Notes for Meeting 1. Introduction and Background

What is artificial intelligence?

- This course is about the field of artificial intelligence (AI). What does that phrase mean?
- Artificial intelligence is the computational study of structures and processes that support intelligent behavior.
- The name accurately reflects the fact that most work in the area involves the creation of computational artifacts.
- But other names - complex information processing, computational intelligence, cognitive systems - have been proposed.

What does artificial intelligence study?

- Any scientific field must commit to some phenomena it aims to understand and explain.
- Any engineering field must commit to some capabilities that it aims to produce and support.
- Artificial intelligence focuses on phenomena and abilities that, informally, we recognize as intelligent in humans.
- However, this includes many (apparently) distinct functionalities.

What are some examples of intelligent behavior?

- Can you carry out transitive reasoning?
  [Joe shorter Sam, Bob taller Tom, Joe taller Bob, who is tallest?]
- Can you prove that two triangles are similar?
- Can you draw plausible inferences from a sentence like:
  "John needed money. He got his gun. He drove to the bank."
  Would you make different inferences from a sentence like:
  "John needed money. He got his gun. He drove to the pawn shop."
- Can you compute the sum of 647 and 359?
- Can you give directions to a familiar place?
- Can you follow such directions to reach the place? Can you adapt them if a road is closed?
- Can you solve crossword or sudoku puzzles?
- Can you carry out a dialogue to support joint activity?

Some questions about intelligence

- Can animals handle any of the tasks we just discussed?
- Can a young child handle any of the tasks just discussed?
- Are some aspects of intelligence learned?
- Could intelligence be what makes us distinctively human?
- Should AI focus on these distinctive capabilities?

Why is it important to study intelligence?

- Basic science
  - Understanding the nature of mind is a grand challenge equivalent to understanding the universe, matter, life, and society.
  - Computational models of mental abilities can give insights into human psychology.
- Engineering
  - AI can support more effective educational processes
  - AI can automate or assist in handling complex tasks
  - AI can provide more interactive entertainment

A brief history of artificial intelligence

- AI officially launched at 1956 Dartmouth Conference, which recently celebrated its 50th anniversary.
- That year also saw the first AI system and the first programming language for symbolic processing (the same year as Fortran).
- AI has been so successful that its influence is seldom recognized; its ideas underlie:
  - spreadsheets and tax software
  - digital road maps and route finders
  - data mining and recommender systems
  - many other common computer applications
- However, the past 20 years has seen AI diverge from some themes that guided early research, and that we will emphasize in class.

Artificial intelligence and high-level cognition

- Early AI research revolved around the study of high-level cognition.
- When we say that humans exhibit intelligence, we do not mean they can perceive objects, recognize concepts, or execute motor skills.
- Rather, we mean they can carry out multi-step reasoning, solve novel problems, and use natural language to interact with others.
- Recent AI work has moved away from these challenges or, at best, focused on greatly reduced versions of them.
- This shift has produced short-term gains in narrow areas, but it has not brought us closer to understanding the mind.

Symbolic representations and processes

- The central insight that launched AI was that computers are not just number crunchers, but rather general symbol manipulators.
- Intelligence is enabled by, and depends on, symbolic processing.
- This activity requires ways to represent symbolic structures, to interpret them, and to manipulate them.
- Much of AI's five decades of progress has relied on advances in symbolic representations and mechanisms that operate on them.
- Recent excitement in AI about statistical approaches has not made this insight any less valid or important.

The role of knowledge in intelligent behavior

- Another important assumption in early AI was that knowledge plays a key role in high-level cognition.
- Studies of human experts revealed that their abilities came from domain-specific content acquired over time.
- Techniques for representing and utilizing such knowledge led to the subfield of expert systems.
- This produced many knowledge-based systems that business and government groups adopted to replace or augment human experts.
- These systems typically encoded knowledge in terms of modular symbolic structures.

Satisficing and heuristics in complex cognition

- Early AI research recognized that, when confronted with complex problems, a cognitive system cannot find optimal solutions.
- Any resource-limited agent must satisfice, that is, be content with finding acceptable solutions in reasonable time.
- Satisficing results naturally from the use of heuristic methods rather than algorithmic ones.
- Psychological studies have shown that people satisfice, and it is the only practical way for AI to approach complex cognitive tasks.
- The field's increasing focus on finding optimal solutions in restricted settings does not counter these facts.

Artificial intelligence and cognitive psychology

- Much initial AI work was motivated by, and incorporated insights from, studies of human cognition.
- Examining high-level cognition in humans suggests both challenge tasks for AI and promising methods for machine cognition.
- Moreover, AI systems can serve as computational models of human cognitive processing.
- Most modern AI researchers have abandoned this connection, as have most cognitive psychologists.
- But the current divide between AI and psychology does not lessen the potential benefits of further interaction.

Complete cognitive systems

- Although early AI focused on components of intelligence, its aim was to develop complete cognitive systems.
- Many initial efforts focused on general mechanisms designed to support such intelligent agents.
- Research on cognitive architectures began to explore ways to combine these components into unified theories of cognition.
- However, over the past 20 years, AI has become increasingly fragmented, with most studying components in isolation.
- This piecemeal emphasis, even when leading to commercial success, does not make AI's original vision any less important.

High-level languages for intelligent systems

- Early AI research also favored developing high-level programming languages that could speed creation of intelligent systems.
- Some encodings, like list structures and rules, arise repeatedly in many aspects of cognition.
- Certain mechanisms, like relational pattern matching and heuristic search, also occur in many cognitive tasks.
- Embedding these ideas in high-level programming constructs supports compact code, rapid programming, and effective debugging.
- The current emphasis on Java and C, and their associated libraries, does not make these insights any less viable.

Summary of main themes

- Early research on artificial intelligence often favored:
  1. focusing on challenging, high-level cognitive tasks
  2. incorporating symbolic representations and processes
  3. developing knowledge-based systems with domain expertise
  4. using heuristic methods that satisfice on complex tasks
  5. drawing on results from psychology for tasks and mechanisms
  6. attempting to construct complete cognitive systems
  7. using high-level programming languages to support development
- Recent AI work has often diverged from these themes, but they are still valid and we will emphasize them in this course.

Metaphors for research in AI and cognitive systems

- Artificial intelligence researchers draw on a number of distinct metaphors for mental processing:
  - Thinking as logical reasoning
  - Thinking as search through a problem space
  - Thinking as retrieval from memory
- Each metaphor suggests different ways of understanding the mind, emphasizes different tasks, and has different communities.
- These metaphors are not mutually exclusive and can be combined.
- All share the view that cognition involves symbolic processing, and they appear in both AI and cognitive psychology.

The mind and the brain

- Many identify the mind with the brain, but they are not the same.
- We can specify theories of the mind that are independent of the specific hardware or wetware on which they operate.
- A given program can run on many different types of computers; the same holds for cognitive systems.
- Chemistry predated quantum physics by 100 years, and the former still offers a useful level of description.
- This does not mean that brain science cannot suggest insights about the mind, but they are not required.

Emphases of the course

- Although the course focuses on integrated intelligent systems, we must still examine their component mechanisms.
- We will start with processes for inference, execution, and problem solving, then consider their use in high-level abilities.
- Exercises will embed these ideas in the context of a cognitive architecture for complex cognition.
- The course has a historical flavor, so many readings will involve classic papers and survey articles.
- Many early results on AI and cognitive systems are still valid and more important than recent developments.
- Exercises and project will emphasize the creation of intelligent agents that operate in virtual environments.

Details of course Operation

- Assignments include reading papers, preparing comments or questions, and discussing them; this is not a traditional lecture class.
- One simple Lisp exercise, six programming exercises using the Icarus language, Prolog, or Lisp, and one class project.
- Grading policy
  - Each student has a total of four free late days.
  - Once exhausted, 20 percent off per (fractional) day late.
- Auditors are welcome but should limit their class discussion.
- Offices and office hours
  - Pat Langley, BYENG 446
    - Mondays 6:15PM-7:15PM or by appointment
  - Glen Hunt, BYENG 443 AB

Assignments for meeting 2
Frameworks and metaphors for intelligent systems

Read the article:

- Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. Communications of the ACM, 19, 113-126.
- This paper introduces two key ideas:
  - The notion of physical symbol systems
  - The idea of heuristic search through a problem space
- Newell and Simon claim both are key features of intelligent systems.