EEL5840: Elements of Machine Intelligence

Announcements

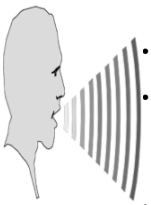

- Today's Handouts:
  - > Outline Class 2
- Web Site
  - > www.mil.ufl.edu/5840
  - > Software and Notes
  - > Programming Assignment Format
- Reading Assignment:
  - > Nilsson Chapter 2 & 3
  - > LISP Chapters 1-4
- Written Assignment
  - > Homework 1 Exercises 2.1-2.6 Due Thu. 9/1/09 in class

EEL5840: Elements of Machine Intelligence

Today's Menu

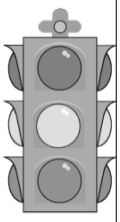

- Approach used in the Nilsson Text
- An example of a Classical AI Problem
  - > N-Queens Problem
- An Example of a Modern Machine Intelligence Problem
  - > Q-Learning: Learning to Push a Box
- Stimulus-Response (*SR*) Agents

EEL5840: Elements of Machine Intelligence

Approach Used in the Text

- Ideas are presented in the context of ever more capable and complex "agents" in *grid-space world*.
- Ideas then are easy to describe—yet a variety of enhancements makes the world sufficiently rich to demand intelligence out of its inhabiting agents.
- A typical grid-space world is the 3-D world of TJs in our lab—Nilsson's floor is conveniently demarcated by two-dimensional grid of "cells" or tiles on the floor. Objects must be on the floor or supported by a stack of objects resting on the floor.
- There may be wall-like boundaries between sets of cells. The agents are confined to the floor and move from cell to cell.

EEL5840: Elements of Machine Intelligence

Approach Used in the Text

- The first set of agents are called *reactive* agents: agents that have various means of sensing their worlds and acting in them.
- More complex reactive agents will have the ability to remember properties and to store internal models of the world.
- The actions taken by these agents are functions of the current and past states of their worlds—as they are sensed and remembered.
- Reactive agents may (and often do) have quite complex perceptual and motor processes.

MILECE

EEL5840: Elements of Machine Intelligence
Approach Used in the Text

- Most AI systems use some sort of model or representation of their world and task.
- A model is a symbolic structure and set of computations on it that correlate sufficiently with the world in that the computations yield information about the world useful to the agent. Information may be about present or future states.
- Iconic models: the use of data structures and computations that simulate aspects of an agent's environment and the effect of agent actions upon that environment. Example: n-queens, 8-puzzle
- Feature-Based models: use declarative descriptions of the environment.

MILECE

EEL5840: Elements of Machine Intelligence
Approach Used in the Text

- The second series of agents will have the ability to anticipate the effects of their actions and take those that are expected to lead toward their goals—agents that make plans.
- Grid-space worlds will have implicit constraints that are analogous to properties of real worlds, e.g., two objects cannot occupy the same grid at the same time. Agents that can take these and other constraints into account are said to "reason" and to "deduce" properties of their world that are only implicit in their constraints.
- The final set of agents live in a world inhabited by other agents —agent communication is required.

MILECE

EEL5840: Elements of Machine Intelligence
Classical AI Example

- Comprehensive Example: N-Queens Problem
- DEF: HEURISTIC - A rule of thumb, strategy, method, intuitive rule or trick used to improve the efficiency of a system which tries to discover the solution of complex problems. From the Greek "EUREKA," meaning 'serving to discover.'
- Problem: Place N Queens on an N x N chess board so that no two can attack one another. Choose a suitable representation and derive a solution. Can we device a suitable heuristic or a strategy?

MILECE

EEL5840: Elements of Machine Intelligence

EEL5840: Elements of Machine Intelligence

Classical AI Example

MILECE

- Problem: Place N Queens on an N x N chess board so that no two can attack one another. Choose a suitable representation and derive a solution. Can we devise a suitable heuristic or a strategy?
- Choose n-tuples to represent the data (x_1 , x_2 , x_3 , x_4) Let each x_i represent the queen in row *i*, i.e., x_1 =2 means queen in row 1 column 2. Clearly *x* can be {1,2,3,4}. The solution is (2,4,1,3) or (3,1,4,2)

```
_ Q _ _        _ _ Q _
_ _ _ Q        Q _ _ _
Q _ _ _        _ _ _ Q
_ _ Q _        _ Q _ _
```


EEL5840: Elements of Machine Intelligence

Classical AI Example

MILECE

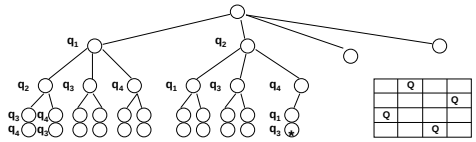

EEL5840: Elements of Machine Intelligence

Classical AI Example

MILECE

- HEURISTIC: If we have the current queen in column i then do not place the next queen in column i+1 or i-1

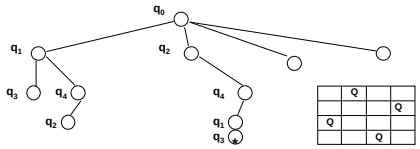

EEL5840: Elements of Machine Intelligence

Machine Intelligence Example

MILECE

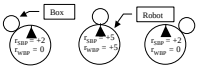

- MI Example

Q-TABLE Characteristics for All Experiments

Qty	Sensor	Input States	Factor
3	IR	3 (close, midrange, far)	3^3
2	IR	Combined (none, detect)	2^1
		Total Number of States in Q-Table	$3^3\ 2^1 = 54$

EEL5840: Elements of Machine Intelligence

MILECE

Machine Intelligence Example

Bumper used only as a negative reward generator in collision avoidance. All other reinforcement rewards are positive. This has an impact on sequential learning.

The various box surfaces, red blotches on white, blue and white stripes, brown cardboard.

- **Learning Behaviors Investigated**

Algorithm	Intuitive Description
Collision Avoidance	**Don't bump into anything massive enough to trigger the bumper.**
Weak Box Pushing	**Get close to objects in front and move forward.**
Strong Box Pushing	**Get close to objects in front (best), or on either side, and move fwd**

EEL5840: Elements of Machine Intelligence

MILECE

Machine Intelligence Example

Learning Parameters

Algorithm	α	γ	Run Time	ΣQ + 1000	Reinforcement Schedule
Collision Avoidance	0.5	0.8	30 minutes	22	r = +5 forward motion, -9 bump
Weak Box Pushing	0.5	0.8	40 minutes	27	r = +5 object close in front, +0.5 forward motion.
Strong Box Pushing	0.5	0.8	15 minutes	30	r =+5 object close in front, +2 object close on right, +2 object close on left, +0.5 forward motion

EEL5840: Elements of Machine Intelligence

MILECE

"Learning is an important part of autonomy. A system is autonomous to the extent that its behaviour is determined by its immediate inputs and past experience, rather by its designer's. Agents are usually designed for a class of environments, where each member of the class is consistent with what the designer knows about what the real environment might hold in store for the agent. Truly autonomous systems should be able to operate successfully in any environment, given sufficient time to adapt. The system's internal knowledge structures should therefore be constructible, in principle, from its experience of the world…"

The End!