TDDE13 - Multi-Agent Learning

Johan Källström, Industrial PhD Student at Saab and LiU/AIICS Johan.Kallstrom@liu.se





Multi-Agent Learning

Motivation and Learning Theories.





Motivation

- In multi-agent systems it may be beneficial to learn
 - Environment dynamics
 - Reward functions
 - Other agents' strategies
- Can be useful for cooperative as well as competitive scenarios
- Finding solutions by hand may be difficult and time-consuming
 - Learning may help
 - Also lets the agent adapt to changes in the environment





Challenges

- Non-stationary environment when multiple agents learn
 - Simultaneous learning and teaching for individual agent
- Lack of observability (e.g., actions and rewards of other agents)





Model-based and model-free learning

- Model-based algorithms
 - Try to model the behavior of other agents
 - More specific models may improve performance, but at a loss of generality
- Model-free algorithms
 - Do not use models
 - May perform worse or take longer to converge





Learning Theories

- Descriptive Theories
 - Model learning in real life
 - Human behavior
 - Animal behavior
- Prescriptive Theories
 - A study of how agents should learn to meet some performance requirements





Evaluation Metrics

- Descriptive Theories
 - Realism
 - There should be a good match between the learning theory and the realworld phenomenon
 - Convergence
 - The learning should converge to some solution concept of the game being played
 - The learning should converge to an "interesting state" that provides insight





Evaluation Metrics

- Prescriptive Theories
 - Learning performance
 - Opponent behavior affects performance, may be difficult to find a procedure that is optimal for all opponents
 - Evaluate against different types of agents
 - Evaluation approaches
 - Ask which learning rules are in equilibrium with each other
 - Ask whether a learning strategy achieves payoffs that are "high enough"





Evaluation Metrics

- Prescriptive Theories
 - Payoffs that are "high enough"
 - Safety: Should achieve maxmin payoff
 - Rationality: Should give best response to stationary strategy
 - No-regret: Should achieve payoff that is no less than pure strategy





www.liu.se

