



Coordination Game Tournament

Human-Human Mean = .68



Human-Computer Mean = .77



Quiz questions

- Observation: 3 worst had disputed answers
- Too many questions, so new constraints:
 - Wait if you've had a question selected
 - Select only 1 question/day

Architectures & Ethics for Robots


COGS300.002

Peter Danielson

Jan 14 2016


Tele-Robotics

- Nikola Tesla remote controlled boat 1898
 - In 1906 Lee de Forest invented the Audion, the first triode vacuum tube and the first electrical device which could amplify a weak electrical signal and make it stronger.
- “a borrowed mind”
 - Mind control “moral panic”
- When a New York Times writer suggested that Tesla could make the boat submerge and carry dynamite as a weapon of war, the inventor himself exploded. Tesla quickly corrected the reporter: "You do not see there a wireless torpedo, you see there the first of a race of robots, mechanical men which will do the laborious work of the human race."



Tele-Robotics

- 1940 Version




Feedback Control - 1

- Sperry Gyro-Pilot
- Test 1923: “No Human Hand Touching Wheel in 12,000-Mile Voyage”
 - Spokane Daily Chronicle
Sept 3, 1923, p 19
- Current AirBus
- “With the introduction of “Metal Mike” the former man at the wheel may draw up an easy chair, sit down and study navigation... The gyro-pilot relieves him of the responsibility of holding his head on the compass needle and his hand on the wheel.”

Autonomous Robots


- First 1948-9
- W. Grey Walter
“Machina speculatrix”
 - Here: Elsie
- Unpredictable



Feedback Control - 2

- Proportional-integral-derivative controller (PID controller)
- “In the absence of knowledge of the underlying process, .. historically been considered to be the best controller”
- Complex tuning
 - 3 parameters

Closed (Feedback) Loop



Programmed Industrial Robots

- Unimate (1961)
- First Industrial Robot
- Car, CRT assembly
- ~4k 3-d waypoints
“programmed” by
example
- 4,000 pounds
- Protected space



Robot Ethics Baseline


- What is “epsilon” for/ I wouldn’t let X into my home w/o supervision; X =
 1. Roomba (video)
 2. Dog + cat
 - Agent constraints (provable?)
 - Environmental constraints (not universal)
 3. Human caretaker

GOGFAI & GOFAIR

- G(ood) O(ld) F(ashioned) AI
 - Newell & Simon: Search & Means Ends Reasoning
- STRIPS & robotics

Shakey & STRIPS

- 1966 – 72 Stanford Research Institute (Problem Solver)
- STRIPS planner
 - Means end planning
 - <http://www.aispace.org/planning/version4.0.0/launcher.html>
- Simplified room environment
- Speed: 2 meters/hour!



Q1

Which of the following is NOT a meta-assumption of GOFAIR?

- (a) Perception is needed only to determine initial world state
- (b) Requires a mostly observable, open world
- (c) A single agent
- (d) Serial action execution order
- (e) All of the above are meta-assumptions of GOFAIR

-Aurelia

Evaluate Q1

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Q2

In a Sudoku puzzle, "every row must contain every digit from 1 to 9" is an example of what?

- a) Variable
- b) **Constraint**
- c) Constraint satisfaction problem
- d) Domain
- e) Value

Chloe

Evaluate Q2

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Constraint Satisfaction

- <http://aispace.org/constraint/version4.6.1/launcher.html>

E.g. Clueless Crossword

A1/D1	D2	D3
A2		
A3		

B	O	A
E	A	R
E	F	T

“Clue” for A1: add, ado, age, ago, aid, ail, aim, air, and, any, ape, apt, arc, are, ark, arm, art, ash, ask, auk, awe, awl, aye, bad, bag, ban, bat, bee, boa, ear, eel, eft, far, fat, fit, lee, oaf, rat, tar, tie

GOFAIR vs. Soccer as Task

	Sudoku	Soccer
Number of agents	1	23
Competition	No	Yes
Collaboration	No	Yes
Real time	No	Yes
Dynamics	Minimal	Yes
Chance	No	Yes
Online	No	Yes
Planning Horizons	No	Yes
Situated Perception	No	Yes
Partially Observable	No	Yes
Open World	No	Yes
Learning	Some	Yes

Q3

Constraint Satisfaction problem methodology excels at problem solving due to its ability to

React to multiple unforeseen variables in the real world.

Act in a closed environment where the result of actions is well understood.

Understand the impact that its actions have on a real world environment.

Prioritize desired states and planning for the long term.

A. 1 and 3

B. 4 and 1

C. 3 and 2

D. 4 and 2

E. 1 and 2

Tyler

Evaluate Q3

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Horizontal Architecture

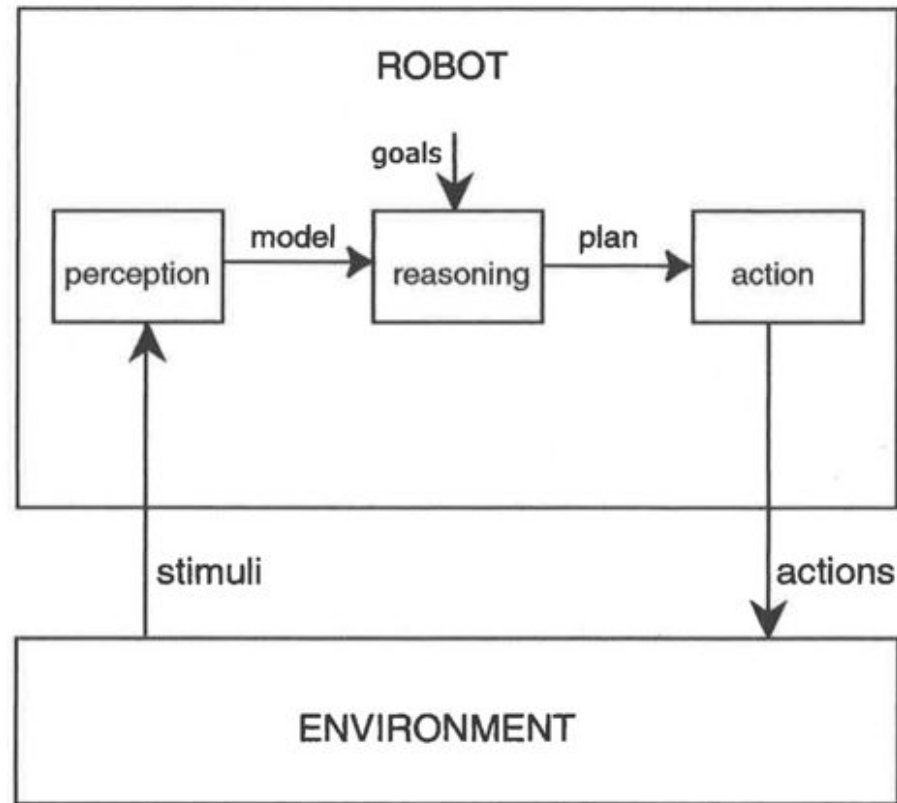
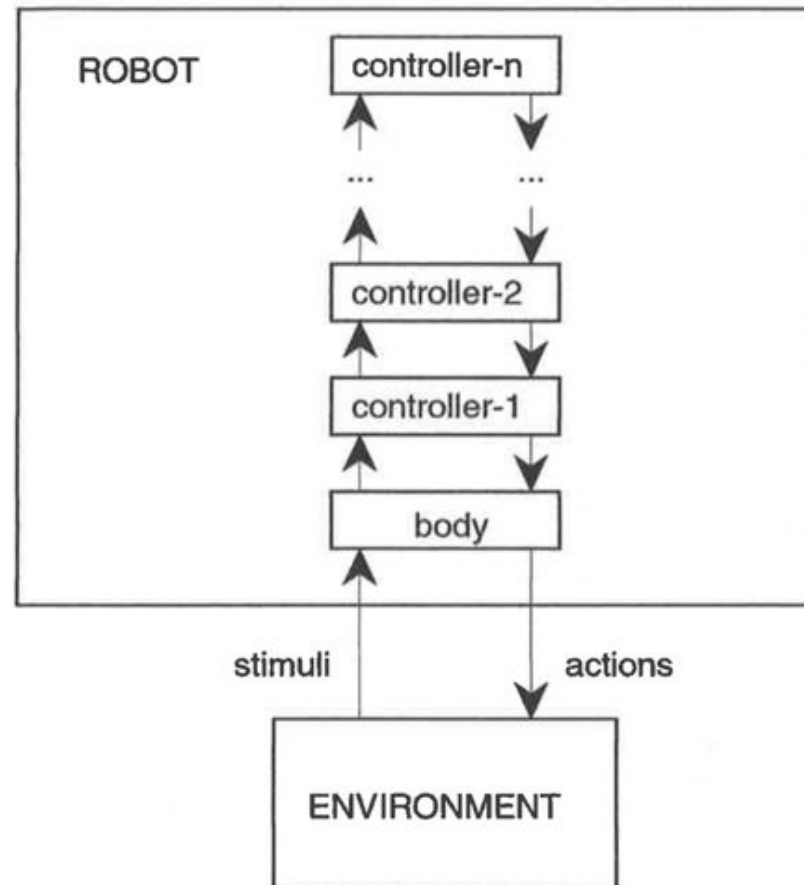


Figure 4. A Horizontal Architecture for a GOFAIR Robot

Wedding Cake Architecture



Q4

Which of the following does a Vertical Robotic System Architecture allow that a classical Horizontal Architecture does not?

- A. Prioritized constraints
- B. Reaction to real-time changes
- C. Perception of environmental stimuli
- D. **A and B**
- E. A and C

-- Trevor

Evaluate Q4

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Q5

The prioritized constraint order in Mackworth's soccer robot are :

1. Ball in image
2. Robot at ball
3. Base heading pan
4. Ball in center

- A. 1->2->3->4
- B. 4->3->2->1
- C. 3->4->2->1
- D. 1->4->3->2
- E. Any order as long as all are satisfied

Winson

Evaluate Q5

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Q 6

Which of the following is NOT part of Asimov's Laws of Robotics:

- a) A robot may decide, as it develops autonomy and free will, what to do ethically should it not receive orders from a human being.
- b) A robot may not harm a human being or allow a human being to come to harm through its inaction.
- c) A robot must protect own existence, as long as such protection does not conflict with the other Asimov's Laws.
- d) A robot must obey the orders given to it by a human being unless it conflicts with the other Asimov's Laws.
- e) None of the above.

-Sophia

Evaluate Q6

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Asimov & Constraints

- Asimov's Laws as constraints
 - Goal-based alternative (Utilitarianism: Next)
- For a general purpose robot
 - Constrain task domain
- Mackworth argues that robot soccer is a good task domain for situated robotics. What would be a good task domain for robot ethics?

Discussion

Ignoring for the moment the presupposition of technological capability, can you see any problems with Asimov's Laws of Robotics? That is, do you think just these three main laws sufficiently cover potential issues that may arise in robot ethics?

- Linda

Evaluate D1

- A. Excellent question
- B. Very Good question
- C. Good question
- D. Acceptable question
- E. Poor question

Next

Please take the Robot Ethics Survey

http://your-views.org/D7/Robot_Ethics_Welcome

- Class code: Robbie
- By Tuesday, Jan 19
- Voluntary participation

- Grau, C. (2011) There is No 'I' in 'Robot': Robots and Utilitarianism (In Machine Ethics, (Eds, Susan, A. & Michael, A.) Cambridge University Press, pp. 451 – 463.

[https://](https://www.academia.edu/208540/There_is_no_I_in_Robot_Robots_and_Utilitarianism)

[www.academia.edu/208540/](https://www.academia.edu/208540/There_is_no_I_in_Robot_Robots_and_Utilitarianism)

[There is no I in Robot Robots and Utilitarianism](https://www.academia.edu/208540/There_is_no_I_in_Robot_Robots_and_Utilitarianism)