

# ARTIFICIAL INTELLIGENCE

## An Overview

Jeffrey A. Barnett

Albuquerque Retired Lispers

# What is AI?

# What is AI?

A spectrum along a line bounded by

# What is AI?



A spectrum along a line bounded by

- Strong AI — Competence Criteria
  - Model and/or develop theory of artifacts and intelligence that show competence at or near human levels of accomplishment.

# What is AI?



## A spectrum along a line bounded by

- Strong AI — Competence Criteria
  - Model and/or develop theory of artifacts and intelligence that show competence at or near human levels of accomplishment.
- Applied AI — Performance Criteria
  - Develop systems in problem-solving domains that assume intelligent processing of data and stimuli, in packages that exhibit useful performance.

# What is AI?



## A spectrum along a line bounded by

- Strong AI — Competence Criteria
  - Model and/or develop theory of artifacts and intelligence that show competence at or near human levels of accomplishment.
- Applied AI — Performance Criteria
  - Develop systems in problem-solving domains that assume intelligent processing of data and stimuli, in packages that exhibit useful performance.

## Chomsky Says

Chomsky makes a similar distinction between competence and performance grammars.

# Is AI a Science?

# Is AI a Science?

There are two popular definitions of science:



# Is AI a Science?



There are two popular definitions of science:

- 1 Science is an endeavor that **accumulates** and **archives** knowledge with goals of insight and enhanced understanding.
  - AI is a science by this criterion—investigations have added to knowledge in many application areas as well as AI itself.

# Is AI a Science?



There are two popular definitions of science:

- 1 Science is an endeavor that **accumulates** and **archives** knowledge with goals of insight and enhanced understanding.
  - AI is a science by this criterion—investigations have added to knowledge in many application areas as well as AI itself.
- 2 Science is an endeavor that produces testable and falsifiable hypotheses about the world then performs the experiments.
  - One must define **the world** and **testable**.
  - The tests for competence and performance AI are different.

# Is AI a Science?



There are two popular definitions of science:

- 1 Science is an endeavor that **accumulates** and **archives** knowledge with goals of insight and enhanced understanding.
  - AI is a science by this criterion—investigations have added to knowledge in many application areas as well as AI itself.
- 2 Science is an endeavor that produces testable and falsifiable hypotheses about the world then performs the experiments.
  - One must define **the world** and **testable**.
  - The tests for competence and performance AI are different.
  - To be fair, we don't know enough today to define **intelligence** or test for it, thus, psychology, educations, neurology, etc., are in the same boat as AI.

# Associated Science Areas

Psychology	education, cognitive, hermeneutics
Philosophy	epistemology, ontology, logic
Education	learning, organization of knowledge
Operations Research	algorithms, goal formation, search
Mathematics	models, proof theory, symbolization
Linguistics	communications, models, representations
Mechanical Engineering	effectors, adaptation, etc.
Control Theory	feedback, adaptation, approximation
Physiology	sensors, effectors, infrastructure
Neurology	natural models, fusion, organizations
Computer Science	everything

# Tools of the Trade — What Practitioners Do/Use

# Tools of the Trade — What Practitioners Do/Use

- Representations — things to manipulate inside the computer  
— models, knowledge, state, operations

# Tools of the Trade — What Practioners Do/Use

- Representations — things to manipulate inside the computer  
— models, knowledge, state, operations
- Search — methods to find/generate/evaluate alternative hypotheses (global)

# Tools of the Trade — What Practioners Do/Use

- Representations — things to manipulate inside the computer  
— models, knowledge, state, operations
- Search — methods to find/generate/evaluate alternative hypotheses (global)
- Deduction/Induction — deduce/form hypotheses (local)



# Tools of the Trade — What Practioners Do/Use

- Representations — things to manipulate inside the computer — models, knowledge, state, operations
- Search — methods to find/generate/evaluate alternative hypotheses (global)
- Deduction/Induction — deduce/form hypotheses (local)
- Uncertain Reasoning — represent and combine bodies of evidence

# Tools of the Trade — What Practioners Do/Use

- Representations — things to manipulate inside the computer — models, knowledge, state, operations
- Search — methods to find/generate/evaluate alternative hypotheses (global)
- Deduction/Induction — deduce/form hypotheses (local)
- Uncertain Reasoning — represent and combine bodies of evidence
- Simulation — examine states and consequences (learn by doing and observing)

# Tools of the Trade — What Practioners Do/Use

- Representations — things to manipulate inside the computer — models, knowledge, state, operations
- Search — methods to find/generate/evaluate alternative hypotheses (global)
- Deduction/Induction — deduce/form hypotheses (local)
- Uncertain Reasoning — represent and combine bodies of evidence
- Simulation — examine states and consequences (learn by doing and observing)
- Dynamic Languages — specialized programming tools (Lisp, Scheme, Prolog)

# Tools of the Trade — What Practitioners Do/Use

- Representations — things to manipulate inside the computer — models, knowledge, state, operations
- Search — methods to find/generate/evaluate alternative hypotheses (global)
- Deduction/Induction — deduce/form hypotheses (local)
- Uncertain Reasoning — represent and combine bodies of evidence
- Simulation — examine states and consequences (learn by doing and observing)
- Dynamic Languages — specialized programming tools (Lisp, Scheme, Prolog)
- Practical Stuff — sensors, effectors, etc.

# Some Limited Successes

Areas where AI has done well:

- Computer Vision
- Speech Understanding
- Diagnostic Systems
- Robotics
- Game Playing
- Symbol Manipulation
- Expert Systems
- Scheduling/Automation
- Language Understanding
- Chemical Structures

## Status

AI did not solve these problems; Rather it provided new insights, methods, and novel tools that are used in other's homerooms were progress continues.

# Some Limited Successes

Areas where AI has done well:

- Computer Vision
- Speech Understanding
- Diagnostic Systems
- Robotics
- Game Playing
- Symbol Manipulation
- Expert Systems
- Scheduling/Automation
- Language Understanding
- Chemical Structures

## Status

AI did not solve these problems; Rather it provided new insights, methods, and novel tools that are used in other's homerooms were progress continues.

## The Aler's Lament


When we finally know how to do it better, it's no longer AI.

# Representation Makes a Big Difference

# Representation Makes a Big Difference

## A Tough Problem

Remove two diagonally opposed corners from an  $8 \times 8$  board. Can it be completely covered by  $1 \times 2$  and  $2 \times 1$  tiles?






# Representation Makes a Big Difference

## A Tough Problem

Remove two diagonally opposed corners from an  $8 \times 8$  board. Can it be completely covered by  $1 \times 2$  and  $2 \times 1$  tiles?




## Insight

Color it as a checker board—alternating red and black squares.

# Representation Makes a Big Difference

## A Tough Problem

Remove two diagonally opposed corners from an  $8 \times 8$  board. Can it be completely covered by  $1 \times 2$  and  $2 \times 1$  tiles?



## Insight

Color it as a checker board—alternating red and black squares.


## Solution

Each tile covers exactly one red and one black square. The removed corners are both red or both black. So there is no tiling.

# Representation Makes a Big Difference

## A Tough Problem

Remove two diagonally opposed corners from an  $8 \times 8$  board. Can it be completely covered by  $1 \times 2$  and  $2 \times 1$  tiles?



## Insight

Color it as a checker board—alternating red and black squares.

## Solution

Each tile covers exactly one red and one black square. The removed corners are both red or both black. So there is no tiling.

## Conclusion

Insights about proper representations can make huge differences in the ease with which problems in a domain are solved.

# System Goals

## Optimization

We have a natural desire to find best/optimal solutions.

# System Goals

## Optimization

We have a natural desire to find best/optimal solutions.

## Ignorance Runs Rampant

But we don't always have a definition of optimal. What is **your** tradeoff between engine smoothness, mileage, and acceleration when you get a tuneup?

# System Goals

## Optimization

We have a natural desire to find best/optimal solutions.

## Ignorance Runs Rampant

But we don't always have a definition of optimal. What is **your** tradeoff between engine smoothness, mileage, and acceleration when you get a tuneup?

## Satisficing search finds solutions that are good enough

Consider finding sharp needles in a haystack:

# System Goals

## Optimization

We have a natural desire to find best/optimal solutions.

## Ignorance Runs Rampant

But we don't always have a definition of optimal. What is **your** tradeoff between engine smoothness, mileage, and acceleration when you get a tuneup?

## Satisficing search finds solutions that are good enough

Consider finding sharp needles in a haystack:

### Optimization

Find sharpest needle:  
Need statistics on upper sharpness bound in areas of the haystack



# System Goals

## Optimization

We have a natural desire to find best/optimal solutions.

## Ignorance Runs Rampant

But we don't always have a definition of optimal. What is **your** tradeoff between engine smoothness, mileage, and acceleration when you get a tuneup?

## Satisficing search finds solutions that are good enough

Consider finding sharp needles in a haystack:

### Optimization

Find sharpest needle:  
Need statistics on upper sharpness bound in areas of the haystack

### Satisficing

Find sharp enough needle:  
Need statistics on density of sharp enough needles in areas of the haystack

# Attribution of Intelligence

# Attribution of Intelligence

## Observation

An ant moves along the beach on a complex path. His progress is slow but apparent. Shall we ascribe intelligence to the ant based on our observation of his behavior?

# Attribution of Intelligence

## Observation

An ant moves along the beach on a complex path. His progress is slow but apparent. Shall we ascribe intelligence to the ant based on our observation of his behavior?

## Simon Says

Consider a mechanism that reacts to local contours to manage its energy expenditure. This simple being—our ant—exhibits complex behavior because of interactions with a complex environment—the ripple pattern in the sand.

# Attribution of Intelligence

## Observation

An ant moves along the beach on a complex path. His progress is slow but apparent. Shall we ascribe intelligence to the ant based on our observation of his behavior?

## Simon Says

Consider a mechanism that reacts to local contours to manage its energy expenditure. This simple being—our ant—exhibits complex behavior because of interactions with a complex environment—the ripple pattern in the sand.

## Question/Agenda

Does complex behavior, even that which we believe requires intelligence, need an esoteric explanation or will simple rules applied in a complex environment provide an adequate explanation?

# What is Knowledge?

# What is Knowledge?

Philosophers have worried this question for a long time:

- There is no good easy (or hard) answer to this question.
- A better question might be does agent  $A$  know  $K$ ?

# What is Knowledge?

Philosophers have worried this question for a long time:

- There is no good easy (or hard) answer to this question.
- A better question might be does agent  $A$  know  $K$ ?

## Newell Says

Agent  $A$  knows  $K$  if it applies  $K$  and its consequences whenever it is appropriate. Thus, one might talk about to what degree does  $A$  know  $K$ .



# What is Knowledge?

Philosophers have worried this question for a long time:

- There is no good easy (or hard) answer to this question.
- A better question might be does agent  $A$  **know**  $K$ ?

## Newell Says

Agent  $A$  knows  $K$  if it applies  $K$  and its consequences whenever it is appropriate. Thus, one might talk about to what degree does  $A$  know  $K$ .

## Reality Check

To what degree do **you** know calculus, politics, English grammar, chess, or company policy? Would you require more/less of an artifact to attribute knowledge?

# A View of Computation

Every computation systems manipulates entities in an interior model—a sort of simulation.

That model may be explicit (knowledge, objects, operators) or highly compiled in its code and data structures.

# A View of Computation

Every computation systems manipulates entities in an interior model—a sort of simulation.

That model may be explicit (knowledge, objects, operators) or highly compiled in its code and data structures.

Payroll Package	Question Answering
Knowledge of employees, tax laws, organizations, salary, etc.	Knowledge of grammar, meaning, domain rules, etc.
Objects to represent the above and current/next state.	Objects to represent language, parses, entities, domain rules.

Both examples are systems that represent external objects, perform operations on the interior objects inside the computation, then communicate facts about the interior representation—basically, a simulation.

# Almost Hierarchical Organizations

## Simon Says

A watch maker assembles a watch from 1000 parts. When the phone rings he answers it and must start over. Another watch maker designs a watch where each 10-part assemble is stable. He can build watches despite the phone!

# Almost Hierarchical Organizations

## Simon Says

A watch maker assembles a watch from 1000 parts. When the phone rings he answers it and must start over. Another watch maker designs a watch where each 10-part assemble is stable. He can build watches despite the phone!

- Almost all natural and artificial systems have a nearly hierarchical structure: animals, vegetables, corporations, airplanes, watches, even knowledge repositories.

# Almost Hierarchical Organizations

## Simon Says

A watch maker assembles a watch from 1000 parts. When the phone rings he answers it and must start over. Another watch maker designs a watch where each 10-part assemble is stable. He can build watches despite the phone!

- Almost all natural and artificial systems have a nearly hierarchical structure: animals, vegetables, corporations, airplanes, watches, even knowledge repositories.
- It is necessary to deviate from strict hierarchy to provide efficient functionality. (Parable of the new pencil.)

# Almost Hierarchical Organizations

## Simon Says

A watch maker assembles a watch from 1000 parts. When the phone rings he answers it and must start over. Another watch maker designs a watch where each 10-part assemble is stable. He can build watches despite the phone!

- Almost all natural and artificial systems have a nearly hierarchical structure: animals, vegetables, corporations, airplanes, watches, even knowledge repositories.
- It is necessary to deviate from strict hierarchy to provide efficient functionality. (Parable of the new pencil.)
- AI provides a large repertoire of tools to build almost hierarchical systems. This is one of its most interesting contributions.

# Learning

An intelligent being acquires knowledge from a mentor or from the school of hard knocks.

- A truly intelligent artifact could be educated.
- Can we be its mentor/evaluator?



An intelligent being acquires knowledge from a mentor or from the school of hard knocks.

- A truly intelligent artifact could be educated.
- Can we be its mentor/evaluator?

## Human

Our Vision

Heuristic: Where there's  
smoke there's fire.

# Learning

An intelligent being acquires knowledge from a mentor or from the school of hard knocks.

- A truly intelligent artifact could be educated.
- Can we be its mentor/evaluator?

## Human

Our Vision

Heuristic: Where there's smoke there's fire.

## Artifact/Alien Species

IR Vision

Heuristic: Where there's an intense hot spot there's fire.

# Learning

An intelligent being acquires knowledge from a mentor or from the school of hard knocks.

- A truly intelligent artifact could be educated.
- Can we be its mentor/evaluator?

## Human

Our Vision

Heuristic: Where there's smoke there's fire.

## Artifact/Alien Species

IR Vision

Heuristic: Where there's an intense hot spot there's fire.

- We can't teach the latter heuristic, we can't evaluate its application, and, in short, we can't appreciate it.

An intelligent being acquires knowledge from a mentor or from the school of hard knocks.

- A truly intelligent artifact could be educated.
- Can we be its mentor/evaluator?

## Human

Our Vision

Heuristic: Where there's smoke there's fire.

## Artifact/Alien Species

IR Vision

Heuristic: Where there's an intense hot spot there's fire.

- We can't teach the latter heuristic, we can't evaluate its application, and, in short, we can't appreciate it.
- Nothing about a good definition of intelligence/learning should involve specific facts about human wet ware.

An intelligent being acquires knowledge from a mentor or from the school of hard knocks.

- A truly intelligent artifact could be educated.
- Can we be its mentor/evaluator?

## Human

Our Vision

Heuristic: Where there's smoke there's fire.

## Artifact/Alien Species

IR Vision

Heuristic: Where there's an intense hot spot there's fire.

- We can't teach the latter heuristic, we can't evaluate its application, and, in short, we can't appreciate it.
- Nothing about a good definition of intelligence/learning should involve specific facts about human wet ware.

Moore asked, "Who is Merlin's mommy?"

# A Few Great Hacks

# A Few Great Hacks

## Lenat — Associate Mathematician

Starts with simple knowledge of “element of” relation for multisets and heuristics about what makes concepts more or less interesting. Searches for interesting concepts.

# A Few Great Hacks

## Lenat — Associate Mathematician

Starts with simple knowledge of “element of” relation for multisets and heuristics about what makes concepts more or less interesting. Searches for interesting concepts.

## Winograd — Shrdlu

English interactions to query, command, and define events in a virtual blocks world.



# A Few Great Hacks

## Lenat — Associate Mathematician

Starts with simple knowledge of “element of” relation for multisets and heuristics about what makes concepts more or less interesting. Searches for interesting concepts.

## Winograd — Shrdlu

English interactions to query, command, and define events in a virtual blocks world.

## Swartout — Explanation Subsystem

Expert system technology so that a system—a digitalis advisor—can describe its own behavior and defend it.

# A Few Great Hacks

## Lenat — Associate Mathematician

Starts with simple knowledge of “element of” relation for multisets and heuristics about what makes concepts more or less interesting. Searches for interesting concepts.

## Winograd — Shrdlu

English interactions to query, command, and define events in a virtual blocks world.

## Swartout — Explanation Subsystem

Expert system technology so that a system—a digitalis advisor—can describe its own behavior and defend it.

## McCarthy — Lisp

Systems that can inspect and reason about their own state.

# A Few More Great Hacks

# A Few More Great Hacks

Sussman — Electronic Circuit Repair Instructor

Teach repair skills interactively.

# A Few More Great Hacks

Sussman — Electronic Circuit Repair Instructor

Teach repair skills interactively.

Feigenbaum — Dendral

Determines molecular structure from mass spectrometry data using expert human rules of thumb.

# A Few More Great Hacks

## Sussman — Electronic Circuit Repair Instructor

Teach repair skills interactively.

## Feigenbaum — Dendral

Determines molecular structure from mass spectrometry data using expert human rules of thumb.

## Shank — Conceptual Information Processing

A broad sketch of a complete theory of language understanding, synthesis, and natural inference.

# A Few More Great Hacks

## Sussman — Electronic Circuit Repair Instructor

Teach repair skills interactively.

## Feigenbaum — Dendral

Determines molecular structure from mass spectrometry data using expert human rules of thumb.

## Shank — Conceptual Information Processing

A broad sketch of a complete theory of language understanding, synthesis, and natural inference.

## Berliner — Chess and Backgammon

The father of computer chess and creator of an automation that won an international backgammon tournament.

## Early Days

Strong AI  
Ignore limitations  
Everything in play

- Early days where characterized by visionary funding agencies, researchers, and kids who wanted to prove they could make progress on the hardest problems imaginable.



# AI Time Line and Focus

## Early Days

Strong AI  
Ignore limitations  
Everything in play

## Middle Period

Performance AI  
Fit what you have  
Balanced systems

- Early days where characterized by visionary funding agencies, researchers, and kids who wanted to prove they could make progress on the hardest problems imaginable.
- Since that time, universities have become more dependent on external funding and those funding sources have went from visionary, to mission oriented, to short term = short sighted.

# AI Time Line and Focus

## Early Days

Strong AI  
Ignore limitations  
Everything in play

## Middle Period

Performance AI  
Fit what you have  
Balanced systems

## Today

Military relevance  
Power,  $\rightarrow$  knowledge  
Math based

- Early days where characterized by visionary funding agencies, researchers, and kids who wanted to prove they could make progress on the hardest problems imaginable.
- Since that time, universities have become more dependent on external funding and those funding sources have went from visionary, to mission oriented, to short term = short sighted.
- Much if not most of current progress is based on cycle exploitation, not theoretical breakthroughs. It's been a long time since a single dissertation has made a large competence statement. How long since you've read a dissertation?

# The Final(?) Word

Dan Cohen Says

AI is a very promising technology. It always has been and always will be!

Jeff Says

Amen! Oh, wait a minute.