

Digital Readiness: AI/ML


The thinking system quest

Gregg Vesonder
Stevens Institute of Technology
<http://vesonder.com>

Three Talks

- Digital Readiness: AI/ML, The thinking system quest.
 - Artificial Intelligence and Machine Learning (AI/ML) have had a fascinating evolution from 1950 to the present. This talk sketches the main themes of AI and machine learning, tracing the evolution of the field since its beginning in the 1950s and explaining some of its main concepts. These eras are characterized as “from knowledge is power” to “data is king”.
- Digital Readiness: AI/ML, Finding a doing machine.
 - In the last decade Machine Learning had a remarkable success record. We will review reasons for that success, review the technology, examine areas of need and explore what happened to the rest of AI, GOFAI (Good Old Fashion AI).
- Digital Readiness: AI/ML, Common Sense prevails?
 - Will there be another AI Winter? We will explore some clues to where the current AI/ML may reunite with GOFAI (Good Old Fashioned AI) and hopefully expand the utility of both. This will include extrapolating on the necessary melding of AI with engineering, particularly systems engineering.

Current State



Check out aiindex.org

GOFAI

Roadmap Talk 1

- Approaches to AI
- AI the early years, first summer
- Can computers “think”
- AI first winter
- AI second summer – Business
- AI second winter – why
- AI’s influence on computing in the 90s


Role of sci-fi

Winter is Coming?

- First Summer: Irrational Exuberance (1948 – 1966)
- First Winter (1967 – 1977)
- Second Summer: Knowledge is Power (1978 – 1987)
- Second Winter (1988 – 2011)
- Third Summer (2012 – ?)
- Why there might not be a third winter!

Henry Kautz – Engelmore Lecture

- Amazon \$159.95
- 1136 pages
- Fourth Edition



Approaches to AI

Thinking Humanly "The exciting new effort to make computers think ... <i>machines with minds</i> , in the full and literal sense" Haugeland, 1985	Thinking Rationally "The study of computations that make it possible to perceive, reason and act." Winston, 1985
Acting Humanly "The art of creating machines that perform functions that require intelligence when performed by people." Kurzweil, 1991	Acting Rationally "AI ... is concerned with intelligent behavior in artifacts." Nilsson, 1998

p. 2 AI : A modern approach: 3rd edition


Scruffy vs Neat

What's AI

- Artificial Intelligence not only tries to understand intelligence but also build intelligent entities (Russell & Norvig)

Cognition	Systems that think like humans	Systems that think rationally
Behavior	Systems that act like humans -> empirical science	Systems that act rationally -> math and engineering

Very Simplified Agent



Classifying the Agent Task Environments

- Fully vs. Partially observable -can agent detect all relevant aspects?
- Deterministic vs. Stochastic - is next state of environment solely determined by current state and agent's actions
- Episodic vs. Sequential - does current decision affect future decisions?
- Static vs. Dynamic - can the environment change while the agent is deliberating?
- Discrete vs. continuous applies to state, time and to percepts & actions of the agents
- Single agent vs. Multi-agent - competitive or cooperative? Communication?
- You're in trouble with partially observable, stochastic, sequential , dynamic, continuous and multi-agent

AI & Agents

- **Agent:** something that acts, operates under autonomous control, perceiving the environment, persisting over long time periods, adapting, being able to take on another's goals
- A **rational agent** achieves the best outcome or, given uncertainty, the best expected outcome
- **Perfect** rationality - always doing the right thing, is not feasible in complicated environments
- **Limited** rationality - acting appropriately when there is not enough time (or ability or feasibility) to do all the calculations one might like

Types of Agents

- Simple Reflex Agents
 - Simple rule bases
- Model based reflex agent
 - Model state reflects part of the “percept history”
 - Information needed:
 - How world evolves independent of agent
 - How agent’s actions affect world



Types of Agents - 2

- Goal based agent
 - Agents use goals to select actions
 - Search and Planning devoted to action sequences achieving agent's goals
- Utility based
 - If one world state preferred to another then it has higher utility for agent
 - Utility functions are the measure
 - Enables tradeoffs on conflicting goals
 - Provides a way to measure likelihood of success against importance of goals

Learning and Agents


- Four components
 - Learning element
 - Performance element - selects external actions (The agent)
 - Critic - provides feedback
 - Problem generation - suggests actions leading to new informative experiences (exploration)

AI Beginnings

Dartmouth 1956



Herb Simon



John McCarthy



Marvin Minsky

Can computers think?

- Most of the research explored whether computation could accomplish a task or set of tasks
 - Address “A computer can’t do X”
- Demonstration experiments in a sense
- Did these tasks indicate that machines could think
- The Turing Test was in the back of their minds

Eliza

Talk to Eliza

* Hello I am Gregg
> Did you come to me because you are Gregg?
* No I came to you because I am sad
> How long have you been sad?
* Since I offered to give three talks
> What does that suggest to you?
* That I am normal
> Do you enjoy being normal?
* very much so
> Tell me more...|

Input:

JOSEPH WEIZENBAUM

COMPUTER POWER AND HUMAN REASON

FROM JUDGMENT
TO CALCULATION

<http://psych.fullerton.edu/mbirnbaum/psych101/Eliza.htm>

AI Summer 1 ~1948 to 1966 Can Computers Think?




Table of Contents: Computers and Thought

- AI
 - Can a machine think?
 - Computing machinery and intelligence - Turing
 - Machines that play games
 - Chess playing programs and the problem of complexity – Newell, Shaw and Simon
 - *Some studies in Machine learning using the game of checkers – A.L. Samuel*
 - Machines that prove mathematical theorems
 - Two important applications*
 - Question answering machines
 - Pattern recognition
- Simulation of Cognitive Processes
 - *GPS: A program that simulates human thought - Newell & Simon*
 - The simulation of verbal learning behavior
 - Programming a model of human concept formation
 - Simulation of behavior in the binary choice experiment
 - Model of the trust investment process
 - Computer model of elementary social behavior

*line balancing & symbolic integration

nim – some simple context

- Nim
- “learning program”
- Start with 21 sticks (standard game)
- Players selects 1, 2 or 3 sticks
- One of the players is the machine
- Player selecting last stick loses nim

Elements of nim learning design

Sticks left	Take 1	Take 2	Take 3
1	100	-99	-99
2	0	0	-99
3	0	0	0
4	0	0	0
5	0	0	0
6	0	0	0
...			

The original program for nim was written in basic in 1974 and was 19 lines long

Learn from game play

Record moves of both players
Take moves of winner and increment by 1
Take moves of loser and decrement by 1

Sticks left	Take 1	Take 2	Take 3
1	100	-99	-99
2	57	-57	-99
3	-57	57	-57
4
5
6
...			

Samuel's checker playing program

- Used Claude Shannon's min-max procedure for chess
- "look ahead" procedure
- Played games against itself
- Checkers more reasonable, chess a bit worse 10^{120} continuations
- Last version achieved "respectable" amateur status"

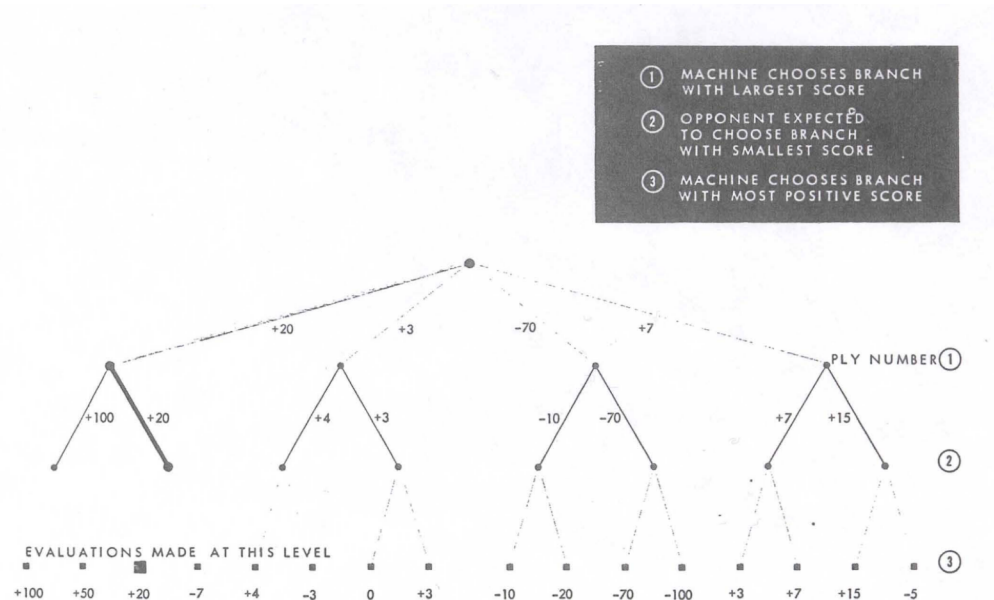


Figure 2 Simplified diagram showing how the evaluations are backed-up through the "tree" of possible moves to arrive at the best next move. The evaluation process starts at ③.

GPS

- GPS – General Problem Solver – not humble!
 - Cut to the end it wasn't a general problem solver!
- Goal was to solve lots of programs using same general reasoning mechanism
- Used means ends analysis. Initial states and goal states
- Class of systems that attempts to solve problems as a human does
- Initial state, current state, goal state driven by distance to the goal and operators to traverse states
- Hobbits and Orcs is an example of the type of problem it could solve

State space

- State set
- Start state
- Goal state test
- Operators (and costs)

Planning graphs

THE HOBBITS-AND-ORCS PROBLEM

Try solving this problem. (The answer is at the end of the chapter.)

Three Hobbits and three Orcs arrive at a river bank, and they all wish to cross to the other side. Fortunately, there is a boat, but unfortunately, the boat can hold only two creatures at one time. There is another problem. Orcs are vicious creatures, and whenever there are more Orcs than Hobbits on one side of the river, the Orcs will immediately attack the Hobbits and eat them up. Consequently, you should be certain that you never leave more Orcs than Hobbits on any river bank. How should the problem be solved? (It must be added that the Orcs, though vicious, can be trusted to bring the boat back!)



Hobbits and Orcs solution

Send 2 orcs across

1 orc

2 orcs

1 orc

2 hobbits

1 hobbit and 1 orc

2 hobbits

1 orc

[at this point, all hobbits are
on the goal side and all orcs
are on the start side]

2 orcs


1 orc

2 orcs Done!

River bank	River boat	Other bank
3H, 3O		
3H, 1O	2O>	
3H, 1O	<1O	1O
3H,	2O>	1O
3H,1O	<1O	2O
1H,1O	2H>	2O
1H,1O	<1H,1O	1H,1O
2O	2H>	1H,1O
2O	<1O	3H
1O	2O>	3H
1O	<1O	3H, 1O
	2O>	3H, 1O
		3H, 3O

Robotics

- Shakey – Stanford
- STRIPS



Thanks wikipedia

STRIPS

- An initial state;
- The specification of the goal states – situations which the planner is trying to reach;
- A set of actions. For each action, the following are included:
 - preconditions (what must be established before the action is performed);
 - postconditions (what is established after the action is performed).

STRIPS: monkey and bananas

Initial state: At(A), Level(low), BoxAt(C), BananasAt(B)

Goal state: Have(bananas)

Actions: // move from X to Y
Move(X, Y)
Preconditions: At(X), Level(low)
Postconditions: not At(X), At(Y)
// climb up on the box
ClimbUp(Location)
Preconditions: At(Location), BoxAt(Location), Level(low)
Postconditions: Level(high), not Level(low)
// climb down from the box
ClimbDown(Location)
Preconditions: At(Location), BoxAt(Location), Level(high)
Postconditions: Level(low), not Level(high)
// move monkey and box from X to Y
MoveBox(X, Y)
Preconditions: At(X), BoxAt(X), Level(low)
Postconditions: BoxAt(Y), not BoxAt(X), At(Y), not At(X)
// take the bananas
TakeBananas(Location)
Preconditions: At(Location), BananasAt(Location), Level(high)
Postconditions: Have(bananas)

AI Winter 1

- DARPA funding cuts
 - Early failure of practical speech recognition and autonomous tank project
 - Mansfield Amendment (1969) mandated that DARPA stop funding basic undirected research
- United Kingdom: Lighthill Report (1973)
 - Criticized failure of AI outside of toy micro-worlds
 - Claimed AI could never tame combinatorial explosion of real-world domains
 - Complete dismantling of AI research in UK


Second Summer (1966-87)

- Knowledge is power
 - Knowledge Engineering
 - Expert Systems
 - MYCIN
 - R1
 - DENDRAL
 - ACE
 - Expert System Environments
 - Lisp Machine and their programming environments

Knowledge Engineering Process

1. Review institutionalized knowledge
2. Gather problems
3. Construct rules with current knowledge
4. Solve real problems in parallel with experts – John Henry technique
5. Highlight differences
6. Discover new knowledge or correct knowledge
7. Attained satisfactory level of performance – “satisfiability” – exit
8. Goto 3

Rule Based System




- Data driven applications
- Originally used for expert systems
- Can act as a blackboard
WM blackboard, rules independent modules

1. Rule Base Operation

1. Match rules
2. *Apply conflict resolution*
3. Fire rule (changes working memory + side effects)
4. Goto 1

Blackboard Model



Note memory is blackboard
and computation occurs at
all independent KS
elements

rule base

Panel 3. The Stoplight Example

Working Memory Elements

```
traffic_light color
car_motion    starting_point
               my_car
               current_direction
               desired_direction
               ending_point
```

Working Memory Example

```
traffic_light color=red
car_motion    starting_point=home
               my_car=me
               current_direction=approach_intersection
               desired_direction=straight
               ending_point=airport
```

Rules

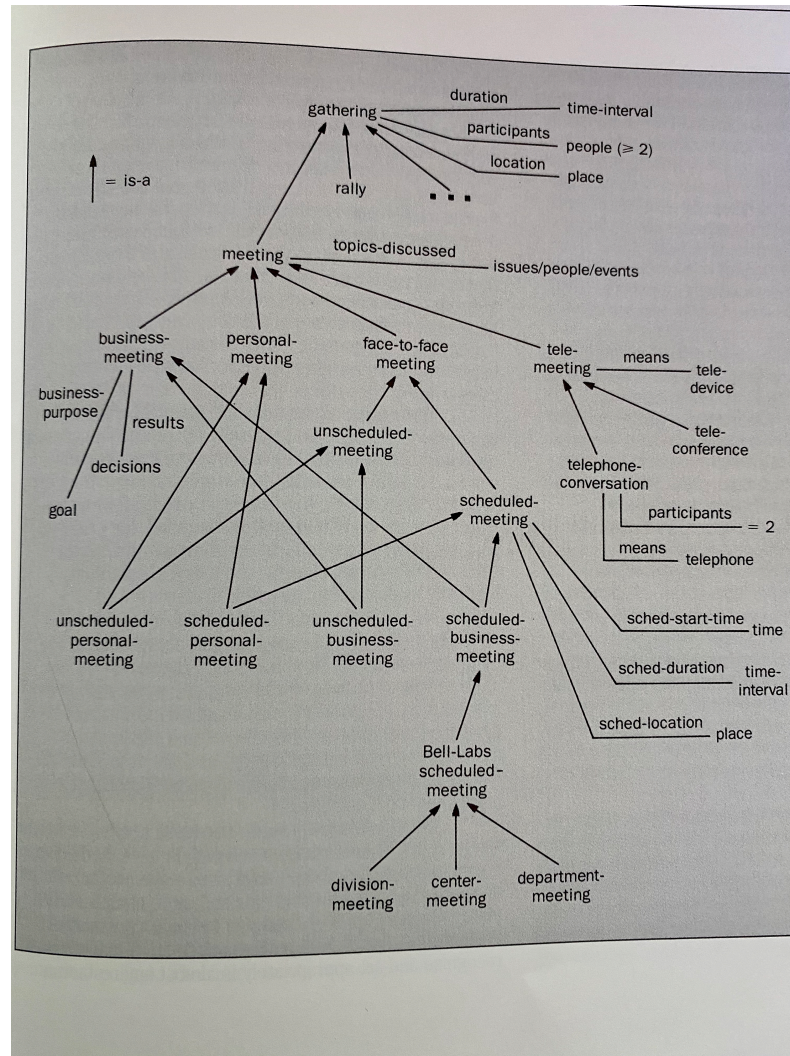
Rule 1:

```
IF car_motion    my_car=me
                  current_direction=approach_intersection
                  desired_direction=(member-of straight, left)
and traffic_light color=red
THEN change current_direction to equal stopped
```

Rule 2:

```
IF car_motion    my_car=me
                  current_direction=stopped
                  desired_direction=straight
and traffic_light color=green
THEN change current_direction to equal straight
```

semantic network



frame


Panel 3. A Frame for a Department Meeting

```
11257-DEPARTMENT-MEETING
is-a:      MEETING
attendee:   at-least 3 11257-DEPARTMENT-MEMBER
              (require: RON)
location:   a ROOM
              (default: 3D-473)
              (if-added: (reserve location))
date:       a DATE
              day: {Mon,Tues,Wed,Thurs,Fri}
agenda:     a list of
              TOPIC
                name:      a TOPIC-NAME
                speaker:    a PERSON
expected-duration: a TIME-INTERVAL
              (prefer: (less-than ONE-HOUR))
actual-duration:  a TIME-INTERVAL
              (default: (same-as expected-duration))
start-time:       a TIME
              (prefer: (and (before 4:30-pm) (after 1:30-pm)))
expected-end-time: a TIME
              (if-needed: (plus expected-duration start-time))
```

Tools

- KEE – Knowledge Engineering Environment, Intellicorp
- ART – Automated Reasoning Tool
- Knowledge Craft
- OPS4, 5, 83
- InterLisp, Lisp machine lisp (Maclisp), Franz Lisp
- Common Lisp

Lisp Machines




- XEROX: Dorado, Dolphin, Dandelion
- SYMBOLICS
- LMI
- Texas Instruments Explorer

Biologically Inspired Computing

- aka Natural computing:
 - Genetic algorithms
 - Cellular automata
 - Swarm intelligence
 - Ant algorithms (pheromones)
- Explore Alife too
- Simon's ant on the beach

BOIDS

- Craig Reynolds
 - Separation: steer to avoid crowding
 - Alignment: steer towards average heading
 - Cohesion: steer to move toward the average position



Ants, Evolution and Pheromones

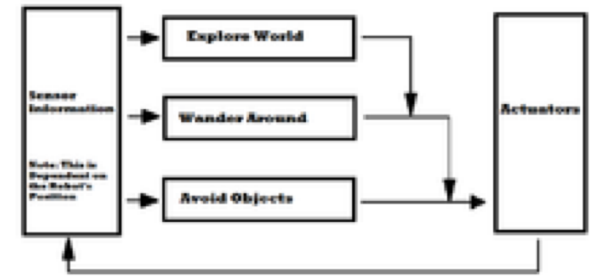
- An ant finding food marks trail when returning to colony, other ants follow reinforcing trail, until food exhausted
- Sniffing square ahead, moving ahead or turning left or right – Jefferson and Taylor
- Traveling salesman problem

Genetic Algorithm Loop

1. Generate initial population
2. Calculate fitness of each organism
3. Select a pair of parents
4. Crossover with probability x
5. If no crossover parents duplicated in population
6. Mutate each gene with certain probability
7. Potential parents available go to 3, else go to 2 and start next generation

Subsumption Architecture

- Situatedness – directly interact with the world rather than model it (robots contain no internal model of the world)
- Embodied agent – real robots, linking sense data to meaningful actions
- “Intelligence is determined by the dynamics of interacting with the world”
- Emergence – only determined by observing the agent and its environment



Thanks wikipedia

Rodney Brooks - 2

- Subsumption architecture –behavior based programming
 - When it finds dirt a specific behavior kicks in
 - If it is stuck algorithms to get itself unstuck kicks in
 - If battery is low it goes into search mode looking for the guiding beacon and guiding behavior takes it to the dock



<https://www.createdigital.org.au/rethink-robotics-collaborative-robots-pioneer/>

AI Winter 2

- Knowledge-engineered expert systems proved costly to maintain
- Collapse of market for specialized AI workstations
- Failure of Japan's Fifth-Generation AI effort based on Prolog hardware and software, Europe's Esprit and US Strategic Computing

Influence

- Rapid prototyping
- Programming environments
- Requirements discovery
- Declarative programming
- System engineering
- Agile
- Closures, functional programming, object oriented programming

Next Week

- Digital Readiness: AI/ML, Finding a doing machine.
 - In the last decade Machine Learning had a remarkable success record. We will review reasons for that success, review the technology, examine areas of need and explore what happened to the rest of AI, GOFAI (Good Old Fashion AI).