

PS452
Intelligent Behaviour

**Lecture 6: Artificial
Intelligence - Where Next?**

Maxwell J Roberts
Department of Psychology
University of Essex
www.tubemapcentral.com

version date: 20/02/2018

Part 2: Intelligent Behaviour in Machines

- **Lecture 3: What is Artificial Intelligence?**
The science of intelligent machines
 - Computer scientists define intelligence?
 - What are their tools and assumptions?
- **Lecture 4: Artificial Intelligence landmarks**
Famous programs and findings
 - Which programs have defined the field
 - What can they do, what can't they do?

Part 2: Intelligent Behaviour in Machines

- **Lecture 5: Artificial Intelligence observations, objections**
Why have goals not been achieved?
 - Acerbic critics and ardent philosophers
 - The five fundamental barriers to Artificial Intelligence
- **Lecture 6: Artificial Intelligence - where next?**
Beyond traditional approaches
 - Can breaking with tradition move us forwards?

Lecture 6: AI – Where Next?

- **6.1 Traditional AI: Bigger, Faster, Better?**
 - Deep Blue, world chess champion
 - CYC and common sense: vision
 - CYC and common sense: reality
- **6.2 Alternative Approaches for AI**
 - Connectionism: artificial brains?
 - Connectionism: reality
 - Artificial Life and Situated Cognition
- **6.3 Final thoughts on the Symbolic Search Space Paradigm**

6.1 Traditional AI: Bigger, Faster, Better?

- Is the SSSP fundamentally flawed wrong or has computer power been insufficient for it to shine?
- Huge increases in computer power since early AI attempts
- How much more can SSSP achieve?

Deep Blue, World Chess Champion

- 1997: Defeated world champion:
Gary Kasparov 2½ – Deep Blue 3½
 - 256 processors
200,000,000 moves per second
Position scoring heuristics
Two levels of scoring to prune state space
Standard move sequences programmed
700,000 grandmaster games
Look ahead around 14, up to 40 moves
 - Today, efficient software on personal computers can defeat champions
- ➔ Classic brute force methods can take AI to victory

Deep Blue, World Chess Champion

- Corporate response by IBM
 - Comparing Best human with Deep Blue

	Best Human	Deep Blue
Moves/sec	3	200,000,000
Knowledge	very high	high
Calculation	low	high
Intuition	yes	no
Adaptive learning	yes	no
Independence	yes	no

➔ Deep Blue is **NOT** Artificial Intelligence

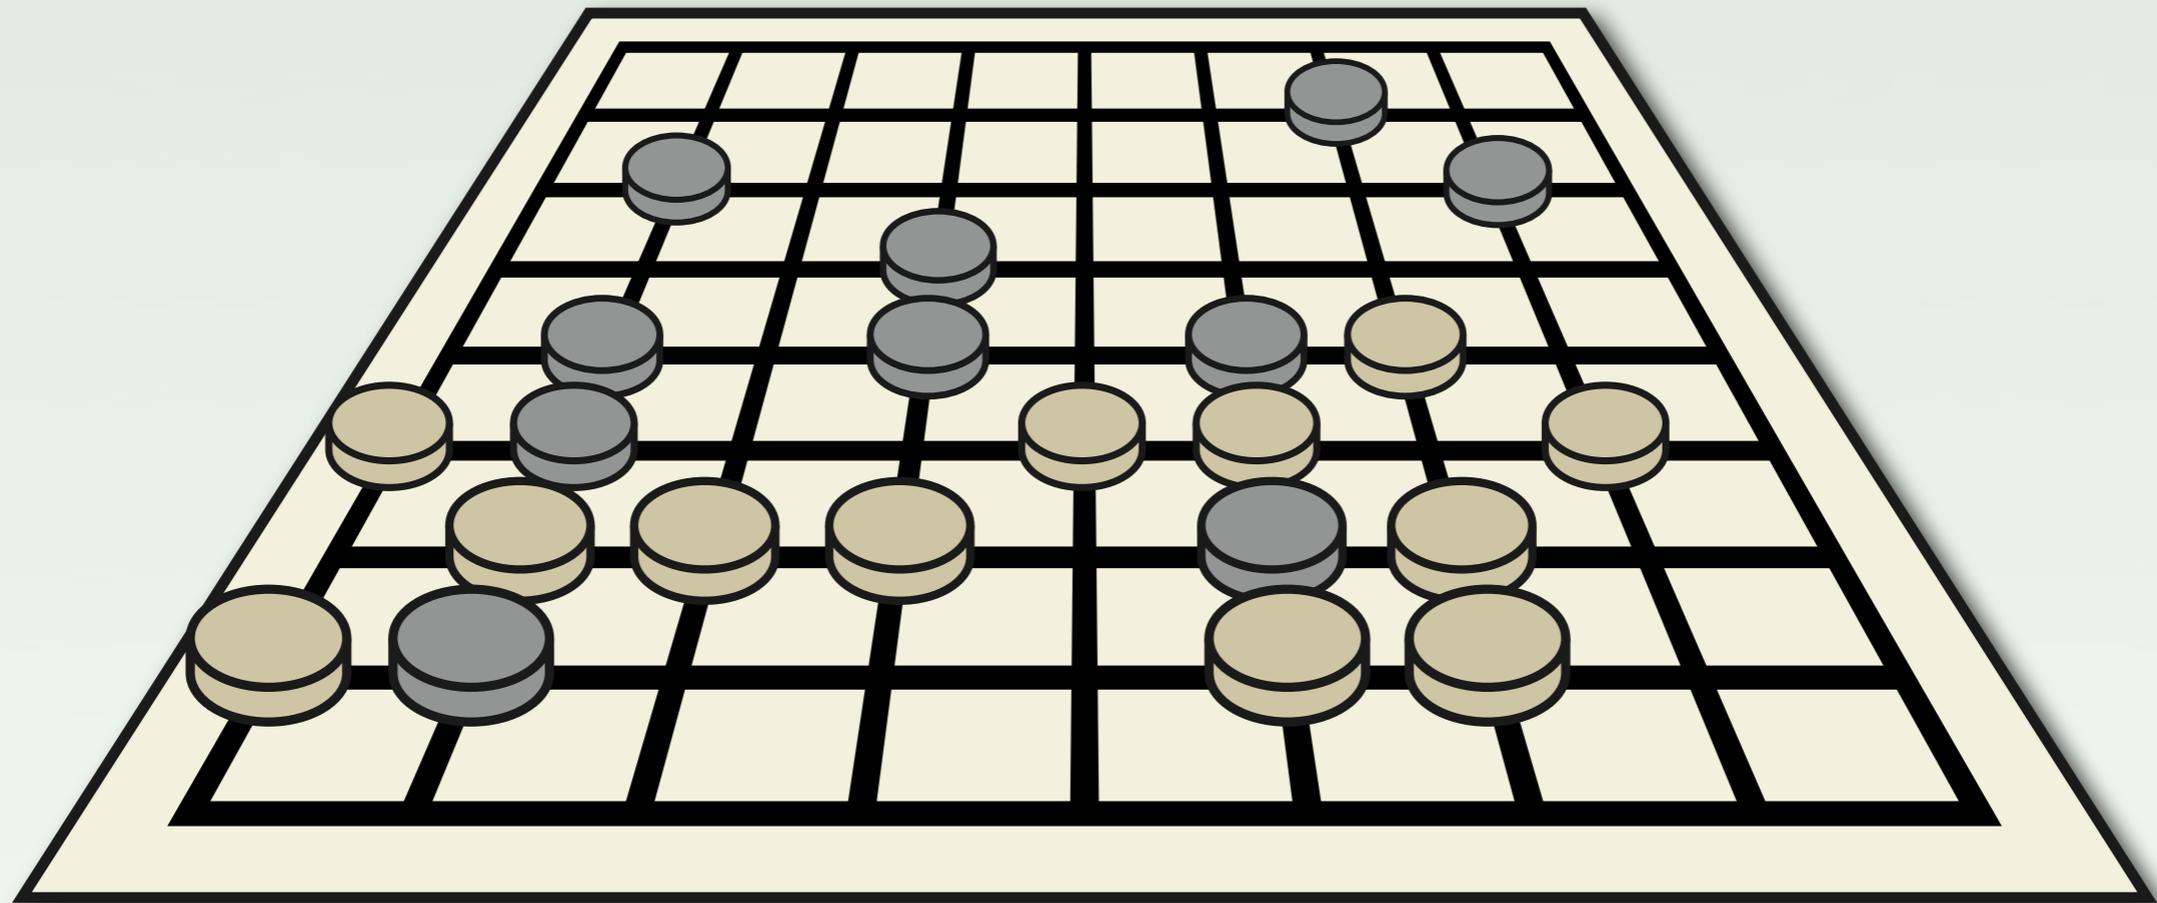
Note, IBM rejects task-based definition of AI

Deep Blue, World Chess Champion

- Nilsson (2010)
 - Combine human/machine strengths for devastating intelligence
 - Ekbida (2008)
 - Machines can beat humans when humans try to mimic machines
- ➔ Human intelligence = finding ways to play chess well despite cognitive limitations

Deep Blue, World Chess Champion

- Lessons learnt
 - Computer superiority for well-defined state space problems now demonstrated, *even for Go* (2016)
 - ➔ Computer power increases can/will tame state spaces



Deep Blue, World Chess Champion

- Lessons learnt
 - Computers perform well to the extent that humans can formalise the world
 - ➔ The rest of the world's challenges?
 - No sign of any emergent properties such as meaning, understanding, or common-sense reasoning
 - ➔ Still quantitative shortcomings, or qualitative ones too?

CYC and Common Sense: Vision

- **CYC** project (initiated by Lenat, 1984+)
 - SAM defeated by simple non-script questions:
Was your left foot in the restaurant?
 - Expert system errors must be trapped manually
 - ➔ AI stalled because lack of common sense?
- Lenat and Feigenbaum (1991) Critiques: Copeland (1993), Smith (1991), Dreyfus (1993) and Ekbria (2008)

CYC and Common Sense: Vision

- **CYC** project *cont.*
 - Lack of common sense = lack of basic knowledge?
 - Give computers knowledge to support common-sense?
 - Will accelerate learning?
 - Overcome the problems of brittleness and scalability?
 - ➔ Modular foundation for future?

CYC and Common Sense: Vision

- **CYC** project *cont.*
 - Assumption: common sense derived from:
 - Explicitly coded facts
 - Logical inferences
 - Lenat and Feigenbaum (1991)

“No sophisticated as-yet-unknown control structure is required for intelligent behaviour”

L2: control structure is both known and sophisticated!?

- Explicitly NOT trying to model humans
- Mimic human common sense without mimicking human cognitive processes???

CYC and Common Sense: Vision

- **CYC** project *cont.*
 - Resources, methodology and targets (early)
 - \$50,000,000 budget
 - Aim: to complement an encyclopaedia; knowledge required to understand one
 - Hence en-**CYC**-lopaedia
 - By 1990: over 1,000,000 facts added manually (0.1% of consensus reality)

CYC and Common Sense: Vision

- **CYC** project *cont.*
 - Resources, methodology and targets (early)
 - From newspapers, magazines, adverts
 - Explicit information + information that the articles assumed the reader knows
 - Living things get diseases
 - Water flows downhill
 - Food prevents starvation etc. etc.
 - 100,000,000 facts/principles = final target

CYC and Common Sense: Vision

- **CYC** project *cont.*
 - Methodology and targets (*cont.*)
 - 1994: 30% to 50% consensus reality
 - Quicker for system to acquire facts and ask questions than to type in material
 - 2001: Human level understanding
 - Lenat: 60% chance of success
 - Last standing for SSSP?
 - What if sufficient knowledge but no common-sense
 - What could the next big SSSP project try instead?

CYC and Common Sense: Reality

- Copeland (1993)
 - Large knowledge base must be searched
 - Needs sound logic, ontology, epistemology
 - **Ontology:** theory of classification;
What categories of objects exist in the world and what are their properties
 - **Epistemology:** theory of knowledge;
What can we know about the world?
What should we believe about the world?
 - Without these, insufficient foundations for common sense

CYC and Common Sense: Reality

- Copeland (1993) *cont.*
 - ***Classification (ontology)***
Entities and their properties
 - What types of object?
Can two objects occupy the same place?
 - CYC has five types of object:
tangible, intangible, composites, events and mathematical

The clouds swept past me

versus

The unicorn swept past me

➔ Enough? Mutually exclusive? Well defined?

CYC and Common Sense: Reality

- Copeland (1993) *cont.*

- ***Causality***

- How does one event cause another?
How is this inferred
- CYC tags pairs of events together if causal

It was raining so the river flooded

versus

It was raining so we went to museum

➔ Insufficient to distinguish complexities/subtleties

E.g. ***X determines Y*** versus ***X led to Y***

CYC and Common Sense: Reality

- Copeland (1993) *cont.*
 - **Logic**
 - Inferences about an indeterminate world
 - CYC attempts non-monotonic logic
 - Retracts inferences if found false
 - Identifies erroneous inference chain
 - Rejects/modifies knowledge if necessary

PS452 sounds like an interesting course

I enjoy taking interesting courses

If I take PS452, I will enjoy it

I did not enjoy PS452

CYC and Common Sense: Reality

- Copeland (1993) *cont.*
 - ***Logic cont.***
 - 100,000,000 facts = slow process to ensure consistent and up to date
 - Which knowledge to reject/modify?
 - No reliable heuristics available
- ➔ *Logic Problem/Frame Problem* contained, not resolved

CYC and Common Sense: Reality

- Copeland (1993) *cont.*
 - **Relevance**
 - Not all knowledge needed for all tasks
 - CYC encodes temporal, informational, spatial proximity
- Where would be a good place to go on holiday this year?***
- May not narrow down sufficiently
 - May miss relevant details
 - ➔ Many causes distant in time and space

CYC and Common Sense: Reality

- Copeland (1993) *cont.*

- ***Consistency***

- What if conflicting facts or inferences?
- CYC divides knowledge into internally consistent micro-theories

IF my pet is in trouble THEN give it assistance

versus

IF the sea is cold and rough THEN do not enter it

- Micro-theories cannot be mutually exclusive unless there are as many of these as facts?
- ➔ Inconsistencies insufficiently trapped and can be fatal

CYC and Common Sense: Reality

- Ekbia (2008)
 - 15 years on, CYC is now a commercial product
 - Multi-context inference engine
 - [Glorified internet search engine?]
 - ➔ Further than ever from scrutiny
 - Example of successful inference?
 - Table 1: Name, job title, employer
 - Table 2: Employer, location
 - ▶ List advanced degree holders in New England
 - Searches for doctors, lawyers, professors
 - whose employer is in New England
 - ➔ Common sense inference possible?

CYC and Common Sense: Reality

- Ekbia (2008) *cont.*
 - Example of failure (from a 2002 visit)
 - Bill & Martha are parents of Bobby & Karen
Is Bill older than Bobby? CYC: ??????????
 - Bill and Bobby are both alive
Is Bill older than Bobby? CYC: ??????????
 - Bill is the father of Bobby
Is Bill older than Bobby? CYC: ??????????
 - Bill is the biological father of Bobby
Is Bill older than Bobby? CYC: YES!
 - ▶ Two hours of searching inference chains to diagnose!
 - ▶ Failures to many other obvious questions

CYC and Common Sense: Reality

- Ekbia (2008) *cont.*

“Even if existent, this ‘knowledge’ is buried under a tangled and cluttered pile of propositions, axioms, and inference rules”

- Emphasis on explicit knowledge wrong
 - How many facts and principles needed?
 - How much of knowledge is tacit?
- ➔ A brittle expert system whose domain is common sense

CYC and Common Sense: Reality

- Overall Impressions
 - Copeland (1993): naive solutions

"[Achieving goals within predicted time is] nothing less than absurd"
 - Ekbia (2008): fatal neglect of control

"crumbling under its own weight"

"remedial work is merely better camouflage of problems with fundamental assumptions"

"brittleness may not be due to lack of knowledge"
- ➔ CYC has not delivered real world robust common sense
- ➔ Time to move on from *Symbolic Search Space Paradigm?*

6.2 Alternative Approaches for AI

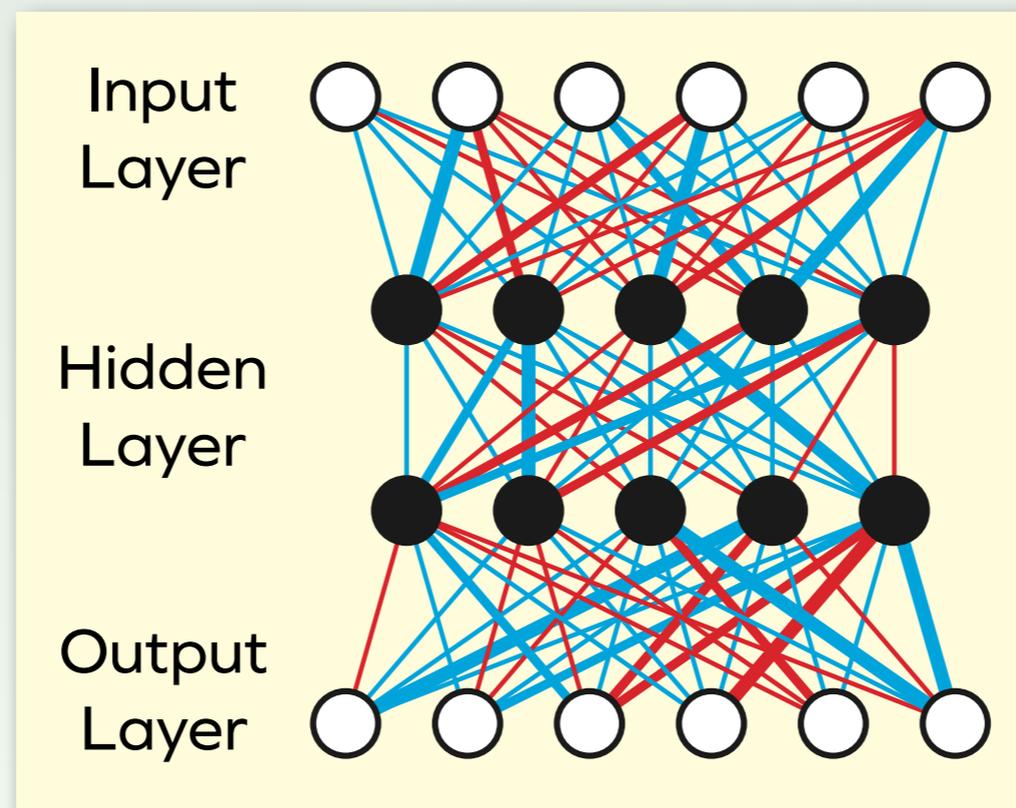
- Connectionism: Artificial Brains
 - Can learning be achieved by making computers more brain-like?
- Artificial Life & Situated Cognition
 - Can high level intelligence be attained by lowering ambitions and incrementally adding new features?

Connectionism: Artificial Brains?

- Highly interconnected networks, various names
 - Connectionism
 - Parallel Distributed Processing (PDP)
 - Neural Networks
- Bermúdez (2010), Copeland (1993), Ekbia (2008), Johnson-Laird (1988) for discussions

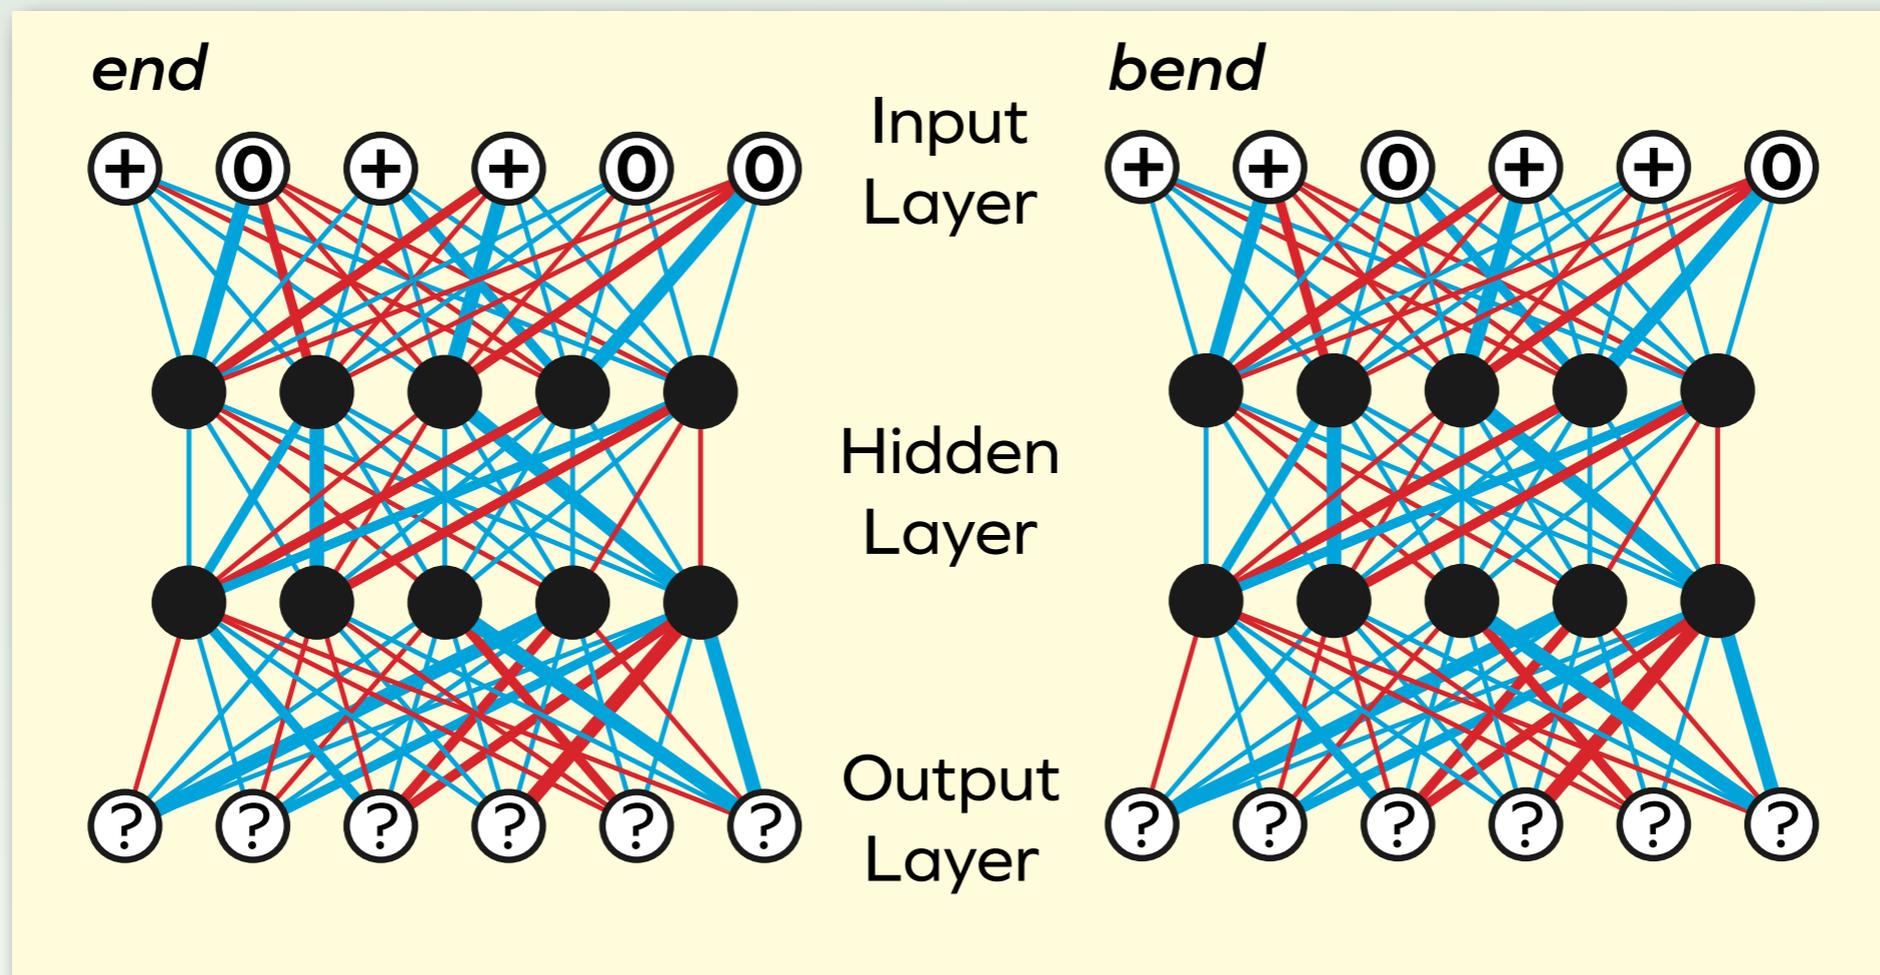
Connectionism: Artificial Brains?

- Thousands of nodes organised in layers
 - Input units
 - Hidden units
 - Output units
- Nodes individually inter-connected by links of variable strength: strongly excitatory to strongly inhibitory



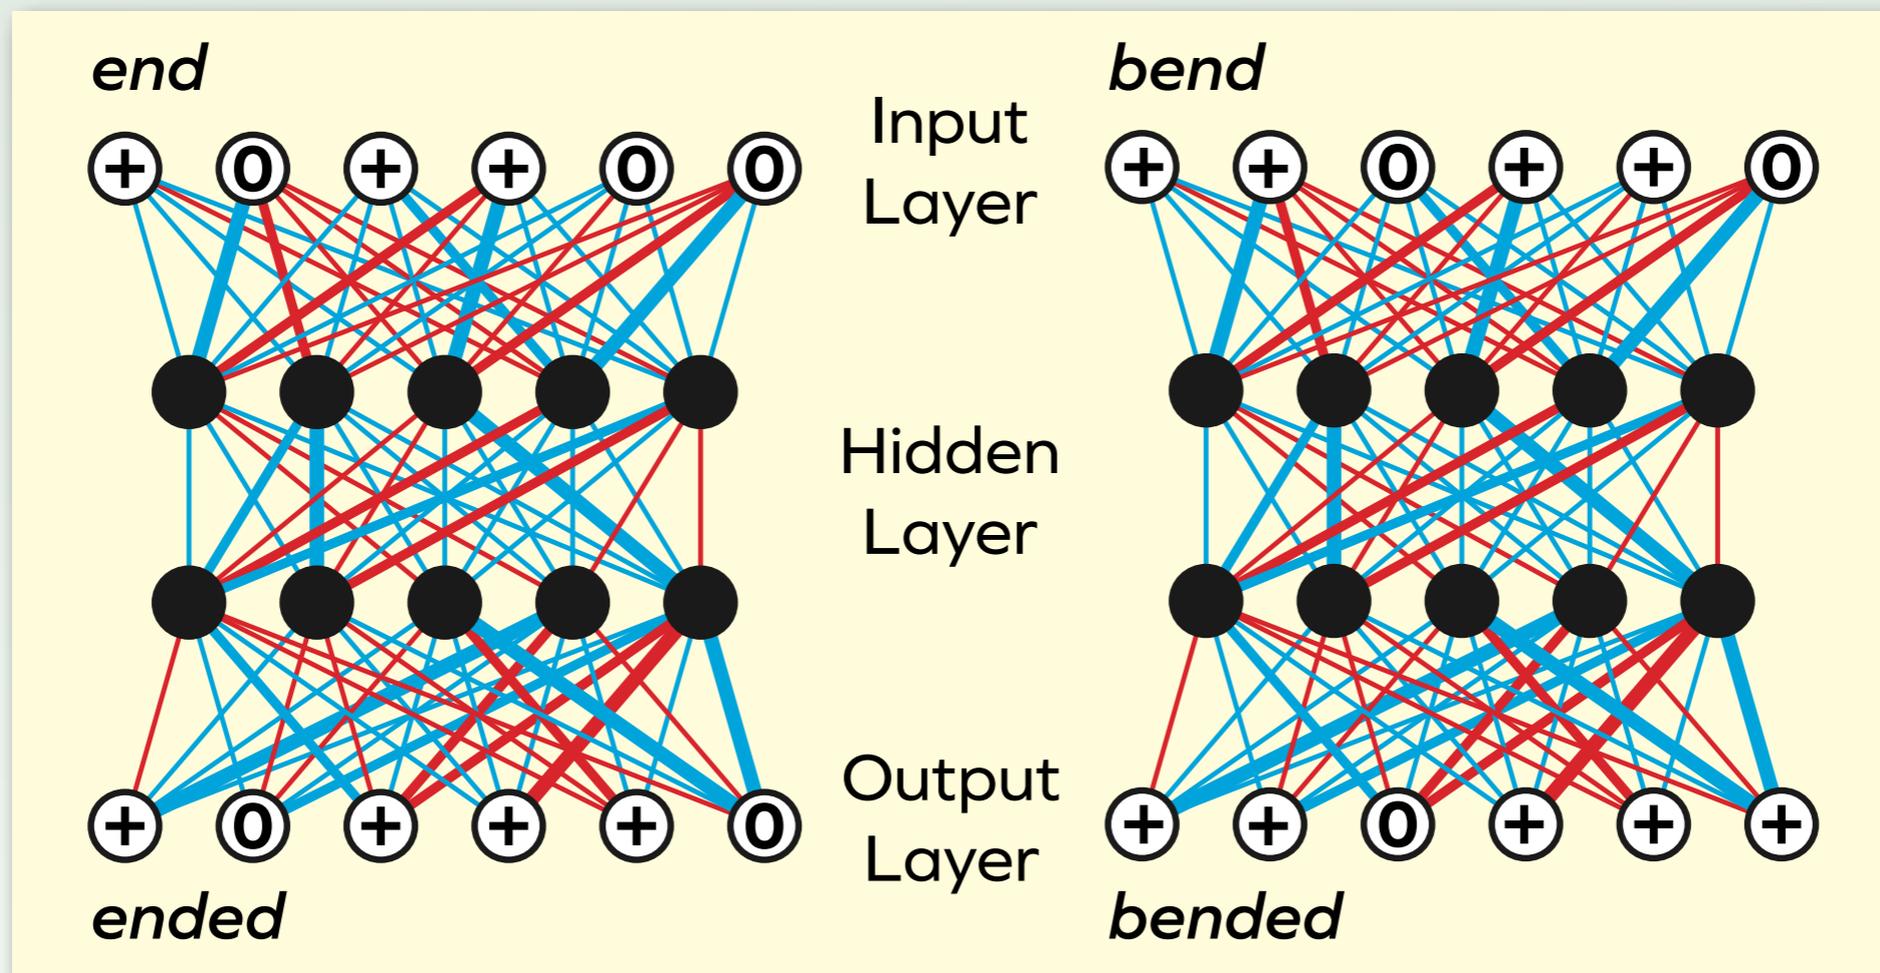
Connectionism: Artificial Brains?

- Each node has two states, active/inactive
- Input layer node status depends on how information is manually encoded
- Only active nodes feed excitation/inhibition to their connected neighbours



Connectionism: Artificial Brains?

- Activity of subsequent nodes depends on sum total of receipts from active nodes
- Eventually, output nodes reach stable states = response to input information
- Output is decoded to identify response

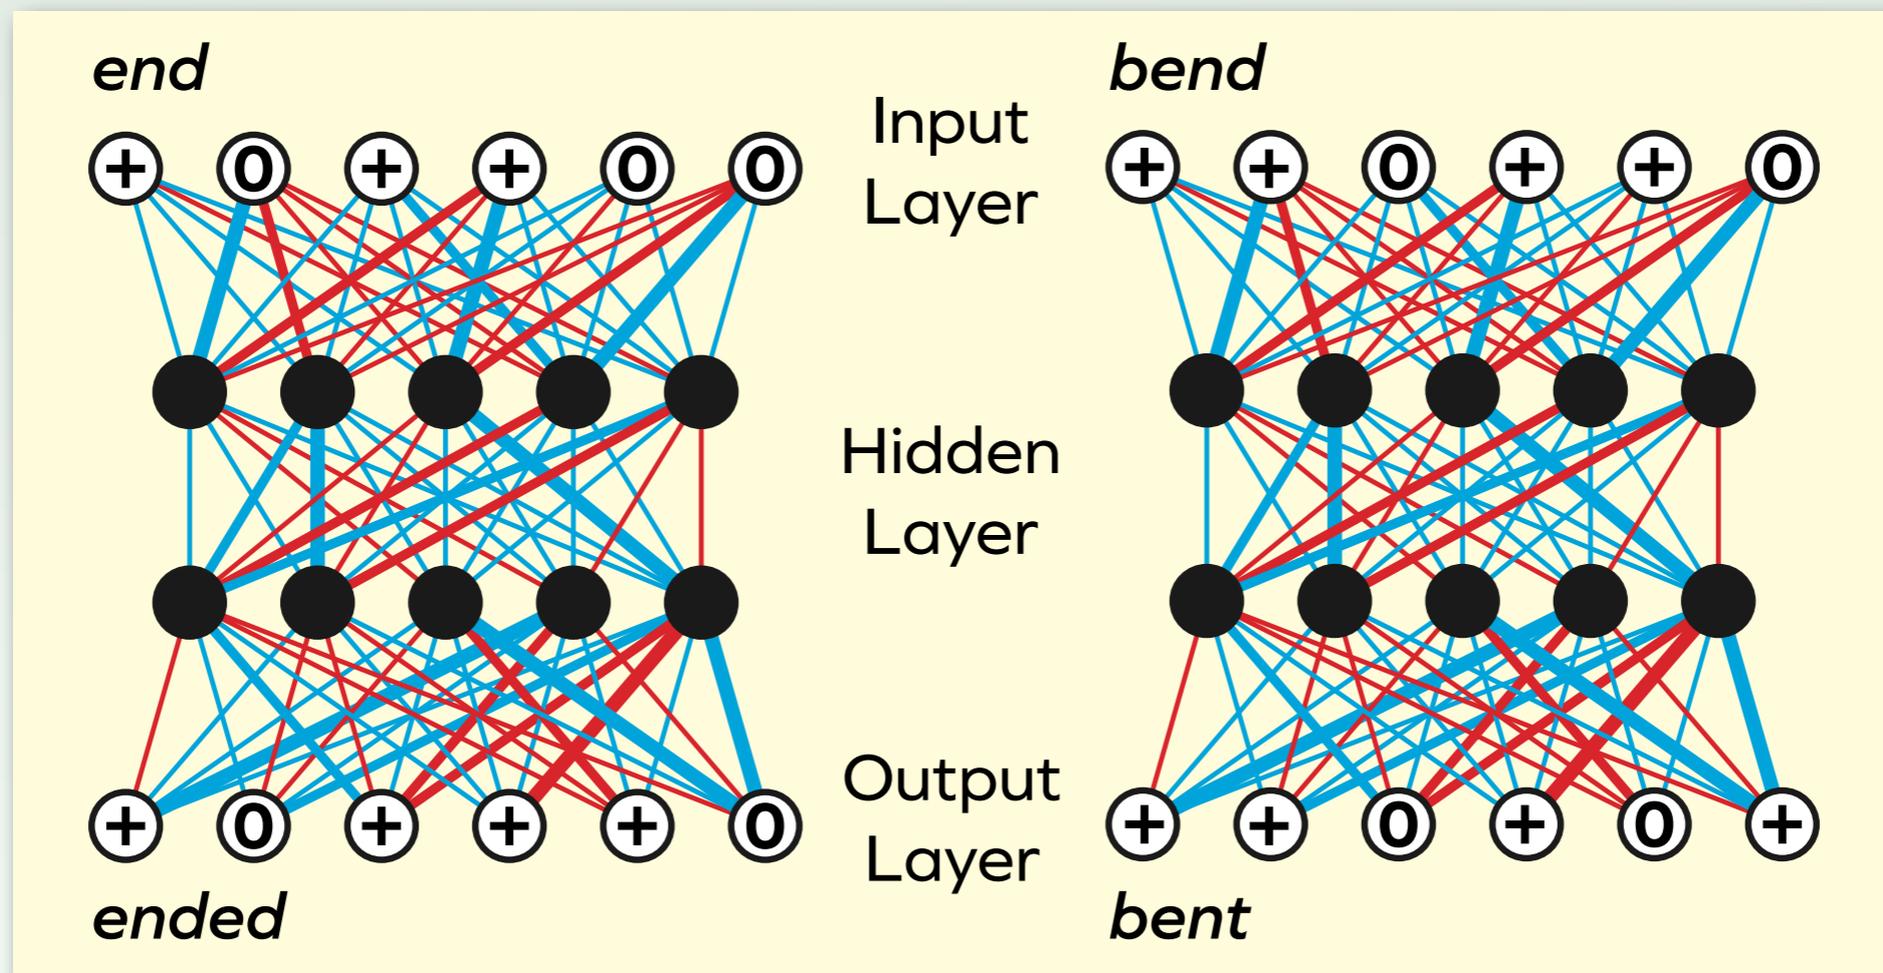


Connectionism: Artificial Brains?

- Only supervised learning is currently practical
- Back propagation: correction applied to each output
 - (1) Difference between desired output and actual output is taken from output units
 - (2) Algorithm adjusts strengths of links for hidden units
 - (3) Cycle many times for many items
- Output gets progressively closer to the desired result for the training set

Connectionism: Artificial Brains?

- Famous findings
 - McClelland & Rumelhart (1985a)
 - Processing verb past tense *end* → *ended*; *go* → *went*
 - ▶ *go* → *went* THEN *go* → *goed* THEN *go* → *went*

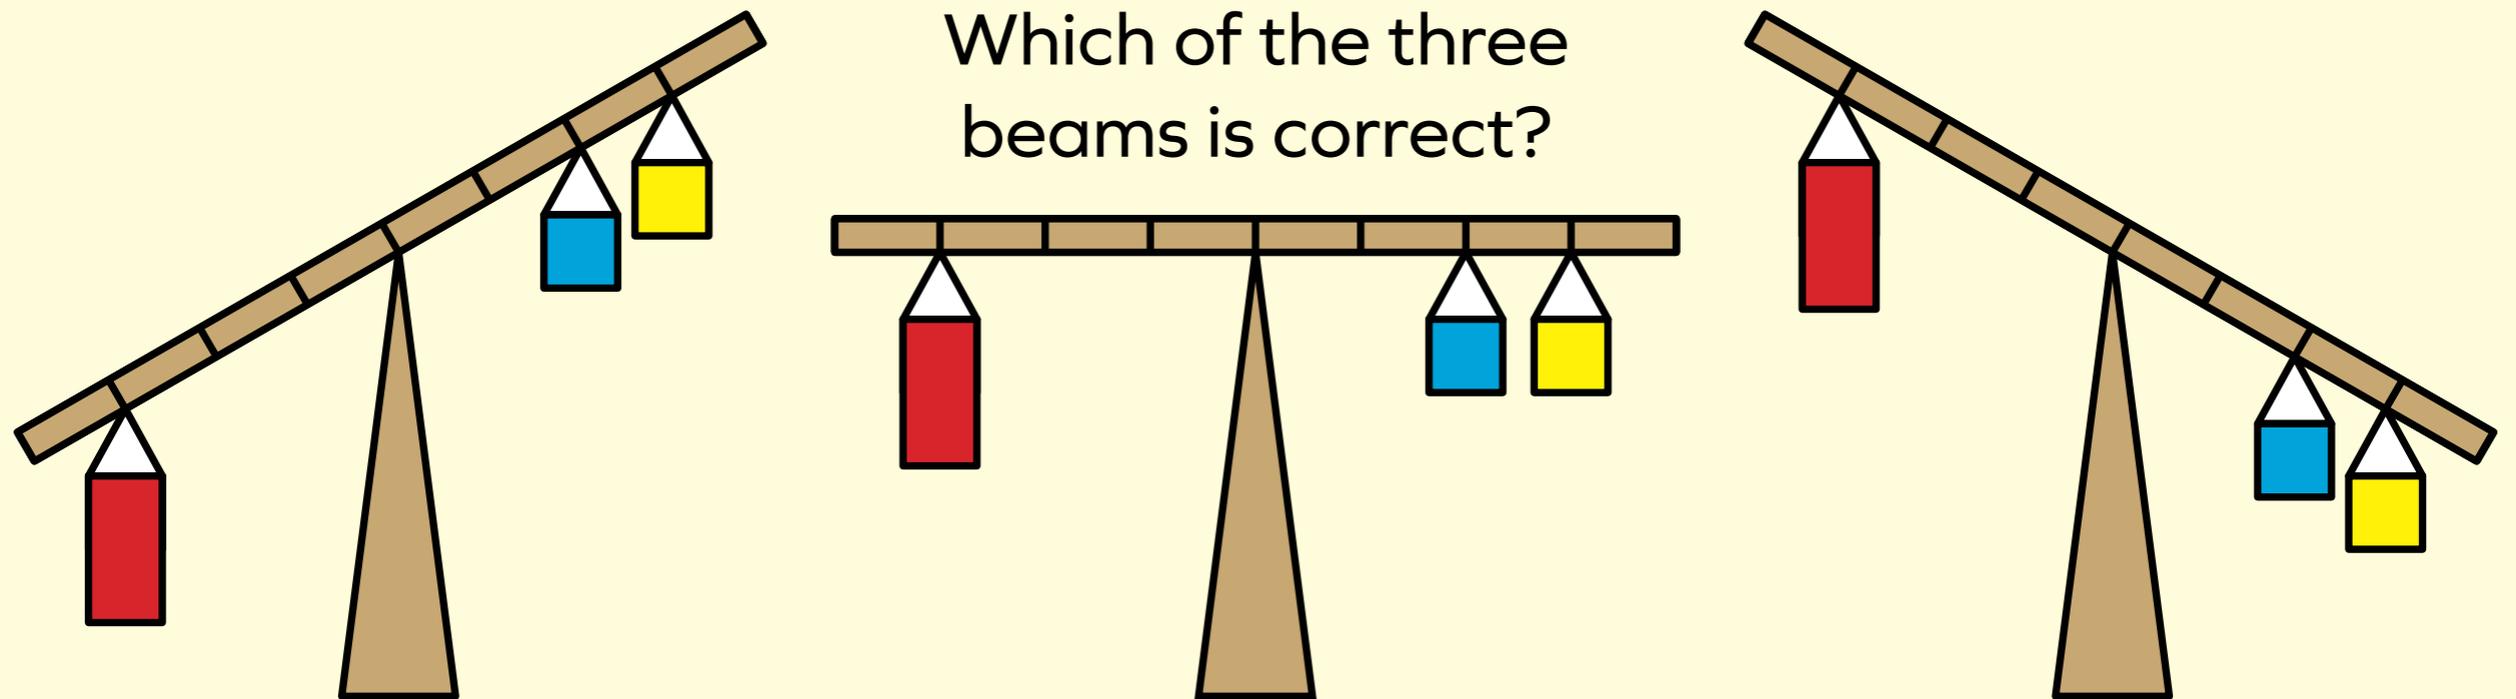


Connectionism: Artificial Brains?

- Famous findings
 - McClelland & Rumelhart (1985a)
 - Processing verb past tense
 - ▶ Learns individual cases
 - Learns rule but overgeneralises to exceptions
 - Learns to apply rule except for special exceptions
 - ▶ Generalises to rule followers not present in training set
 - ➡ Human patterns of grammar
 - learning from blank-slate network?

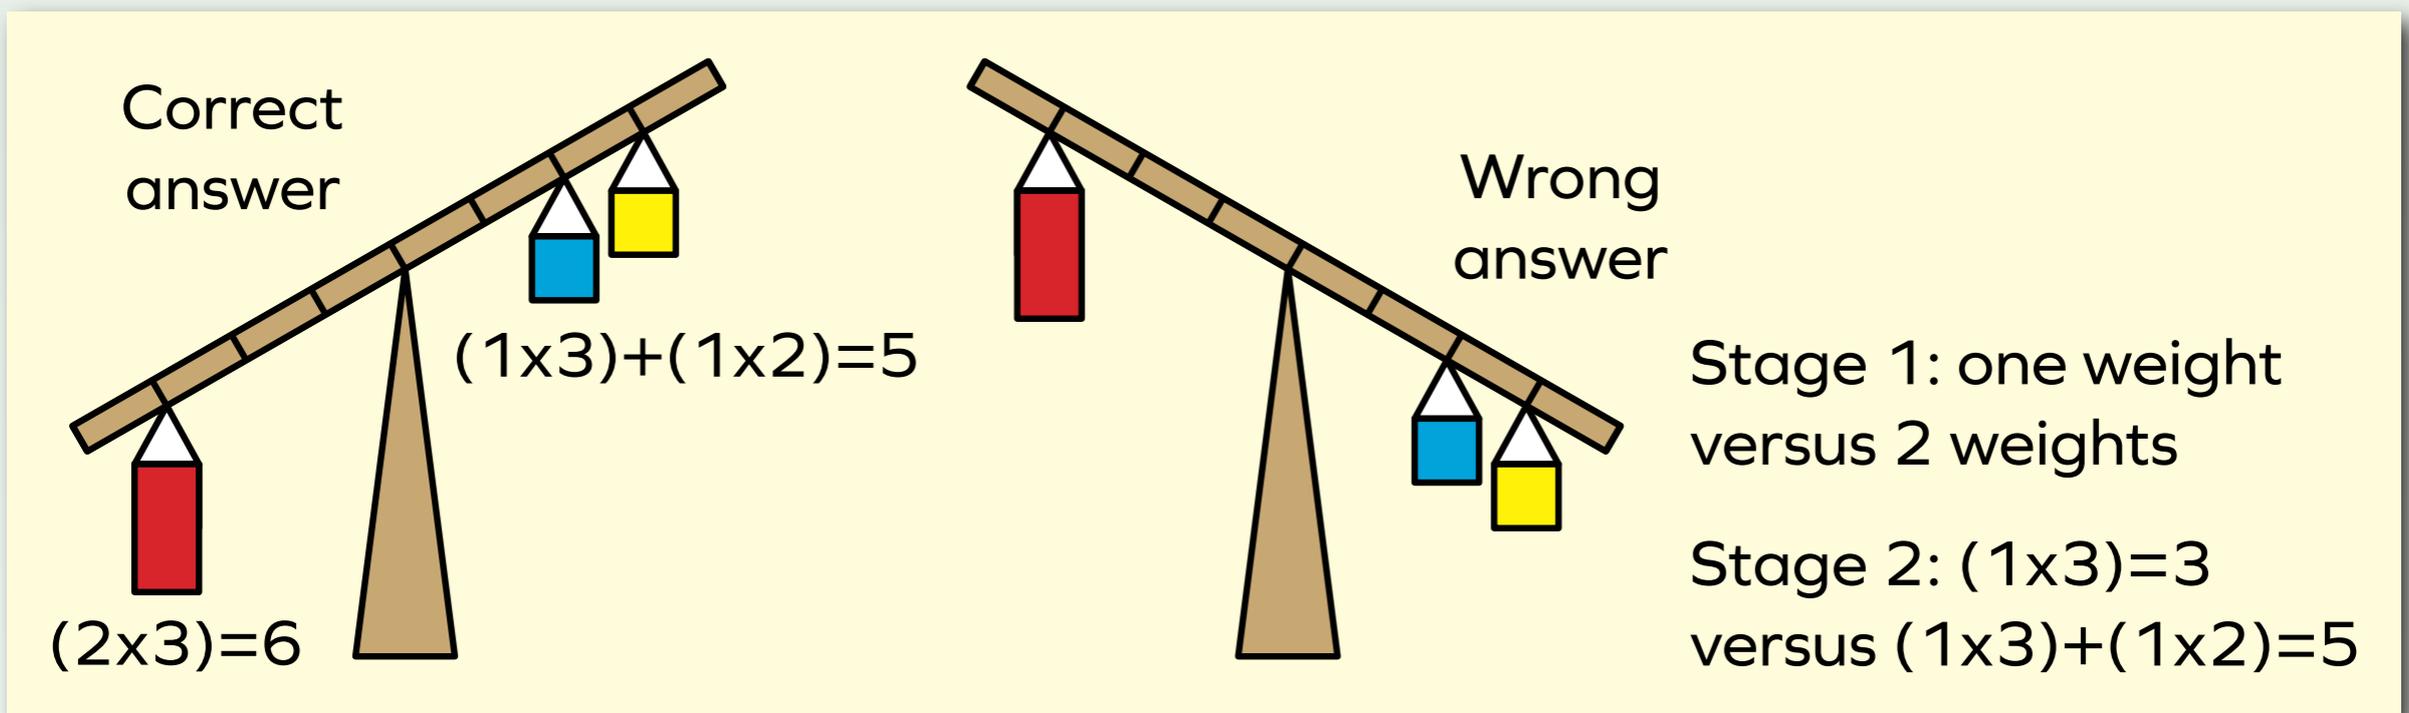
Connectionism: Artificial Brains?

- Famous findings
 - McClelland & Jenkins (1991):
 - Beam balance judgment



Connectionism: Artificial Brains?

- Famous findings
 - McClelland & Jenkins (1991):
 - Beam balance judgment
 - Human sequence
 - S1: Judge by total weights, not their distance
 - S2: Judge by distance if equal weights
 - S3: Judge by distance OR weight OR number
 - S4: Judge by distance AND weight AND number



Connectionism: Artificial Brains?

- Famous findings
 - McClelland & Jenkins (1991):
 - Beam balance judgment
 - Human sequence
 - S1: Judge by total weights, not their distance
 - S2: Judge by distance if equal weights
 - S3: Judge by distance OR weight OR number
 - S4: Judge by distance AND weight AND number
 - ▶ Network displays same sequence, not always reaching S4
 - ➡ Complex rule-based concept acquired, same as human sequence

Connectionism: Artificial Brains?

- ➔ Connectionist networks good for:
 - 'One-shot' low level tasks
 - Perception
 - Categorisation
 - Simple decisions (words, faces)
 - 'Cleaning' data
 - Noise removal
 - Partial matching
- ➔ Bad for sequences
- ➔ Sound basis for explaining and implementing perception, memory, learning?

Connectionism: Artificial Brains?

- Why are connectionist networks interesting?
 - Network looks/acts like a brain?
 - Knowledge distributed through network
 - Less brittle: if 'lesioned', knowledge not necessarily lost
 - Programmer takes a 'hands off' approach
 - ➔ ***Ecological validity?***
 - Human patterns of performance modelled
 - ➔ ***Answers questions in psychology?***

Connectionism: Artificial Brains?

- Why are connectionist networks interesting? *cont.*
 - Enables *Knowledge Problem* to be by-passed
 - No need for explicit rules/knowledge
 - Network can learn/make inductions
 - Memory may be content addressable
- ➔ ***Solves persistent AI problems?***
- Information is sub-symbolic
- ➔ ***Viable alternative to SSSP?***

Connectionism: Reality

- ***Not ecologically valid?***
 - Supervised learning procedures, e.g. back-propagation, not psychologically plausible
 - Unsupervised procedures limited, under development
 - ➔ Learning is possible, but better mechanisms required
- Basic structural differences with brain:
neuronal variety, quantity, connectivity, operation
- Physical similarities are superficial
- ➔ Hard to evaluate importance of differences without a theory of networks

Connectionism: Reality

- ***What are the answers to questions in psychology?***
 - Inter-connected networks of nodes with little or no pre-specification of their connections can generate human-like patterns of learning
 - ➔ 'Empiricist' cognitive development, no 'nativist' innate grammar, number, innate anything
 - ➔ Behaviour might appear rule driven, but these are mere external descriptions, not present in system
 - ➔ Cognitive processes are sub-symbolic; no language of thought, humans are not symbol systems

Connectionism: Reality

- *Incomplete answers*
 - How do symbols and rules emerge from an asymbolic system?
- *Wrong answers*
 - Pinker & Prince (1988)
 - E.g. verb generalisation results suspect
 - ▶ jump → no response
 - ▶ mail → membled
 - ➡ Extracting statistical regularities from training set **NOT** the same as modelling complexities of human behaviour

Connectionism: Reality

- ***Atheoretical answers***
 - Ekbia (2008)
 - Connectionism research =
Can piecemeal aspects of human cognition be mimicked?
 - Findings influenced by network architecture, encoding of stimuli, training task set, but there is no theory of ...
 - What networks can/can't do
 - What happens in a network when it learns
 - How training sets influence findings
 - How stimulus inputting influences findings
 - ➔ Without a theory of networks, research comprises arbitrary statistical modelling, no serious theory of cognition on offer

Connectionism: Reality

- *Only solves some AI issues and also creates new ones*
 - Ekbia (2008)
 - Networks = statistical knowledge filters
 - No understanding, meaning ascribed to output by humans
 - How does meaning emerge from diffuse statistical associations within a network?
 - ➔ *Symbol Grounding Problem* even more serious
 - *Binding Problem*
 - Long red pen on round blue table extremely difficult to encode reliably
 - *[See next section]*

Connectionism: Reality

- ***No alternative to SSSP?***
 - SSSP =
Symbol System Hypothesis + State Space Approach
 - Avoids State Space Approach,
what about Symbol System Hypothesis?
 - Two opposing factions
 - (1) PDP is just an implementation/recreation
of low-level symbols (Broadbent, 1985)
 - (2) PDP: symbol free, qualitatively different from symbol based
approach (McClelland and Rumelhart 1985a and 1985b)

Connectionism: Reality

- ***No alternative to SSSP?***
 - Fodor & Pylyshyn (1988):
 - No separable elements, just patterns of activations
 - ***John loves Mary*** hard to distinguish from ***Mary loves John***
 - *Binding Problem*
 - ▶ No evidence for combinatorial properties of syntax/ semantics expected from genuine symbolic language
 - ➔ Without them, not modelled language
With them, have created a symbol system

Connectionism: Reality

- ***No alternative to SSSP?***
 - Ekbia (2008)
 - How does statistical learning lead to symbolic approximation?
 - Recreation of symbol system from bottom up, or something genuinely different?
 - ➡ System opacity/lack of theory of networks prevents resolution

Connectionism: Reality

- ***No alternative to SSSP?***
 - All connectionism research takes place using digital computers, which simulate them perfectly
 - If the Symbol System Hypothesis is wrong
THEN
 - A digital computer cannot support intelligence/mind
AND
 - A connectionist network that a digital computer can simulate cannot support intelligence/mind
- ➔ If objections to SSH are hardware based, then connectionism is not the answer

Connectionism: Reality

- Connectionism Evaluation
 - Copeland (1993)
"intriguing but trivial results"
 - Dreyfus (attrib)
"the next 20 years of research are unlikely to produce anything with the intellectual ability of a slug"

Connectionism: Reality

- Connectionism Evaluation
 - Advantages
 - Tackles learning issue head on
 - Could complement, rather than replace the SSSP
 - Disadvantages
 - Solves a few issues perhaps, but not all
 - Without a theory of networks, no theoretical insights
 - No exit route from Symbol System Hypothesis

Artificial Life and Situated Cognition

- The science of (simple) self-organising systems
- Attempt to understand fundamental processes of behaviour organisation using computer simulation
- Two sub-disciplines relevant
 - (1) Simulation of worlds in which organisms compete/cooperate for resources (Dyer)
 - (2) Construction of robots that can function in the real world (Steels, Brooks)
- Levy (1992) for a popular account, also Bedau (2003), Brooks (1991, 1992), Dyer (1994), Ekbida (2008), Floreano & Mattiussi (2008), Johnston (2008), Maes (1994), Steels (1994)

Artificial Life and Situated Cognition

- ***Avoid symbols, embody intelligence***
 - Symbols created by programmers are arbitrary and meaningless to a computer

SSSP = learning language from dictionary
 - ➔ Create robots whose behaviour is grounded in real world activities
- ***Avoid cognitive cost of detailed representations***
 - World provides rich data, too complicated to rebuild it inside entities, sample it
 - ➔ Exploit structure in the world to link perception directly to action (c.f. Gibson)

Artificial Life and Situated Cognition

- *Seek minimal criteria, exploit emergence*
 - Analyse task: seek efficient ways to achieve basic goals
 - Focus on simple behaviours and ability to learn/evolve
- ➔ Capitalise on instances where simple rules lead to complex structured behaviour

Artificial Life and Situated Cognition

- Brooks (1991): Artificial Intelligence versus Artificial Life
 - How traditional AI (via SSSP) attempts to create intelligent behaviour
 - Find out how the outside world should be represented inside the organism
 - Identify the cognitive processes that work on these representations
 - Every decision/action will require extensive processing

➔ Knowledge-orientated, top down: build a mind

"... Intelligence is too complex and little understood to be correctly decomposed (...). Furthermore, we will never understand how to decompose human level intelligence until we've had a lot of practice with simpler level intelligences"

Artificial Life and Situated Cognition

- Brooks (1991) *cont.*
 - How ALife attempts to create intelligent behaviour
 - Identify simplest way to create required actions
 - Program entity minimally to achieve exactly this
 - Add more advanced behaviour only if required by the environment or sanctioned by circumstances
 - ➔ Behaviour-oriented, bottom up: build a basic organism
 - "... mobility, acute vision and the ability to carry out survival related tasks in a dynamic environment provide a basis for ... true intelligence"*

Artificial Life and Situated Cognition

- Steels (1994)

“Behaviour is intelligent if it maximises preservation of the system in its environment”

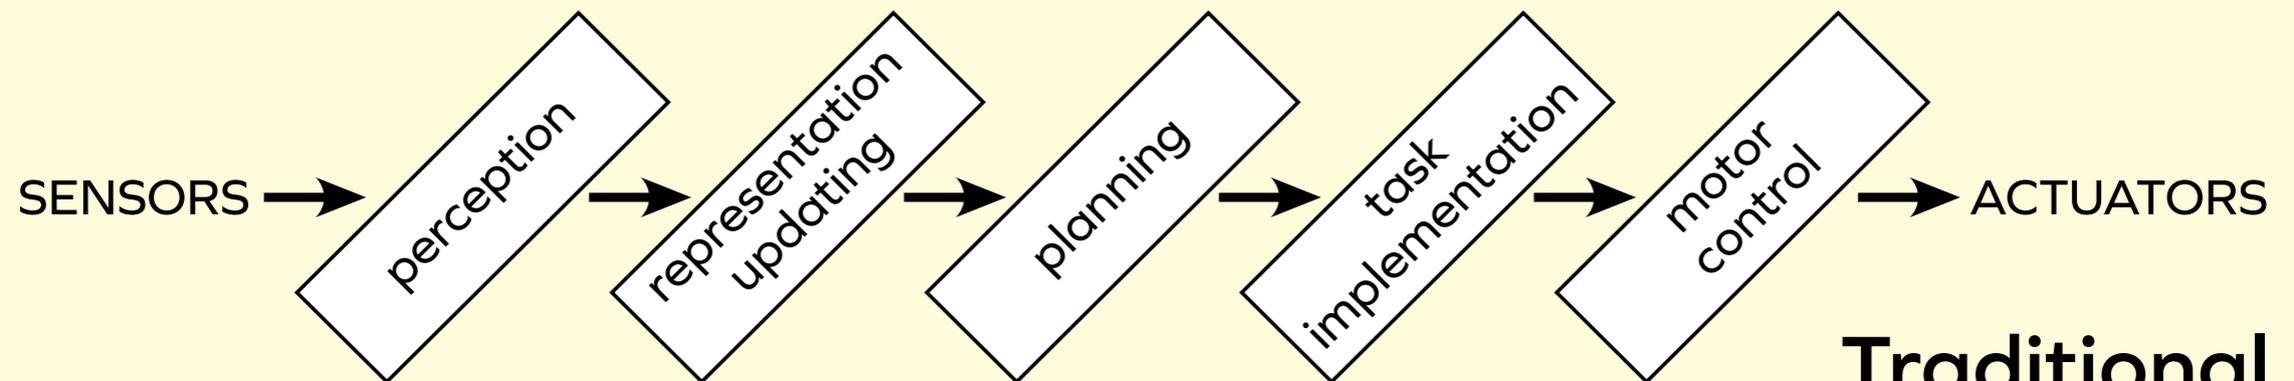
- Brooks (1991)

- Survival behaviour must real time and robust
 - Complex rules slow decision making
 - Robot with a rich internal representation and theory of obstacles will be too slow to survive
 - Luxury of complex behaviour requires basic survival plus hierarchical goals

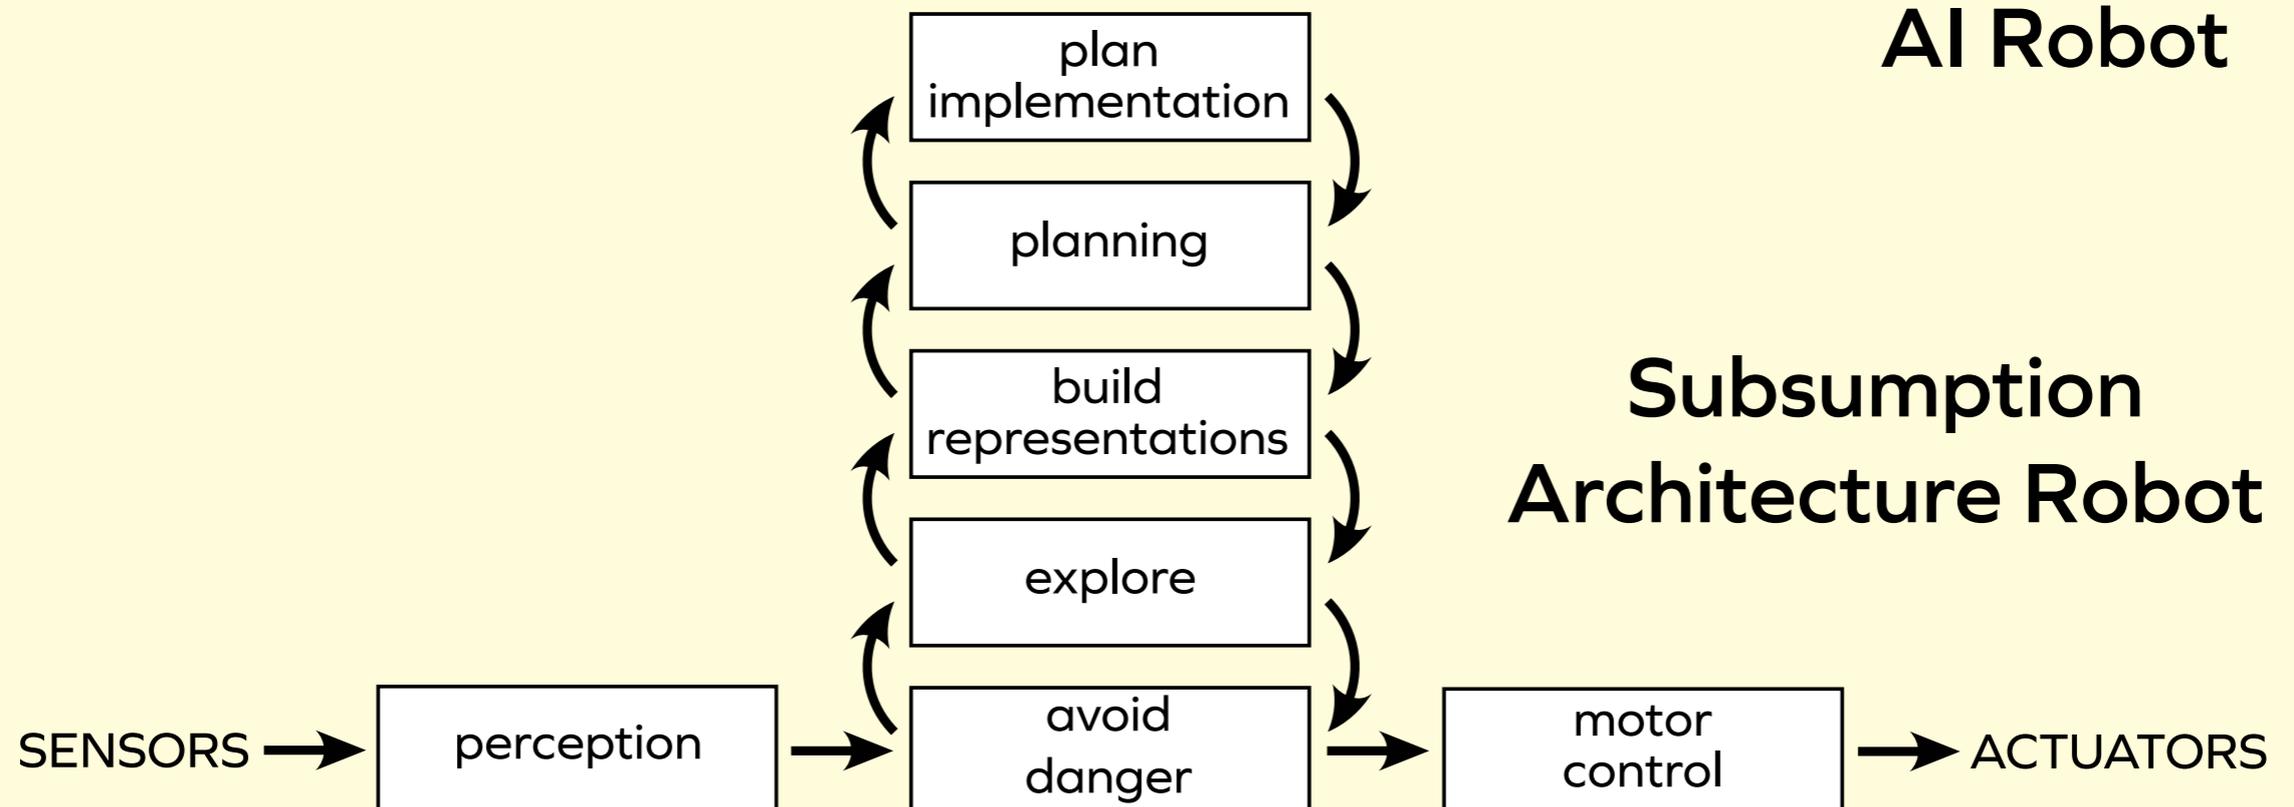
➔ Layered ***subsumption architecture*** permits complexity, avoids processing bottlenecks

Artificial Life and Situated Cognition

- Traditional AI versus ALife *Subsumption Architecture*



Traditional AI Robot



Subsumption Architecture Robot

Artificial Life and Situated Cognition

- Achievements
 - ▶ Complex emergent behaviours from computer simulations and robots (e.g., see Floreano & Mattiussi, 2008, Johnston 2008)
 - Simulated birds with 'nearest neighbours rules' show flocking behaviour
 - A robot programmed to seek walls and avoid collisions follows the walls
 - ▶ Brooks (1991): Ghengis 'stalks' sources of infra-red radiation can steer, track, climb over obstacles
 - ▶ Steels (1994): Simulated robots collect ore, individual message trails co-ordinate group (cf. ant foraging)

Artificial Life and Situated Cognition

- ALife goals
 - Incrementally add complex behaviour
 - (1) Get survival right ✓
 - (2) Next permit exploration ✓
 - (3) Next seek to manipulate environment ✓
 - Solved problems have useful generalisations
 - Less brittle
 - Easy to scale from laboratory to real world
 - Modular robot behavioural mechanism?

Artificial Life and Situated Cognition

- ALife goals
 - ➔ Brooks (1991): robots have insect level intelligence
 - Incrementally add complex behaviour ...
 - (?) Human intelligence in the distant future?
 - No! Brooks (2002):

“I had been thinking that after completing the current round of insect robots, we might move on to reptilelike robots. Then perhaps a small mammal, then a larger mammal, and finally a primate. At this rate (...), it was starting to look as though I might be remembered as the guy who built the best artificial cat. Somehow this legacy just did not quite fit my self-image.”
 - ➔ Refocused on creating expressive human-like robots

Artificial Life and Situated Cognition

- Artificial Life Evaluation
 - Bedau (2003)
 - Same problem as traditional AI
 - Rapid initial gains
 - Stalling, same brick wall problem
 - Ekbia (2008)
 - Now we have a dichotomy
 - Complex behaviour in simple domains (micro worlds)
 - Simple behaviour in complex domains
 - No obvious way to scale up
 - ➔ Combined approaches needed?

Artificial Life and Situated Cognition

- Artificial Life Evaluation
 - Advantages
 - Trying to create real time, real behaviour
 - Easy to attain, realistic goals
 - Seeks understanding of how environment enabled us to become intelligent
 - Disadvantages
 - Reflects lowered aspirations
 - High-profile research programs stalled
 - How far can emergence go?
 - Has the problem of human intelligence been solved or postponed

6.3 Final Thoughts on the SSSP

- CYC tackled the key AI problem head on
- What does its failure mean for the SSSP?
 - (1) The particular approach of CYC to AI is wrong?
 - What knowledge-based approaches could be tried next?
 - Who will resource the time and expenditure?
 - (2) The knowledge approach to AI is wrong?
 - Intelligence = a control problem, not a knowledge problem
 - But general approaches to AI not much more successful

6.3 Final Thoughts on the SSSP

- CYC tackled the key AI problem head on
- What does its failure mean for the SSSP?

(3) The SSSP approach to AI is wrong?

- Intelligence IS NOT a logic driven search of a state space constrained by knowledge and heuristics
- A digital computer is sufficient for AI, but alternatives to brute force state space search are not obvious

(4) Or is Symbol System Hypothesis AI is wrong?

- Intelligent behaviour cannot be implemented on a digital computer
- Nowhere near off-the-shelf non-von-Neuman machines

6.3 Final Thoughts on the SSSP

- Lack of CYC success somewhat uninformative
 - State Space Approach: supplement or replace?
 - Nothing takes us away from the SSH
 - Fresh approaches risk echoing SSSP lameness while they continue to utilise digital computers
 - Repeating pattern likely, every new paradigm gets so far then hits the brick wall
 - Computer speed and memory improvements can offer quantitative enhancement, not qualitative
- Despite setbacks, could a major breakthrough solve the (1 to 5) barriers to AI and save the SSSP?
- Cognitive psychologists: where are the insights, how much do we really know about human cognition?

Major Sources

- Copeland, B.J. (1993). *Artificial Intelligence: A Philosophical Introduction*. Oxford: Blackwell.
- Partridge, D. (1991). *A New Guide to Artificial Intelligence*. Norwood, NJ: Ablex.

For more recent updates:

- Bermudez, J.L. (2010). *Cognitive Science*. Cambridge: Cambridge University Press.
- Ekbia, H.R. (2008). *Artificial Dreams*. Cambridge: Cambridge University Press.
- Floreano, D., & Mattiussi, C. (2008). *Bio-inspired artificial intelligence*. Cambridge: MA, MIT Press
- Johnston, J. (2008). *The allure of machinic life*. Cambridge: MA, MIT Press.
- Nilsson, N.J. (2010). *The Quest for Artificial Intelligence*. Cambridge: Cambridge University Press.

References

- Bedau, M.A. (2003). Artificial life: organization, adaptation and complexity from the bottom up. *Trends in Cognitive Sciences*, 7, 505-512.
- Broadbent, D. (1985). A question of levels: comment on McClelland and Rumelhart. *Journal of Experimental Psychology, General*, 114, 189-192.
- Brooks, R.A. (1992). Intelligence without representation. *Artificial Intelligence*, 47, 139-159.
- Brooks, R.A. (1992). *Robot*. London: Penguin.
- Dreyfus, H.L. (1993). *What Computers Still Cannot Do*. Cambridge, MA: MIT Press.
- Dyer, M.G. (1994). Toward synthesizing artificial neural networks that exhibit cooperative intelligent behavior: some open issues in artificial life. *Artificial Life*, 1, 111-134.
- Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture. *Cognition*, 9, 136-196.
- Harnad, S. (1994). Levels of functioning equivalence in reverse bioengineering. *Artificial Life*, 3, 293-302.
- Johnson-Laird, P.N. (1988). *The Computer and the Mind*. London: Fontana.
- Lenat, D.B., & Feigenbaum, E.A. (1991). On the thresholds of knowledge. *Artificial Intelligence*, 47, 231-250.
- Levy, S. (1993). *Artificial Life*. London: Penguin.
- Maes, P. (1994). Modelling adaptive autonomous agents. *Artificial Life*, 1, 136-162.

References

- McClelland, J.L., & Jenkins, E. (1991). Nature, nurture, and connectionism: Implications for connectionist models of development. In K. VanLehn (ed.), *Architectures for intelligence: The 22nd (1988) Carnegie Symposium on Cognition*. Hillsdale, NJ: Lawrence Erlbaum.
- McClelland, J.L., & Rumelhart, D.E. (1985a). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology, General*, 114, 159-188.
- McClelland, J.L., & Rumelhart, D.E. (1985b). Levels indeed! A response to Broadbent. *Journal of Experimental Psychology, General*, 114, 193-197.
- Pinker, S., & Prince, A. (1988). On language and connectionism: analysis of a parallel distributed processing model of language acquisition. *Cognition*, 28, 73-193.
- Smith, B.C. The owl and the electric encyclopaedia. *Artificial Intelligence*, 47, 251-288.
- Steels, L. (1994). The artificial life roots of artificial intelligence. *Artificial Life*, 1, 75-110.

Copyright Notice

- The text of and organisation of this presentation is copyright ©Maxwell J Roberts, 2018. These slides may be distributed in unaltered form, but must not be reused or reformatted for presentations or teaching, whatever the purpose, and they must not be rehosted for downloading at any other web site, without the express permission of the copyright holder.
- The following images are copyright ©Maxwell J Roberts, 2018, and may not be reused for any purpose except for fair-use educational/illustrative purposes. In particular, they may not be used for any commercial purpose (e.g., textbook, non-open-access academic journal) without the express permission of the copyright holder.
- Slides 9, 30-32 & 34, 36/37, 59
- The web page from which this presentation was downloaded was www.tubemapcentral.com
- All subjective evaluations expressed are the personal opinions of the author
- *This slide must not be deleted from this presentation*