Can Artificial Life Engender Real Understanding?

Bruce MacLennan
Dept. of Computer Science
www.cs.utk.edu/~mclennan

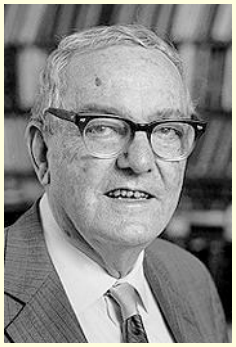

"Perhaps the greatest significance of the computer lies in its impact on Man's view of himself… [T]he computer aids him to obey, for the first time, the ancient injunction *Know thyself*."

—Herbert Simon
(Nobel Laur., 1978)

I. Disembodied Reasoning

Historical Background

- Reason & language as characteristic human abilities
- Cartesian dualism
- Thought as computation
  - "By *ratiocination* I mean *computation*." (Hobbes)
- Mechanized logic
  - Leibniz, Boole, Jevons, …

Development of Cognitive Science

- Convergence of scientific & technological developments c. 1960
- Behaviorism inadequate for explaining cognitive processes
- Computer models of cognition provide an alternative
- More powerful computers permit testing the hypothesis that thought is computation

The Cognitive Sciences

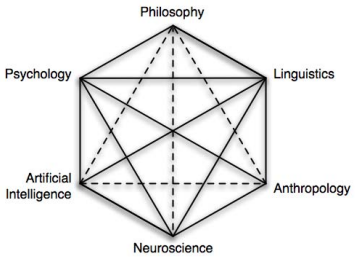

(based on Gardner, 1985)

Traditional Definition of Artificial Intelligence

- "*Artificial Intelligence* (AI) is the part of computer science concerned with designing intelligent computer systems,
- that is, systems that exhibit the characteristics we associate with intelligence in human behavior —
- understanding language, learning, reasoning, solving problems, and so on."

— *Handbook of Artif. Intell.*, vol. I, p. 3

2005-10-20 7

Traditional AI

- Long-term goal: equaling or surpassing human intelligence
- Approach: attempt to simulate "highest" human faculties:
  - language, discursive reason, mathematics, abstract problem solving
- Cartesian assumption: our essential humanness resides in our reasoning minds, not our bodies
  - *Cogito, ergo sum.*

2005-10-20 8

Formal Knowledge-Representation Language

• Spot is a dog	• dog(Spot)
• Spot is brown	• brown(Spot)
• Every dog has four legs	• $(\forall x)(\text{dog}(x) \rightarrow \text{four-legged}(x))$
• Every dog has a tail	• $(\forall x)(\text{dog}(x) \rightarrow \text{tail}(x))$
• Every dog is a mammal	• $(\forall x)(\text{dog}(x) \rightarrow \text{mammal}(x))$
• Every mammal is warm-blooded	• $(\forall x)(\text{mammal}(x) \rightarrow \text{warm-blooded}(x))$

2005-10-20 9

Graphical Representation (Semantic Net)

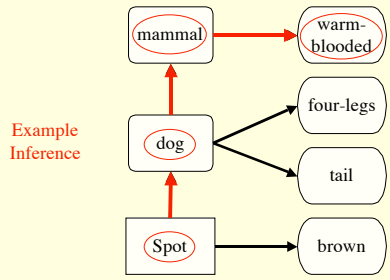

2005-10-20 10

Example of Propositional Knowledge Representation

IF

1) the infection is primary-bacteremia, and
2) the site of the culture is one of the sterile sites, and
3) the suspected portal of entry of the organism is the gastrointestinal tract,

THEN

there is suggestive evidence (.7) that the identity of the organism is bacteroides.

2005-10-20 11

Limitations of Traditional AI

- Brittleness of expert systems
- Combinatorial explosion
- Context-sensitivity & relevance
- Non-classical concepts
- Ungrounded symbols
- Common-sense knowledge
- Non-verbal cognition
- The "cognitive inversion"

2005-10-20 12

Five Stages of Skill Acquisition

1. Novice
   - learns facts & rules to apply to simple "context-free" features
2. Advanced Beginner
   - through experience, learns to recognize similar situations
3. Competence
   - uses developing sense of relevance to deal with volume of facts
4. Proficiency
   - analytical thinking is supplemented by intuitive organization & understanding
5. Expertise
   - skillful behavior is automatic, involved, intuitive, and fluent.

2005-10-20 13

The Cognitive Inversion

- Computers can do some things very well that are difficult for people — **abstract skills**
  - e.g., arithmetic calculations
  - playing chess & other abstract games
  - doing proofs in formal logic & mathematics
  - handling large amounts of data precisely
- But computers are very bad at some things that are easy for people (and even some animals) — **concrete skills**
  - e.g., face recognition & general object recognition
  - autonomous locomotion
  - sensory-motor coordination
- Conclusion: *brains work very differently from digital computers*

2005-10-20 14

The 100-Step Rule

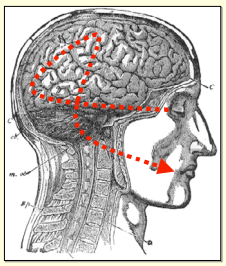

- Typical recognition tasks take less than one second
- Neurons take several milliseconds to fire
- Therefore then can be at most about 100 sequential processing steps

2005-10-20 15

"The New AI"

- A new paradigm that emerged in mid-80s
- Convergence of developments in:
  - philosophy
  - cognitive science
  - artificial intelligence
- Non-propositional knowledge representation
  - imagistic representation & processing
  - propositional knowledge as emergent
- Neural information processing
  - connectionism (implicit vs. explicit representation)
  - critical dependence on physical computation

2005-10-20 16

II. The Embodied Mind

2005-10-20 17

The Embodied Mind

- Brain
  - the brain matters
- Embodiment
  - the body matters
- Situatedness
  - the world matters

2005-10-20 18

How Dependent is Intelligence on its Hardware? Traditional View

- Brain is no more powerful than Turing machine
- Human intelligence is a result of the program running on our brains (Cartesian dualism)
- The same program could be run on any Universal TM
- In particular, it could run on a digital computer and make it artificially intelligent
- Ignores "performance" (as opposed to "competence")

2005-10-20 19

Connectionist View

- Information processing on digital computers (*hardware*) is fundamentally different from that in brains (*wetware*)
- The flexible, context-sensitive cognition we associate with human intelligence depends on the physical properties of biological neurons
- Therefore, true artificial intelligence requires sufficiently brain-like computers (*neurocomputers*)

2005-10-20 20

Neural Information Processing

- *100-Step Rule* & *Cognitive Inversion* show brains operate on different principles from digital computers
  - "wide & shallow" vs. "narrow & deep"
- How do brains do it?
- Can we make *neurocomputers*?

2005-10-20 21

Neural Density in Cortex

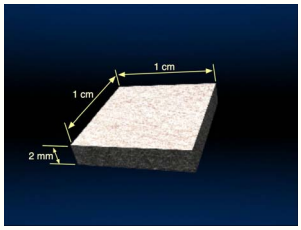

- 148 000 neurons / sq. mm
- Hence, about 15 million / sq. cm

2005-10-20 22

Relative Cortical Areas

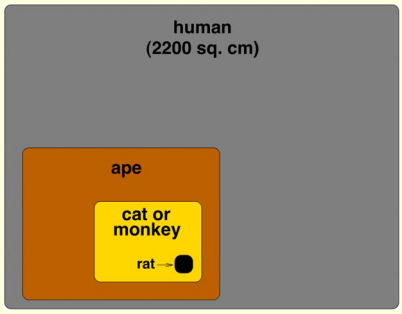

2005-10-20 23

Macaque Visual System

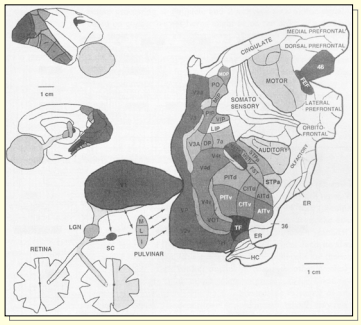

2005-10-20 (fig. from Van Essen & al. 1992) 24

Hierarchy of Macaque Visual Areas

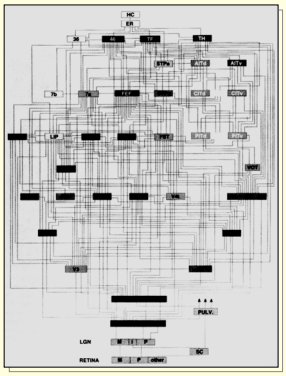

2005-10-20 25

(fig. from Van Essen & al. 1992)

Bat Auditory Cortex

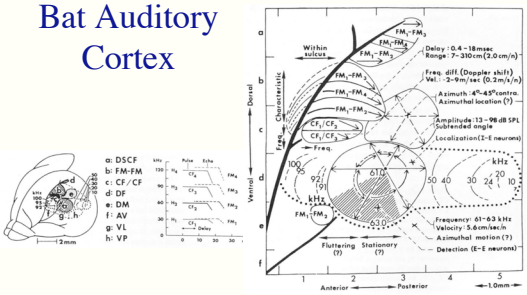

2005-10-20 26

(figs. from Suga, 1985)

Neurocomputing

- Artificial Neural Networks
  - implemented in software on conventional computers
  - are trained, not programmed
  - "second-best way of doing anything"
  - poor match between HW & SW
- Neurocomputers
  - goal: design HW better suited to neurocomputing
  - massively-parallel, low-precision, analog computation
  - electronic? optical? chemical? biological?

2005-10-20 27

Imagistic Representation

- Much information is implicit in an image
- But can be extracted when needed
- Humans have prototype images for each basic category
- Brains use a kind of analog computing for image manipulation

2005-10-20 28

Multiple Intelligences
(Howard Gardner)

- linguistic
- logico-mathematical
- spatial
- musical
- bodily-kinesthetic
- naturalistic
- intrapersonal
- interpersonal
- existential

2005-10-20 29

Artificial Emotions?

- Have been neglected (in cognitive science & AI) due to Cartesian bias
- Importance of "emotional intelligence" now recognized
- Emotions "tag" information with indicators of relevance to us
- Emotions serve important purposes in
  - motivating & directing behavior
  - modulating information processing
- Artificial emotions will be essential for truly autonomous robotics

2005-10-20 30

Propositional Knowledge as Emergent & Approximate

- System may only *appear* to be following rules
  - a spectrum of rule-like behavior
- Recognition of situation can be fuzzy & context-sensitive
- Extraction of relevant elements can be context-sensitive
- May explain subtlety & sensitivity of rule-like behavior in humans & other animals

2005-10-20 31

Natural Computation

- Computation occurring in nature or inspired by computation in nature
- Characteristics:
  - Tolerance to noise, error, faults, damage
  - Generality of response
  - Flexible response to novelty
  - Adaptability
  - Real-time response
  - Optimality is secondary

2005-10-20 32

Being in the World

2005-10-20 33

Importance of Embodied Intelligence

- Traditional (dualist) view: mind is essentially independent of the body
  - in principle, could have an intelligent "brain in a vat"
- Now we understand that much of our knowledge is implicit in the fact that we have a body
- Also, our body teaches us about the world
- Structure of body is foundation for structure of knowledge
- A "disembodied intelligence" is a contradiction in terms?

2005-10-20 34

Structure of Embodied Intelligence

- Representational primitives are skills, not concepts
- Higher-level skills are built on lower-level
- Lowest-level skills are grounded in the body

2005-10-20 35

Embodied & Situated Artificial Intelligence

- Therefore a genuine AI must be:
  - embedded in a body (*embodied*)
  - capable of interacting significantly with its world (*situated*)
- Intelligence develops as consequence of interaction of body with environment, including other agents
- How can we investigate embodied, situated intelligence?

2005-10-20 36

Artificial Life

"Genghis" from Brooks' lab (MIT)

2005-10-20 37

Definition of Artificial Life

- *Artificial Life* is "the study of man-made systems that exhibit behaviors characteristic of natural living systems" (Langton)
- "ALife" includes:
  - synthetic self-reproducing chemical systems, etc.
  - some autonomous robots
  - electronic life forms "living" in a computer's memory

2005-10-20 38

Interactions with Other Agents

- Being situated includes interactions with other agents
- Cooperative interactions:
  - robots with robots
  - robots with humans
- Competitive interactions:
  - robots against robots
  - robots against humans
  - robots against animals

"Robonaut"

2005-10-20 39

"Mind Reading" and Other Social Skills

- Need to understand other agents' mental states & processes
- Need to communicate (or misrepresent) one's own mental state & processes
- Non-verbal communication: gesture, eye contact, gaze
- Imitation as basis of learning & social understanding

2005-10-20 40

Shared Cooperative Activities

- Commitment to joint activity & mutual support
- Joint intention theory
- Simulation theory
- Ability to take perspective of other agent

2005-10-20 41

Leonardo

- Cynthia Breazeal's Lab, MIT
- "Sociable Robots" Project
- Vehicle for exploring socially guided learning & cooperative activity

(video < Breazeal's Lab)

2005-10-20 42

Socially Guided Learning

- Leo is taught to "turn on all the lights"
- Leo generalizes to new situation
- Leo displays commitment to joint activity in spite of incorrect action

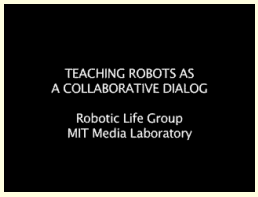

(video < Breazeal's Lab)

2005-10-20 43

Truly Autonomous Robots

- The ultimate test of intelligence is to be able to function effectively in a complex natural environment
- Natural environments do not come parsed into context-free categories
- Natural environments are characterized by complexity, unpredictability, uncertainty, openness, & genuine novelty
- There is also a practical need for *truly autonomous robots*

2005-10-20 44

Starting Small

- In science, it's generally considered prudent to start by studying the *simplest* instances of a phenomenon
- Perhaps it is premature to attempt human-scale embodied artificial intelligence
- It may be more fruitful to try to understand the simplest instances of embodied intelligent behavior

2005-10-20 45

Collective Intelligence

2005-10-20 46

Mound Building by *Macrotermes* Termites

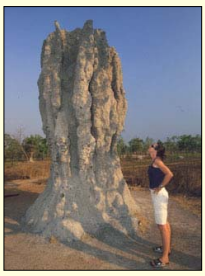

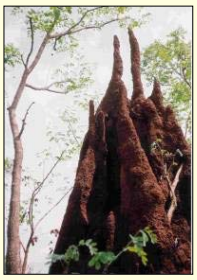

2005-10-20 47

Structure of Mound

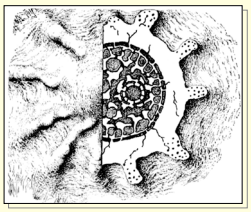

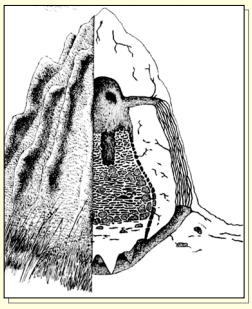

2005-10-20 figs. from Lüscher (1961) 48

Fungus Cultivator Ants

- "Cultivate" fungi underground
- Construct "gardens"
- Plant spores
- Weed out competing fungi
- Fertilize with compost from chewed leaves

2005-10-20 49

Harvester Ants

- Find shortest path to food
- Prioritize food sources based on distance & ease of access
- Adjust number involved in foraging based on:
  - colony size
  - amount of food stored
  - amount of food in area
  - presence of other colonies
  - etc.
- Collective decision making can be as accurate and effective as some individual vertebrate animals

2005-10-20 50

Slime Mold
(*Dictyostelium discoideum*)

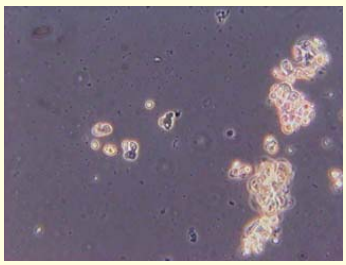

2005-10-20 51

Complete Life Cycle

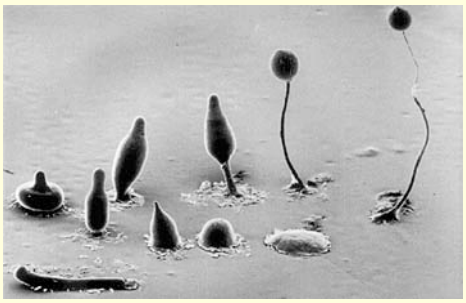

2005-10-20 52

Migration of Slug

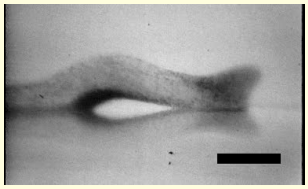

- 1 frame = 20 sec., scale bar = 100 µm ⟷

2005-10-20 (video < Zool. Inst., Univ. München) 53

Early Culmination

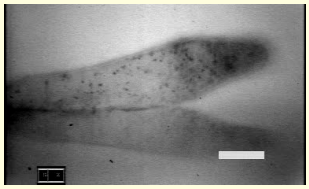

- During early culmination all cell in prestalk rotate
- 1 frame = 25 sec., scale bar = 50 µm

2005-10-20 (video < Zool. Inst., Univ. München) 54

Emergence

- The appearance of *macroscopic* patterns, properties, or behaviors
- that are not simply the "sum" of the *microscopic* properties or behaviors of the components
  - non-linear but not chaotic
- Macroscopic order often described by fewer & different variables than microscopic order
  - e.g. ant trails vs. individual ants
  - *order parameters*

2005-10-20 55

Self-Organization

- Order may be imposed from outside a system
  - to understand, look at the external source of organization
- In *self-organization*, the order emerges from the system itself
  - must look at interactions within system
- In biological systems, the emergent order often has some adaptive purpose
  - e.g., efficient operation of ant colony

2005-10-20 56

Some Principles of Emergence & Self-Organization

- Many non-linearly interacting agents
- Microdecisions lead to macrobehavior
- Circular causality (macro / micro feedback)
- Distributed information storage & processing
- Cooperation + competition
- Diversity
- Amplification of random fluctuations

2005-10-20 57

Adaptation in Artificial Life

- Learning (individual & collective)
- Self-repair (individual & collective)
- Reproduction (individual & collective)
- Artificial evolution

2005-10-20 58

Microrobots

- We don't know enough about human intelligence to reproduce it in a machine,
- but issues of:
  - embodied intelligence
  - autonomous activity
  - social context of intelligence
- may be explored by means of microrobots
- Many potential applications

2005-10-20 59

"Ant" Microrobots Clustering Around "Food"

- "Food" amongst other objects in environment
- First "ant" to encounter food, signals others
- Others cluster at food source
- Brooks' Lab (MIT)

2005-10-20 (video < Brooks' lab, MIT) 60

Nanobots

- How small can we go?
- Viruses & bacteria show how robots could be implemented at micrometer scale
- Genetically engineer:
  - existing organism
  - new organism
- Apply same principles to nonorganic robot

2005-10-20 (video < Hybrid Medical Animation) 61

Computing with Microorganisms

- Bacteria and other microorganisms have have a large amount of "junk DNA"
- Can be genetically engineered to create internal artificial biochemical networks
- GE'd bacteria can cooperate through chemical signals, for:
  - emergent computation
  - microrobotics & nanorobotics

2005-10-20 62

The General-Purpose Bacterial Robot

- An assortment of general genetic circuits
- Ensemble of useful sensors & effectors
- GE to customize operation
- Genetic circuits blocked or enabled by chemical & other means

2005-10-20 63

The Sciences of Complexity

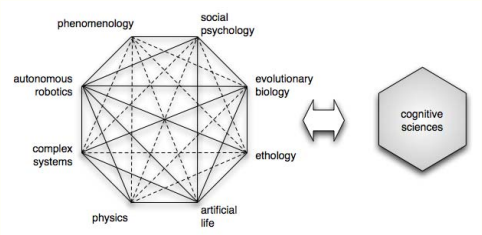

2005-10-20 64

Can artificial life engender real understanding?

- Two senses:
- Can artificial life help us to understand intelligence in humans & other animals?
- Can artificial agents exhibit genuine understanding themselves?

2005-10-20 65

Can ALife help us to understand intelligence?

- Permits embodied, situated testing of theories
- Permits dealing with issues of embodiment & situatedness
- Provides a distinctly different form of "life" for comparison & contrast with ordinary living things

Yes!

2005-10-20 66

Can artificial agents exhibit genuine understanding?

- Symbols are grounded
  - in perceptions, sensorimotor skills, etc.
- Representations are relevant to agent's skillful action in real world
- If they are truly autonomous, then their representations are meaningful *to them*

2005-10-20 67

Thank you!

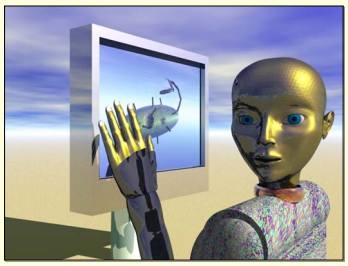

2005-10-20 68