Master Recherche IAC
Robots et agents autonomes

Jamal Atif – Michèle Sebag
TAO
CNRS – INRIA – LRI, Université Paris-Sud

Nov. 30th, 2012

Contents

WHO

- Jamal Atif, vision — TAO, LRI
- Michèle Sebag, machine learning — TAO, LRI

WHAT

1. Introduction
2. Vision
3. Navigation
4. Reinforcement Learning
5. Evolutionary Robotics

WHERE: http://tao.lri.fr/tiki-index.php?page=Courses

Exam

Final: same as for TC2:

- Questions
- Problems

Volunteers

- Some pointers are in the slides
  more ?

  here a paper or url
- Volunteers: read material, write one page, send it (sebag@lri.fr)

Questionaire

Admin: Ouassim Ait El Hara

Debriefing

- What is clear/unclear
- Pre-requisites
- Work organization

Overview

Introduction

The AI roots

Situated robotics

Reactive robotics

- Swarms & Subsumption
- The Darpa Challenge

Principles of Autonomous Agents

Myths

1. Pandora (the box)
2. Golem (Praga)
3. The chess player (The Turc)
   Edgar Allan Poe
4. Robota (still Praga)
5. Movies...

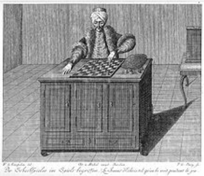

Types of robots: 1. Manufacturing

∗closed world, target behavior known
∗task is decomposed in subtasks
∗subtask: sequence of actions
∗no surprise

Automating Biology Using Robot Scientists

Ross D. King,
University of Manchester, ross.king@manchester.ac.uk

The Concept of a Robot Scientist

Computer systems capable of originating their own experiments, physically executing them, interpreting the results, and then repeating the cycle.

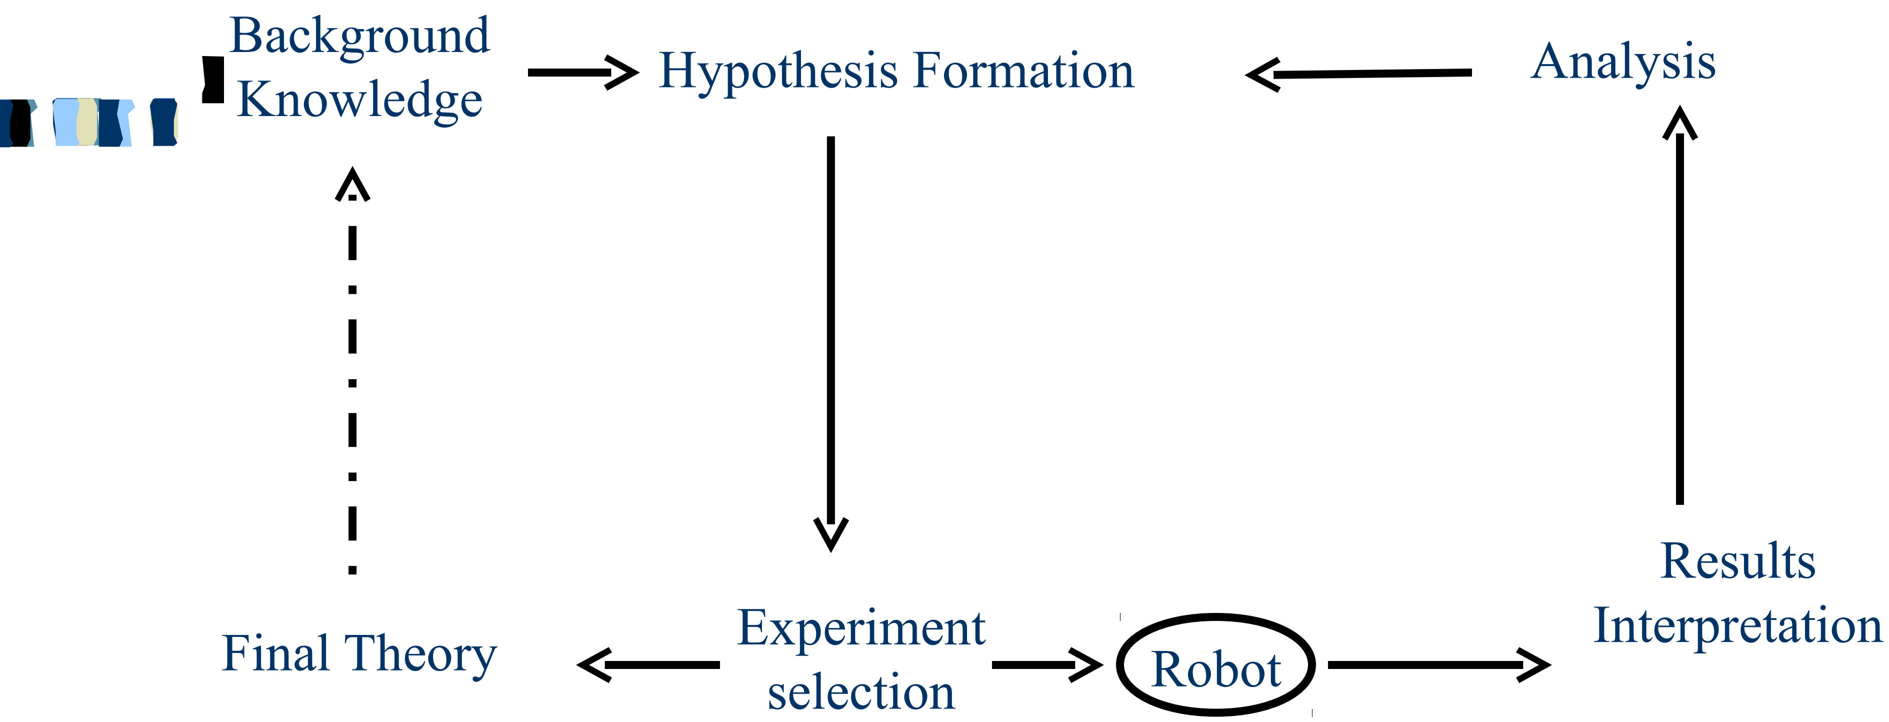

Types of robots: 1, followed

*no adaptation to new situations

Slotine et al., 95

Types of robots: 2. Autonomous vehicles

∗open world
∗task is to navigate
∗action subject to precondition

Types of robots: 2. Autonomous vehicles

∗a wheel chair

∗controlled by voice

∗validation ?

more ?

J. Pineau, R. West, A. Atrash, J. Villemure, F. Routhier. "On the Feasibility of Using a Standardized Test for Evaluating a Speech-Controlled Smart Wheelchair". International Journal of Intelligent Control and Systems. 16(2), pp.121-128, 2011.

Types of robots: 3. Home robots

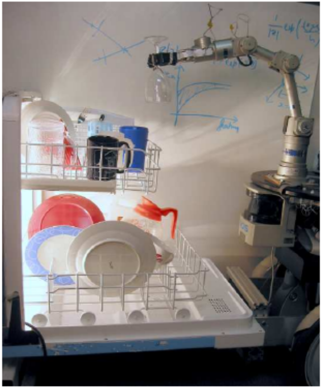

open world

sequence of tasks

each task requires navigation and planning

Vocabulary 1/3

- **State of the robot** set of states $\mathcal{S}$
  A state: all information related to the robot (sensor information; memory)
  Discrete ? continuous ? dimension ?
- **Action of the robot** set of actions $\mathcal{A}$
  values of the robot motors/actuators.
  e.g. a robotic arm with 39 degrees of freedom.
  (possible restrictions: not every action usable in any state).
- **Transition model**: how the state changes depending on the action
  deterministically $tr : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$
  probabilistically or $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$
  Simulator; forward model. deterministic or probabilistic transition.

Vocabulary 2/3

- **Rewards**: any guidance available. $r : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
  How to provide rewards in simulation ? in real-life ?
  What about the robot safety ?

- **Policy**: mapping from states to actions.
  deterministic $\pi : \mathcal{S} \mapsto \mathcal{A}$ or stochastic $\pi : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$
  this is the goal: finding a good policy
  good means:
  ∗reaching the goal
  ∗receiving as many rewards as possible
  ∗as early as possible.

Vocabulary 3/3

Episodic task

- Reaching a goal (playing a game, painting a car, putting something in the dishwasher)
- Do it as soon as possible
- Time horizon is finite

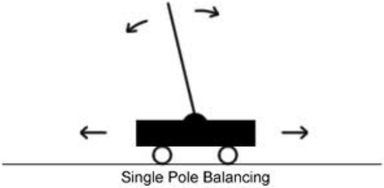

Continual task

- Reaching and keeping a state (pole balancing, car driving)
- Do it as long as you can
- Time horizon is (in principle) infinite

Case 1. Optimal control

Case 1. Optimal control, foll'd

Known dynamics and target behavior

1. state u, action $a \to$ new state u'
2. wanted: sequence of states

Approaches

- Inverse problem
- Optimal control

Challenges

- Model errors, uncertainties
- Stability

Case 2. Reactive behaviors

The 2005 Darpa Challenge

The terrain

The sensors

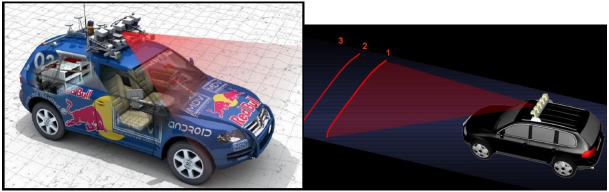

Case 3. Planning

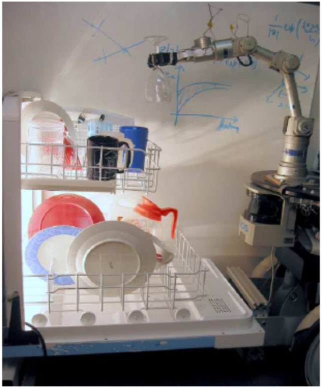

An instance of reinforcement learning / planning problem

1. Solution = sequence of (state,action)
2. In each state, decide the appropriate action
3. ..such that in the end, you reach the goal

Case 3. Planning, foll'd

Approaches

- Reinforcement learning
- Inverse reinforcement learning
- Preference-based RL
- Direct policy search (= optimize the controller)
- Evolutionary robotics

Challenges

- Design the objective function (define the optimization problem)
- Solve the optimization problem
- Assess the validity of the solution

Overview

The AI roots

J. McCarthy 56

We propose a study of artificial intelligence [..]. The study is to proceed on the basis of the conjecture that **every aspect of learning or any other feature of intelligence** can in principle be so precisely described that a machine can be made to simulate it.

Before AI...

Alan Turing

Machine Learning, 1950

by (...) mimicking education, we should hope to modify the machine until it could be relied on to produce definite reactions to certain commands.

Before AI...

Alan Turing

Machine Learning, 1950

by (...) mimicking education, we should hope to modify the machine until it could be relied on to produce definite reactions to certain commands.

How ?

One could carry through the organization of an intelligent machine with only two interfering inputs, one for pleasure or reward, and the other for pain or punishment.

The imitation game

The criterion:

Whether the machine could answer questions in such a way that it will be extremely difficult to guess whether the answers are given by a man, or by the machine

Critical issue

The extent we regard something as behaving in an intelligent manner is determined **as much by our own state of mind and training, as by the properties of the object under consideration**.

Oracle = human being

- Social intelligence matters

The imitation game, 2

So cute !

The imitation game, 2

The uncanny valley

more ?

http://www.androidscience.com/proceedings2005/MacDormanCogSci2005AS.pdf

AI and ML, first era

General Problem Solver

. . . not social intelligence

Focus

- Proof planning and induction
- Combining reasoners and theories

AM and Eurisko Lenat 83, 01

- Generate new concepts
- Assess them

Reasoning and Learning

Lessons Lenat 2001

the promise that *the more you know the more you can learn (..) sounds fine until you think about the inverse, namely, you do not start with very much in the system already. And there is not really that much that you can hope that it will learn completely cut off from the world.*

Interacting with the world is a must-have

Overview

Behavioral robotics

Rodney Brooks, 1990

Elephants don't play chess

- GOFAI: intelligence operates on (a system of) symbols
  ∗symbols (perceptual and sensori primitives) are given
  ∗narrow world, enabling inference (puzzlitis);
  ∗heuristics (monkeys and bananas)

- Nouvelle AI: situated activity
  ∗representations are physically grounded
  ∗mobility, acute vision and survival goals are essential to develop intelligence
  ∗intelligence emerges from functional modules
  ∗perception is an active and task dependent operation.

Milestones

A (shaky) evolutionary argument

Hardness is measured by the time needed for (biological entitities) to master it.

-4.5 MM	Earth
-3.8 MM	Single cells
-2.3 MM	Multicellular life
-550 M	Fish and vertebrates
-370 M	Reptiles
-250 M	Mammals
-120 M	First primates
-2.5 M	Humans
-19,000	Agriculture
-5,000	Writing

Key issues

Efficiency: the innate vs acquired debate

- Some things can be built-in, others are more difficult to be programmed
- Some things must be learned (training methodology ?)

High level vs low-level

- Learn low-level primitives ? (perceptual primitives)
- Learn how to combine elementary skills/concepts ? (planning)
  ?? symbol anchoring

Reactive behaviors

Claims

- The world is its own model
- Perception-action loop
- Reaction – adaptivity

Types of reactive behaviors

- Collective
- Individual

Reactive collective behaviors



Reactive collective behaviors

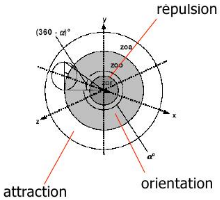

- Not too far from the group — safety
- Not too close — avoid crowding
- Same direction — cohesion

more ?

http://www.red3d.com/cwr/boids/

Intuition

- The noise in the environment
- $+$ the structure of reactions
- $\rightarrow$ emergence of a complex system.

Subsumption architecture

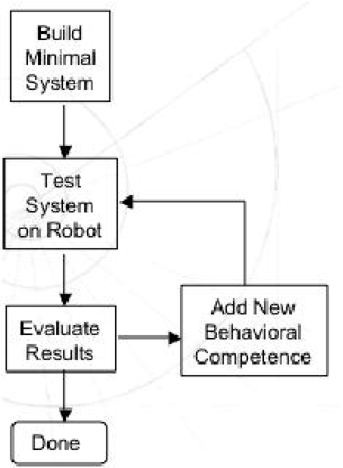

- Modular (∼ routines)
- Bottom-up

Subsumption architecture

Principle

- A finite-state machine
- Layer-wise architecture connecting sensors to motors
- Registers, timers, message sending

PROS

- Modularity (only perception required for the task is achieved)
- Testability hum.

CONS

- Scalability (few layers)
- Control (Action selection)

[same limitations as expert systems...]

Autonomous robotics

Autonomous navigation

Move (part of itself) throughout its operating environment without human assistance.

Interact and learn

Gain information about the environment.

Sustainability

Work for an extended period without human intervention.

Safety

Avoid situations that are harmful to people, property, or itself [unless those are part of its design specifications].

Three laws of Asimov

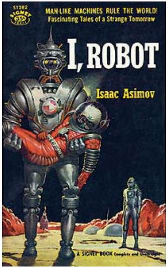

First law

A robot may not injure a human being or, through inaction, allow a human being to come to harm.

Second law

A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.

Third law

A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Overview

Reactive behaviors

Features

- No model of the world
- No reasoning (no planning, no action selection)
- Actuator values = F(sensor values)

Implementation

- Rules (if obstacle on right, go left)
- Built-in: software or hardware

Example: Braitenberg obstacle avoidance

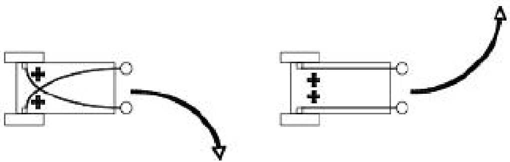

Light

Connexions excitatory, inhibitory

Examples

- Seeking/avoiding light
- Seeking/avoiding obstacles

Remarks

- Single behavior; robust behavior
- Can be misled for intelligence (finding the exit).

The Darpa Challenge

What

∗drive for 175 miles (trajectory known 2 hours before)
∗path defined by landmarks (no planification)
∗no crossing

Goal

∗going as fast as possible
∗avoid obstacles

The Darpa Challenge

Actions

- Direction
- Speed

State

- Position (uncertain)
- Speed
- Lasers, camera

Required

- Is a region navigable ?

Training a reactive controller

Acquiring a training set

1. State = vector of sensor values, camera image
2. States are labelled (region ahead drivable Yes/No)

Exploiting it to build a controller

- Train classifiers: action applicable in a state, yes/no.
- Simple controller (if action applicable, apply it)

Challenges

- From sensations to perceptions
- PERCEPTION biases (your brain constructs what you see)
- Variability

Lifelong learning

Detection from high-definition, low-range camera: accurate

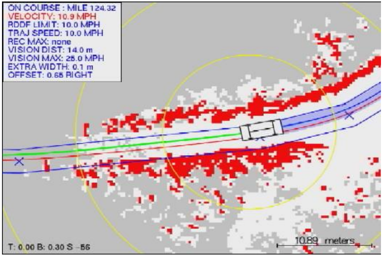

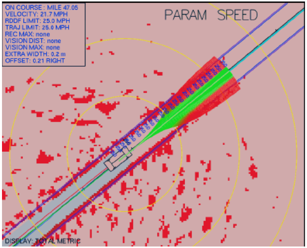

...used to label long-range sensor data

S. Thrun, Burgard and Fox 2005

more ?

http://sss.stanford.edu/coverage/powerpoints/sss-thrun.ppt

Vision

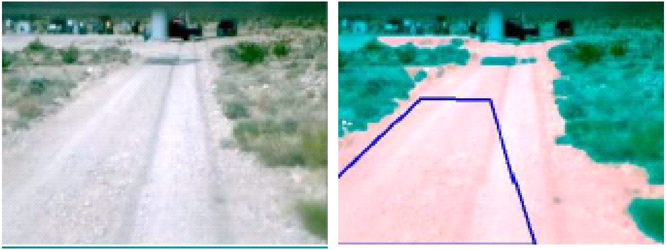

Online learning and Boostrap

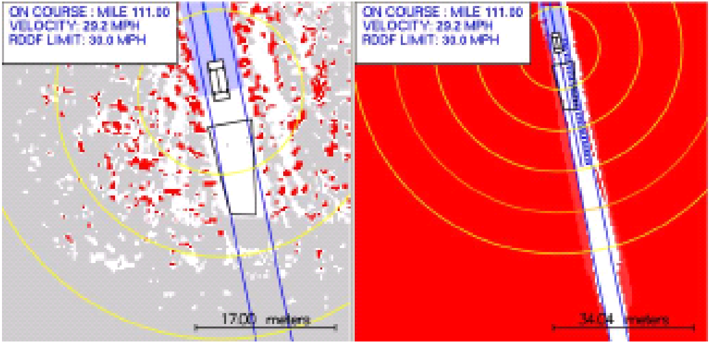

Fig. 2. Real-time generated map of the vehicle vicinity.
The left image shows close-range map as generated by the laser scanners. Red cells are occupied by obstacles, white cells are drivable and grey cells are (as yet) unknown. The black quadrangle is fit into the known empty area and shipped to the computer vision algorithm as training data for drivable road surface.
The right image shows a long-range map as generated by the computer-vision algorithm described in this paper. It can be seen that the road detection range is in this case about 70m, as opposed to only 22m using lasers.

Going fast !

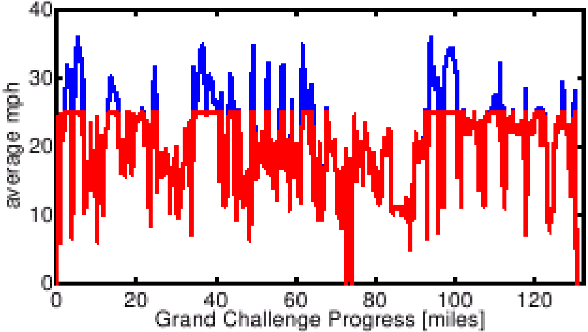

more ?

http://robots.stanford.edu/papers/dahlkamp.adaptvision06.pdf

Results

2004: max. distance travelled 12 miles

2005: 22 robots go farther !

- 5 participants reach the end ($4 < 10$ hours)

6h54 Stanley (Stanford, S. Thrun)

7h04 Sandstorm (CMU, R. Whitaker)

7h14 H1ghlander (Pennsylvania)

7h29 Kat-5 (New Orleans).

2007: Urban Challenge
Idem, + avoid other cars and driving rules.

The CMU revenge...

Follow-on

Google

- hires Sebastian Thrun and part of his team
- Google car appears in 2011
- massive use of Street View
- algorithms ??

Validation

- Safety, regulation
- 3 US states allow driverless cars (2011, 2012)

Complete Agent Principles

Rolf Pfeiffer, Josh Bongard, Max Lungarella,
Jurgen Schmidhuber, Luc Steels, Pierre-Yves Oudeyer...

Situated cognition

Intelligence: a means, not an end

brains are first and foremost control systems for embodied agents, and their most important job is to help such agents flourish.

The agent's goals

- Survival
- Individual priorities *autotelic*
- External duties *standard robotics*

Nouvelle nouvelle AI

Business as usual

- Decompose the problem in sub problems
- Solve them

Bounded rationality

> In complex real-world situations, optimization becomes approximate optimization since the description of the real world is radically simplified until reduced to a degree of complication that the decision maker can handle.
>
> Satisficing seeks simplification in a somewhat different direction, retaining more of the detail of the real-world situation, but settling for a satisfactory, rather than approximate best, decision.
>
> Herbert Simon, 1982

Complete Agent Principles

Rolf Pfeifer, Josh Bongard

more ?

How the Body Shapes the Way We Think: A New View of Intelligence, 07

http://www.agcognition.org/papers/anderson_review2.pdf

Design frame

1 Integrated design of the ecological niche, definition of the desired behaviors and tasks, and design of the agent.

6 There has to be a match between the complexities of the agent's sensory, motor, and neural systems.

The environment helps

2 When designing agents we must think about the complete agent behaving in the real world.

3 If agents are built to exploit the properties of the ecological niche and the characteristics of the interaction with the environment, their design and construction will be much easier, or cheaper.

5 Through sensory-motor coordination structured sensory stimulation is induced.

Complete Agent Principles

Working hypotheses

4 Redundancy : intelligent agents must be designed in such a way that (a) their different subsystems function on the basis of different physical processes and (b) there is partial overlap of functionality between the different subsystems.

7 Intelligence is emergent from a large number of parallel processes that are often coordinated through embodiment, in particular via the embodied interaction with the environment.

8 Intelligent agents are equipped with a value system which constitutes a basic set of assumptions about what is good for the agent.