PS452
Intelligent Behaviour

Lecture 5:
Artificial Intelligence
Observations, Objections

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

version date: 14/02/2018

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 5: AI obsevations, objections

- **5.1 The Turing Test revisited**
  - Lines of attack
  - The need for process-based definitions
  - The simulation objection
  - Turing Test aftermath

Lecture 5: AI obsevations, objections

- **5.2 The five barriers to Artificial Intelligence**
    - The Logic Problem
    - The Knowledge Problem
    - The Frame Problem
    - The Symbol Grounding Problem
    - The Intentionality Problem

- **5.3 Is the brain not equivalent to a digital computer?**
    - The status of the Symbolic Search Space paradigm
    - What is a biological computer?

5.1 The Turing Test Revisited

- 5.1 The Turing Test revisited
  - Extensive commentary
    [e.g. Lieber, 1991; Halpern, 1987; Block, 1981]

  - Objections to the test usually comprise two elements

  (1) Objection to utility
    - A computer that passes the test cannot think because …

  (2) Prediction of impossibility
    - A computer will never be able to pass the test because …

  ➡ Which element dominates?

5.1 The Turing Test Revisited

- Turing (1950)
  - *Can machines think?* is a meaningless question
  - Replace with behaviour-based definition
  - ➡ A machine can think if it can win the imitation game
  - ➡ If achieved, is it reasonable to deny thought

Lines of Attack

- Turing (1950)
  - Discusses anticipated objections
  - Some have not resurfaced as theoretical objections

(1) ***Theological:*** Machines cannot have souls

Presupposes the abilities and intentions of God

(2) ***Fear:*** We must not build such machines

Not an argument against the test

(9) ***Telepathy:*** Computers not telepathic (!)

Turing accepts this, but positive findings from the 1950s have been debunked

Lines of Attack

- Turing (1950) *cont.*
  - Objections: low-level limitations of digital computers

    (6) ***Versatility:*** Machines just follow rules, can never be creative

    Machines can surprise with emergent properties

    (7) ***Complexity:*** Human behaviour cannot be coded as rules

    How can we be so sure?

    (3) ***Mathematical:*** Impossible for computers to do some things [Gödel's Theorem and non-computable functions]

    What if brains can't perform these procedures either?

    (4) ***Architecture:*** Brains too different for computers to mimic

    Too subtle to detect in a conversation?

➡ These ***will*** come back to haunt us

Lines of Attack

- Turing (1950) *cont.*
  - Objections: high-level properties of humans

    (6) ***Versatility:*** Machines just follow rules, can never be creative

    Define creativity to exclude machine creations?

    (4) ***Disablement:*** Things machines will never be able to do [e.g. laugh, taste, love]

    This is an induction/pre-judgement from experience

    (5) ***Qualia:*** Computers cannot have consciousness/emotions [Integral to human thought]

    This is an empirical issue, look at the Turing Test data

➡ If problems of slide 9 solved, these become moot

Lines of Attack

- Computers no nearer passing Turing Test than in 1950
  - All claims of computers winning are in error
  - Humans only fooled by partial versions of the test
- ➡ Low-level computer limitation objections look more potent in the light of six decades of AI programming
- ➡ But could a winning computer sweep aside all rational objections leaving behind only human prejudice?
- ➡ What is the logical meaning of passing/failing test

Lines of Attack

- Copeland (1993)
  - Passing Turing Test requires a very narrow and severe subset of intelligent behaviour
  - Success = deceive in a verbal battle
    - Ability to deceive ➔ Ability to think

      [What if machine refused to do take part?!]
    - Skilled verbal ability ➔ Ability to think
  - Animals deceive, communicate?
  - Animals cannot pass this test
  - ➡ Test failures do not necessarily mean non-thinkers

Lines of Attack

- Copeland (1993) *cont.*
  - Turing Test not a simple index of thinking vs nonthinking
  - It is a threshold for untenable denial
    - Minimum universally agreed behavioural requirements
    - ***OK, SO WHAT DOES MY COMPUTER HAVE TO DO TO CONVINCE YOU THEN***
  - ***BUT*** hidden assumption of test:

    Winning test = ability to think + no cheating possible

  ➡ Could a ‘cheating’ program pass the test?

Lines of Attack

- Copeland (1993) *cont.*
  - Thought experiment
    - Finite number of grammatical, meaningful statements 100 words long
    - Finite number of valid conversations in a finite time
    - Program a super-super-supercomputer with every legal conversation
    - Whatever input, look up a suitable response
  - Computer should pass the test easily
  - ➡ Test is not a minimal behavioural threshold, in theory a non-thinking computer can win
  - ➡ A definition of thinking must be processed-based

The Need for Process-Based Definitions

- Copeland (1993) *cont.*
  - Human intelligence =
    intelligent behaviour + intelligent processes
  - Behaviour-based Turing Test definition is not diagnostic
  - ➡ Behaviour criterion must be supplemented
    with a process-based design criterion

    (1) Program must use identical cognitive processes to humans

    But rules out alien ways of being intelligent

    (2) Program is modular, i.e. can make
    anything intelligent such as robots

    Makes Turing Test even more stringent

➡ Now do we have a threshold for untenable denial

The Simulation Objection

- Hallmarks of human intelligence?
  - Articulateness
  - Ingenuity
  - Versatility
- ***BUT*** can a computer with all of these think?

- ***NO!*** simulation objection has not been resolved
  - A simulation of X is not X itself
  - A simulated storm does not make you wet

➡ Simulated thought is not ‘real’ thought

The Simulation Problem

- Copeland (1993)
  - Simulation objection is simplistic
    - Simulation = imperfect copy (good attempt)

      or
    - Simulation = perfect copy (e.g. artificial vitamins)

  ➡ Simulation need not fail to have properties of original

- Turing Test cannot distinguish between simulation categories
- Nor can any other criterion or definition?

➡ Determined skepticism cannot be overturned

Turing Test Aftermath

- A problematic target for AI researchers
- Nonetheless a useful benchmark and thought-experiment generator?
  - Highlights difficulty of evaluating AI
  - Many philosophical issues raised
  - Makes us think about our prejudices about the capabilities of machines

➡ Would humans ever contemplate ascribing ability to think to anything less than a Turing Test winner?

Turing Test Aftermath

- Is the Simulation Objection ideological?
- Why are some people skeptical?

  (1) Negative evaluations of attempts at AI to date

  (2) Belief that barriers to AI are insurmountable

➡ Tenable belief? Empirically and by definition, current Artificial Intelligence cannot create a human-like mind

5.2 The Five Barriers to Artificial Intelligence

- Why hasn’t the Symbolic Search Space Paradigm yielded human-like intelligence, thinking, etc?

➡ Because there are five barriers to Artificial Intelligence

- The Logic Problem
  How can computers be given precise imprecise instructions?
- The Knowledge Problem
  How can knowledge be stored and retrieved ***efficiently****?*
- The Frame Problem
  How can knowledge be updated safely?
- The Symbol Grounding Problem
  How can abstract symbols acquire meaning?
- The Intentionality Problem
  What would it take for a symbol system to understand?

The Logic Problem

- Copeland (1993) for full discussion
- Logics give standard responses to set situations, e.g.

  Modus Ponens: *If A then B, A✓ ∴ B✓*
  Modus Tollens: *If A then B, B✗ ∴ A✗*

- Advantage: computer will respond consistently in a well-understood way
- Without well-understood logic
  - Hard to understand unexpected computer behaviour
  - Cannot develop or test systems easily

➡ Well-understood logic essential if computers are to control important processes (e.g. aircraft)

The Logic Problem

- ***Monotonic logic***
  - Only supports monotonic inferences …
    - No specification for how to change inferences in the light of new information

Operator:	Polly is a parrot Can Polly fly?
Computer:	ALL parrots fly THEREFORE Polly can fly
Operator:	No, Polly is a rare breed of flightless parrot Can Polly fly?
Computer:	??????

The Logic Problem

- Fully formalised logics ***only*** for deductive reasoning
    - Three types of inference specified
        - Guaranteed definitely correct
        - Possible but not guaranteed correct
        - Definitely wrong

The Logic Problem

- Popular logic choice for AI is First Order Predicate Calculus
  - Formalised meanings of IF, AND, OR, NOT: Evans et al., 1993
- Well understood and predictable
- ***Deductive*** and ***monotonic*** hence;
  - Facts are clear-cut, as stated, consistent
  - Ambiguity/uncertainty is tolerated poorly
  - Not formalised for time (causal learning)

➡ Limited compared with human inference

The Logic Problem

- Real world survival/intelligence needs *inductive reasoning*:
- Inductive reasoning is beyond current formalised logic
  - Need rules for inferences from incomplete information
    - Plausible inferences, which might need further investigation
    - Unlikely but possible inferences
    - Useful inferences even though may be incorrect
  - Must be non-monotonic
    - Rules for revision of inferences from new information
  - Must encompass causal reasoning
    - Rules for distinguishing arbitrary pairings vs causal pairings

➡ Without rules for inductive reasoning, cannot tell (programme) computer what to do in the real world

The Logic Problem

- *Fuzzy logic* supposedly copes with imprecision, uncertainty, partial truth
  - Not fully formalised, poorly understood
  - Computer behaves in a bizarre way, why?
    - Sanctioned by the logic
    - Error by programmer, needs fixing
    - Now computer programming becomes matter of debate
  - Using poorly understood logic defeats purpose of logic
- Many proposals but ***still*** little progress (Nilsson, 2010)

➡ *Max's conjecture: a being cannot create an entity more intelligent than itself, because that requires a more advanced logic than the being can formalise*

The Knowledge Problem

- Human knowledge
  - Retrieved quickly and effortlessly
    - Not obviously slowed down, possibly speeded by quantity
    - Compiled (proceduralised) for speed and accuracy
  - Rapid inferences to fill in gaps
    - Implicit meanings in communication
    - Common-sense questions
- Computer knowledge easy to encode ***BUT***
  - Difficult to retrieve efficiently
    - Locating knowledge intelligently and quickly is hard; it must be structured
    - Adding knowledge slows the search

The Knowledge Problem

- How can a large number of facts be represented in a computer and retrieved quickly?

  When you are visiting the pyramids in Egypt, is your left foot also in Egypt

  (1) Alphabetical encyclopaedia?

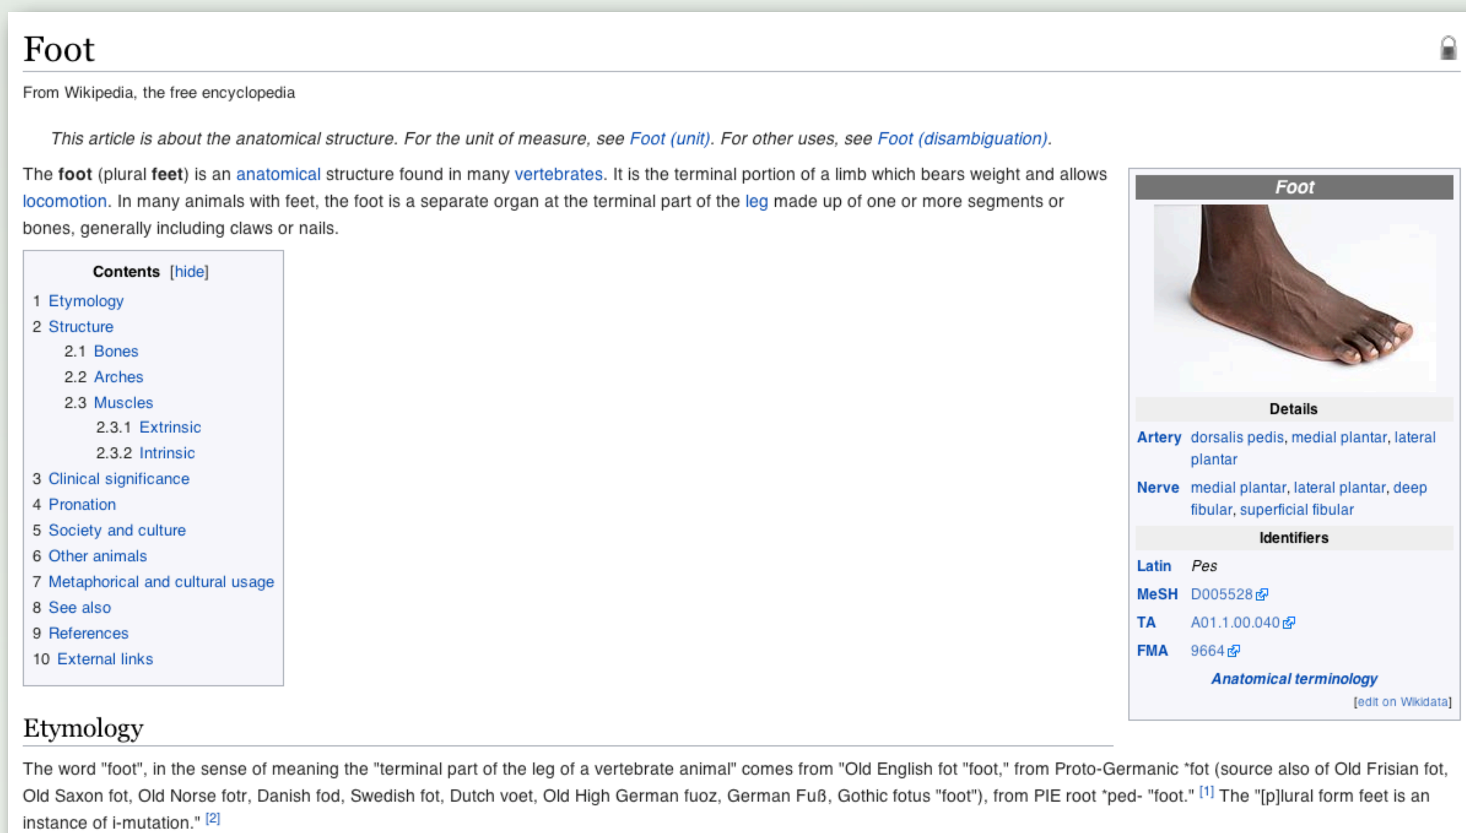

Foot

From Wikipedia, the free encyclopedia

This article is about the anatomical structure. For the unit of measure, see Foot (unit). For other uses, see Foot (disambiguation).

The **foot** (plural **feet**) is an anatomical structure found in many vertebrates. It is the terminal portion of a limb which bears weight and allows locomotion. In many animals with feet, the foot is a separate organ at the terminal part of the leg made up of one or more segments or bones, generally including claws or nails.

Contents [hide]

Foot

Details

Artery	dorsalis pedis, medial plantar, lateral plantar
Nerve	medial plantar, lateral plantar, deep fibular, superficial fibular

Identifiers

Latin	*Pes*
MeSH	D005528
TA	A01.1.00.040
FMA	9664

Anatomical terminology

[edit on Wikidata]

Etymology

The word "foot", in the sense of meaning the "terminal part of the leg of a vertebrate animal" comes from "Old English fot "foot," from Proto-Germanic *fot (source also of Old Frisian fot, Old Saxon fot, Old Norse fotr, Danish fod, Swedish fot, Dutch voet, Old High German fuoz, German Fuß, Gothic fotus "foot"), from PIE root *ped- "foot." [1] The "[p]lural form feet is an instance of i-mutation." [2]

➡ Impossible!

The Knowledge Problem

- How can a large number of facts be represented in a computer and retrieved quickly?

  When you are visiting the pyramids in Egypt, is your left foot also in Egypt

  (2) Production system:
  Huge numbers of propositions and rules?

  IF THERE IS NO REASON TO BELIEVE YOUR FEET ARE DISCONNECTED FROM YOUR BODY THEN YOUR FEET ACCOMPANY YOU AT ALL TIMES

  - Hard to account for all eventualities
  - Very slow to search, combinatorial explosion
  - No obvious way to guide search,
  - Contradictions difficult to detect

➡ Theory of meta-knowledge necessary

The Knowledge Problem

- How can a large number of facts be represented in a computer and retrieved quickly?

When you are visiting the pyramids in Egypt, is your left foot also in Egypt

(3) Semantic network?

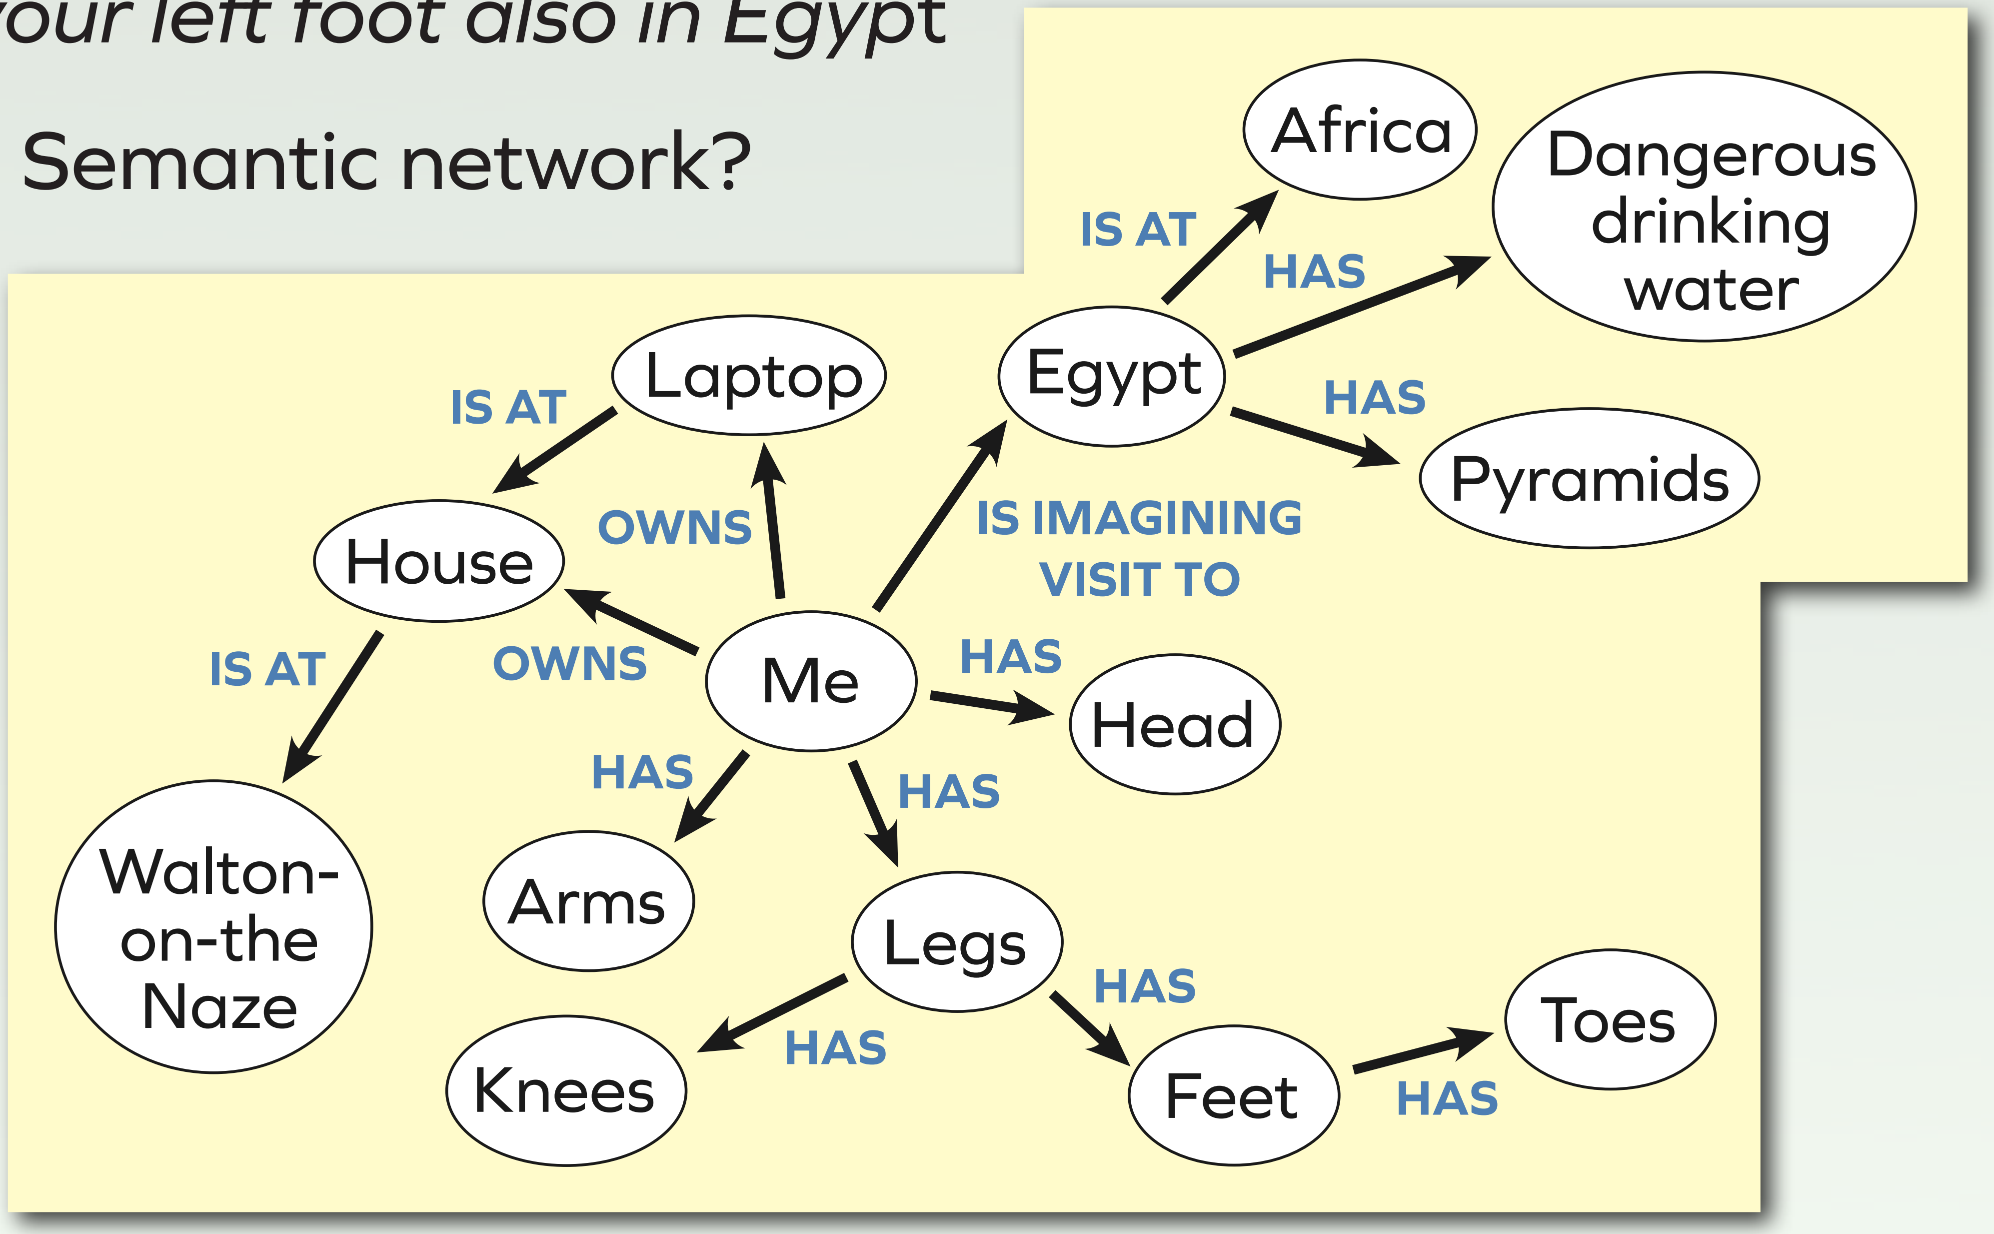

The Knowledge Problem

(3) Semantic network *cont.*

- Difficult to …
  - Implement negation, exception, disjunction
  - Represent vague concepts; usually and probably
    - *How should indeterminacy be interpreted when following links to make inferences?*
  - Organise and update (see *Frame Problem*)
    - *Imagining a visit to Egypt is temporary*

The Knowledge Problem

(3) Semantic network *cont.*

- Partridge (1991)
  - Beware impressive appearances
  - Humans very good at implicit inferences
    - Gives network elements meaning
    - Fills in gaps in network
  - Where is the meaning in abstract network without a human interpreter?
    - Internal consistency = meaning?
    - ➡ Almost certainly not, Symbol Grounding Problem

➡ Semantic networks = satisfying representations for observer rather than prescriptive for programmer

The Knowledge Problem

- Copeland (1993)
  - Memory addressing is the fundamental issue
    - Human memory is content addressable:
      Information activated/retrieved by association
    - Computer memory is compartment addressable:
      Information retrieved from designated locations
      Locations dissociated, cannot cross-activate

      Look in register 1010101 for the name of my pet

      Look in register 1011100 for what animal my pet is
    - Information retrieved only if the exact address known
    - Or else search all registers for required information

    ➡ Nowhere near a practical content addressable memory

The Frame Problem

- Haugeland (1985) for detailed discussion
- If one item of knowledge changed,
  how does the system know what else to update?
- Some rules for a car:

  IF CAR WILL NOT START AND STARTER MOTOR TURNS, HEADLIGHTS BRIGHT THEN CHECK LEADS AND PLUGS

  IF LEADS AND PLUGS OK THEN CHECK DISTRIBUTOR

  - But what if you can't find a distributor?
    - Car has electronic ignition
    - Rules will need updating, some OK, some wrong

  ➡ How do you know which rules to scrutinise?

The Frame Problem

- The world is in a constant state of flux
- Updating knowledge is important
  - Changing one item may affect many others
  - Must add rules, modify or delete others
  - Knock-on means unrelated rules might need updating

➡ Maintenance of a computer knowledge-base not trivial

The Frame Problem

- A particular problem for production rules
  - Isolated statements
  - Difficult to keep a system coherent

(1) Assume that nothing else has changed unless find out otherwise?

➡ Could be dangerous (e.g. expert systems)

The Frame Problem

(2) Check every rule pair for consistency?

- If two rules contradict then at least one must be modified, but which, and how?
- Need to check the modified rules against all others in case inconsistencies introduced

➡ Impossible for a large program, e.g. 500 rules = 124750 pairs to check and resolve as necessary

(3) Organise knowledge into categories, check changes within categories?

- Real world rarely that neat!

➡ Only practical for small systems?

The Frame Problem

(4) Use heuristics to determine which attributes update after particular events

- E.g. moving an object does not usually change its colour

➡ Impossible to construct comprehensive heuristics

The Frame Problem

- Humans better at updating knowledge
  - But still not always easy *(c.f. moving house)*
  - Helped by …
    - Inductive reasoning
    - Non-monotonic reasoning
    - Easy to retrieve associated knowledge
    - Common sense

➡ Solve Logic/Knowledge problems to solve Frame Problem?

The Symbol Grounding Problem

- Harnad (1990)
  - How can an abstract set of symbols acquire meaning?
  - Fundamental problem for SSSP
  - Semantics from syntax: what is needed?
    - Size of database
    - Internal consistency of database
    - Embodiment/links to real world
    - ????
  - Similar problem for human brain?
  - Semantics from synapse: what is needed
    - How can communicating neurones embody meaning

The Symbol Grounding Problem

- ➡ Level of description problem?
    - Know humans representations have meaning, so philosophical difficulties are not attended to
    - More obvious that computer operation is lacking?
- ➡ Invoke *emergent property* argument
    - Even if valid for humans, does not mean that 'more = magic' argument must apply to computers

The Intentionality Problem

- Searle (1980)
    - Suppose we have a room into which we pass sentences written in Chinese
        - Replies are lucid and intelligent
        - Indistinguishable from a native speaker
        - Someone/something in the room understands Chinese
    - Inside the room is a human
        - Books of rules for responding to Chinese
        - Human looks up rules and applies them
        - Input is translated into output
    - Does the human understand Chinese?

➡ No, the human follows rules blindly

The Intentionality Problem

- Searle (1980) *cont.*
  - SSSP uses rules to manipulate symbols
    - Cannot possibly understand their meanings
  - ➡ Strong SSH is WRONG:
    - Human mind content is meaningful, computer program content can never be meaningful
    - Computer symbol systems can never be equivalent to human brains, hardware equivalence is a fallacy
  - ➡ Strong AI is impossible
    - Computer can only simulate a mind, never be a mind

The Intentionality Problem

- Searle (1980) *cont.*
  - Anticipates some rebuttals:

(1) The man is component of a system that collectively understands Chinese

- Suppose the man were to memorise all of the rules, he would still not understand

(2) Suppose the rules were programmed into a walking, talking, seeing robot

- Man could be manipulating rules that control the robot (ignorant homunculus)

(3) Suppose a computer simulates each and every neurone

- Man could be working a set of water pipes and valves which work in the same way

The Intentionality Problem

- Searle (1980) *cont.*
  - Underlying thinking: Human thinking is intentional
    - We think ***about*** things; semantic
  - Computer rules are syntactic
    - No content to them
  - ➡ Semantics cannot be derived from syntax
- Searle (1983)
  - Something special about physical/ chemical properties of brains
  - ➡ Only computers with these properties can have intentional (semantic) thoughts

The Intentionality Problem

- Intentionality problem more general than Symbol Grounding Problem
  - *Understanding* and other intentional concepts border on *qualia*
    - *I understand this text*
    - *This text is easy to understand*
    - *I like this text because it is easy to understand*
  - How can an abstract set of symbols …
    - Acquire meaning?
    - Acquire feeling?

➡ Would solving the Symbol Grounding Problem solve the Intentionality Problem?

The Intentionality Problem

- Responses to the *Chinese Room* thought experiment *cont.*
  - Created a storm of protest
  - Gardner (1987)
    - Intentional position is inadequately explained
  - Hofstadter (Searle, 1983, commentary)
    - Searle has a religious position not a scientific case
  - Pylyshyn (Searle, 1983, commentary)
    - What if each neurone was replaced one by one by a chip that imitated its performance
    - When would understanding/intention cease?

The Intentionality Problem

- Responses to the *Chinese Room* thought experiment *cont.*
  - System Argument is the best rebuttal
  - Copeland (1993)
    - Searle's argument is that:
      - The man in the room cannot understand Chinese
      - Therefore the room cannot understand Chinese
    - And is logically equivalent to:
      - The lecturer at the university does not teach maths
      - Therefore the university does not teach maths

    ➡ The Whole System understands

The Intentionality Problem

- Responses to the *Chinese Room* thought experiment *cont.*
  - Copeland (1993) *cont.*
    - Memorising rules argument = evasion
      - If the man cannot understand Chinese,
      - Then no part of him can understand Chinese
    - Is logically equivalent to:
      - If the man cannot regulate his heartbeat,
      - No part of him can regulate his heartbeat

    ➡ Whole system cannot be dissected

    ➡ Chinese room argument is invalid

    ➡ In theory, the right SSSP program can understand

The Intentionality Problem

- My Chinese Room observations
  - What exactly does the word *understand* mean?
  - How do I know that YOU understand?
  - Behaviour-based/task-based definitions possible, but what would comprise a process-based definition?

  ➡ Hard to define word *understand* such that *Chinese Room* does not understand

    - Negative definition: these particular processes comprise non-understanding
    - Subjective/personal feeling qualia definition, but such feelings can be erroneous

  ➡ Foundation of *Chinese Room* is missing

The Intentionality Problem

- My Chinese Room observations *cont.*
  - Could the *Chinese Room* even exist?
  - Even in a computer, is this system of subjecting content to rule look-up remotely plausible?
  - Natural language understanding (Lecture 2) has not been a resounding success
  - No obvious system of learning, Searle presumably implicitly intends this to be part of the rule books
  - ➡ SSSP cannot deliver a Chinese Room? unless the Frame Problem is solved (+ Knowledge/Logic Problems?)
  - ➡ SSSP that Searle is criticising is an unimaginably different entity from the one that we currently have

5.3 Is the Brain not Equivalent to a Digital computer

- Barriers are closely linked

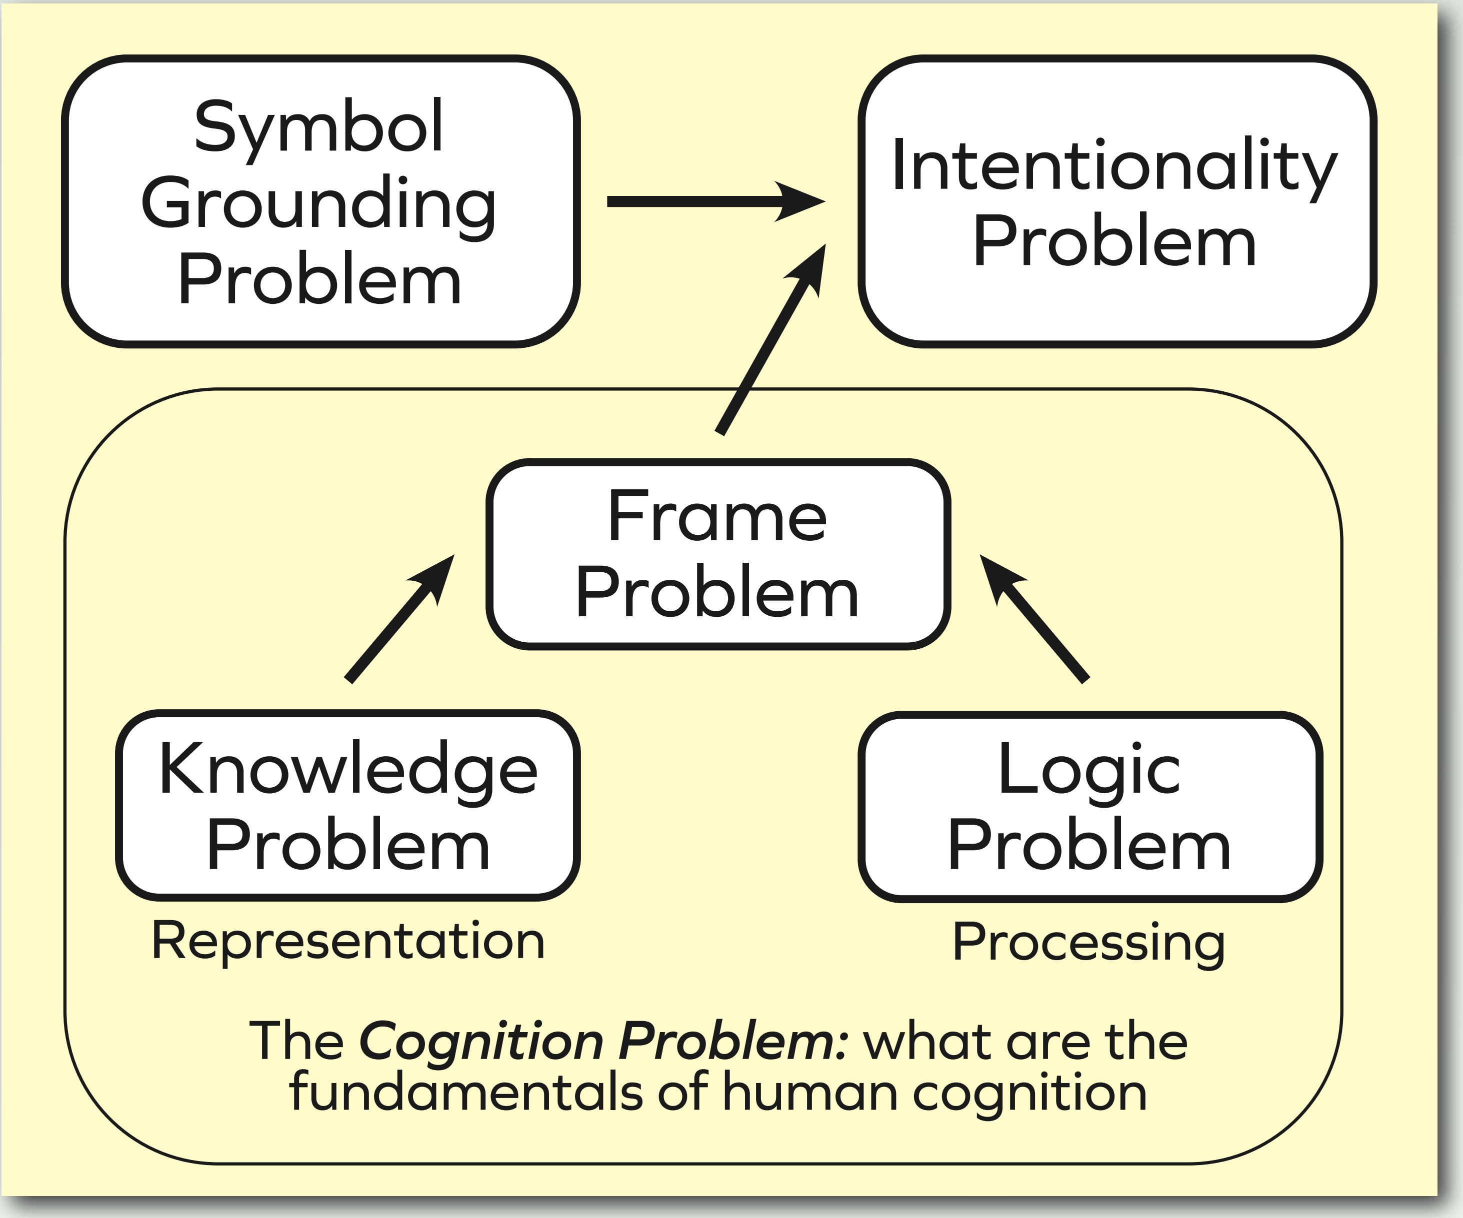

5.3 Is the Brain not Equivalent to a Digital computer

- Solve Logic and Knowledge Problems to solve Frame Problem?
- Solve Frame and Symbol Grounding Problems to solve Intentionality Problem?

➡ Something [one thing?] is fundamentally wrong with the SSSP

  - Fix [one thing?] and Artificial Intelligence happens: SSSP useful and the SSH is plausible
  - Or is that [one thing?] that the SSSP is utterly wrong

➡ Need to look at brain-computer equivalence again

The Status of the Symbolic Search Space Paradigm

- Structural/operational differences between brain and von Neumann machine
  - Brain
    - 100 billion units, highly connected, slow, analogue, parallel, embodied
  - Digital computer
    - A few billion units, poorly connected fast, digital, coordinated serially, usually disembodied

The Status of the Symbolic Search Space Paradigm

- Quantitative differences between brain/digital computer
  - Excellent case for Strong Symbol System Hypotheses

- Qualitative differences between brain/digital computer
  - Symbol System Hypotheses strong/weak/wrong
    - Strong SSH: AI research is a thoroughly worthwhile pursuit
    - Weak SSH: AI might yield useful results but human level intelligence/understanding could be impossible
    - Wrong SSH: AI research is limited, all we can hope for is faster *brute force* in the future

What is a Biological Computer?

- Alternatives to von Neumann machine (digital computer) difficult to identify
- How might a 'brain machine' be created

(1) Searle (1983)

  - Physical/chemical properties make the brain intentional
  - Meaning must be inherent in the system, can't be bolted on

(2) Penrose (1990, 1992)

  - Neurones transmit information via quantum events as well as chemical events
  - Otherwise neurones subject to same restrictions as digital switches

What is a Biological Computer?

- Dreyfus (1993)
  - Brain is an embodied analogue *cybernetic* device
  - Inputs ‘flow’ through the brain, disrupting equilibrium, negative feedback restores equilibrium
- Claims are intriguing but difficult to substantiate

➡ Not many clues for where to go next

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.
- Nilsson, N.J. (2010). *The Quest for Artificial Intelligence*. Cambridge: Cambridge University Press.

References

- Block, N. (1981). Psychologism and behaviorism. *The Philosophical Review*, *90*, 5-43.
- Dreyfus, H.L. (1993). *What Computers Still Cannot Do*. Cambridge, MA: MIT Press.
- Evans, J.St.B.T., Newstead, S.E. and Byrne, R.M.J. (1993). *Human Reasoning: The Psychology of Deduction*. Hove: Lawrence Erlbaum.
- Gardner, H. (1987). *The Mind's New Science*. New York: Basic Books.
- Halpern, M. (1987). Turing's test and the ideology of artificial intelligence. *Artificial Intelligence Review*, *1*, 79-93.
- Harnad, S. (1990). *The symbol grounding problem*. *Physica*, *D42*, 336-346.
- Lieber, J. (1991). *An Invitation to Cognitive Science*. Oxford: Blackwell.
- Penrose, R. (1990). *The Emperors New Mind*. Oxford: OUP.
- Penrose, R. (1992). Précis of the emperors new mind. *Behavioral and Brain Sciences*, *13*, 643-655.
- Searle, J.R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, *3*, 417-457.
- Searle, J.R. (1983). *Intentionality: an Essay in the Philosophy of Mind*. Cambridge: CUP.
- Turing, A. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.

Copyright Notice