

Artificial Intelligence (AI)

What is it?

Recent Examples ?

Artificial Intelligence (AI)

Recent Examples

AI in the News

network security

market trading

game playing

space exploration

off-road driving

Artificial Intelligence (AI)

Definition

what researchers in AI do

what we can not yet do well with computers

acting as humans in demanding situations

generation of intelligent behavior by computers

intelligence \equiv generation of intelligent behavior

Intelligent Behavior

intentional

goal-oriented

input/belief-dependent

adaptive

creative

efficient/in time

others?

self-aware (conscious)

emotional

free will

Dennett

Intentionality

three stances toward behavior description

physical

measurable quantities

functional

component procedures

intentional

goals, beliefs, intentions

intelligence

an attributed property of observed behaviors

{in the eye of the beholder}

most efficient/effective description

Artificial Intelligence

Goals

think	x	x
act	x	x

like a human ideally rational

goals tend to range over these four categories

focus of research can be any of these quadrants

Artificial Intelligence

act like a human

Turing Test

can not tell difference

between computer and human

limited to a typed/indirect interaction

Loebner Prize 1990 ...

think like a human

cognitive modeling/simulation

same sequence of thought

match errors or response times

Artificial Intelligence

Approaches to Research

task/application

functionality

usability

demonstration

theoretical

complexity

completeness

experimentation

cognitive modeling

validity

simulation

Artificial Intelligence

What makes a computer program
a piece of artificial intelligence?

task

requires intelligence

technique

general

declarative

adaptive

Where's the AI?

Artificial Intelligence

Tasks

mundane

perception

language processing

commonsense/ everyday planning

formal

games and puzzles

algebra problems

expert

engineering design

scientific discovery

medical diagnosis

Artificial Intelligence

Arguments Against Possibility

only do what programmed to do

serial (computer) -vs- parallel (brain)

lack of success

free will

by definition

(function of human brain)

Artificial Intelligence

History

- 50's first symbolic programs
 neural networks
- 60's weak method definition
 micro-world successes
- 70's reality check, applications
 importance of knowledge
- 80's expert systems
 industrialization
- 90's learning
 formalization
- 00's agents/robots integration
 learning applications/data mining

Artificial Intelligence

Physical Symbol System Hypothesis

1980 Newell and Simon

intelligent behavior is the product
of symbolic computation
(necessary and sufficient)

"strong AI position"

"cognition is computation"

Hume, Boole, Frege,

Russell, Turing

Physical Symbol System Hypothesis

Symbol Structures

coin of the realm, material of thought

translated by interface

manipulated by processor

stored by memory

symbols

type a

token e.g., (eq a a)

relations

binary

n-ary

Physical Symbol System Hypothesis

Symbol Structures

designation

symbols stand for or represent things

environmental elements

internal states and processes

provide access

to other symbols

to relevant, learned behaviors

by acting as keys to memory

Physical Symbol System Hypothesis

Cognitive Architecture

environmental interface

receptors

environmental information

==> symbol structures (perception)

effectors

symbol structures

==> environmental information (action)

processor

interpreter

implements elementary symbolic operations

interacts with an active memory

memory

long-term storage and retrieval

Physical Symbol Systems

primitive operations

symbols

type/token creation

discrimination

association/designation

symbol structures (relations)

construction/selection

copy/alter/delete

read/write

memory

store active memory -> long term memory

retrieve long-term memory -> active memory

Agents and Environments

a general framework for considering
intelligent systems

behavior seen as interaction between an
agent and an environment

we can characterize agents and environments

success of agent of given type depends
to a large extent on type of environment

What does it take to be intelligent
in particular environments?

Environments

properties

observable - partially observable

deterministic - nondeterministic (stochastic)

episodic - sequential

static - dynamic - strategic

discrete - continuous

single agent - multiple agent

values on the left have typified many

toy problems and micro-worlds

values on the right typify many elements of

"real-world" applications

Agents

entities that generate (rational) behavior
in response to an environment

rational behavior

satisfies goals

"maximizes" value of actions

relative to some utility measure

relative to

what is believed in long-term memory

what has been (recently) perceived

Agent

Process Model

while(alive/awake)

 get-environment-input

 update-memory

 select-action

 update-memory

 do-action-in-environment

specific models differ in inclusion,

 complexity, and generality of the above

Agents

elements

architecture

program

types of agents

types of programs

Agents

reflex agents

while(alive/awake)

 get-environment-input

 select-action

 do-action-in-environment

no update of memory

action selection a

function of current input

for each possible input,

 associate an action in a

 (perception,action) table

Agents

reflex agents

rule-based agent

action selection based on
(condition, action) rules

conditions as generalizations
of reflex table entries
allow variables or categories to generalize
more general matching procedure

results of learning

rule-based systems

reactive plans or compiled knowledge

Agents

model based agent

while(alive/awake)

 get-environment-input

 update-perceived-state-in-memory

 select-action

 update-expected-state-in-memory

 do-action-in-environment

 maintain state model

 expected, perceived

 actions can be selected in

 terms of recent inputs

 changes or differences

 cope with partial accessibility

Agents

goal-based agent

compare state to goal

action selected based on

differences between state and goal

explicit representation of desired future

involved in action selection

planning agent

action selection based

on existing or new plan

more flexible

adopt different goals

utility-based agent

policies represented as utility function

action selected to maximize value

Task-Specific Agent

define by **PEAS** analysis

Performance measure

Environment

Actuators (actions)

Sensors (sensory percepts)

call center agent