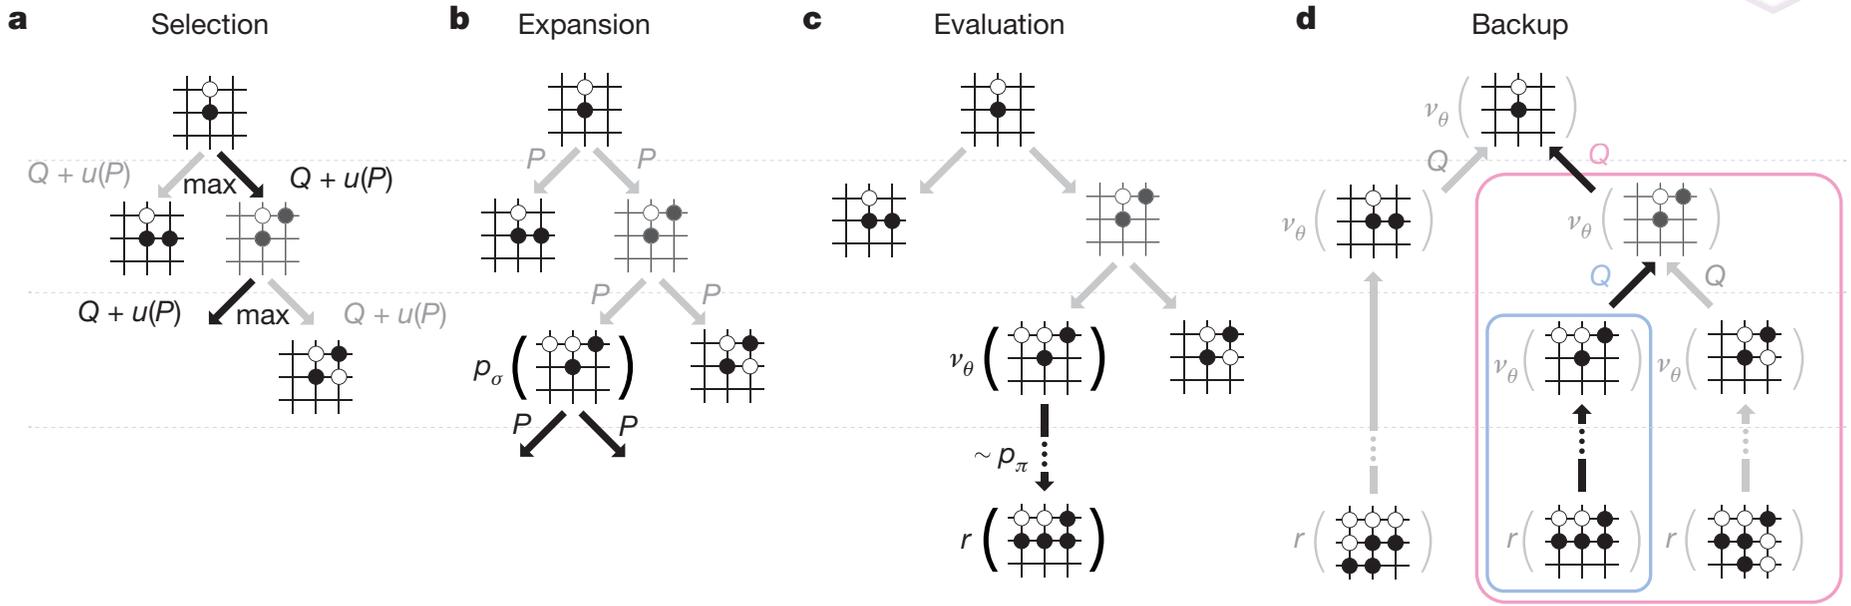


Figure 3 | Monte Carlo tree search in AlphaGo. **a**, Each simulation traverses the tree by selecting the edge with maximum action value Q , plus a bonus $u(P)$ that depends on a stored prior probability P for that edge. **b**, The leaf node may be expanded; the new node is processed once by the policy network p_σ and the output probabilities are stored as prior probabilities P for each action. **c**, At the end of a simulation, the leaf node

is evaluated in two ways: using the value network v_θ ; and by running a rollout to the end of the game with the fast rollout policy p_π , then computing the winner with function r . **d**, Action values Q are updated to track the mean value of all evaluations $r(\cdot)$ and $v_\theta(\cdot)$ in the subtree below that action.

Review on AlphaGo

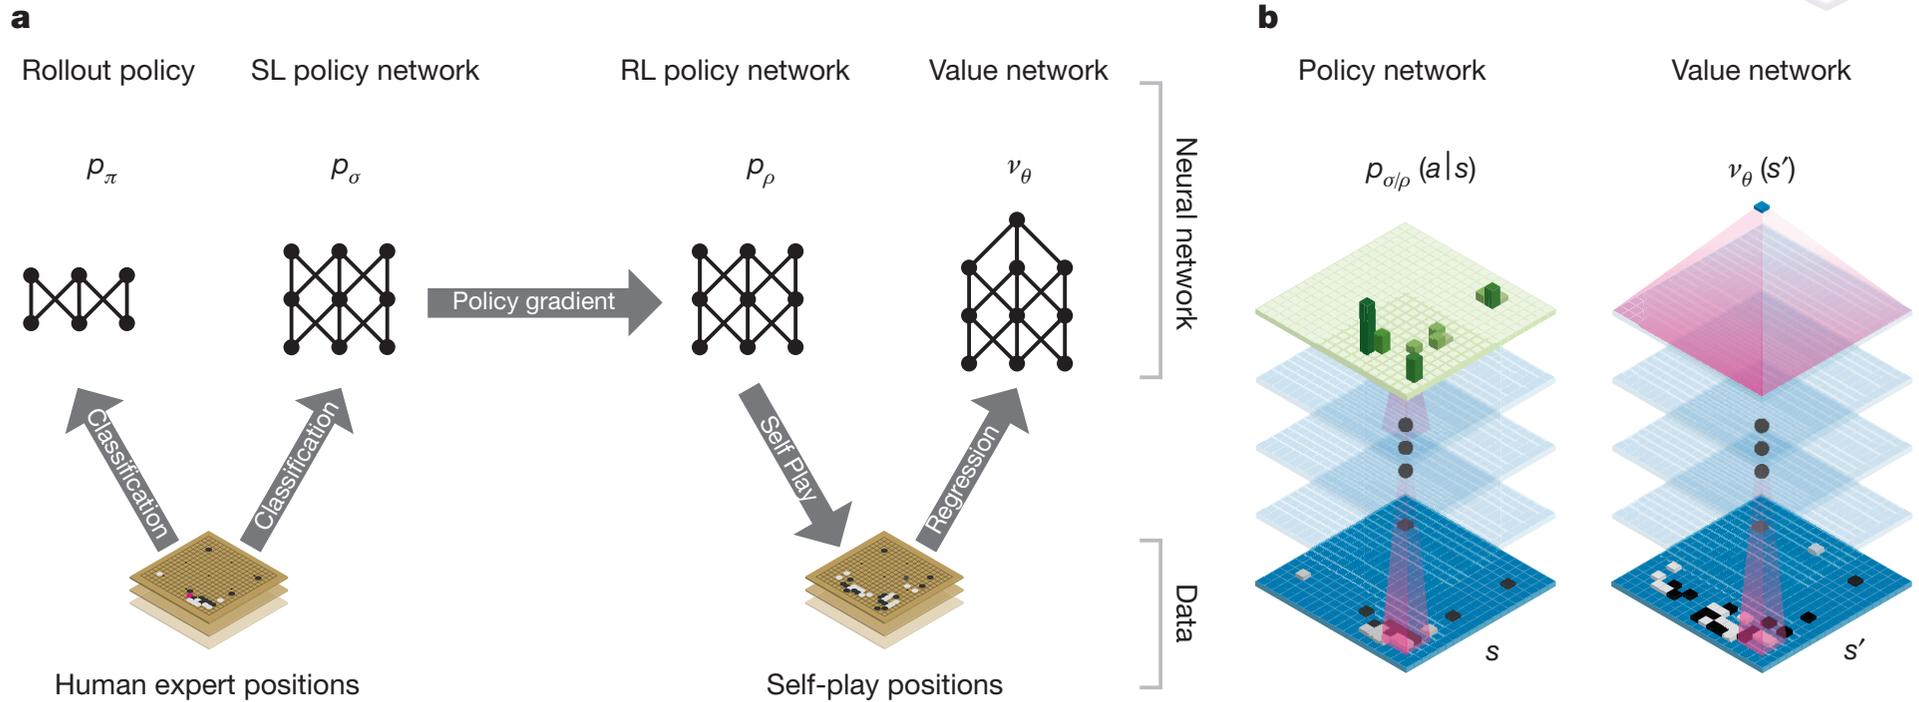
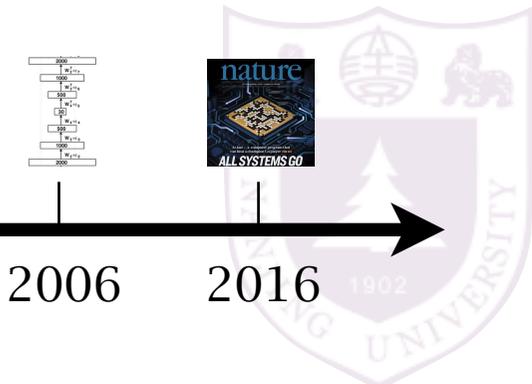


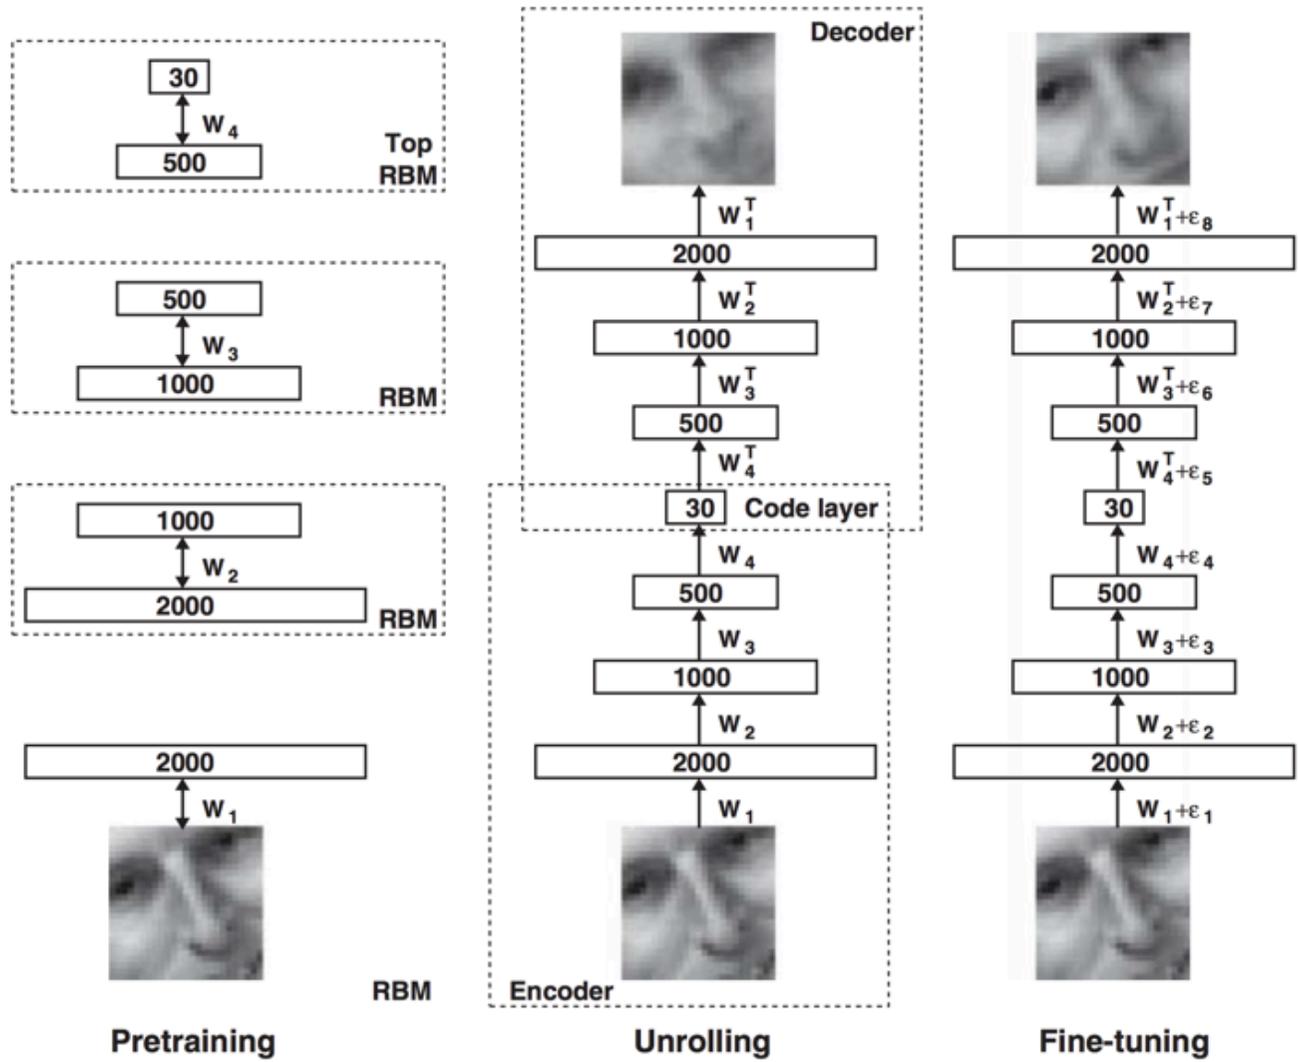
Figure 1 | Neural network training pipeline and architecture. **a**, A fast rollout policy p_π and supervised learning (SL) policy network p_σ are trained to predict human expert moves in a data set of positions. A reinforcement learning (RL) policy network p_ρ is initialized to the SL policy network, and is then improved by policy gradient learning to maximize the outcome (that is, winning more games) against previous versions of the policy network. A new data set is generated by playing games of self-play with the RL policy network. Finally, a value network v_θ is trained by regression to predict the expected outcome (that is, whether

the current player wins) in positions from the self-play data set. **b**, Schematic representation of the neural network architecture used in AlphaGo. The policy network takes a representation of the board position s as its input, passes it through many convolutional layers with parameters σ (SL policy network) or ρ (RL policy network), and outputs a probability distribution $p_\sigma(a|s)$ or $p_\rho(a|s)$ over legal moves a , represented by a probability map over the board. The value network similarly uses many convolutional layers with parameters θ , but outputs a scalar value $v_\theta(s')$ that predicts the expected outcome in position s' .

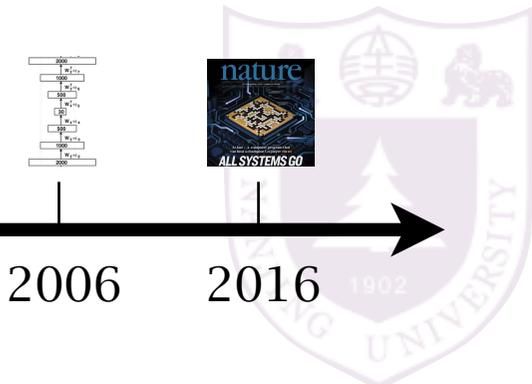
History



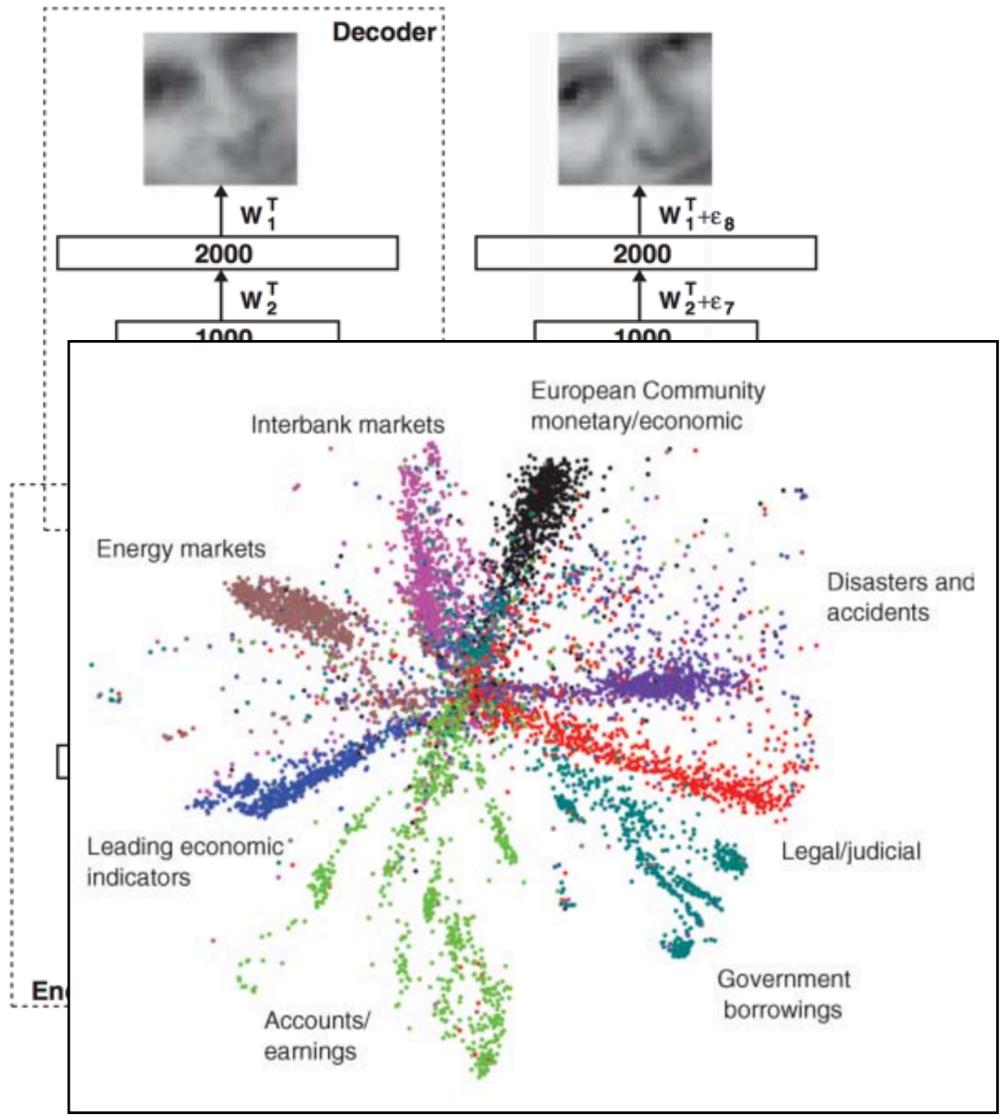
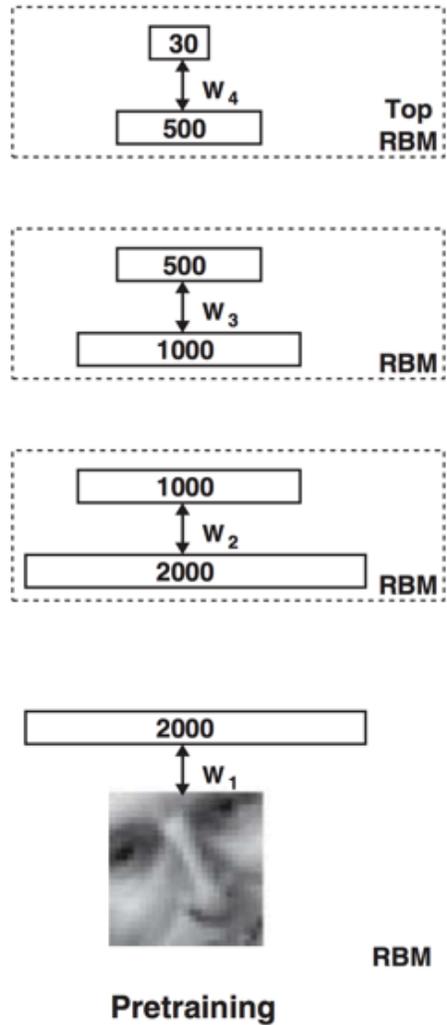
Geoff Hinton



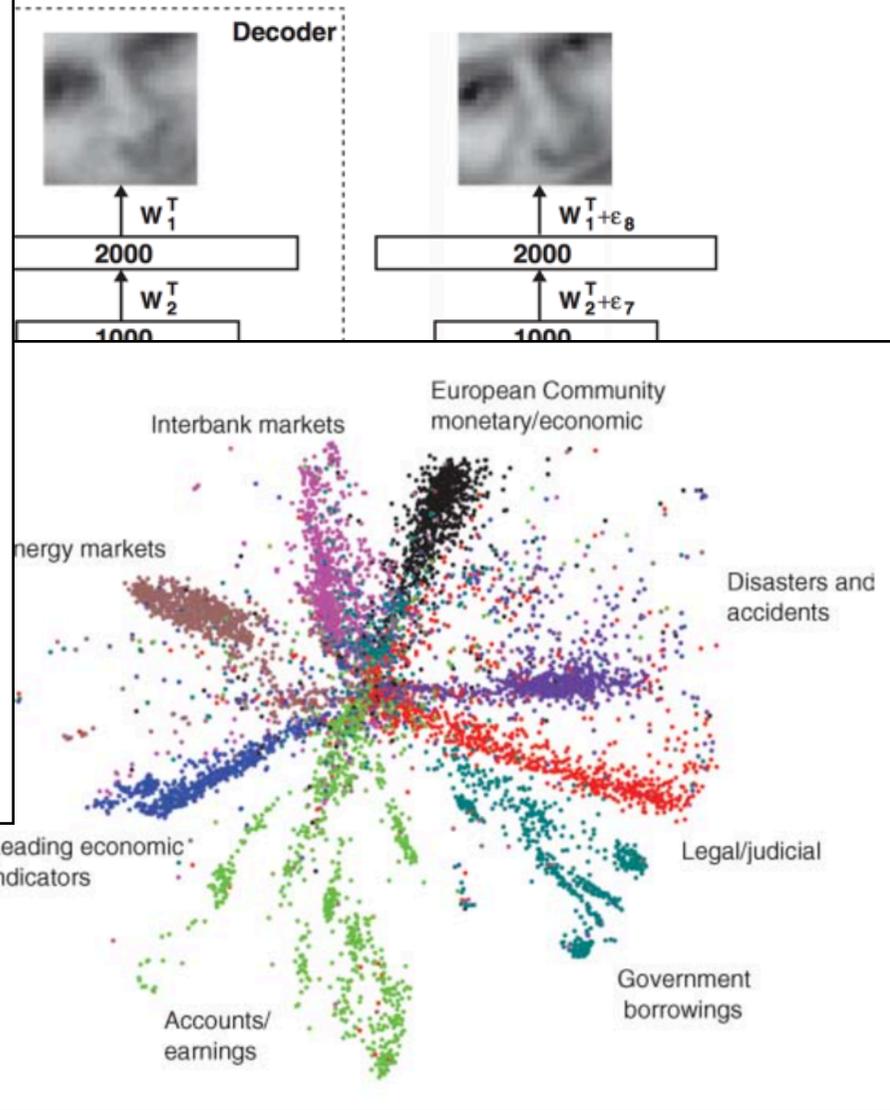
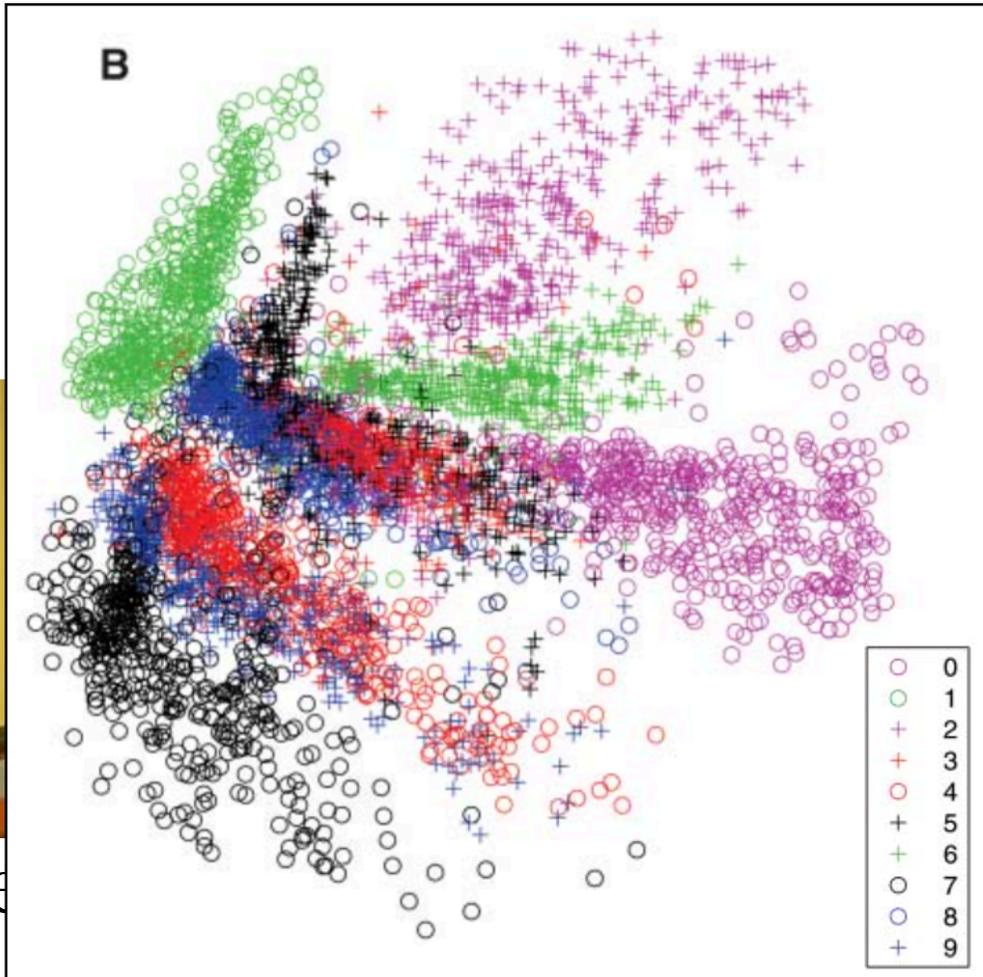
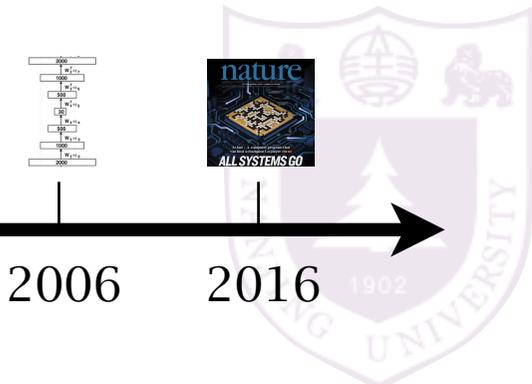
History



Geoff Hinton



History



Pretraining

RBM Encoder



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.



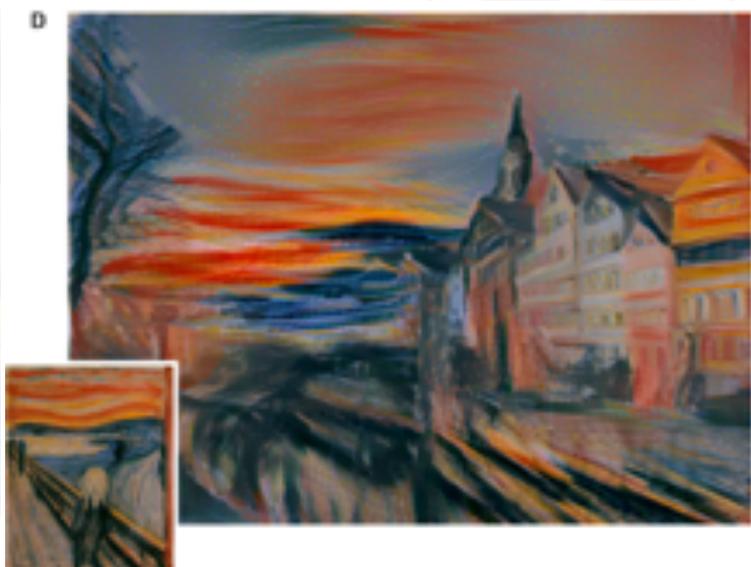
A close up of a cat laying on a couch.

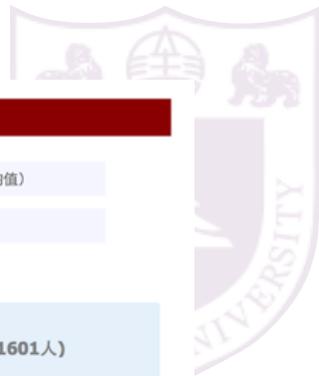


A red motorcycle parked on the side of the road.



A yellow school bus parked in a parking lot.





店铺28天服务情况

纠纷退款率: 0.0036% 小于 0.0064% (行业均值) 退款自主完结率: 99.98% 大于 99.47% (行业均值)

退款完结时长: 0.96天 小于 1.66天 (行业均值)

店铺动态评分: (所属行业: 书籍音像)

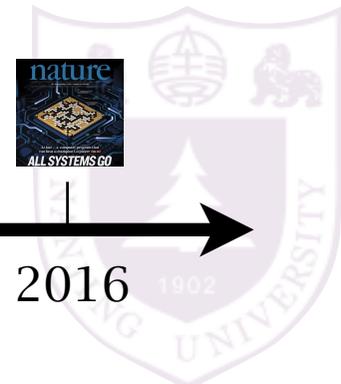
商品与描述相符: 4.8分	比同行业平均水平 低 0.86%	4.8分 ★★★★★ 共34176人
商家的服务态度: 4.7分	比同行业平均水平 低 0.92%	5分: 92.47% (31601人)
商家发货的速度: 4.7分	比同行业平均水平 低 1.09%	4分: 3.99% (1362人)
		3分: 1.61% (551人)
		2分: 0.58% (198人)
		1分: 1.36% (464人)

卖家承诺: 凡使用支付宝服务付款购买本店商品, 若存在质量问题或与描述不符, 本店支持退换货服务并承担来回邮费!

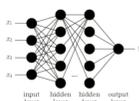
给我推荐

<p>Head First C# (中文版)</p> <p>¥98.6</p>	<p>FUNDAMENTALS AND ALGORITHMS OF FACE RECOGNITION</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 20讲</p> <p>¥42</p>	<p>剑指 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>

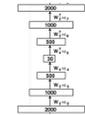
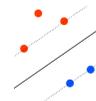




80年代
初期



90年代
中期

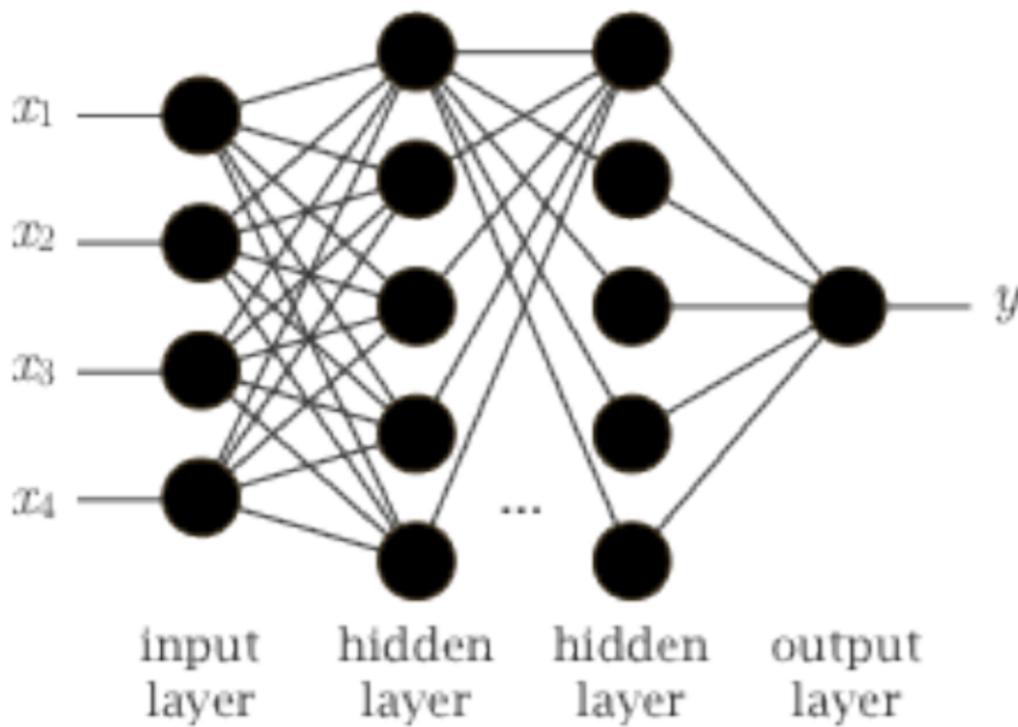


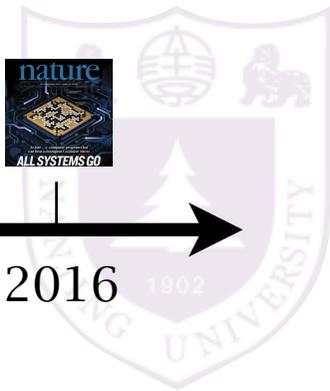
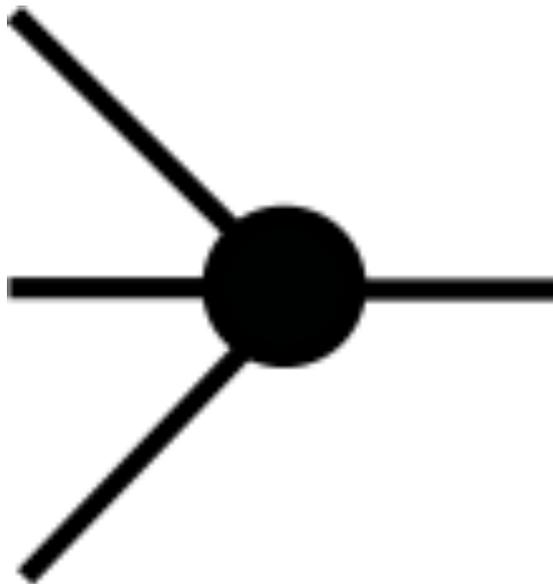
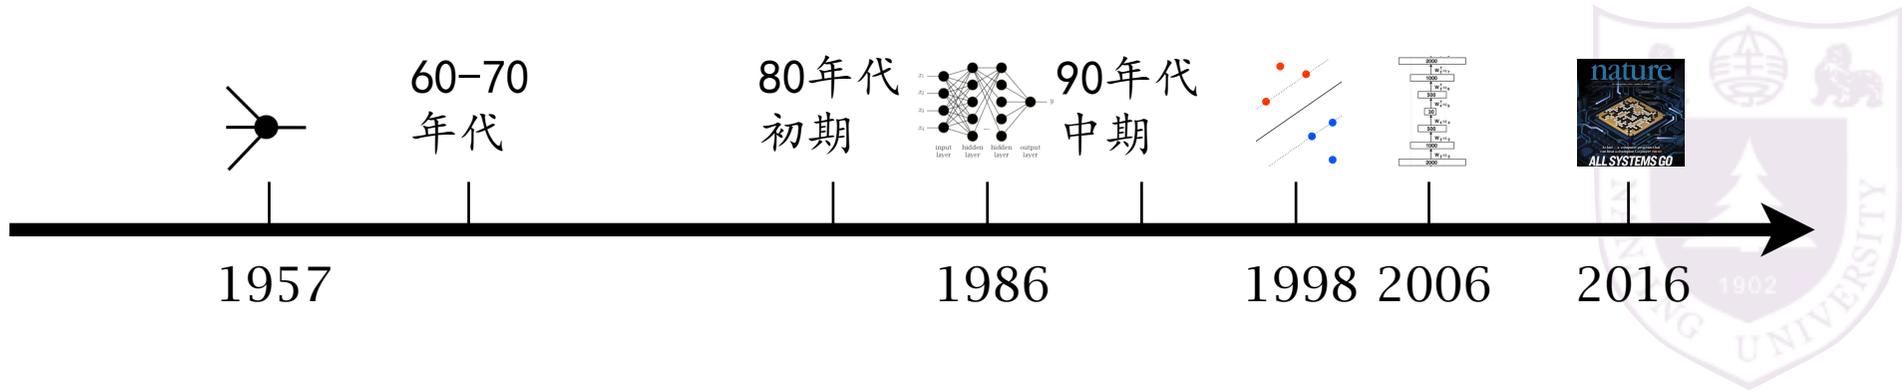
1986

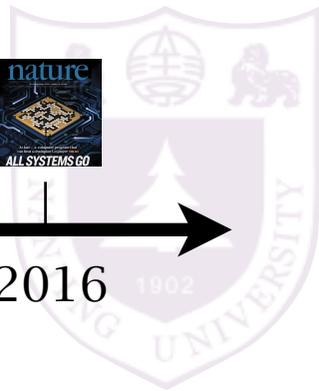
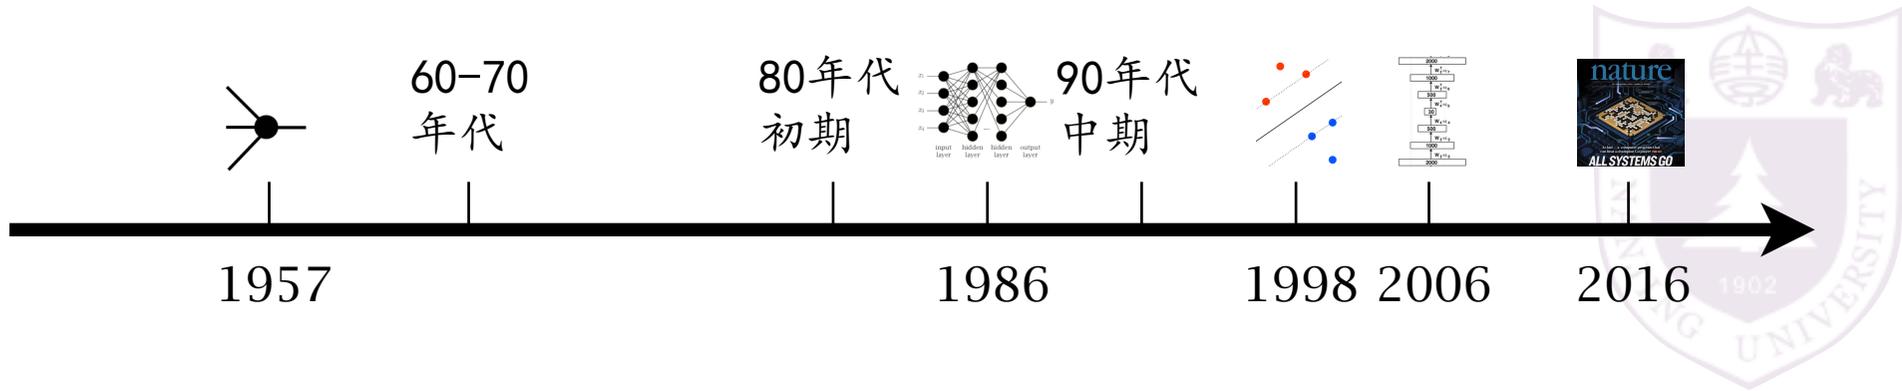
1998

2006

2016





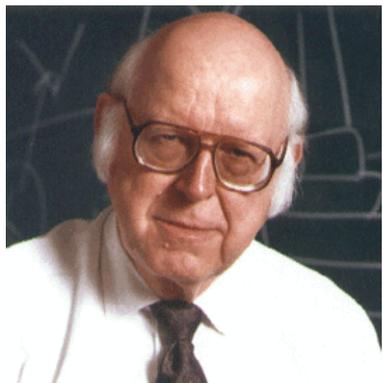
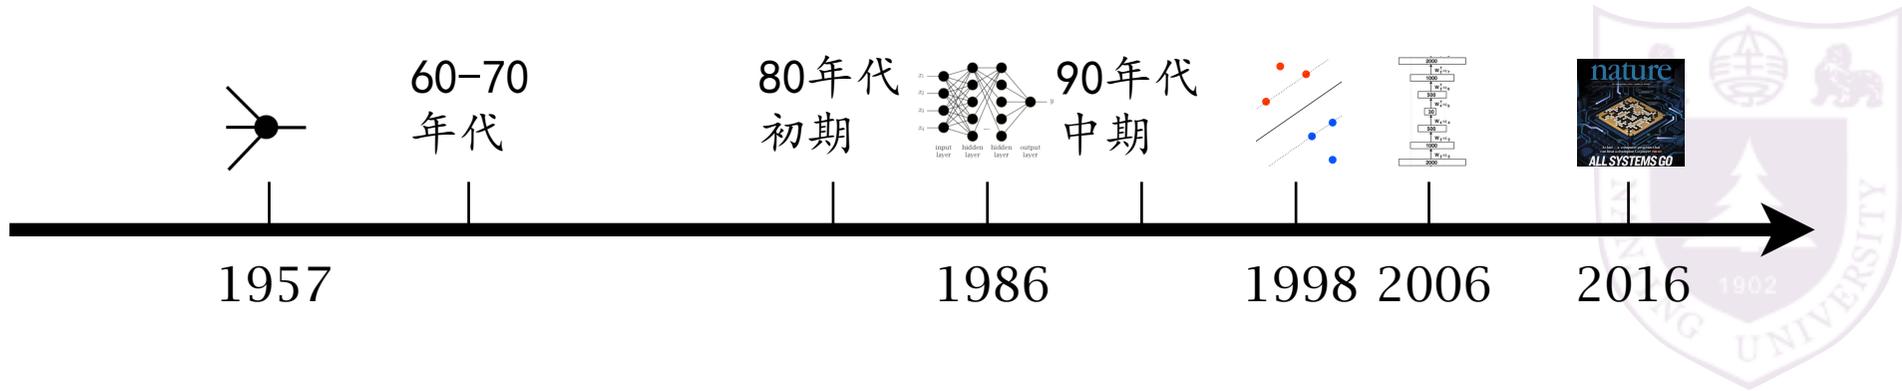


专家系统



Edward Albert Feigenbaum



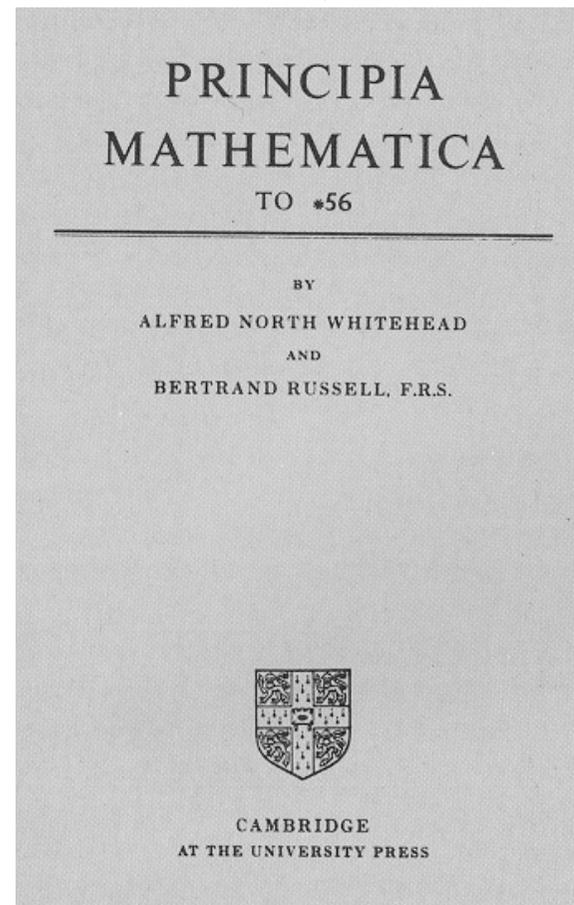


Allen Newell



Herbert Simon

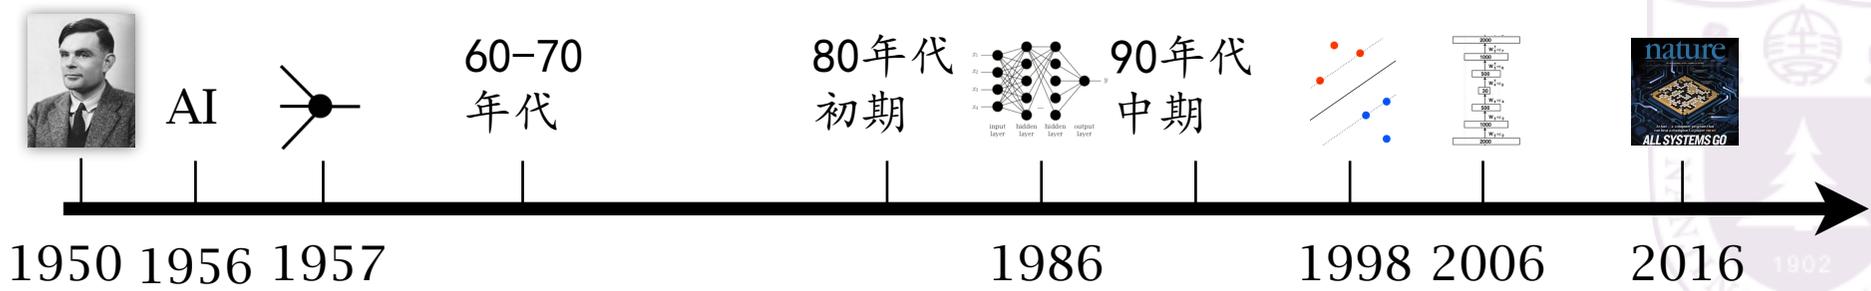
逻辑学家



1956 Dartmouth meeting: “Artificial Intelligence” adopted



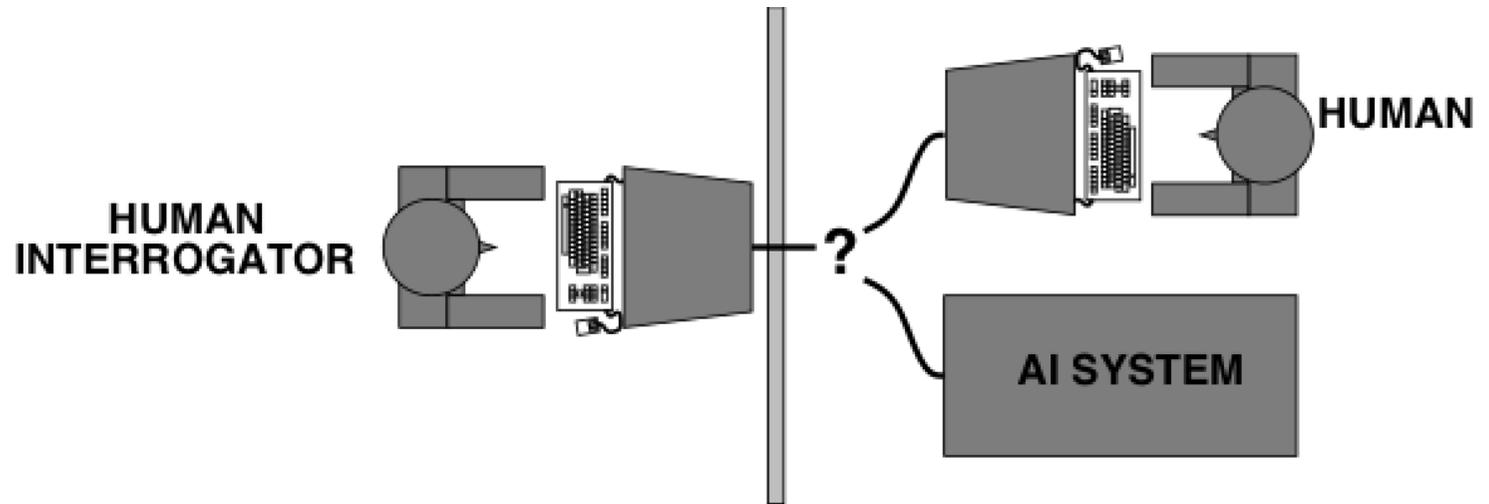
- | | |
|----------------------|--|
| John McCarthy | Turing Award (1971) |
| Marvin Minsky | Turing Award (1969) |
| Claude Shannon | the father of information theory |
| Oliver Selfridge | father of machine perception |
| Herbert A. Simon | Turing Award (1975), Nobel Prize in Economics (1978) |
| Allen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

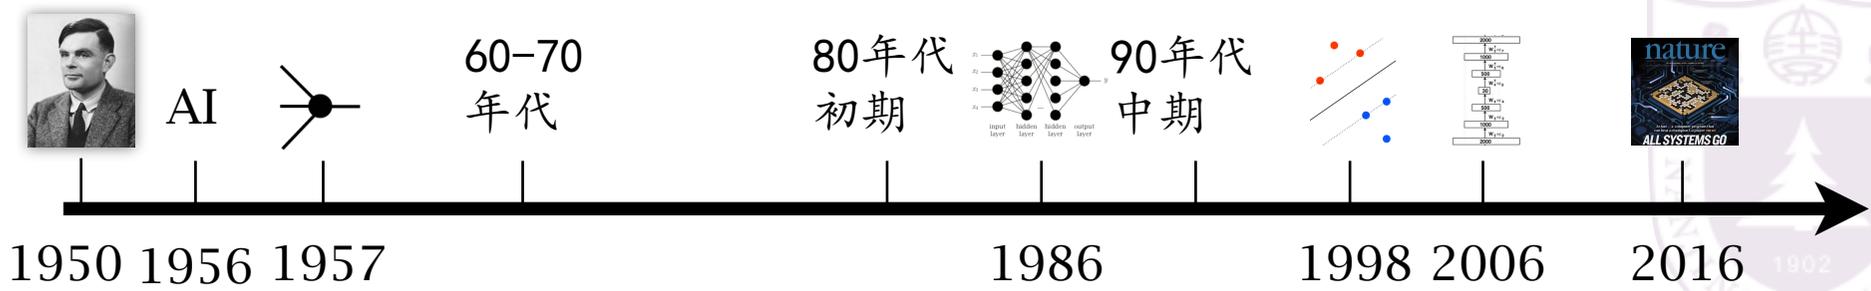


[*Computing machinery and intelligence. Mind* 49: 433-460, 1950.]



Alan Turing
1912-1954



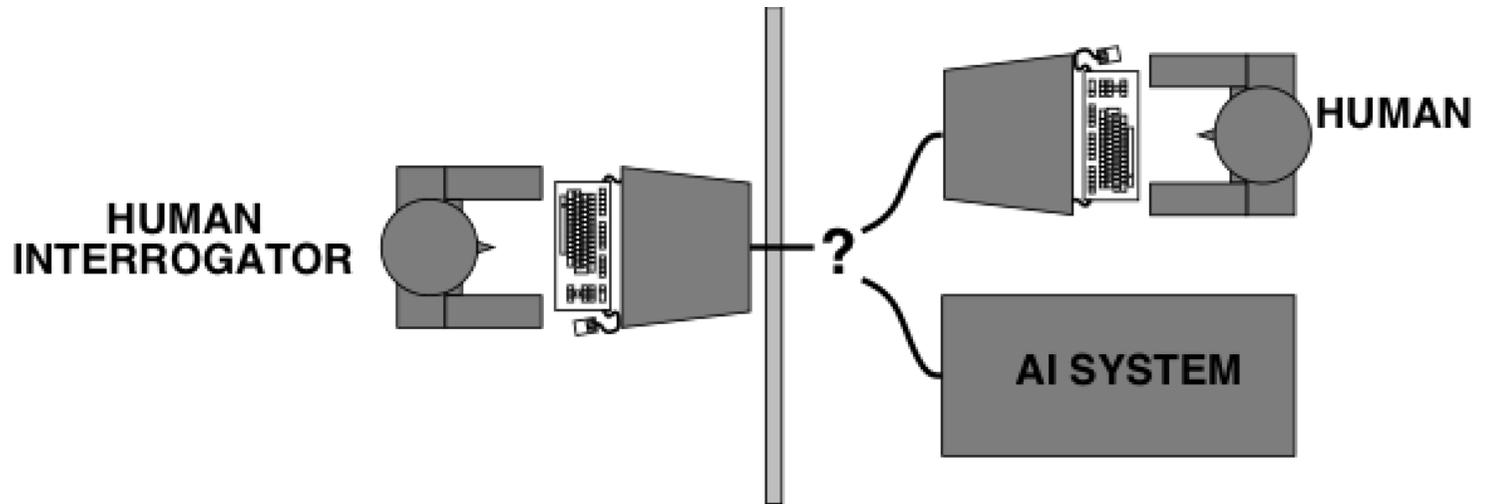


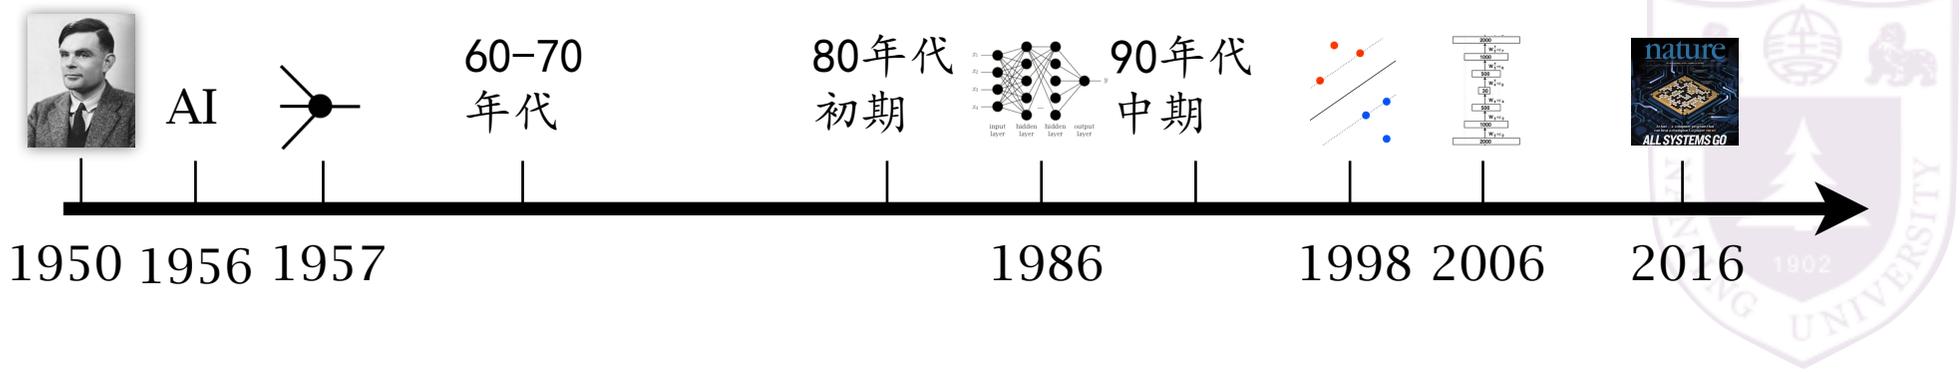
[*Computing machinery and intelligence. Mind* 49: 433-460, 1950.]



Alan Turing
1912-1954

Section 1: Imitation game





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

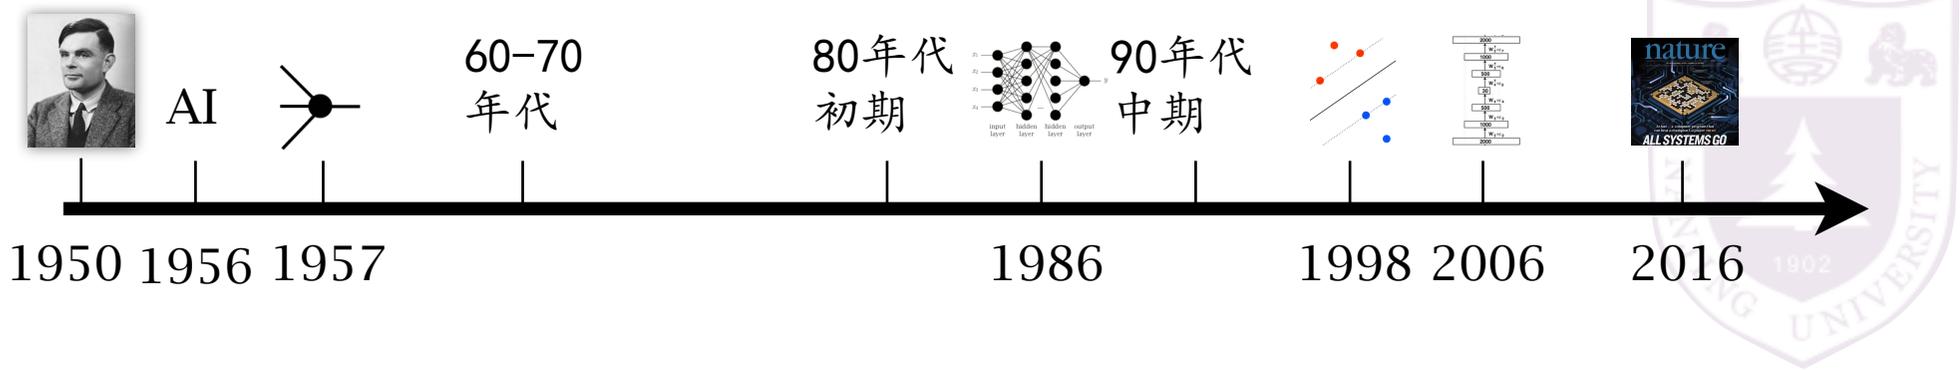


Alan Turing
1912-1954

“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?



[*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

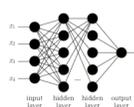


AI

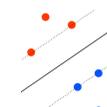


60-70
年代

80年代
初期



90年代
中期



1950 1956 1957

1986

1998 2006

2016



History

学习期



知识期

推理期

AI之冬

AI之冬

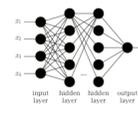


AI

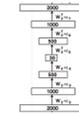
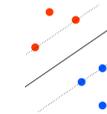


60-70年代

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。”

History

学习期

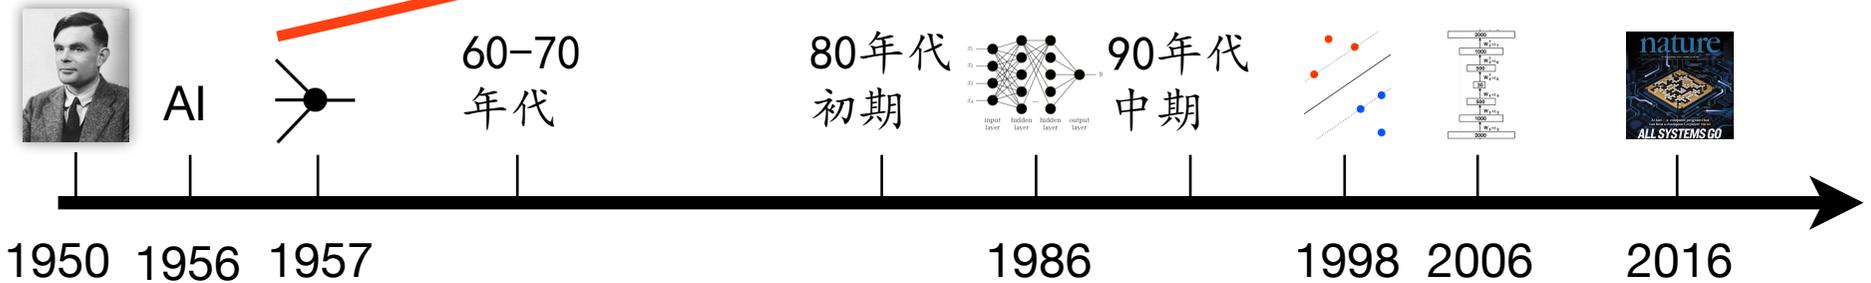


知识期

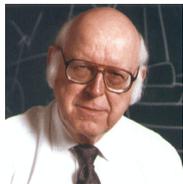
推理期

AI之冬

AI之冬



- “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。”



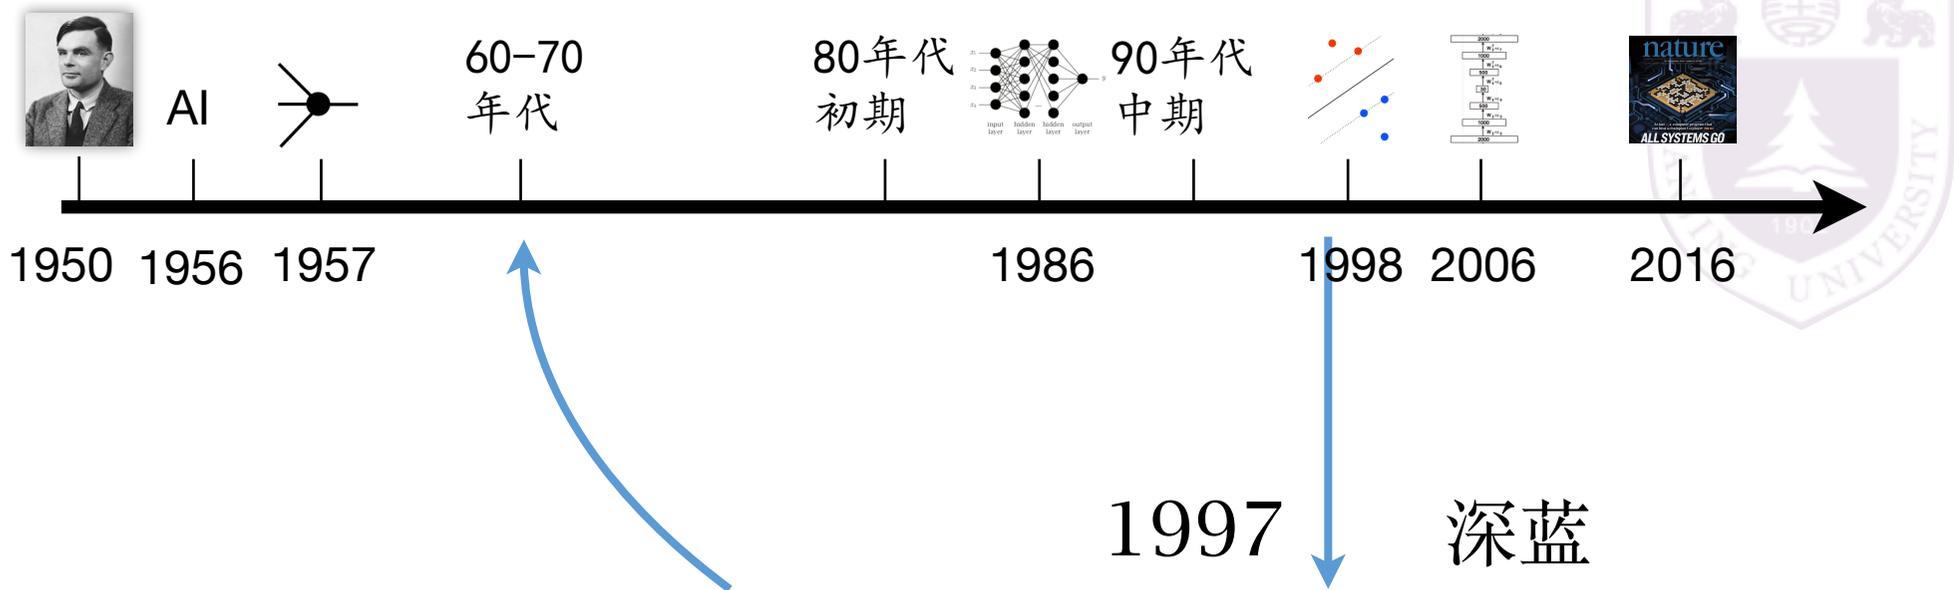
Allen Newell



Herbert Simon



Marvin Minsky



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 artificial intelligence?



1956 Dartmouth meeting: “Artificial Intelligence”

John McCarthy:

“ It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.”



1927-2011

Marvin Minsky:

“ to make computers be capable of doing things that when done by a human, would be thought to require intelligence ”



1927-2016

we will discuss the concept and the history of AI in the last class

What is intelligence?



Intelligence

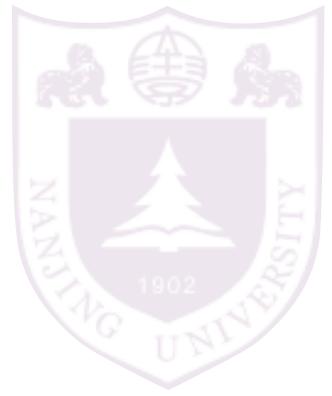
From Wikipedia, the free encyclopedia

For other uses, see [Intelligence \(disambiguation\)](#).

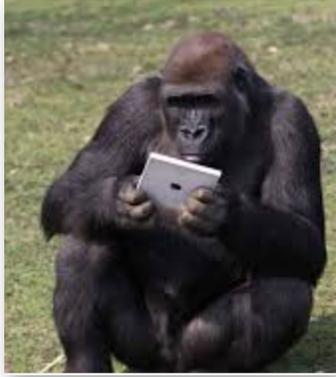
Intelligence has been defined in many different ways such as in terms of one's capacity for [logic](#), [abstract thought](#), [understanding](#), [self-awareness](#), [communication](#), [learning](#), [emotional knowledge](#), [memory](#), [planning](#), [creativity](#) and [problem solving](#). It can also be more generally described as the ability to [perceive](#) and/or retain [knowledge](#) or [information](#) and apply it to itself or other instances of knowledge or information creating referable understanding models of any size, density, or complexity, due to any [conscious](#) or [subconscious](#) imposed [will](#) or instruction to do so.

Intelligence is most widely studied in [humans](#), but has also been observed in non-human animals and in plants. [Artificial intelligence](#) is the simulation of intelligence in machines.

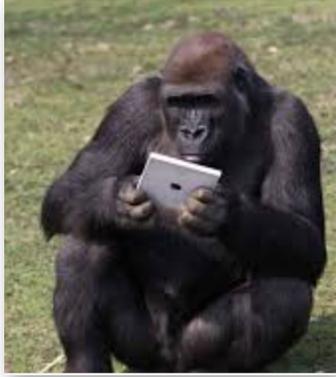
What is intelligence?



What is intelligence?



What is intelligence?



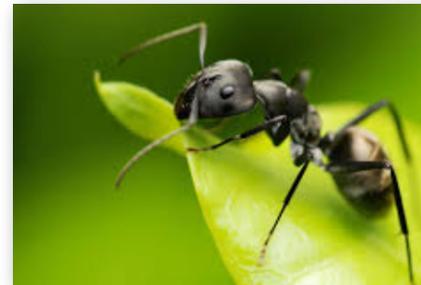
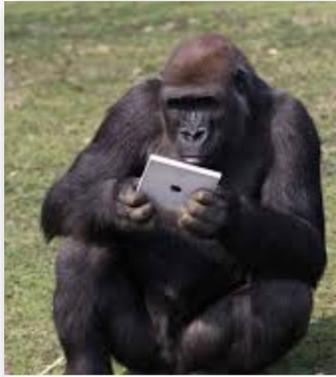
What is intelligence?



What is intelligence?



What is intelligence?



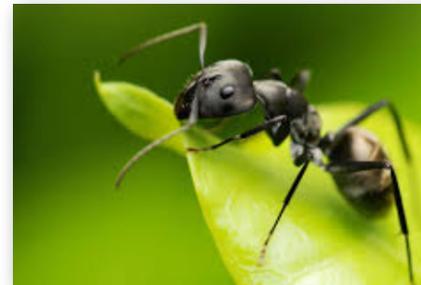
What is intelligence?



What is intelligence?



What is intelligence?



What is intelligence?



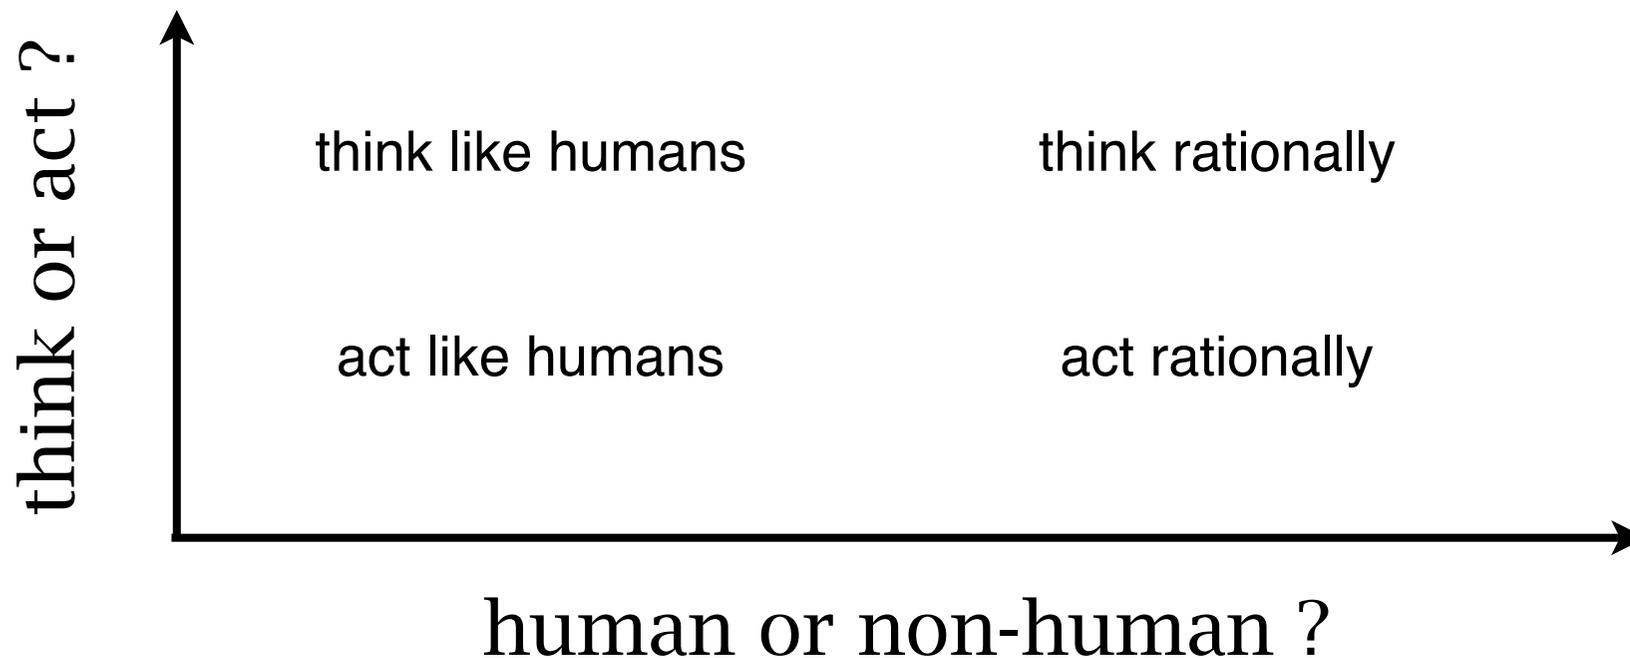
The uncertain about intelligence is a fundamental problem of AI



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

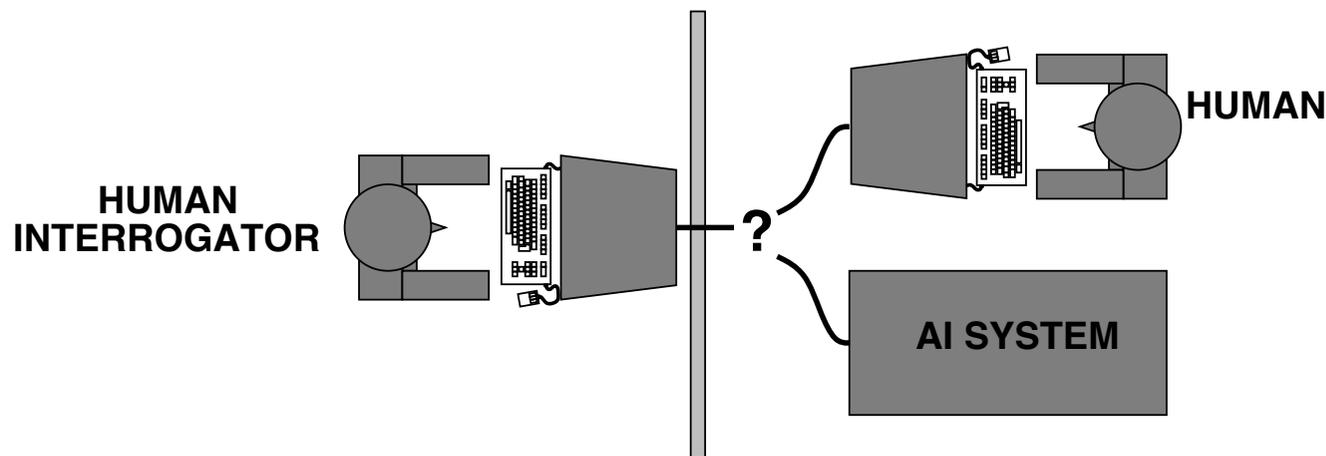
Hence, all three fields share one principal direction!

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

