LeMAS 2016

Bucuresti 21-25 iunie 2016

Multi-Agent Systems Introduction

Adina Magda Florea

AI-MAS Group

Universitatea Politehnica din Bucuresti UPB 1818

Motivations for agents

- Five ongoing trends have marked the history of computing:
  - *ubiquity*;
  - *interconnection*;
  - *intelligence*;
  - *delegation*; and
  - *human-orientation*

Motivations for agents

- Delegation and Intelligence imply the need to build computer systems that can act effectively on our behalf
- This implies:
  - The ability of computer systems to act *independently*
  - The ability of computer systems to act in a way that *represents our best interests* while interacting with other humans or systems

Motivations for agents

- The emergence of a new field in Computer Science
- *Multiagent systems*

Agent definition

- An agent is a computer system that is capable of *independent* action on behalf of its user or owner (figuring out what needs to be done to satisfy design objectives, rather than constantly being told)

Agent = a hardware or (more usually) a software-based computer system that enjoys the following properties:

- **autonomy** - agents operate without the direct intervention of humans or others, and have some kind of control over their actions and internal state;
- **reactivity**: agents perceive their environment and respond in a timely fashion to changes that occur in it;
- **pro-activeness**: agents do not simply act in response to their environment, they are able to exhibit goal-directed behaviour by taking initiative.
- **social ability** - agents interact with other agents (and possibly humans) via some kind of agent-communication language;

Agent definition

Two main streams of definitions

- Define an agent in isolation
- Define an agent in the context of a society of agents ➔ **social dimension** ➔ **MAS**

Two types of definitions

- Does not necessary incorporate intelligence
- Must incorporate a kind of IA behaviour ➔ **intelligent agents**

Agents characteristics

- act on behalf of a user
- be autonomous
- sense the environment and act upon it / reactivity
- purposeful action / pro-activity
- function continuously / persistent software
- mobility

Intelligence

- **Goals, rationality**
- **Reasoning, decision making**
- **Learning/adaptation**
- **Interaction with other agents - social dimension**

Example

- A key air-traffic control system…suddenly fails, leaving flights in the vicinity of the airport with no air-traffic control support. Fortunately, autonomous air-traffic control systems in nearby airports recognize the failure of their peer, and cooperate to track and deal with all affected flights.
- Systems taking the initiative when necessary
- Agents cooperating to solve problems beyond the capabilities of any individual agent

Another example

- An agent searching the internet to achieve a goal for the user.
- The agent would typically be given a task that would require synthesizing pieces of information from various different Internet information sources.
- The agent can cooperate with other agents
- The agent must plan, arrange, buy, negotiate – carry out arrangements of all sorts that would normally be done by its human user

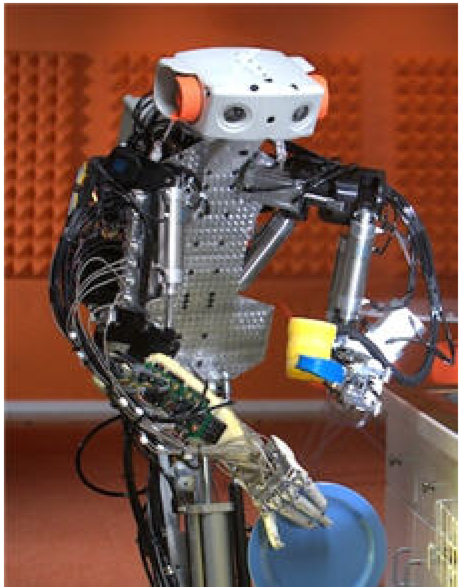

Intelligent agents?

Cognitive agents

The model of human intelligence and human perspective of the world ➔ characterise an intelligent agent using symbolic representations and ***mentalistic notions***:

- knowledge - John knows humans are mortal
- beliefs - John took his umbrella because he believed it was going to rain
- desires, goals - John wants to possess a PhD
- intentions - John intends to work hard in order to have a PhD

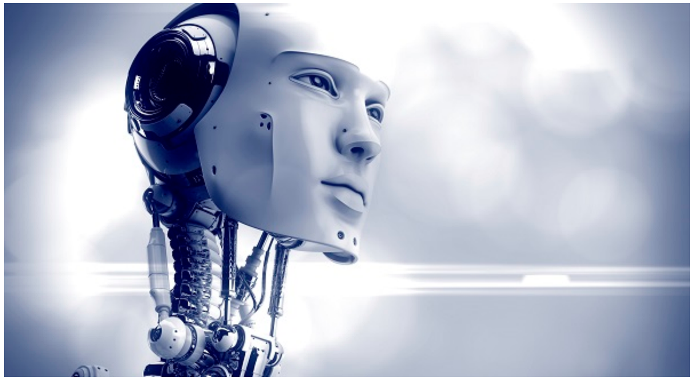

Reactive agents

- Simple processing units that perceive and react to changes in their environment.
- Do not have a symbolic representation of the world and do not use complex symbolic reasoning.
- Intelligence is not a property of the active entity but it is distributed in the system
- Steams as the result of the interaction between the many entities and the environment.

Some typical agent problems

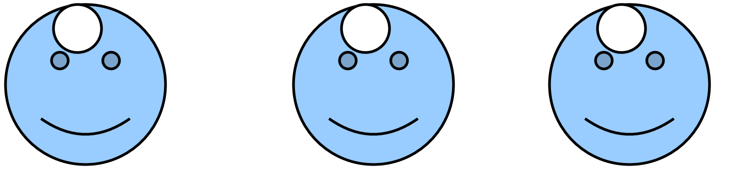

The wise men problem

A king wishing to know which of his three wise men is the wisest, paints a white spot on each of their foreheads, tells them at least one spot is white, and asks each to determine the color of his spot. After a while the smartest announces that his spot is white

The problem of pray and predators

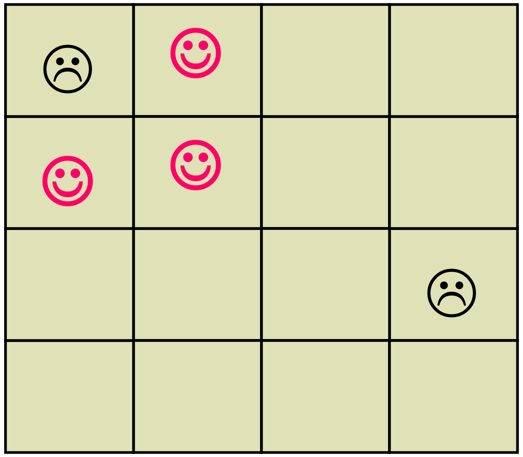

Cognitive agents

- detection of prey animals
- predator agents have goals
- coordination is distributed
- necessity for communication and for coordination

The problem of pray and predators

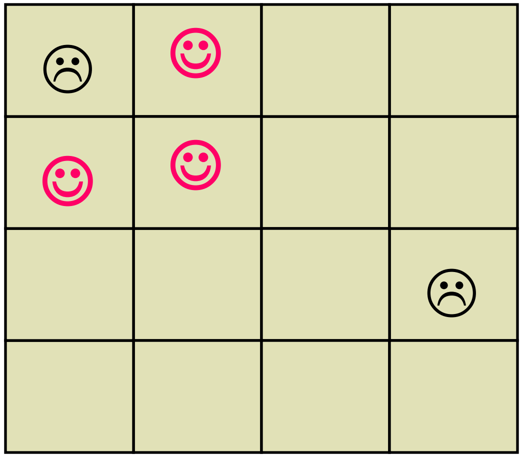

Reactive approach

- The preys emit a signal whose intensity decreases in proportion to distance - plays the role of attractor for the predators
- Hunters emit a signal which acts as a weak repellent for other hunters, so as not to find themselves at the same place
- Each hunter is attracted by the pray and (weakly) repelled by the other hunters

Prisoner’s dilemma

Tanya and Cinque have been arrested for robbing the Hibernia Savings Bank and placed in separate isolation cells. Both care much more about their personal freedom than about the welfare of their accomplice.

A clever prosecutor makes the following offer to each. “You may choose to confess or remain silent. If you confess and your accomplice remains silent I will drop all charges against you and use your testimony to ensure that your accomplice does serious time. Likewise, if your accomplice confesses while you remain silent, he will go free while you do the time. If you both confess I get two convictions, but I'll see to it that you both get early parole. If you both remain silent, I'll have to settle for a heavy charge. If you wish to confess, you must leave a note with the jailer before my return tomorrow morning.”

Utility theory

- A dominant approach to model the agent's interests is utility theory
- **Utility theory** = every state of the world has a degree of usefulness – **utility** or **payoff**, to an agent, and that agent will prefer states with higher utility
- **Decision theory** = an agent is rational if and only if it chooses the actions that yields the highest expected utility, averaged over all possible outcomes of actions

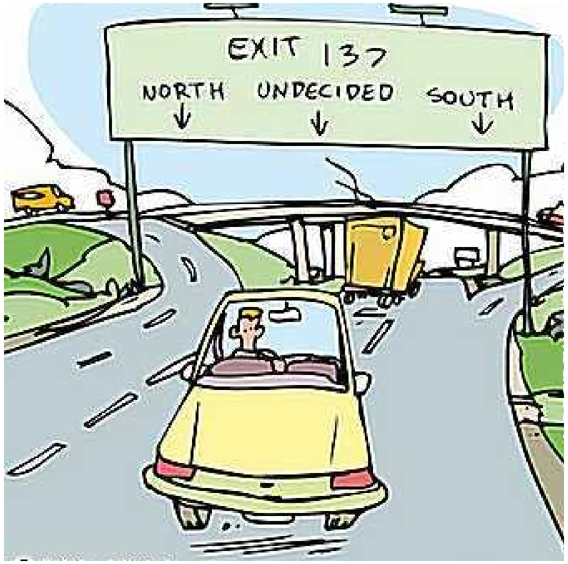

Example

- To show how utility theory functions can be used as a basis for making decisions

			Bob	Carol
Alice:	**Club (c)**	100	-90	1.5 factor
	Movie (m)	50	-40	1.5 factor
	Home (h)	50		

Bob prefers the Club 60% of time, 40% Movie

Carol – 25% Club, 75% Movie

Which is Alice best course of action?

We list Alice’s utilities for each possible state of the world

		Bob	Carol	
Alice:	**Club (c)**	100	-90	1.5 factor
	Movie (m)	50	-40	1.5 factor
	Home (h)	50		

A=c	B=c (60%)	B=m (40%)
C=c (25%)	15	150
C=m (75%)	10	100

A=m	B=c (60%)	B=m (40%)
C=c (25%)	50	10
C=m (75%)	75	15

U(A,c) = 0.25*(0.6*15+0.4*150)+0.75*(0.6*10+0.4*100) = 51.75

U(A,m) = 0.25*(0.6*50+0.4*75)+0.75*(0.6*75+0.4*15) = 46.75

U(A,h) = 50

Game theory

- Agents have utility functions they want to maximize
- As long as the outcomes and their probabilities are known to the agent, it can decide how to act optimally
- Agents need to choose the course of action that maximize the expected utility
- But **when there are 2 or more agents whose actions can affect each other's utility** we turn to **game theory**

Definition of games in normal form

- The **normal form** , also known as the strategic form, is the most familiar representation of strategic interactions in game theory.
- Represent every player’s utility for every state of the world in the special case when states depend only on the player’s combined actions
- Settings in which the states of the world depend also on the randomness of the environment are subject to Bayesian games – but they can be reduces to normal form games

Prisoner's dilema

		Column	player
		Cooperate	Defect
Row player	Cooperate	3, 3	0, 5
	Defect	5, 0	1, 1

- Cooperate (not confessing)
- Defect (confessing)

TCP user's game

- You and one of your colleagues are the only one using Internet
- Backoff mechanism of TCP
- You have 2 possible strategies: use Correct implementation or use a Defective implementation
- If both you and your colleague adopt C – average packet delay is 1ms
- If you both adopt D then the delay is 3ms
- If one adopts D and the other C then no delay for the first and 4ms delay for the second

		Column	**player**
		Correct	Defect
Row player	Correct	-1, -1	-4, 0
	Defect	0, -4	-3, -3

Common payoff games

Coordination game

		Column player	
		Left	Right
Row player	Left	1, 1	0, 0
	Right	0, 0	1, 1

- A **common payoff game** is a game in which for all actions profiles and for any pair of agents ***i***, ***j*** it is the case that $u_i(a)=u_j(a)$
- They are also called **pure coordination games** or **team games**
- Agents have no conflicting interests, their aim is to coordinate on an action that is maximally beneficial to all

Zero sum games

Matching pennies

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

- Properly called constant sum games
- A game is a **constant sum game** if there exists a constant c such that for each strategy profile $a \in A_1 \times A_2$ it is the case that $\mathbf{u_1(a)+u_2(a)=c}$
- If c=0 then **zero-sum games**
- Represent pure competition

Strategies in normal form games

- Agents are supposed to behave rationally
- **Rational behavior** = an agent prefers a greater utility (payoff) over a smaller one
- **Payoff maximization**: what to maximize? From what point of view?

Strategies in normal form games

Outside observer: can some outcomes be considered better than some others?

- **Social welfare**
  - The sum of agents' utilities (payoffs) in a given solution.
  - Measures the global good of the agents
  - Problem: how to compare utilities
- **Pareto optimal solutions**

From the individual agent's point of view

- **Nash equilibrium**

Pareto optimality

- A solution **u**, i.e., a **utility vector $\mathbf{u(a_1), \ldots, u(a_n)}$**, is **Pareto efficient**, i.e., **Pareto optimal**, if there is no other solution **u'** such that at least one agent is better off in **u'** than in **u** and no agent is worst off in **u'** than in **u**.
- Measures global good, does not require utility comparison

Nash equilibrium

Definition

- Two strategies, S_1 of agent A and S_2 of agent B are in a **Nash equilibrium** if:
  - in case agent A follows S_1 agent B can not do better than using S_2 **and**
  - in case agent B follows S_2 agent A can not do better than using S_1.

John Forbes Nash

Examples

Prisoners' Dilemma

		Column	**player**
		Cooperate	Defect
Row player	Cooperate	3, 3	0, 5
	Defect	5, 0	1, 1

Cooperate = not confessing
Defect = confessing

Coordination game (Common payoff)

		Column	**player**
		Left	Right
Row player	Left	1, 1	0, 0
	Right	0, 0	1, 1

Examples

Matching Pennies (Zero-sum)

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

BoS

	Football	Movie
Football	2, 1	0, 0
Movie	0, 0	1, 2

Conclusions

- There is a wealth of research topics in MAS
- MAS combines artificial intelligence, distributed systems, machine learning and many other subjects and disciplines

AI-MAS Group

"AIM AS high as you can"

Computer Science & Engineering Department

HOME | PROJECTS | PEOPLE | EVENTS | PUBLICATIONS | COOPERATION | COURSES | LINKS

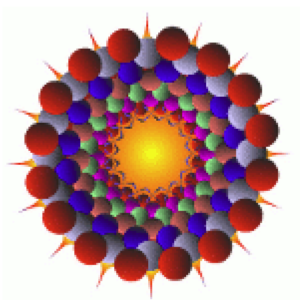

Highlights

[Master of Science in Artificial Intelligence]

[AI-MAS Winter Olympics]

[ERRIC FP7 Project]

Home

AIM

At the **Laboratory of Artificial Intelligence and Multi-Agent Systems** our research is focused on the development of models and architectures for single- and multi-agent systems endowed with artificially intelligent behavior, on the one hand, and on ambient intelligence, smart web services, and web agents, on the other hand.

We believe that techniques and technologies from the field of artificial intelligence can improve the experience of the user and the performance of distributed systems, learning platforms, web services and applications in general.

The Laboratory of Artificial Intelligence and Multi-Agent Systems is part of the Department of Computer Science, Faculty of Control and Computers at University Politehnica of Bucharest.
The AI-MAS Lab is directed by **Professor** [Adina Magda Florea].

RESEARCH

We investigate issues related to coordination mechanisms, automated negotiation, agent learning, and affective agents. Recent endeavors relate to self-organizing systems and swarm intelligence.

A key aspect of the research is the development of technologies that enable agent models to be used in different areas of applications, such as e-commerce, e-learning, supply chain management, and ambient intelligence.

In the vast field of Ambient Intelligence, we work in the application of multi-agent systems in AmI, posture, activity and habit recognition, Ambient Assisted Living, and automating the user's tasks.

Recently, we have started to investigate issues related to self-organization of complex systems and how agents can contribute to the development of such systems.

The connection between web services, semantic representation, and agents on the Web is also one of our concerns. We direct our efforts towards developing agents that offer semantic web services to support different aspects of business processing and enterprise interoperability.

AmILab

AI-MAS Brochure

AI-MAS Pictures

Search

search...

Contact

Splaiul Independentei
Nr. 313 Sector 6
Bucharest 06042
Romania

info@aimas.cs.pub.ro

Login Form

Username

Password

Login

Ongoing Research Projects and Student Positions

Internal projects

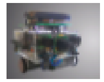

AIR Swarm
Ant Intelligent Robotic Swarm
https://www.youtube.com/channel/UCL6_X14h_mxyOQ05AJTlVhQ

AmIciTy
Ambient Intelligence for Collaborative Integration of Your Tasks
http://aimas.cs.pub.ro/amicity/

Mobile@Old
Assistant for elderly people based on mobility patterns
http://aimas.cs.pub.ro/mobile-old/

Nao Evolution
A tool abstraction framework for Nao
http://aimas.cs.pub.ro/nao-evolution/

- more information on the Projects page.

Wiki AI-MAS

We support **ARIA - Asociația Română pentru Inteligență Artificială** - The Romanian Association for Artificial Intelligence

The AI-MAS Laboratory is proud to pe part of IFAAMAS and COST Action IS0801 AGREEMENT TECHNOLOGIES.

IFAAMAS

AGREEMENT TECHNOLOGIES
COST Action IS0801