ECLab - Summer Lecture Series, 2003

Lecture 2: Conducting Empirical Investigations in Multiagent Systems and Parallel Models

Lecturers:

Adrian Grajdeanu & Liviu Panait

Multiagent Systems

Distributed Artificial Intelligence

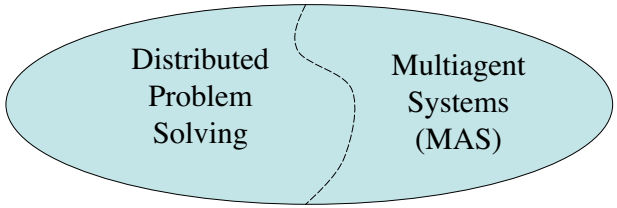

Agents: (Wooldridge and Jennings, '95)
- weak agents
  - autonomy
  - social ability
  - reactivity
  - pro-activeness
- strong agents
  - beliefs, desires, intentions
  - knowledge, commitments, trust

Multiagent Systems: (Jennings et al, '98)
- agents with incomplete information
- decentralized control
- decentralized data
- asynchronous computation

Multiagent Learning

- Learning = improving performance with accumulated experience, as indicated by a metric measure (Mitchell, '97)
- Multiagent Learning = improving the performance of individual agents or of teams of agents in a MAS setting
- We assume learning affects more than one agent
- Throughout the lecture, EC is THE learning technique

MAS Evolutionary Learning

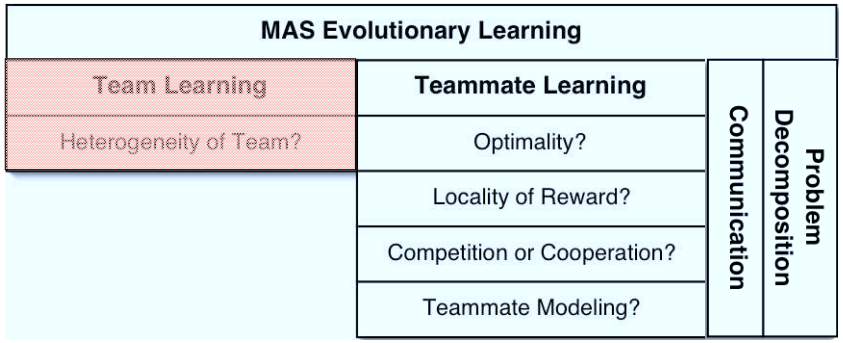

Team Learning

- EC particularly suited for team learning
- An individual codes for the behavior of an entire team
- Relatively similar to standard EC
- Team composition
  - domain specific (soccer)
  - scalable to larger teams (MAV)

Team Learning Approaches

- Homogeneous Team Learning
  - an individual contains a single behavior used for all agents
  - fast, scalable, possible suboptimal results
- Heterogeneous Team Learning
  - an individual contains one behavior for each of the agents
  - can potentially code for any homogeneous team
  - slower, potentially non-scalable, allows agent specialization
  - restricted inter-breeding may be better (Luke and Spector, '96)
- Hybrid Team Learning
  - an individual codes for team behaviors composed of heterogeneous groups of homogeneous behaviors
  - usually requires additional parameters for coding and manipulating the hybrid teams
  - breeding for hybrid teams?
  - learning team decomposition (Hara and Nagao, '99)

Heterogeneity and Performance

- Adding heterogeneity increases performance IF ENOUGH TIME IS AVAILABLE and
  - in domains that require task specialization (Balch, '98)
  - in inherently decomposable domains (Bongard, '00)
  - in domains that require increased number of different skills (Potter et al, '01)

Empirical Investigations in Team Learning

- Empirical investigations in team learning are very similar to those in standard EC
  - analysis of performance is straightforward
  - best-so-far curves
  - standard statistics/visualization tools
- How to measure scalability?
  - plot learning curves for different numbers of agents
  - plot performance versus number of agents
- How to measure heterogeneity?
  - problematic in GP
- How to measure and quantify the relation between heterogeneity and domain features?

Teammate Learning

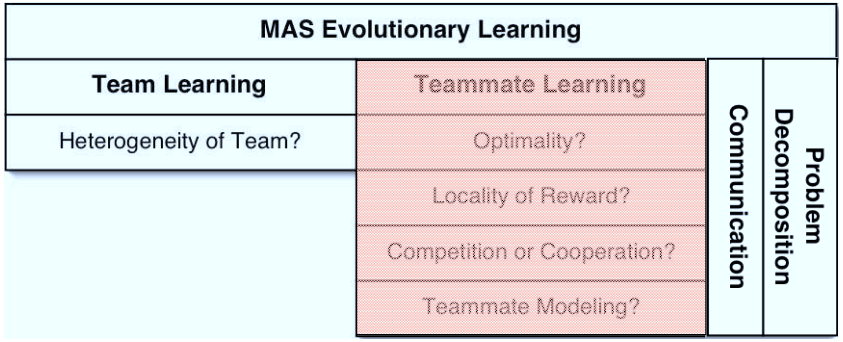

Teammate Learning

- Introduction
- Research Directions
- Issues
- Conduction Empirical Investigations

Teammate Learning

- Agents performing own learning processes
  - decentralized learning
  - closer to the concept of MAS
- Teammate learning better than team learning
  (Iba, '96, '98)
- Teammate learning worse than team learning
  (Miconi, '03)
- Theoretical comparisons
  (Jansen and Wiegand, '03)

Coevolution for MAS learning

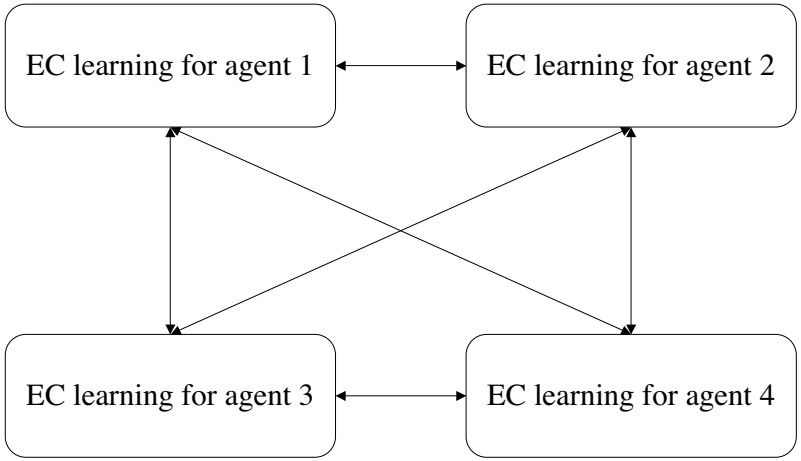

Teammate Learning

- Introduction
- Research Directions
- Issues
- Conduction Empirical Investigations

Teammate Learning

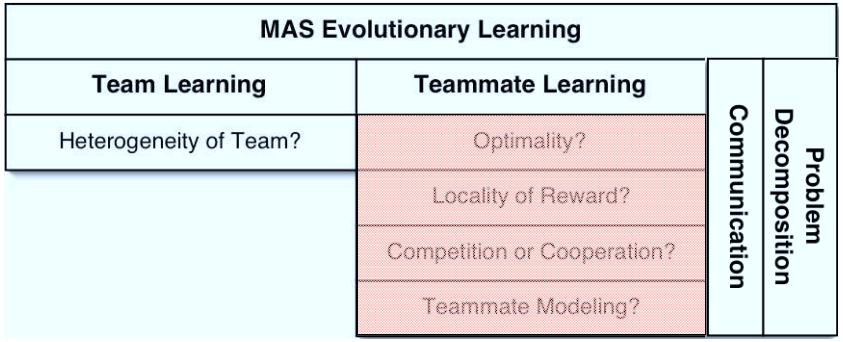

Optimality

- Search influenced by *performance* and *balance*
  (Panait et al, '03)
- Cooperative tasks with joint reward functions
  - standard algorithms not guaranteed to find optima, even with 'relaxed' settings
  - robustness of solutions?
    - better when teamed with optimal collaborator
    - better when teamed with many other collaborators
- Competitive tasks
  - what is optimal?
    - duel methodology
    - renaissance-man methodology
- Good news: fertile area for future research

Locality of Reward

- Influences performance
- Influences heterogeneity

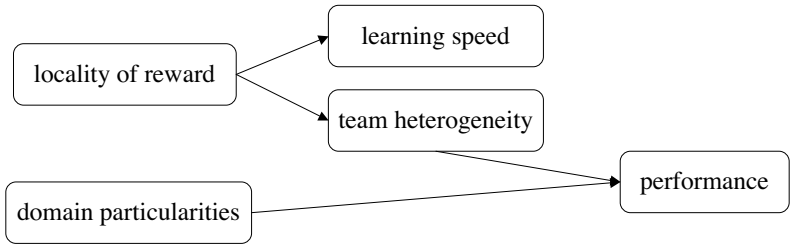

- Future research opportunities: automatic adjustment of locality

Cooperation or Competition

- No clear relation among agents, relations might change over time
- Learning opportunities
  - manipulation
  - exploitation of other agents' faults
  - mutual trust
  - reciprocity

Teammate Modeling

- Recursive modeling
- Flavors
  - single focus of learning
  - modeling combined with learning
- Initial beliefs are VERY important

Teammate Learning

- Introduction
- Research Directions
- Issues
- Conduction Empirical Investigations

Issues in Team Learning

- Search Space
- Red Queen Effect
- Exploration
- Credit Assignment
- Learning Cycles
- Loss of Gradient
- Mediocre Stability

Search Space for Team Learning

Agent 1

Agent 2	a_1	a_2	a_3
b_1	R_{11}	R_{12}	R_{13}
b_2	R_{21}	R_{22}	R_{23}
b_3	R_{31}	R_{32}	R_{33}

Search Space for Teammate Learning

Agent 1

a_1	a_2	a_3
R_{a1}	R_{a2}	R_{a3}

Agent 2

b_1	b_2	b_3
R_{b1}	R_{b2}	R_{b3}

Search Space for Teammate Learning

Agent 1

a_1	a_2	a_3
R^0_{a1}	R^0_{a2}	R^0_{a2}
R^1_{a1}	R^1_{a2}	R^1_{a2}
R^2_{a1}	R^2_{a2}	R^2_{a2}
R^3_{a1}	R^3_{a2}	R^3_{a2}

Agent 2

b_1	b_2	b_3
R^0_{b1}	R^0_{b2}	R^0_{b2}
R^1_{b1}	R^1_{b2}	R^1_{b2}
R^2_{b1}	R^2_{b2}	R^2_{b2}
R^3_{b1}	R^3_{b2}	R^3_{b2}

Search Space for Teammate Learning

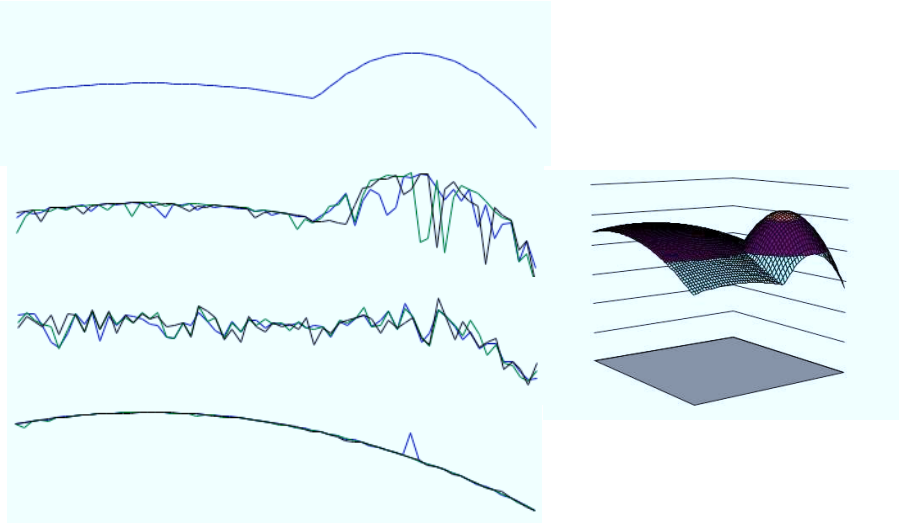

Red Queen Effect

- "Change in a moving landscape may go unnoticed"
- Individuals are evaluated in the context of other individuals
- Subjective performance metrics may hide progress, stagnation, or learning cycles

Exploration

- An agent's exploration process affects the learning processes of other agents, with later repercussions on the agent's learning process
- Similar to an agent learning in a dynamic environment, where the dynamicity is directly related to the agent's behavior

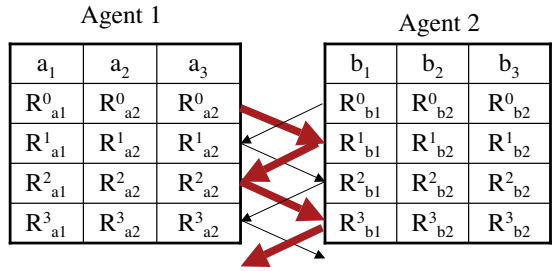

Credit Assignment

- Inter- and intra- agent credit assignment
- Individual reinforcement information may influence agents to learn greedy strategies focused on individual, rather than team, performance

Learning Cycles

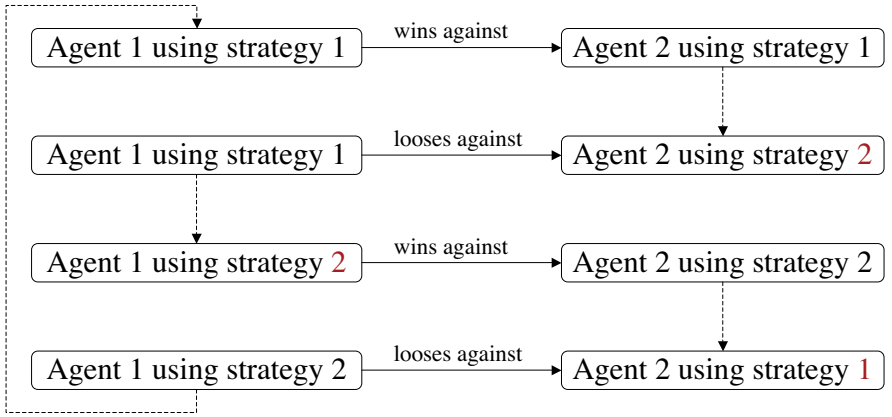

Loss of Gradient

Mediocre Stability

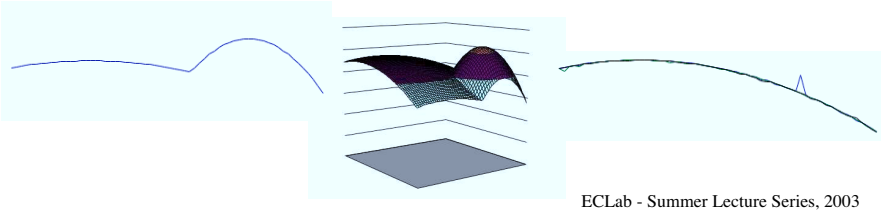

Teammate Learning

- Introduction
- Research Directions
- Issues
- Conduction Empirical Investigations

Empirical Investigations in Teammate Learning

- What is being measured?
- Performance is subjective (Red Queen Effect)
  - possible solutions
    - choose domains where objective performance measure is available (Panait and Luke, '02), (Bucci and Pollack, '03)
    - use benchmarks
    - dominance tournament (Stanley and Miikkulainen, '02)
    - hall of fame? (Rosin and Belew, '97)
  - measure for team heterogeneity?
  - measure for sizes of basins of attractions?
- What is meant by 'better' or 'best'?
  - (Panait and Luke, '02)
    - duel methodology
    - renaissance-man methodology

Empirical Investigations in Teammate Learning

- What are the assumptions of the experiments?
  - global information does not guarantee optimality
  - recommendations to restrict assumptions about other agents when their behaviors are unknown
  - coevolution may be improved when assuming other agents are competing or cooperating
- How to select problem domains?
  - "my method is better than your method" stage of investigation
  - for theoretical analysis, use very simple domains (game matrixes)
  - pay attention to assumptions

Empirical Investigations in Teammate Learning

- Visualization
  - visualization needs to capture the relation among different coevolutionary algorithms
  - plot the trajectories of the search process
  - search driven by *balance* and *performance*
    - visualization of search space: basins of attraction?
    - assess difficulty of domain based on sizes of basins of attraction for suboptimal peaks
- Statistical methods
  - because performance assessment is subjective, the results of statistical tests will depend on the other 'components'
  - co-adaptation and learning cycles
    - time may be an especially important characteristic
  - assess performance based on final results for all agents

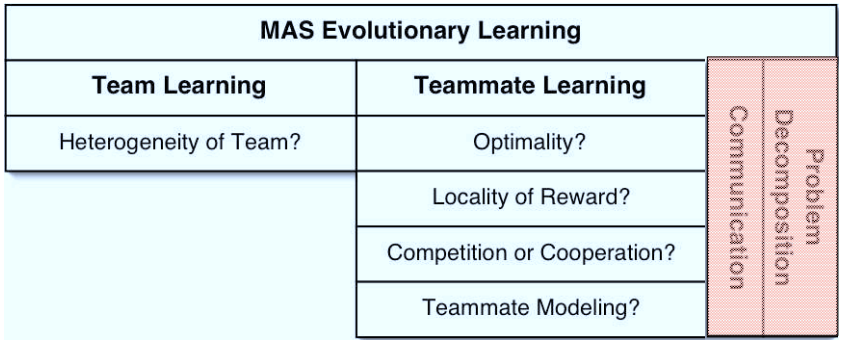

Problem Decomposition

- Flavors
  - task decomposition
  - behavior decomposition
  - layered learning
  - shaping
- Questions:
  - automatic problem decomposition
  - decentralized problem decomposition

Communication

- MAS + unrestricted communication = centralized system (Stone and Veloso, '00)

- Via rapidly decaying information
  - may increase the search space
  - may improve performance
  - emergent vocabularies
- Via slowly decaying information (example: pheromones)
  - long-lasting shared information
- Via embodiment

Empirical Investigations and Communication

- What are the ***assumptions***?
- Additional parameters to tune
  - range, bandwidth
  - evaporation and diffusion rates
  - communication topologies
- How to measure relation between learning algorithm and communication?
- Emergent vocabularies?
- Test communication via embodiment?

Conclusions

- Empirical investigations in team learning
  - pretty much straightforward
  - analysis of heterogeneity and scalability
- Empirical investigations in teammate learning
  - subjective evaluation → no clear performance criteria
  - visualize and measure balance and its relation to performance as the components driving the search process
  - assumptions about other agents are very important
- Empirical investigations and problem decomposition
  - representations
- Empirical investigations and communication
  - assumptions
  - test of emergent vocabularies
  - test of communication via embodiment

Conclusions

M. Wooldridge and N. Jennings, (1995), "Intelligent Agents: Theory and Practice", *The Knowledge Engineering Review*, 10(2):115-152

N.R. Jennings, K. Sycara, and M. Wooldridge, (1998), "A Roadmap of Agents Research and Development", *Autonomous Agents and Multi-Agent Systems*, 1:7-38

T. Mitchell, (1997), *Machine Learning*, McGraw-Hill

S. Luke and L. Spector, (1996), "Evolving Teamwork and Coordination with Genetic Programming", In *Genetic Programming 1996: Proceedings of the First Annual Conference*, MIT Press

A. Hara and T. Nagao, (1999), "Emergence of cooperative behavior using ADG; Automaticaly Defined Groups", In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*

T. Balch, (1998), *Behavioral Diversity in Learning Robot Teams*, PhD Thesis, College of Computing, Georgia Institute of Technology

J. Bongard, (2000), "The Legion System: A Novel Approach to Evolving Heterogeneity for Collective Problem Solving", In *Genetic Programming: Proceedings of EuroGP'2000*, Springer-Verlag

M. Potter, L. Meeden, and A. Schultz, (2001), "Heterogeneity in the Coevolved Behaviors of Mobile Robots: The Emergence of Specialists", In *Proceedings of The Seventeenth International Conference on Artificial Intelligence (IJCAI-2001)*

H. Iba, (1996), "Emergent Cooperation for Multiple Agents Using Genetic Programming", In *Parallel Problem Solving from Nature IV: Proceedings of the International Conference on Evolutionary Computation*

H. Iba, (1998), "Evolutionary Learning of Communicating Agents", *Information Sciences*, Volume 108

T. Miconi, (2003), "When Evolving Populations is Better than Coevolving Individuals: The Blind Mice Problem", In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*

T. Jansen and R. P. Wiegand, (2003), "Exploring the Explorative Advantage of the Cooperative Coevolutionary (1+1) EA", In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2003)*

L. Panait, R. P. Wiegand, and S. Luke, (2003), "Improving Coevolutionary Search for Optimal Multiagent Behaviors", In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-2003)*

L. Panait and S. Luke, (2002), "A Comparison of Two Competitive Fitness Functions", In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2002)*

A. Bucci and J. Pollack, (2003), "Focusing versus Intransitivity: Geometrical Aspects of Co-evolution", In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2003)*

K. O. Stanley and R. Miikkulainen, (2002), "The Dominance Tournament Method of Monitoring Progress in Coevolution", In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2002) Workshop Program*

C. Rosin and R. Belew, (1997), "New Methods for Competitive Coevolution", *Evolutionary Computation Journal*, 5(1):1-29

P. Stone and M. Veloso, (2000), "Multiagent Systems: A Survey from a Machine Learning Perspective", *Autonomous Robots*, 8(3):345-383