COMPSCI 111 / 111G

Mastering Cyberspace:
An Introduction to Practical Computing

Artificial Intelligence

Outline of Lecture

- What is Artificial Intelligence?
  - Definition and Aims
  - AI in Fiction, Fact, and History
  - Common Themes
- Metaphors for Artificial Intelligence
  - Complex Reasoning
  - Search Through a Maze
  - Retrieving Knowledge from Memory
- Integrated Intelligent Systems
- Relation to Other Fields / Summary

Introductory Remarks

What is Artificial Intelligence?

Artificial intelligence is the *computational study of structures and processes that support intelligent behavior*.

The name reflects a focus on creating computational *artifacts*.

- Early researchers hoped to create systems with broad, general skills for reasoning, problem solving, and language use.
- In recent years, most AI practitioners have adopted narrower goals with shorter-term payoffs.

However, both approaches share many assumptions about the principles that underlie intelligent behavior.

Why Study Artificial Intelligence?

Artificial intelligence has multiple objectives. Some concern basic *scientific* questions:

- Understanding the mind is a grand challenge analogous to understanding the universe, matter, life, and society;
- Computer models of mental abilities can give us insights into *human* cognition.

Still other AI objectives revolve around *engineering* concerns:

- Tools for automating / assisting in complex cognitive tasks;
- More effective computer-based education and entertainment.

Scientific and engineering aims are not mutually exclusive, as progress in one arena can aid the other.

What is Intelligence?

When we say that humans are *intelligent*, we mean they exhibit certain high-level cognitive abilities, including:

- Carrying out complex reasoning
  - E.g., solving physics problems, proving mathematical theorems
- Drawing plausible inferences
  - E.g., diagnosing automobile faults, solving murder cases
- Using natural language
  - E.g., reading stories, carrying out extended conversations
- Solving novel, complex problems
  - E.g., completing puzzles, generating plans, designing artifacts

We do *not* mean that people can recognize familiar objects or execute motor skills, abilities they share with dogs and cats.

Degrees of Intelligence

However, it is misguided to ask if a person, animal, or program *is* intelligent; there are *degrees* of intelligence:

- Dogs and cats are more intelligent than rats or mice;
- Adult humans are more intelligent than infant humans;
- Some adult humans seem to be more intelligent than others;
- Some computer programs are more intelligent than others.

Also, there appear to be different *facets* of intelligence, as many IQ tests attempt to measure.

We should ask not whether a computer program *is* intelligent, but about its *level* of intelligence on given dimensions.

Artificial Intelligence in Fiction

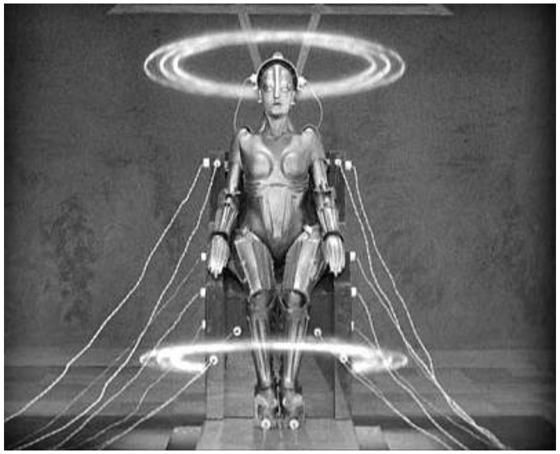

Maria (1927)

Robby (1956)

HAL (1968)

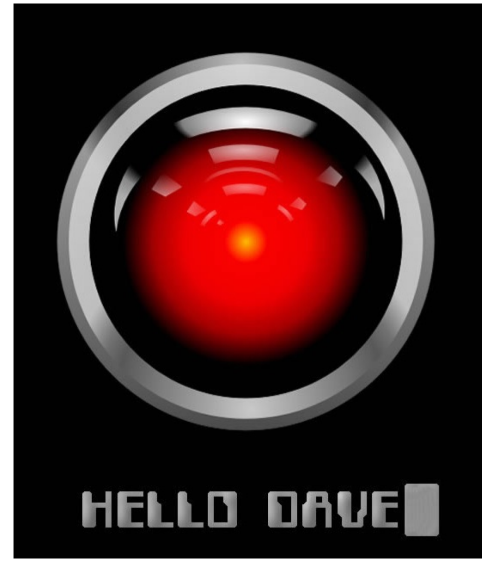

Data (1987)

R2D2 (1977)

Sonny (2004)

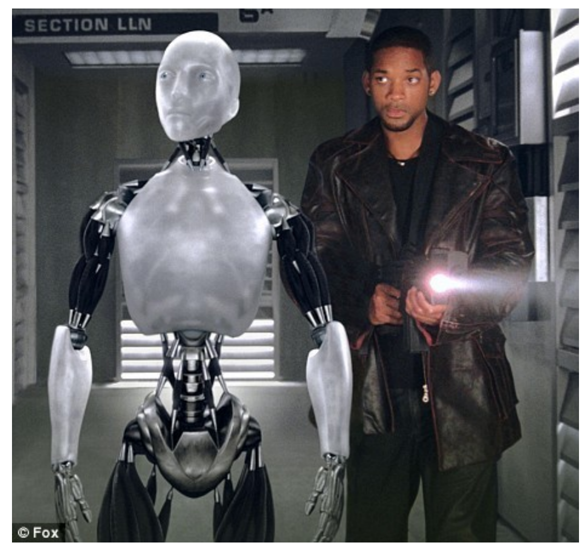

Artificial Intelligence in Fact

The reality of AI is very different from the image presented in popular fiction.

- The field has made great strides, but we are still far from the intelligent systems in movies.
- AI is a scientific and engineering field focused on *research*, with widespread but narrow *applications*.
- Intelligence is a complex set of phenomena that will take decades or centuries to understand.

We may be able to create Robby or HAL, but it could be some time before that happens.

History of Artificial Intelligence

We can divide the history of AI into a number of periods with different concerns and styles:

- Dartmouth conference / first running AI program (1956)
- Audacious period / general methods (1956–late 1970s)
- Early applications period (late 1970s–late 1980s)
- AI winter / doubts about potential (late 1980s–late 1990s)
- Narrowed research and applications (late 1990s–present)

AI has been diverse and productive during all periods, but these labels reflect trends and attitudes about the field.

Early AI was only loosely connected to computer science, and sometimes seen as a distinct discipline.

Symbol Structures and Processing

The key insight that launched AI was that computers are not just number crunchers, but rather *general symbol manipulators*.

- Intelligence is enabled by, and depends on, symbol processing;
- This requires ways to *represent* symbol structures, to *interpret* such structures, and to *manipulate* them;
- These often take the form of *list structures* that can encode logic and logic-like relations;
- Much of AI's five decades of progress has relied on advances in symbolic notations and mechanisms that operate on them.

Recent excitement about using statistical techniques in AI has not made this insight any less valid or important.

Rule-Based Processing

Many AI systems are implemented in rule-based programming languages that:

- Specify behavior entirely in terms of *if-then rules*;
- Emphasize the *conditional* nature of behavior;
- Often utilize *list structures* and *pattern matching*; and
- Support coding of highly *flexible* behaviors.

Rule-based formalisms have many practical applications, but they come from studies of human thinking and logic.

Two such frameworks – *production systems* and *logic programs* – play important roles in AI research and applications.

Metaphors for Artificial Intelligence

Three Metaphors for Intelligence

Most research on AI adopts one of three distinct metaphors for thinking about the mind:

- Intelligence as multi-step *reasoning*
- Intelligence as *search* through a maze
- Intelligence as retrieval of *knowledge* from memory

None of these metaphors is right or wrong, but each one offers important insights into the nature of mental processing.

Metaphor 1: Reasoning and Intelligence

One metaphor for intelligence emphasizes the human ability to carry out complex *reasoning*, including:

- *Deduction*, which draws logical conclusions from givens;
- *Abduction*, which finds plausible accounts of observations;
- *Analogy*, which maps new situations onto known ones.

AI researchers have developed systems that exhibit each of these cognitive capabilities.

Despite differences in operation, they all operate over structured, relational encodings of content.

Automated Reasoning

Some of the earliest AI research dealt with automated reasoning:

- Newell, Shaw, Simon's (1957) *Logic Theorist*: The first AI system, based on studies of human reasoning, proved theorems in logic;
- Gelernter's (1959) *Geometry Theorem Machine*: Proved geometry theorems with diagrams, some requiring constructions;
- Slagle's (1963) *SAINT*: Solved final exam problems in symbolic integration from a first-year Calculus course.

Automated reasoning methods are now mature enough to appear in software like Maple and to be used in hardware verification.

Propagation of values in Excel is an even more common use of such reasoning methods.

The SHRDLU System

SHRDLU (1970) was an early AI system that answered questions in English and executed complex commands.

Although limited in scope, it reproduced many aspects of human reasoning.

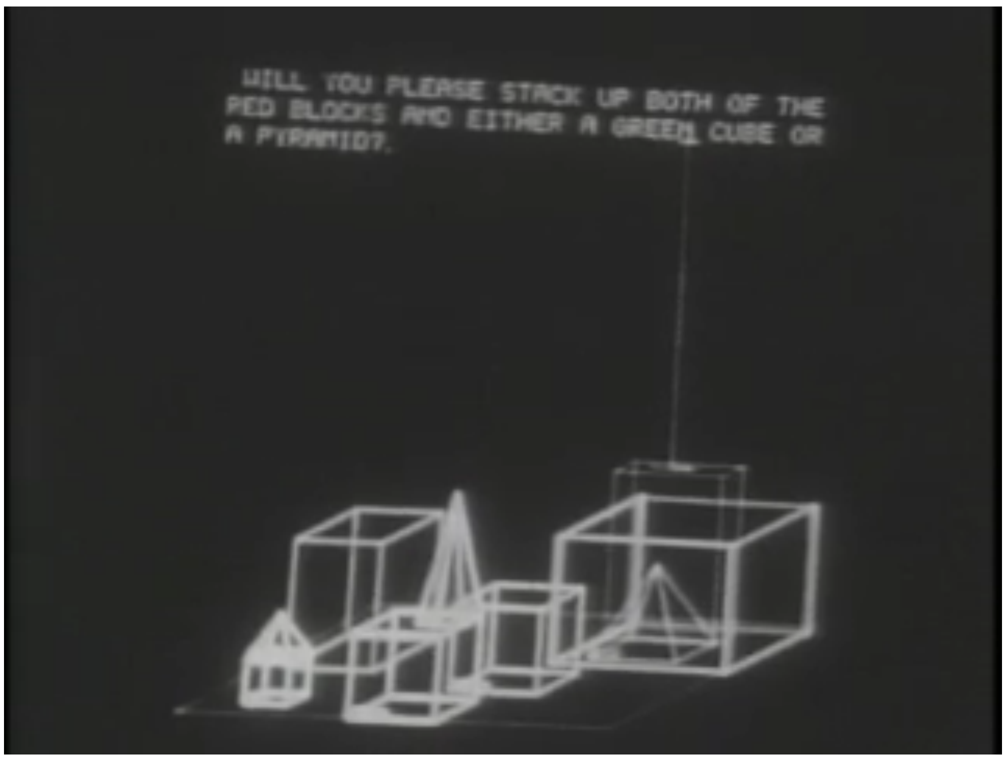

Person: Pick up a big red block.
Computer: *OK.*
Person: Grasp the pyramid.
Computer: *I don't understand which pyramid you mean.*
Person (changing his mind): Find a block which is taller than the one you are holding and put it into the box.
Computer: *OK.*
Person: What does the box contain?
Computer: *The blue pyramid and the blue block.*
Person: What is the pyramid supported by?
Computer: *The box.*
Person: How many blocks are not in the box?
Computer: *Four of them.*
Person: Is at least one of them narrower than the one which I told you to pick up?
Computer: *Yes, the red cube.*

Metaphor 2: Intelligence and Search

A second metaphor for intelligence rests on an analogy with finding one's way through a physical maze.

This view is commonly adopted for *problem solving* tasks.

These require one to find some sequences of steps that solves an unfamiliar problem.

Instances include generating new plans, schedules, and designs.

The search view is not antithetical to the reasoning metaphor, but they emphasize different aspects of intelligence.

Search and the Tower of Hanoi

Consider a puzzle that is known as the *Tower of Hanoi*, which involves three pegs and N disks.

The task is to move all disks from the initial peg to another peg, which involves problem solving and search.

The Notion of a Problem Space

Viewing problem solving as maze traversal leads naturally to the notion of a *problem space* that involves:

- *States* in the problem space (places in the maze), including:
  - The *initial* state (the maze entrance); and
  - The *goal* state (the maze exit or center);
- *Operators* that change states (steps in the maze); and
- *Solution path(s)* for the problem (way through the maze).

These elements of a problem space are central to both AI and psychological analyses of problem solving.

They offer a means to specify a space *implicitly* and to explore it selectively through search.

Problem Space for Tower of Hanoi

Starting from a single state, we can apply operators repeatedly to generate the entire problem space, or we can explore it more selectively.

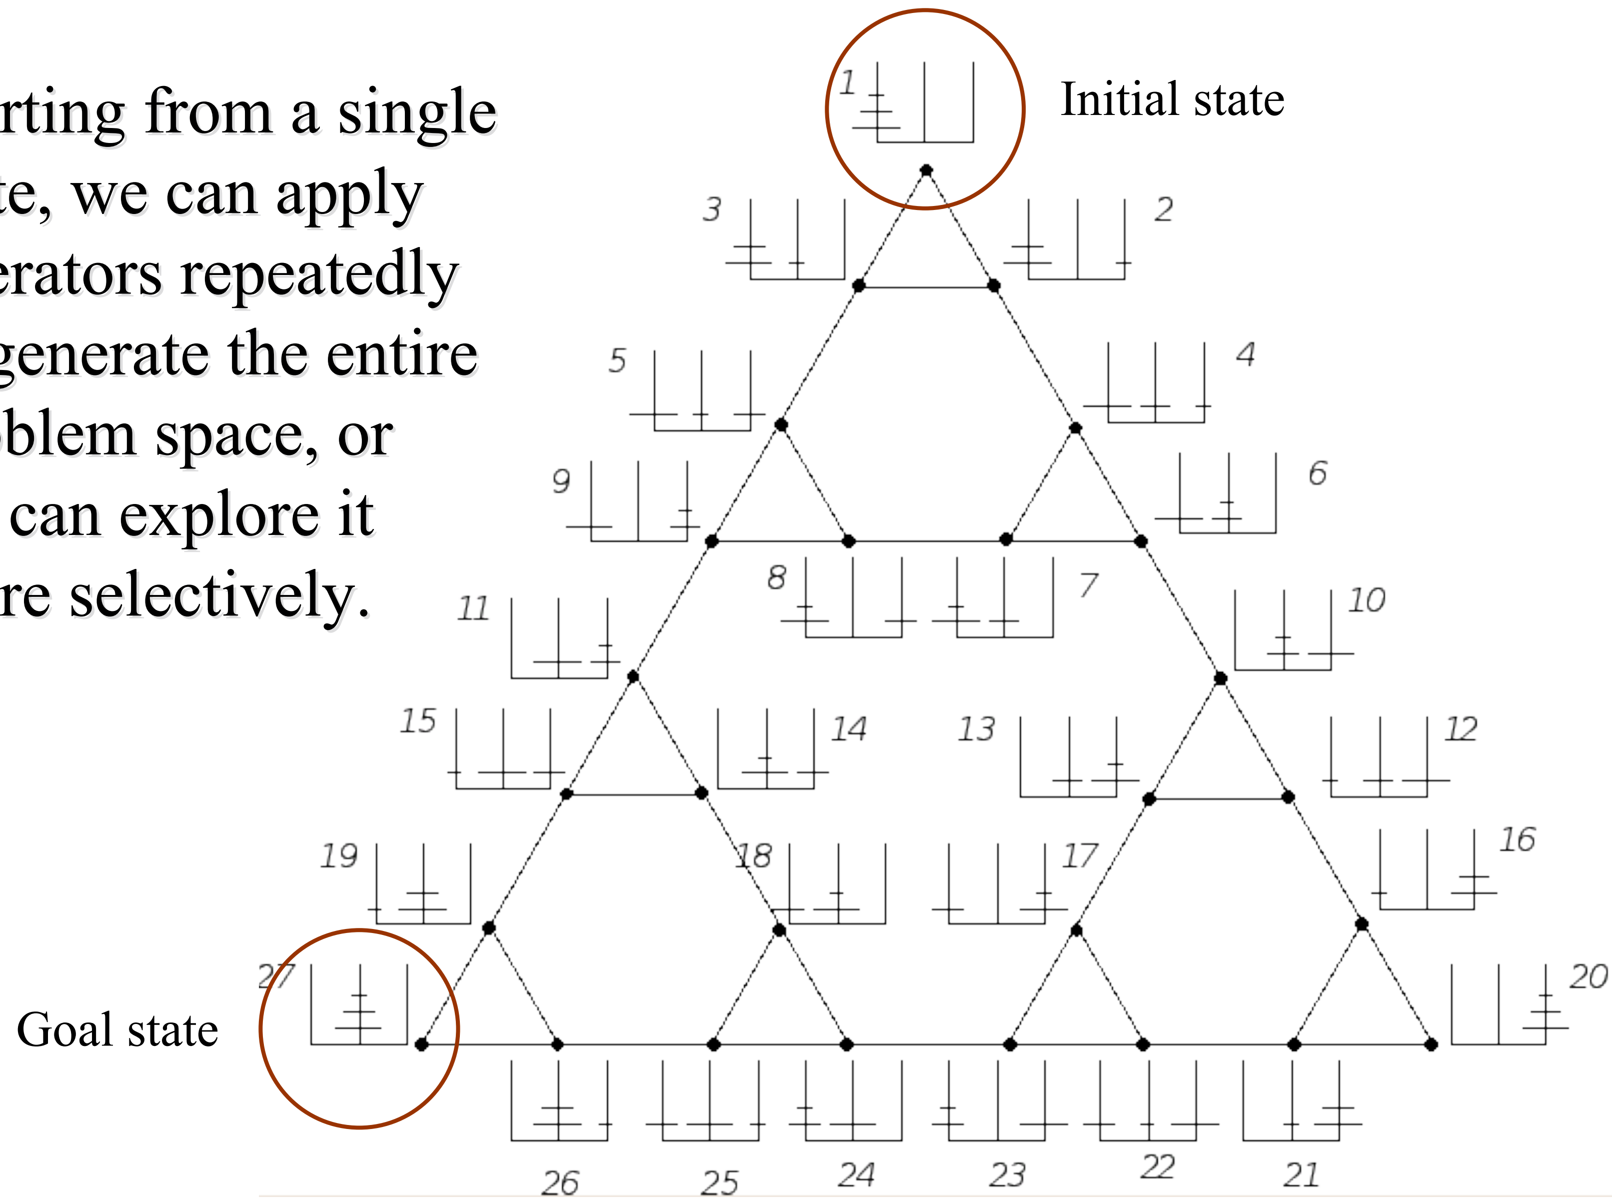

AI Planning Systems

One important use of search mechanisms is to generate a *plan* for achieving some goal.

AI planning methods have become mature enough to support a variety of applied systems, including:

- Automobile navigation systems that generate routes to follow;
- Orbitz and other travel sites, which propose airline itineraries;
- DART, which generated logistical plans for the US military;
- The Hubble space telescope, Mars rovers, and even copiers.

These systems are much less flexible than human planners, but they find solutions that people might overlook.

Game-Playing Programs

Early work on games like chess, one of AI's original challenge problems, led to many insights about representation and search.

- In 1957, Herbert Simon predicted that ‘‘*within ten years, a computer will be the world's chess champion*''.
- Simon's prediction was too optimistic, but only his timing was off.
- In 1997, *Deep Blue*, a hardware-aided chess program that searched deeper than humans, beat Gary Kasparov, the human champion.

Similar results have occurred for checkers, backgammon, and other complex games that humans find challenging.

However, these systems are tuned for specific games and lack the generalized abilities of human players.

Metaphor 3: Knowledge and Intelligence

A third metaphor for intelligence emphasizes people's ability to draw upon large amounts of *knowledge* in memory about:

- The structure and behavior of physical environments;
- Topics from science, engineering, arts, and humanities; and
- Relations and events that arise in social interactions.

AI researchers have developed systems that store, retrieve, and use knowledge in each of these areas.

This paradigm led to the first widespread applications of AI technology in commerce and industry.

Expert Systems

The *expert systems* movement built upon insights about the role of knowledge in human expertise.

This paradigm encodes knowledge as *if-then rules* or as *stored cases* that are matched to draw conclusions.

- Hundreds of such expert systems have been deployed since the 1980s, making or saving substantial money.
- However, they are costly to maintain as conditions change, and their reasoning is often routine and shallow.

Successes of machine learning and data mining are best viewed as the *automated construction* of expert systems.

The TurboTax System

One widely adopted expert system is TurboTax, which helps its users prepare their US tax returns.

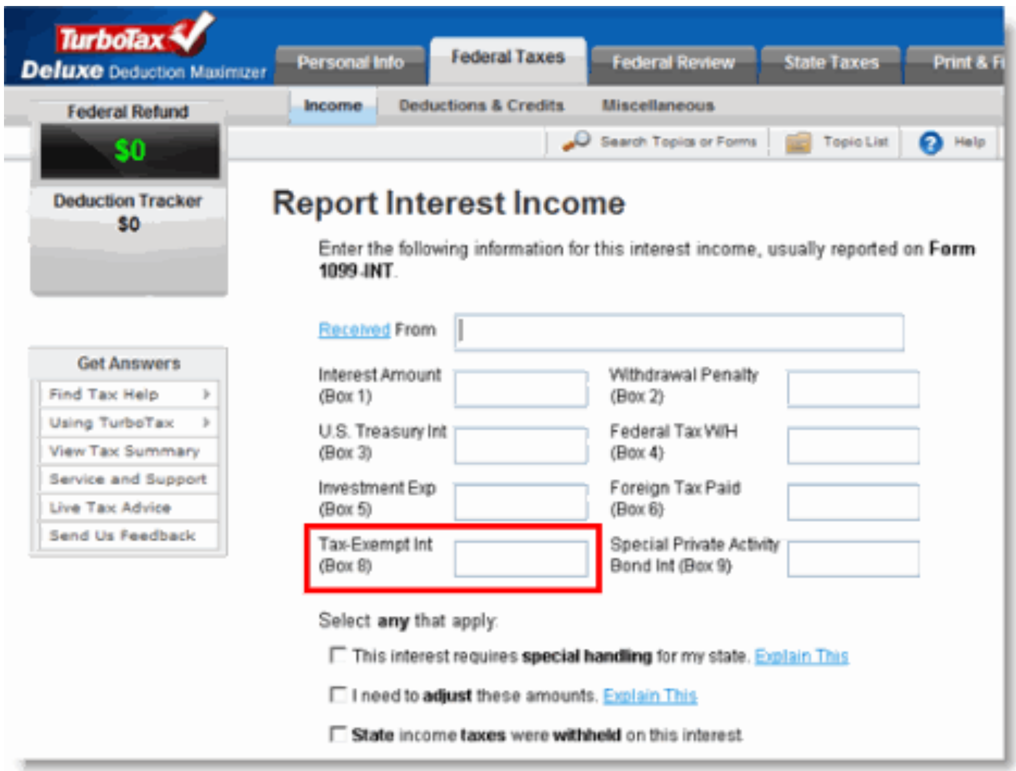

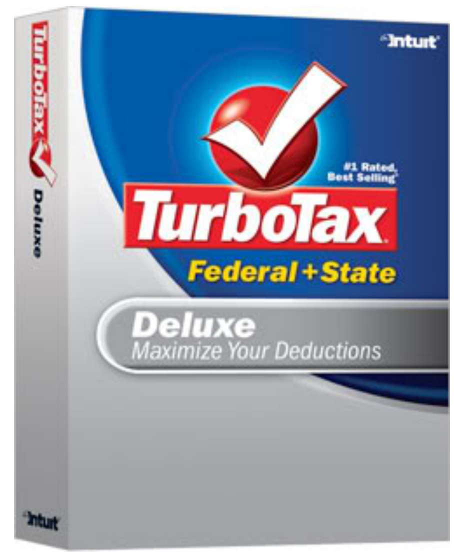

TurboTax has been purchased by millions of customers, reducing their time and effort to file taxes.

Recommender Systems

In the 1990s, increasing availability of Web content, including products on the Web, led to *recommender systems* that:

- Propose, rank, or otherwise present selected items that they predict will interest the user.
- Utilize profiles for users that describe their preferences, which are often learned from explicit or implicit feedback.
- Amazon, Tivo, and many other companies use techniques of this sort to increase sales and customer satisfaction.

Recommender systems have shallow models of users' tastes, but they are used widely and successfully.

Integrated Intelligent Systems

Integrated Intelligent Systems

We have seen that intelligence is a complex phenomenon with many different facets.

Most AI efforts focus on isolated aspects of thinking, but we can also learn much by:

- Developing integrated systems that combine different abilities;
- Clarify interactions and constraints among components; and
- Exhibit abilities closer to those observed in human intelligence.

Most AI programs are *idiot savants* that perform well in narrow areas, but there has been progress on integrated systems.

A Deployed Algebra Tutor (1999–present)

Carnegie Learning's algebra tutor encodes knowledge as rules, infers student models, and provides personalized instruction.

The tutor is based on computer models of human learning.

Hundreds of US schools have now adopted the system.

Studies suggest that it improves student learning by 75%.

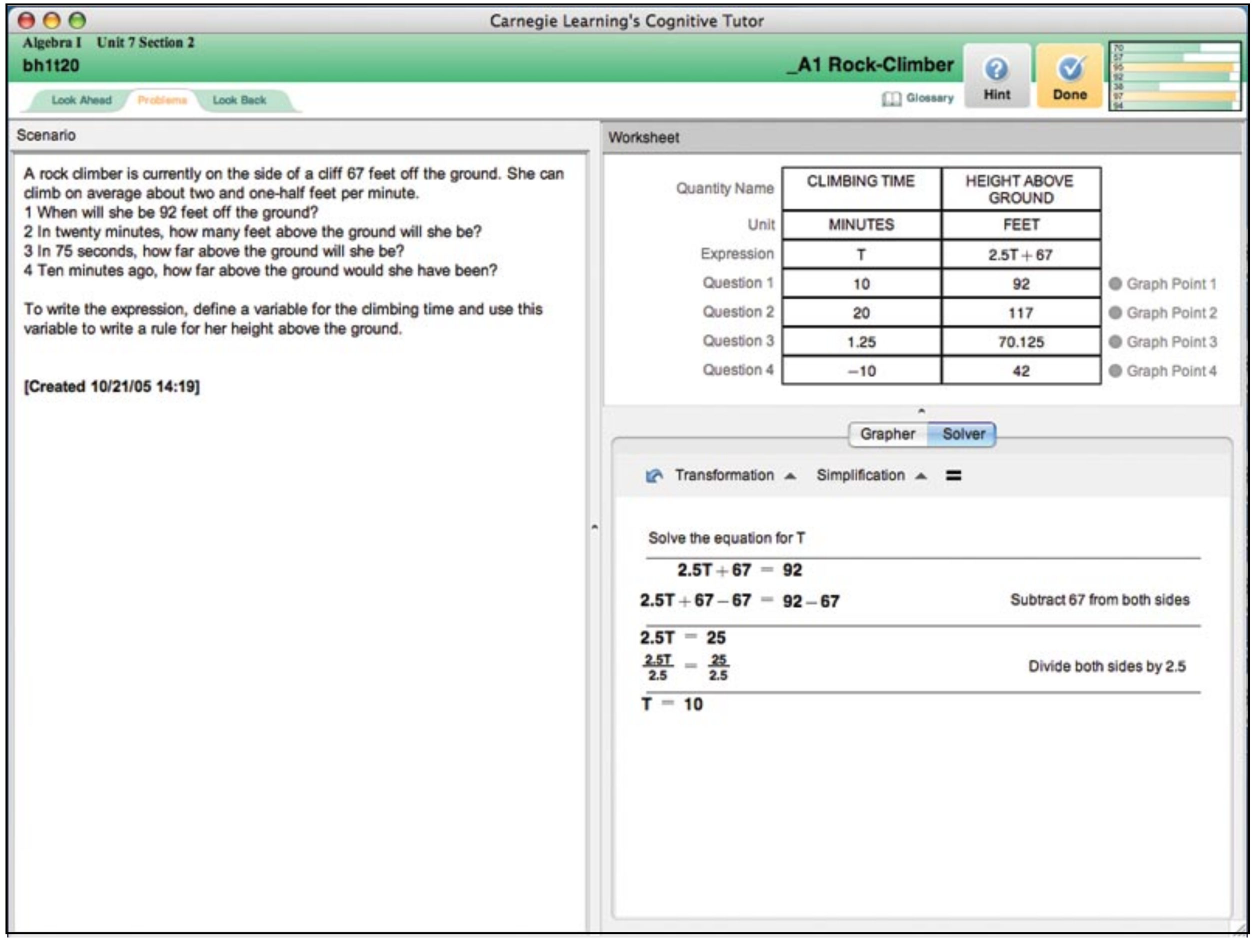

TacAir-Soar (1997)

The TacAir-Soar system reproduces pilot behavior in tactical air combat.

It combines abilities for spatio-temporal reasoning, plan generation / recognition, language, and coordination.

The system flew 722 missions during the STOW-97 simulated training exercise.

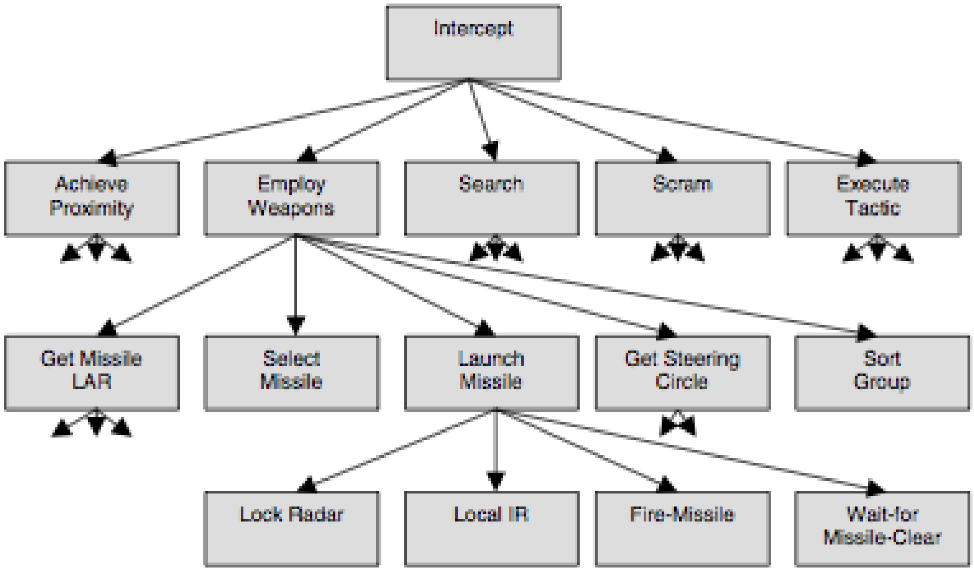

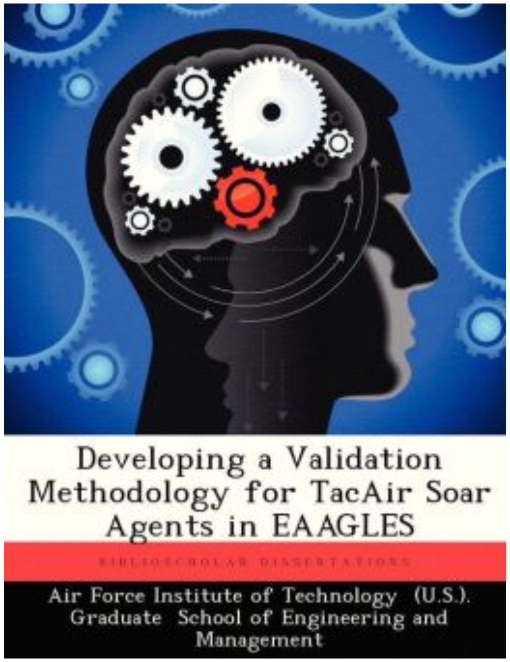

The Watson System (2011)

Watson is an AI system that plays *Jeopardy!*, where contestants are given *answers* and must provide the *question*.

The program constructs its knowledge base by analyzing many online text sources and databases. For a given task, it then:

- Hypothesizes responses with a mixture of deep/shallow methods;
- Filters out poor hypotheses, gathers evidence for remaining ones;
- Merges similar hypotheses, ranks them based on overall support.

Watson combines a variety of language and reasoning techniques to outperform the best human players.

However, it is an expert at the *Jeopardy!* game, not at general language processing.

The Aaron System

Aaron – developed by Harold Cohen continuously from 1973 – composes, draws, and colors novel paintings.

The system produces creative works of art, utilizing knowledge about different facets of artistic composition.

Closing Remarks

Arguments Against AI

Many philosophers have claimed that artificial intelligence is impossible, using arguments like:

- Only carbon-based life forms can exhibit intelligence, or at least consciousness (the *neovitalist* position);
- Only embodied agents that exist in a physical environment can exhibit intelligence (the *situated cognition* position);
- Only systems that operate in a parallel manner, like the human brain, can be intelligent (the *parallel processing* position).

They ignore the many computer programs that already exhibit key facets of intelligence.

But if such arguments are spurious, then where are human-level intelligent systems?

AI's Rate of Progress

Early AI researchers were optimistic about the pace at which their new field would make progress.

- Many people are disappointed that we do not yet have human-level intelligent systems like those seen in movies.
- But do we expect *biology*, a much older discipline, to produce *complete synthetic life forms*?
- If not, then why expect AI, after only 55 years, to have created Robby or HAL by now?

AI has made steady strides, gaining genuine insights into the mind and creating many valuable applications.

We may eventually produce human-level systems, but intelligence is a complex phenomenon that will take time to understand.

AI and Robotics

AI is often linked to robotics, which develops embodied artifacts that operate in the physical world, but:

- Robotics is centrally concerned with sensori-motor behavior, which humans share with other animals;
- Many robotic systems exhibit little or no intelligence in the sense discussed earlier;
- Most applications (vacuum cleaners, walking robots, self-driving cars) focus on low-level control tasks.

There remains potential for interactions between AI and robotics, but they have remained largely separate.

AI and Neuroscience

Many people identify the *mind* with the *brain*, then assume that we cannot understand the former without the latter.

- But theories of the mind can be independent of the hardware or wetware on which they operate.
- Note that the same computer program run on different computer architectures and operating systems.
- Quantum physics may underlie chemistry, yet chemists seldom use it in theory or practice.

Neuroscience may provide insights about the mind, but AI has made great progress without it, and this will continue.

AI and Cognitive Psychology

Much early AI research was inspired by, and drew upon insights from, studies of human thinking.

- Examining high-level cognition in humans suggests both challenge tasks and ideas for machine intelligence.
- Moreover, AI systems can serve as computational models of human cognitive processing.
- This link has produced many of the most powerful ideas in artificial intelligence, including the search metaphor.

Most modern AI researchers have abandoned this connection, but it still has much to offer.

Review of Key Ideas

Artificial intelligence is the *computational study of structures and processes that support intelligent behavior*, a field that:

- Emphasizes symbolic representations and views computers as *general symbol manipulators*;
- Often uses *rule-based languages* that support pattern matching and conditional behavior;
- Draws on metaphors for intelligence that emphasize *reasoning*, *search*, and *knowledge*.

AI has produced many scientific insights about the mind and many applications in business and elsewhere.

But intelligence is a complex phenomenon, and the AI reality remains far from its popular image.