Sequential Decision Making for Cooperative Agents

July 15 2014, 3:15PM – 7:15PM, TU6 - Room 2041

EASSS-14, Chania, Crete

Program for the Theoretical Section of the Tutorial

Part 0: Introductions and Overview of Tutorial

Part 1: An Introduction to Decision Theory
- What is planning, why planning, why decision-theoretic planning;
- States, actions, observations, rewards;
- Planning and Learning;
- An example MDP model abstracted from real environment;
- Value functions;
- Value iteration;
- Policy iteration;
- Partial observability and POMDPs;
- Introduction to belief states;
- Dynamic programming: piecewise linear value functions;
- Point-based methods;
- Factored Models.

Part 2: Multiagent Decision-Making under Uncertainty
- Decentralized actions;
- A straightforward, though nontrivial extension: the Multiagent MDP (MMDP);
- MPOMDP;
- Communication Reduction Methodologies;
- Decentralized observations (non-communicative);
- Introduction of the general formal model: Dec-POMDP.

Part 3: Dec-POMDP Planning Methods
- Dynamic Programming for Dec-POMDPs;
- Heuristic search for Dec-POMDPs;
- Solution methods using best-response computations;
- Graphical multiagent models and methods;
- A brief mention of more advanced topics (e.g., weakly-coupled structure, Dec-POMDP subclasses, communication, hierarchical abstraction, self-interested agents, learning).