

mc375: intro to robotics learning.

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

1

adaptation vs learning.

- learning produces changes within an organism that, over time, enables it to perform more effectively within its environment
- adaptation in learning through making adjustments in order to be more attuned to its environment.
 - different time scales: acclimatization (slow) vs homeostatis (rapid)
 - different levels: genotypic vs phenotypic

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

2

types of adaptation.

- behavioral -- behaviors are adjusted relative to each other
- evolutionary -- descendents are based on ancestor's performance over long time scales
- sensory -- sensors become more attuned to the environment
- learning as adaptation -- anything else that results in a more ecologically fit agent

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

3

importance of learning.

- learning is more than just adaptation
- learning, the ability to improve one's performance over time, is considered the main hallmark of intelligence, and the greatest challenge of AI
- learning is particularly difficult to achieve in physical robots, for all the reasons that make intelligent behavior in the physical world difficult

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

4

what learning enables.

- introducing new knowledge (facts, behaviors, rules) into the system
- generalizing concepts
- specializing concepts
- reorganizing information
- creating or discovering new concepts
- creating explanations
- reusing past experiences

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

5

why learn in robots?

- providing the robot with the ability to adapt to changes in its task and/or environment
- providing the robot with the ability to improve performance
- providing the robot designer with a way of automating the design and/or programming of the robot

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

6

what can be done?

- automatically design the robot's body
- automatically design a physical network for its processor
- automatically generate its behaviors
- automatically store and re-use its previous executed plans
- automatically improve the way its layers interact
- automatically tune the parameters within the behaviors and many more ...

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

7

what has been done?

- parts of robot bodies, brains (i.e., processors), and programs have been automatically generated (i.e., learned)
- robots given initial programs have used experience and trial & error to improve the programs (from parameter tuning to switching entire behaviors)
- robots programmed for a task have adapted to changes in the environment (e.g., new obstacles, new maps, heavier loads, new goals)

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

8

challenges in robots

- situatedness in the world: noise, occlusion, dynamics, hard to model, etc.
- real time constraints: slow learners die easily in the real world
- simultaneous and multi-modal: multiple goals & tasks, need to walk and talk (not just either or)

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

9

challenges in learning.

- credit assignment: who is to credit/blame for outcome?
- saliency: what is relevant right now?
- new term: when should a new concept/representation be created?
- indexing: how to organize the memory?
- utility: what should be forgotten?

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

10

levels of learning.

- within a behavior
 - suitable responses
 - suitable stimulus
 - suitable behavioral function mapping
 - magnitude of response (gain)
 - whole new behavior
- within a behavior assemblage
 - suitable set of behaviors
 - relative strengths
 - suitable coordination function

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

11

learning methods.

- reinforcement learning
- neural network (connectionist) learning
- evolutionary learning
- co-evolutionary learning
- learning from experience
 - memory-based
 - case-based
- inductive learning
- explanation-based learning
- multistrategy learning

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

12

types of learning.

- numeric or symbolic
 - numeric: manipulate numeric functions
 - symbolic: manipulate symbolic representations
- inductive or deductive
 - inductive: generalize from examples
 - deductive: optimize what is known
- continuous or batch
 - continuous: during interaction with the world
 - batch: after interaction, all at once

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

13

some terminology.

- reward/punishment
 - positive/negative feedback
- cost function/performance metric
 - scalar (usually) goodness measure
- induction
 - generating a function (a hypothesis) that approximates the observed examples
- teacher/critic
 - provides feedback

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

14

more terminology.

- plant/model
 - system/agent that we want to train
- convergence
 - reaching a desired (or steady) state
- credit assignment problem
 - who should get the credit/blame?
 - hard to tell over time
 - hard to tell in multi-robot systems

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

15

reinforcement learning (RL).

- currently the most popular approach to learning on mobile robots
- inspired by conditioning in psychology
- Law of effect: applying a reward immediately after the response increases its probability of reoccurring while providing punishment after the response will decrease the probability [Thorndike 1911]
- translated to robotics: some combinations of stimuli (i.e., sensor readings and/or state) and responses (i.e., actions/behaviors) are coupled with subsequent reward in order to increase their probability of future use

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

16

reinforcement learning, cont.

- desirable responses of outcomes are positively reinforced (rewarded) and thus strengthened, while undesirable ones are negatively reinforced (punished) and thus weakened
- this very general notion can be translated into a variety of specific reinforcement learning algorithms

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

17

challenges in RL.

- learning from delayed rewards: the problem is difficult if the feedback is not immediate
- credit assignment problem: when something good or bad happens, what exact state/condition-action/behavior should be rewarded or punished?
- common approach: use the expected value of exponentially weighted past/future reinforcement

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

18

RL algorithms.

- proscribe exact mathematical functions that associate states/situations, actions/behaviors, reinforcement and various associated parameters
- some of these algorithms (Q-learning, TD learning, etc.) have well-understood convergence properties; they are guaranteed to make the robot learn the optimal solution

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

19

optimality in learning.

- optimality depends on strong assumptions: the robot must be given infinitely many trials of each state/action combination, must know what state it is in and what action it has executed, and the world must not change too quickly
- but: there is not enough time for infinite trials, outcomes of actions are uncertain, the world can change
- thus: optimality is impractical

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

20

observability.

- the world to a robot is partially observable, the robot does not know exactly what state/situation it is in at all times
- learning algorithms have been developed for partially observable environments, and while they are more realistic, they require even more time to converge
- in general, managing uncertainty in learning is very hard

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

21

unsupervised learning.

- RL is a form of unsupervised learning
- RL allows a robot to learn on its own, using its own experiences (and some built-in notions of desirable and undesirable situations, associated with reward and punishment)
- the designer can also provide reinforcement (reward/punishment) directly, to influence the robot
- the robot is never told what to do

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

22

RL function and critic.

- RL systems contain a reinforcement function, which determines when and how much positive or negative reinforcement to deliver to the robot
- the critic is the part of the RL system which provides the reinforcement, i.e., executes the reinforcement function

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

23

types of critic.

- the critic can be
 - external: if the user provides the reinforcement
 - internal: if the system itself provides the reinforcement
- in both cases the approach is unsupervised in that the answer is never given explicitly by the critic

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

24

associative learning.

- learning new behaviors by associating sensors and actions into rules
- example:
 - 6-legged walking (Edinburgh)
 - whisker sensors first, IR and light later
 - 3 actions: left, right, ahead
 - user provided feedback -- *shaping*
 - learned avoidance, pushing, wall-following, light seeking

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

31

neural network examples.


- robot motion planning, articulation/manipulation
- control of complex plants: robots, aircraft
- control and coordination of multiple vehicles
- some domains and tasks lend themselves very well to supervised NN learning
- the answer to any given situation is well known and can be trained
- best example is robot motion planning for articulation/manipulation; NN's widely used for learning inverse kinematics

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

32

evolution.




4/21/2001 19:20

Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

33

co-evolution.



4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

34

evolutionary methods.

- genetic/evolutionary approaches are based on the evolutionary search metaphor
- states/situations and actions/behaviors are represented as "genes"
- different combinations are tried in different "individuals" in a "population"
- individuals with the highest "fitness", that perform the best, are kept as "survivors"; the others are discarded -- this is the *selection* process

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

35

evolutionary methods, cont.

- the survivors' genes are mutated, crossed-over, etc., and new individuals are so formed, which are then tested and scored
- in effect, the evolutionary process is searching through the space of solutions to find the one with the highest fitness
- often used for solving optimization problems
- trick is proper definition of representation, operators and fitness function

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

36

evolutionary methods: applications.

- tuning parameters (such as gains in a control system)
- developing controllers (policies) for individual robots
- developing group strategies for multi-robot systems (by testing groups as populations)


4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

37

genetic algorithms.

[Holland, 1975]



4/21/2001 19:20

Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

38

GA's vs GP's.

- GA = genetic algorithm
 - representation is strings of genes
- GP = genetic programming
 - representation is pieces of executable programs
- GP operates at a higher level of abstraction than GA's
- examples:
 - classifier systems
 - evolving structure/behavior [Sims]

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

39

fuzzy logic.

- fuzzy control produces actions using a set of fuzzy rules based on fuzzy logic
- this includes:
 - fuzzifying: mapping sensor readings into a set of fuzzy inputs
 - fuzzy rule base: a set of IF-THEN rules
 - fuzzy inference: maps fuzzy sets onto other fuzzy sets using membership functions
 - defuzzifying: mapping a set of fuzzy outputs onto a set of crisp output commands

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

40

fuzzy control.

- fuzzy logic allows for specifying behaviors as fuzzy rules
- such behaviors can be smoothly blended together
- fuzzy rules can be learned

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

41

memory-based learning (MBL).

- involves storing the exact parameters of a situation/action
- important for systems with a great deal of parameters, where blind search through parameter space would take too long
- takes a lot of space for data storage
- a trade-off against computation time
- assumed it is best to use space to remember rather than time to compute

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

42

symbolic learning.

- induction on decision trees
- learning new rules given a set of predefined rules
- e.g., by inference or other logical techniques
- example: RoboSOAR
 - logical reasoning system
 - production system based on implications or rules
 - uses chunking
 - conflict resolution mechanism when different rules apply
 - continuous conversion of deliberative information to efficient representations

4/21/2001 19:20

course notes adapted from
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

43

multi-strategy learning.

- some learning methods are effectively combined
- NN's and RL is common
 - use NN's to generalize states
 - then apply RL in a reduced state space
- NN's and MBL
 - use NN's to interpolate between instances
- NN's and fuzzy is also popular
- NN's and GA's is also useful

4/21/2001 19:20

course notes adapted from
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

44

learning to play games.

- learning from humans
 - Checkers [Samuels, 1959]
 - too many games are needed
 - humans are noisy
 - humans are learning
- learning from computers
 - Backgammon [Tesauro, 1992]
 - lack of generalization
 - deceptive landscape

4/21/2001 19:20

course notes adapted from
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

45

Tron.

- learning from humans on the Internet
[Funes, Sklar, Juille & Pollack, 1998]
- human users *collectively supply fitness function*
- software agents *collectively supply intelligent opponent*




4/21/2001 19:20

course notes adapted from
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

46

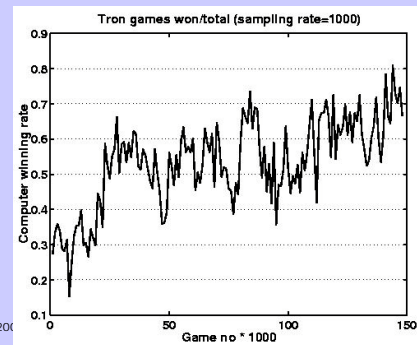
Tron architecture: two levels of co-evolution.



47

Internet LAN
Autonomous Systems © Andreas Birk, VUB

Tron results: computer learning.



4/21/2001

48

state of the art.

- there are many learning techniques
 - use dependant on application
- robotics is a particularly challenging domain for learning
 - very few successful examples
 - active area of research
- research journals
 - Machine Learning, Artificial Intelligence, Fuzzy Computing, Neural Networks, Adaptive Computing

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

49

state of the art, cont.

- supervised learning has been used to learn controllers for complex architectures, e.g., many DOF, as well as for navigating specific paths and performing articulated skills (e.g., juggling, pole-balancing)
- MBL has been used to learn controllers for very similar manipulator tasks (as above)
- RL has been used to learn controllers for individual navigating robots, groups of foraging robots, map-learning and maze-learning robots
- evolutionary learning has been used to learn controllers for individual navigating robots, maze-solving robots, herding robots and foraging robots

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

50

reading.

- Reward Functions for Accelerated Learning [Mataric]
(in course pack I I)
- Issues in Evolutionary Robotics [Harvey, Husbands and Cliff]
(in course pack I I)

4/21/2001 19:20

course notes adapted from:
Introduction to Robotics © Maja Mataric, USC
Autonomous Systems © Andreas Birk, VUB

51