

Artificial Intelligence

Chapter 1: Introduction

What Is AI?

The Foundations of Artificial Intelligence

The History of Artificial Intelligence

What is Artificial Intelligence ?

THOUGHT	Systems that think like humans	Systems that think rationally
BEHAVIOUR	Systems that act like humans	Systems that act rationally
	HUMAN	RATIONAL

WHAT IS AI?

Thinking Humanly

- “The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.” (Haugeland, 1985)
- “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)

Thinking Rationally

- “The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985)
- “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)

DEFINITIONS


Acting Humanly

- “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)
- “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)


Acting Rationally

- “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998)
- “AI . . . is concerned with intelligent behaviour in artifacts.” (Nilsson, 1998)

“The science and engineering of making intelligent machines, especially intelligent computer programs”.
-John McCarthy



Acting humanly: The Turing Test approach



Systems that act like humans

The Turing Test approach

- a human questioner cannot tell if
 - there is a computer or a human answering his question, via teletype (remote communication)
- The computer must behave intelligently

Intelligent behavior

- to achieve human-level performance in all cognitive tasks

Systems that act like humans

These cognitive tasks include:

- Natural language processing
 - for communication with human
- Knowledge representation
 - to store information effectively & efficiently
- Automated reasoning
 - to retrieve & answer questions using the stored information
- Machine learning
 - to adapt to new circumstances

The total Turing Test

Includes two more issues:

- Computer vision
 - to perceive objects (seeing)
- Robotics
 - to move objects (acting)

Systems that act like humans

- The Turing Test approach
 - a human questioner cannot tell if
 - there is a computer or a human answering his question, via teletype (remote communication)
 - The computer must behave intelligently
- Intelligent behavior
 - to achieve human-level performance in all cognitive tasks

Systems that act like humans

- These cognitive tasks include:
 - Natural language processing
 - for communication with human
 - Knowledge representation
 - to store information effectively & efficiently
 - Automated reasoning
 - to retrieve & answer questions using the stored information
 - Machine learning
 - to adapt to new circumstances

The total Turing Test

- Includes two more issues:
 - Computer vision
 - to perceive objects (seeing)
 - Robotics
 - to move objects (acting)

What is Artificial Intelligence ?

THOUGHT	Systems that think like humans	Systems that think rationally
BEHAVIOUR	Systems that act like humans	Systems that act rationally
	HUMAN	RATIONAL

Systems that think like humans: cognitive modeling

- Humans as observed from 'inside'
- How do we know how humans think?
 - Introspection vs. psychological experiments
- Cognitive Science
- “The exciting new effort to make computers think ... machines with minds in the full and literal sense” (Haugeland)
- “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...” (Bellman)

What is Artificial Intelligence ?

	THOUGHT		
	BEHAVIOUR		
	Systems that think like humans	Systems that think rationally	HUMAN
	Systems that act like humans	Systems that act rationally	RATIONAL

Systems that think 'rationally' "laws of thought"

- Humans are not always 'rational'
- Rational - defined in terms of logic?
- Logic can't express everything (e.g. uncertainty)
- Logical approach is often not feasible in terms of computation time (needs 'guidance')
- "The study of mental facilities through the use of computational models" (Charniak and McDermott)
- "The study of the computations that make it possible to perceive, reason, and act" (Winston)

What is Artificial Intelligence ?

	THOUGHT		
	BEHAVIOUR		
		Systems that think like humans	Systems that think rationally
		Systems that act like humans	Systems that act rationally
		HUMAN	RATIONAL

Systems that act rationally: “Rational agent”

- **Rational** behavior: doing the right thing
- **The right thing**: that which is expected to maximize goal achievement, given the available information
- Giving answers to questions is ‘acting’.
- I don't care whether a system:
 - replicates human thought processes
 - makes the same decisions as humans
 - uses purely logical reasoning

Systems that act rationally

- Logic \rightarrow only part of a rational agent, not all of rationality
 - Sometimes logic cannot reason a correct conclusion
 - At that time, some specific (in domain) human knowledge or information is used
- Thus, it covers more generally different situations of problems
 - Compensate the incorrectly reasoned conclusion

Systems that act rationally

- Study AI as rational agent –
- 2 advantages:
 - It is more general than using logic only
 - Because: LOGIC + Domain knowledge
 - It allows extension of the approach with more scientific methodologies

Rational agents

- An agent is an entity that perceives and acts
- Abstractly, an agent is a function from percept histories to actions: $[f: P^* \rightarrow A]$
- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
 - \rightarrow design best program for given machine resource

- Artificial

- Produced by human art or effort, rather than originating naturally.

- Intelligence


is the ability to acquire knowledge and use it"
[Pigford and Baur]

- So AI was defined as:

- AI is the study of ideas that enable computers to be intelligent.
 - AI is the part of computer science concerned with design of computer systems that exhibit human intelligence(From the Concise Oxford Dictionary)

Goals of AI

- To make computers more useful by letting them take over dangerous or tedious tasks from human
- Understand principles of human intelligence



The Foundation of AI

- Philosophy
 - At that time, the study of human intelligence began with no formal expression
 - Initiate the idea of mind as a machine and its internal operations

The Foundation of AI

Mathematics formalizes the three main area of AI: computation, logic, and probability

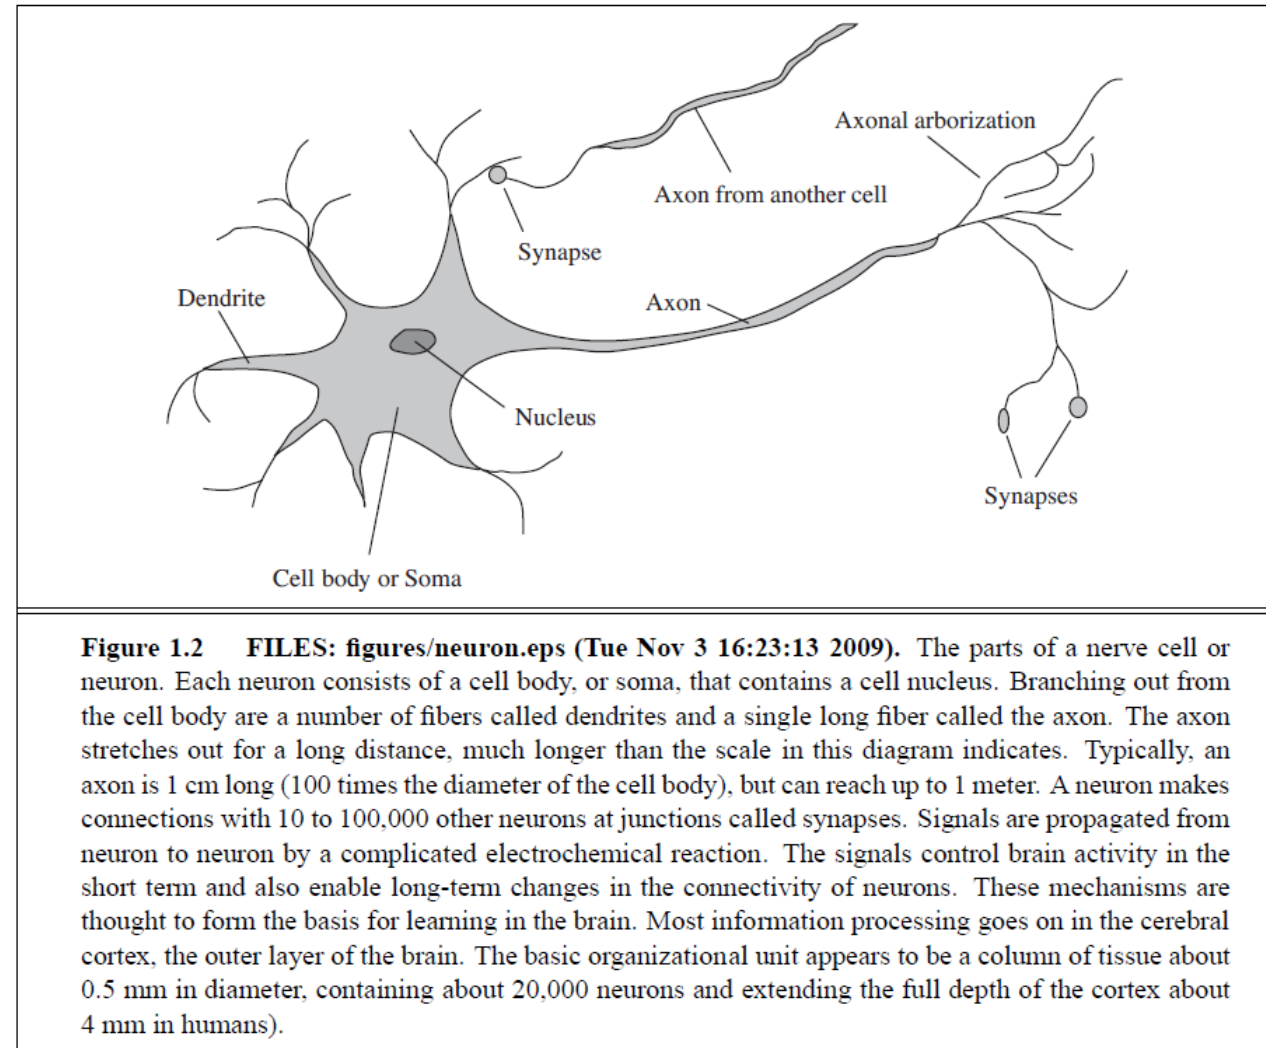
- Computation leads to analysis of the problems that can be computed
 - complexity theory
- Probability contributes the “degree of belief” to handle uncertainty in AI
- Decision theory combines probability theory and utility theory (bias)

The Foundation of AI

Psychology

- How do humans think and act?
- The study of human reasoning and acting
- Provides reasoning models for AI
- Strengthen the ideas
 - humans and other animals can be considered as information processing machines

The Foundation of AI



The Foundation of AI

Computer Engineering

- How to build an efficient computer?
- Provides the artifact that makes AI application possible
- The power of computer makes computation of large and difficult problems more easily
- AI has also contributed its own work to computer science, including: time-sharing, the linked list data type, OOP, etc.

The Foundation of AI

Control theory and Cybernetics

- How can artifacts operate under their own control?
- The artifacts adjust their actions
 - To do better for the environment over time
 - Based on an objective function and feedback from the environment
- Not limited only to linear systems but also other problems
 - as language, vision, and planning, etc.

The Foundation of AI

Linguistics

- For understanding natural languages
 - different approaches has been adopted from the linguistic work
- Formal languages
- Syntactic and semantic analysis
- Knowledge representation

The main topics in AI

Artificial intelligence can be considered under a number of headings:


- Search (includes Game Playing).
- Representing Knowledge and Reasoning with it.
- Planning.
- Learning.
- Natural language processing.
- Expert Systems.
- Interacting with the Environment
(e.g. Vision, Speech recognition, Robotics)

Search

- Search is the fundamental technique of AI.
 - Possible answers, decisions or courses of action are structured into an abstract space, which we then search.
- Search is either "blind" or "informed":
 - blind
 - we move through the space without worrying about what is coming next, but recognising the answer if we see it
 - informed
 - we guess what is ahead, and use that information to decide where to look next.
- We may want to search for the first answer that satisfies our goal, or we may want to keep searching until we find the best answer.

Knowledge Representation & Reasoning

- The second most important concept in AI
- If we are going to act rationally in our environment, then we must have some way of describing that environment and drawing inferences from that representation.
 - how do we describe what we know about the world ?
 - how do we describe it concisely ?
 - how do we describe it so that we can get hold of the right piece of knowledge when we need it ?
 - how do we generate new pieces of knowledge ?
 - how do we deal with uncertain knowledge ?



- Declarative knowledge deals with factoid questions (what is the capital of India? Etc.)
- Procedural knowledge deals with “How”
- Procedural knowledge can be embedded in declarative knowledge

Planning

Given a set of goals, construct a sequence of actions that achieves those goals:

- often very large search space
- but most parts of the world are independent of most other parts
- often start with goals and connect them to actions
- no necessary connection between order of planning and order of execution
- what happens if the world changes as we execute the plan and/or our actions don't produce the expected results?

Learning

If a system is going to act truly appropriately, then it must be able to change its actions in the light of experience:

- how do we generate new facts from old ?
- how do we generate new concepts ?
- how do we learn to distinguish different situations in new environments ?

Interacting with the Environment

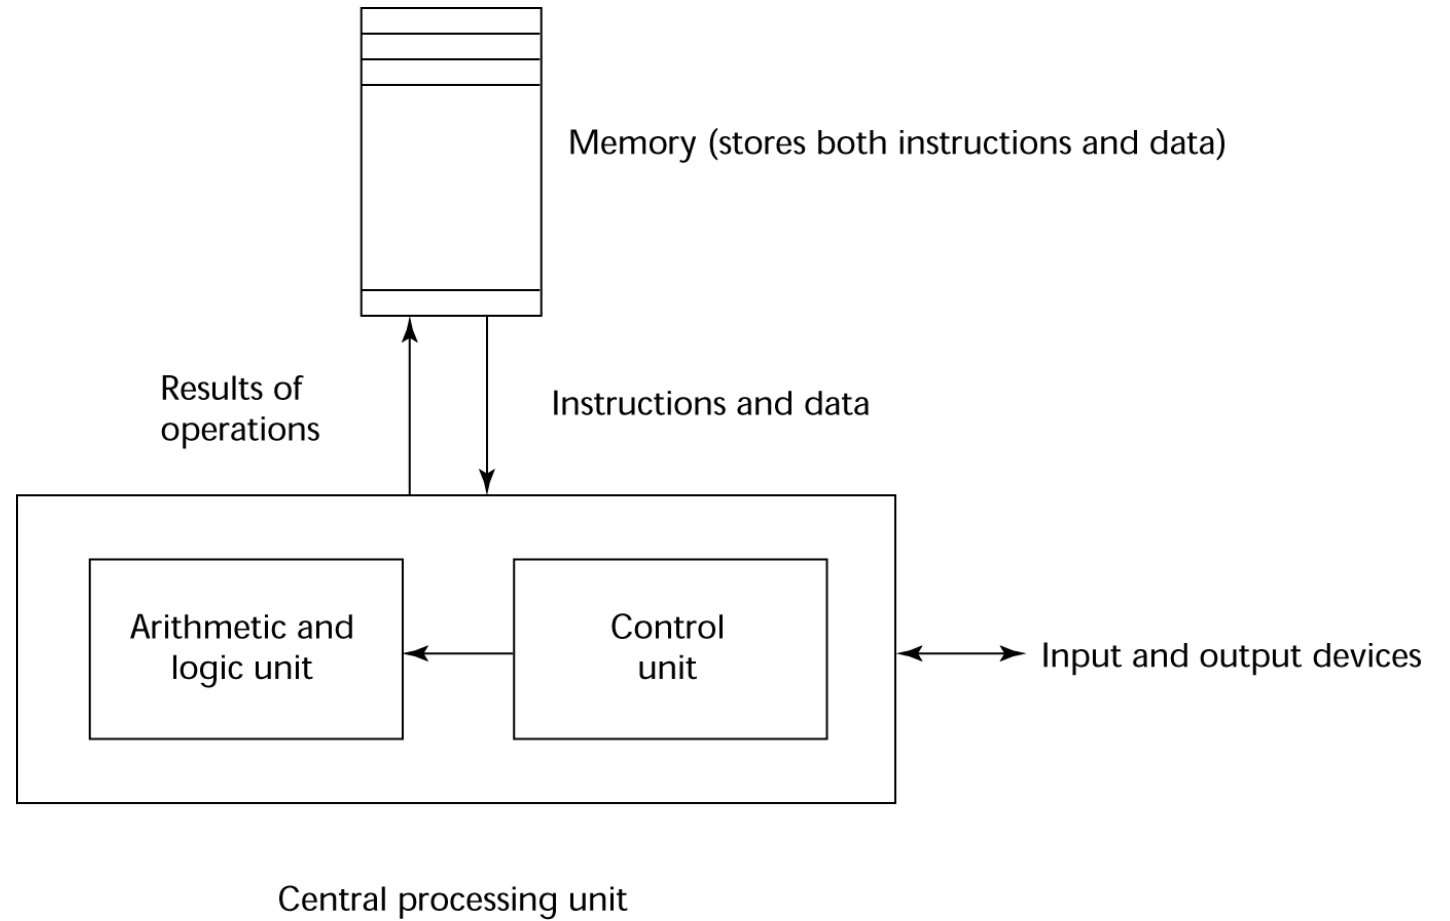
- In order to enable intelligent behaviour, we will have to interact with our environment.
- Properly intelligent systems may be expected to:
 - accept sensory input
 - vision, sound, ...
 - interact with humans
 - understand language, recognise speech,
generate text, speech and graphics, ...
 - modify the environment
 - robotics

History of AI

AI has a long history

- Ancient Greece
 - Aristotle
- Historical Figures Contributed
 - Ramon Lull
 - Al Khowarazmi
 - Leonardo da Vinci
 - David Hume
 - George Boole
 - Charles Babbage
 - John von Neuman
- As old as electronic computers themselves (c1940)

The 'von Neuman' Architecture



History of AI

- Origins
 - The Dartmouth conference: 1956
 - John McCarthy (Stanford)
 - Marvin Minsky (MIT)
 - Herbert Simon (CMU)
 - Allen Newell (CMU)
 - Arthur Samuel (IBM)
- The Turing Test (1950)
- “Machines who Think”
 - By Pamela McCorckindale

Periods in AI

- Early period - 1950's & 60's
 - Game playing
 - brute force (calculate your way out)
 - Theorem proving
 - symbol manipulation
 - Biological models
 - neural nets
- Symbolic application period - 70's
 - Early expert systems, use of knowledge
- Commercial period - 80's
 - boom in knowledge/ rule bases

Periods in AI cont'd

- period - 90's and New Millennium
- Real-world applications, modelling, better evidence, use of theory,
- Topics: data mining, formal models, GA's, fuzzy logic, agents, neural nets, autonomous systems
- Applications
 - visual recognition of traffic
 - medical diagnosis
 - directory enquiries
 - power plant control
 - automatic cars

Fashions in AI


- Progress goes in stages, following funding booms and crises: Some examples:
 1. Machine translation of languages
 - 1950's to 1966 - Syntactic translators
 - 1966 - all US funding cancelled
 - 1980 - commercial translators available
 2. Neural Networks
 - 1943 - first AI work by McCulloch & Pitts
 - 1950's & 60's - Minsky's book on "Perceptrons" stops nearly all work on nets
 - 1986 - rediscovery of solutions leads to massive growth in neural nets research
- The UK had its own funding freeze in 1973 when the Lighthill report reduced AI work severely -Lesson: Don't claim too much for your discipline!!!!
- Look for similar stop/go effects in fields like genetic algorithms and evolutionary computing. This is a very active modern area dating back to the work of Friedberg in 1958.

Symbolic and Sub-symbolic AI

- Symbolic AI is concerned with describing and manipulating our knowledge of the world as explicit symbols, where these symbols have clear relationships to entities in the real world.
- Sub-symbolic AI (e.g. neural-nets) is more concerned with obtaining the correct response to an input stimulus without 'looking inside the box' to see if parts of the mechanism can be associated with discrete real world objects.
- This course is concerned with symbolic AI.


AI Applications

- Autonomous Planning & Scheduling:
 - Autonomous rovers.





AI Applications

- Autonomous Planning & Scheduling:
 - Telescope scheduling





AI Applications

- Autonomous Planning & Scheduling:
 - Analysis of data:






AI Applications

- Medicine:
 - Image guided surgery



AI Applications

- Medicine:
 - Image analysis and enhancement



AI Applications

- Transportation:
 - Autonomous vehicle control:



AI Applications

- Transportation:
 - Pedestrian detection:




AI Applications

Games:



AI Applications

Games:



AI Applications

- Robotic toys:



AI Applications

Other application areas:

- Bioinformatics:
 - Gene expression data analysis
 - Prediction of protein structure
- Text classification, document sorting:
 - Web pages, e-mails
 - Articles in the news
- Video, image classification
- Music composition, picture drawing
- Natural Language Processing .
- Perception.