Multiagent Systems

Brief Description

Multi-agent systems have emerged as one of the most important areas of research and development in information technology. A multi-agent system is one composed of multiple interacting software components known as agents, which are typically capable of cooperating to solve problems that are beyond the abilities of any individual member. Multi-agent systems are important primarily because they have been found to have very wide applicability, in areas as diverse as industrial process control and electronic commerce. This module will begin by introducing the student to the notion of an agent, and will lead them to an understanding of what an agent is, how they can be constructed, how agents can be made to cooperate effectively with one-another to solve problems, and approaches to decision making in multiagent contexts.

Aims

The aims of this module are threefold:

1. to introduce the student to the concept of an agent and multi-agent system, and the main applications for which they are appropriate;
2. to introduce the main issues surrounding the design of intelligent autonomous agents, and the main approaches and techniques for the implementation of such agents
3. to introduce the main issues surrounding the design of multi-agent systems, and the main approaches and techniques for enabling communication and cooperation in such systems;
4. to introduce the main issues surrounding the design of a multi-agent society, and in particular, techniques for automated decision making in multi-agent contexts.

Learning Outcomes

Upon completing this module, a student will:

1. understand the notion of an agent, how agents are distinct from other software paradigms (e.g., objects), and understand the characteristics of applications that lend themselves to an agent-oriented solution;

2. understand the key issues associated with constructing agents capable of intelligent autonomous action, and the main approaches taken to developing such agents;

3. understand the key issues and approaches to high-level communication in multi-agent systems;

4. understand the key issues in designing societies of agents that can effectively cooperate in order to solve problems;

5. understand the main application areas of agent-based solutions;

6. understand the main techniques for automated decision-making in multi-agent systems, including techniques for voting, forming coalitions, allocating scarce resources, and bargaining.

Coursework

The course will be assessed through a written exam.

Outline Syllabus

1. *Introduction*:

   what is an agent: agents and objects; agents and expert systems; agents and distributed systems; typical application areas for agent systems.

2. *Intelligent Agents*:

   - abstract architectures for agents; tasks for agents.
   - the design of intelligent agents: reasoning agents (e.g., Agent0), agents as reactive systems (e.g., subsumption architecture); hybrid agents (e.g., PRS); layered agents (e.g., Interrap).

3. *Multiagent Systems*:

   - ontologies: OWL, KIF, RDF;

- interaction languages and protocols: speech acts, KQML/KIF, the FIPA framework;
- cooperation: cooperative distributed problem solving (CDPS), partial global planning; coherence and coordination;
- applications.

4. *Multiagent Decision-Making*:

   - multi-agent interactions: solution concepts; pure and mixed strategy Nash equlibria; Pareto efficiency; cooperative versus non-cooperative; zero-sum and other interactions; how cooperation occurs — the Prisoner's dilemma and Axelrod's experiments; program equilibria;
   - computational social choice: voting protocols; Arrow's theorem; Gibbard-Satterthwaite theorem; strategic manipulation and the role of computational complexity in preventing manipulation;
   - coalition formation: the core; the Shapley value; representation for coalitional games; coalition structure generation;
   - allocating scarce resources: auction types; auctions for individual items (English, Dutch, Vickrey auctions); combinatorial auctions; winner determination; bidding languages; the VCG mechanism;
   - bargaining: the alternating offers protocol; task-oriented negotiation; resource allocation via bargaining;
   - logical foundations of multi-agent systems: modal logics for epistemic reasoning, reasoning about mental state; cooperation logics; applications of such logics (verification/model checking, direct execution).

Recommended Reading

- M. Wooldridge, *An Introduction to MultiAgent Systems Second Edition.* John Wiley & Sons, 2009.
  ISBN-10: 0470519460
  ISBN-13: 978-0470519462

  (This is the set text, and the course will follow this material closely.)

Other reading:

- Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Gamer-Theoretic, and Logical Foundations*. Cambridge UP, 2008.
- G. Weiss, editor. *Multi-Agent Systems*. The MIT Press, 1999.
- M. Singh and M. Huhns. *Readings in Agents*. Morgan-Kaufmann Publishers, 1997.