Multiagent Systems: A Survey from a Machine Learning Perspective

Peter Stone

An RL and Friends Chat
May 1st, 1996

Taxonomy

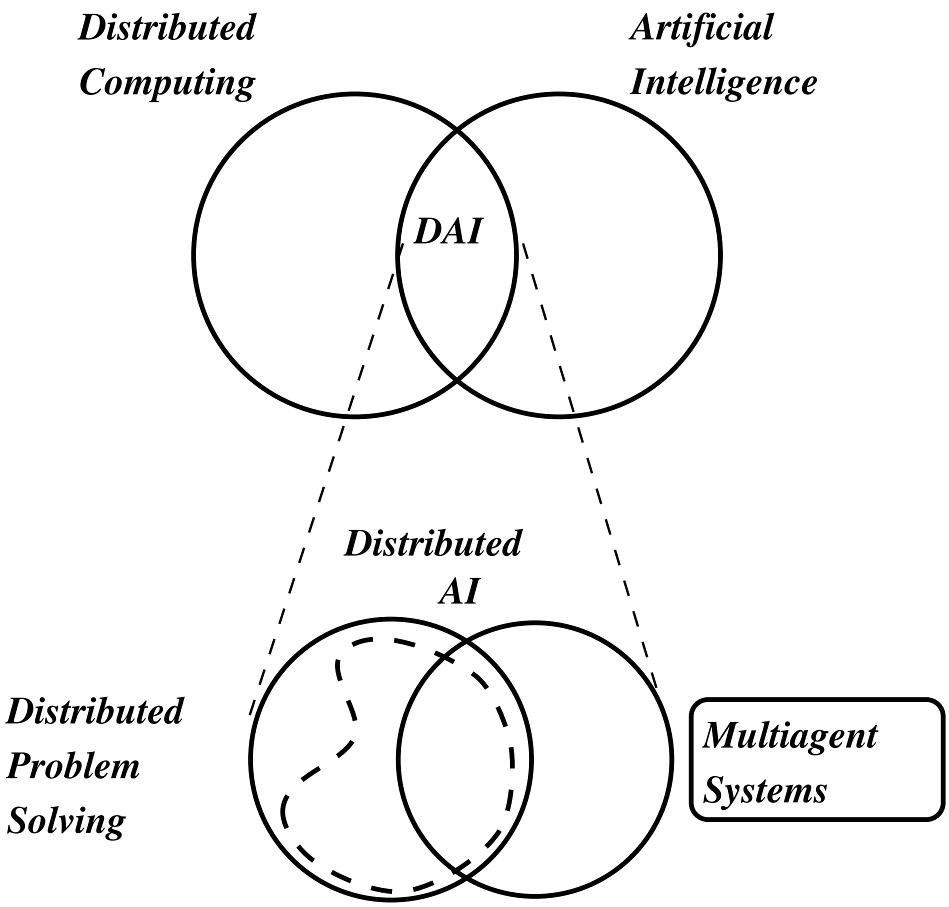

Distributed Artificial Intelligence (DAI) is concerned with the study and design of systems consisting of several interacting entities which are logically and often spatially distributed and in some sense can be called autonomous and intelligent. *Gerhard Weiss*

Definitions

Distributed Computing : Several processors share data, but not control. Focus on low-level parallelization or synchronization issues.

Distributed AI : Control as well as data is distributed. Focus on problem solving, communication, and coordination.

Distributed Problem Solving : Task decomposition (task sharing) and/or solution synthesis (result sharing): information management.

Multiagent Systems : Behavior coordination or behavior management.

- No necessary guarantees about other agents.
- Individual behaviors typically simple relative to interaction issues.

Why Multiagent Systems?

- Some domains require it. (Hospital scheduling)
- Parallelism.
- Robustness.
- Scalability
- Simpler programming.
- “Intelligence is deeply and inevitably coupled with interaction.” – *Gerhard Weiss*

Goals Of This Talk

$\boxed{\textbf{Survey Multiagent Systems}}$

- Present increasingly complex multiagent scenarios (along different dimensions).
  - Degree of agent heterogeneity.
  - Degree of communication.
- Catalog **issues** arising in these scenarios.
- Sample existing **techniques** addressing these issues (ML bias).
- Present some **ML opportunities**.
- Illustrate scenarios with an **example domain**.

The Predator-Prey ("Pursuit") Domain

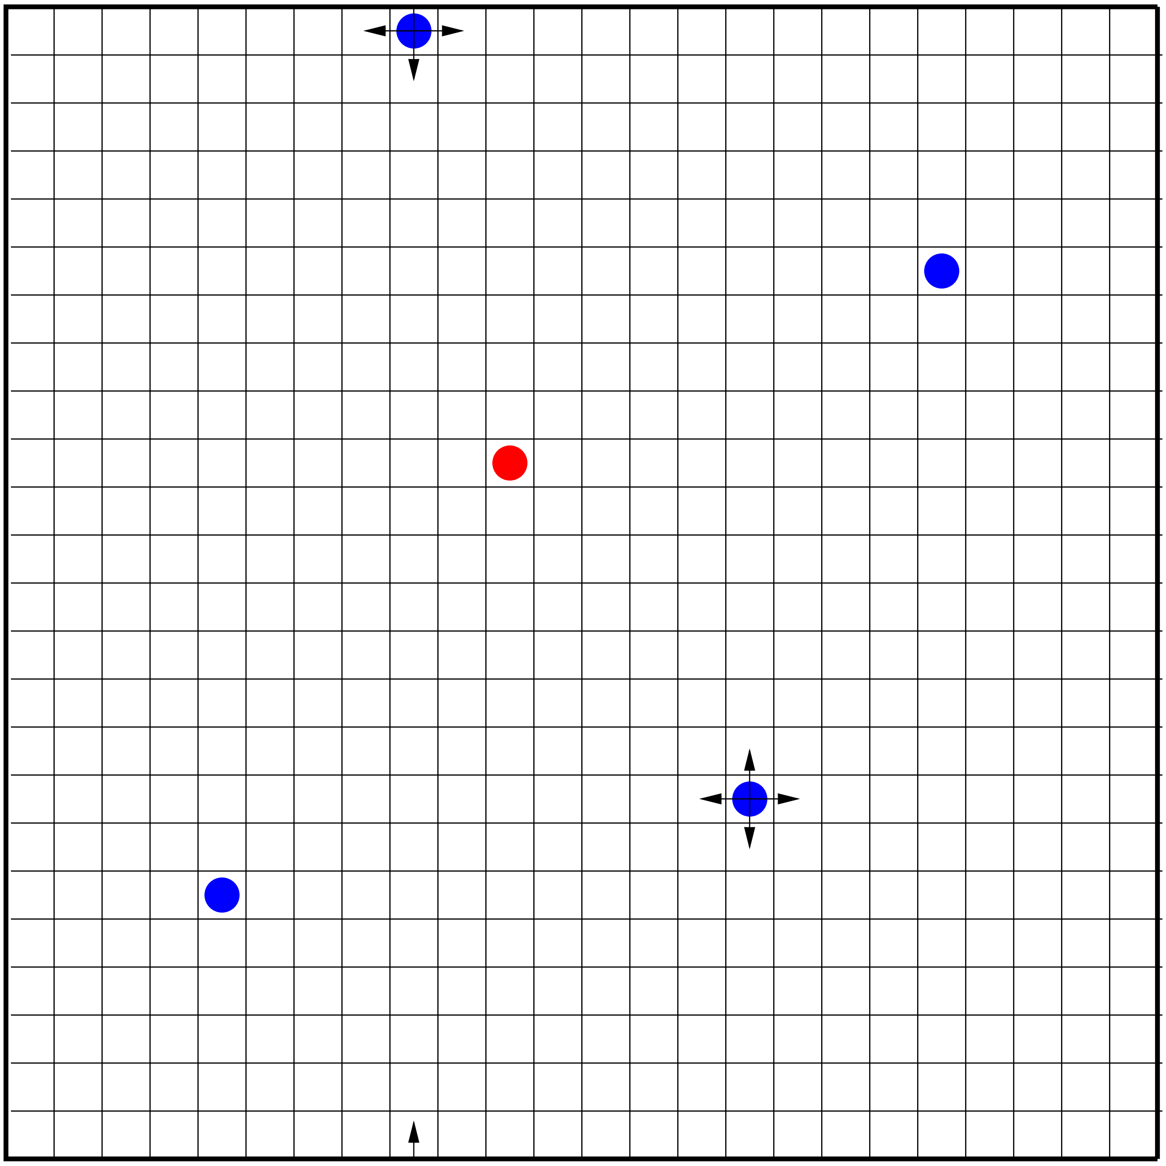

Orthogonal Game in a Toroidal World

Capture

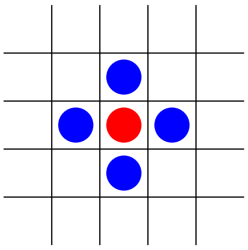

- Predators see each other
- Predators can communicate
- Prey moves randomly
- Prey stays put 10% of time
- Simultaneous movements

Single Agent Pursuit

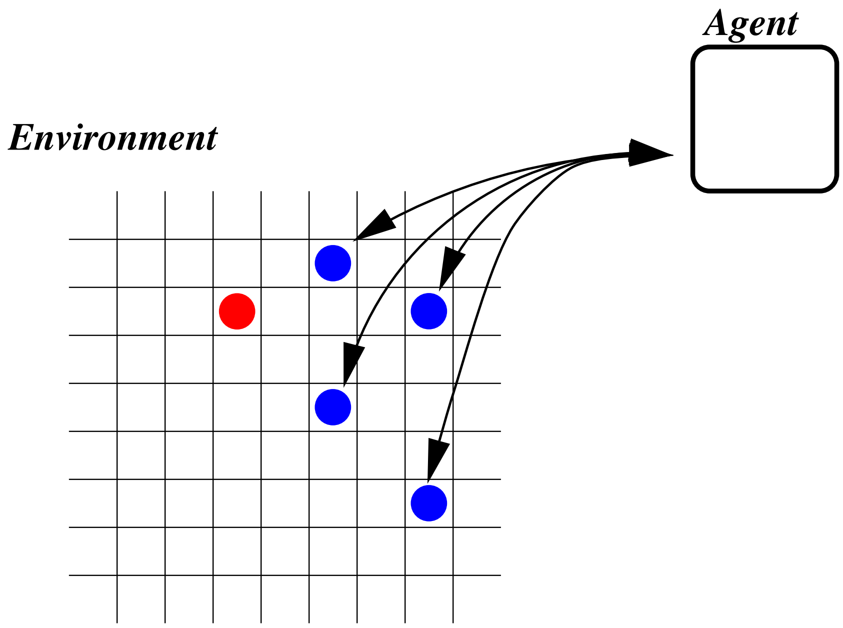

- One agent controls all predators.
- Prey is considered part of the environment.

Single Agent Systems

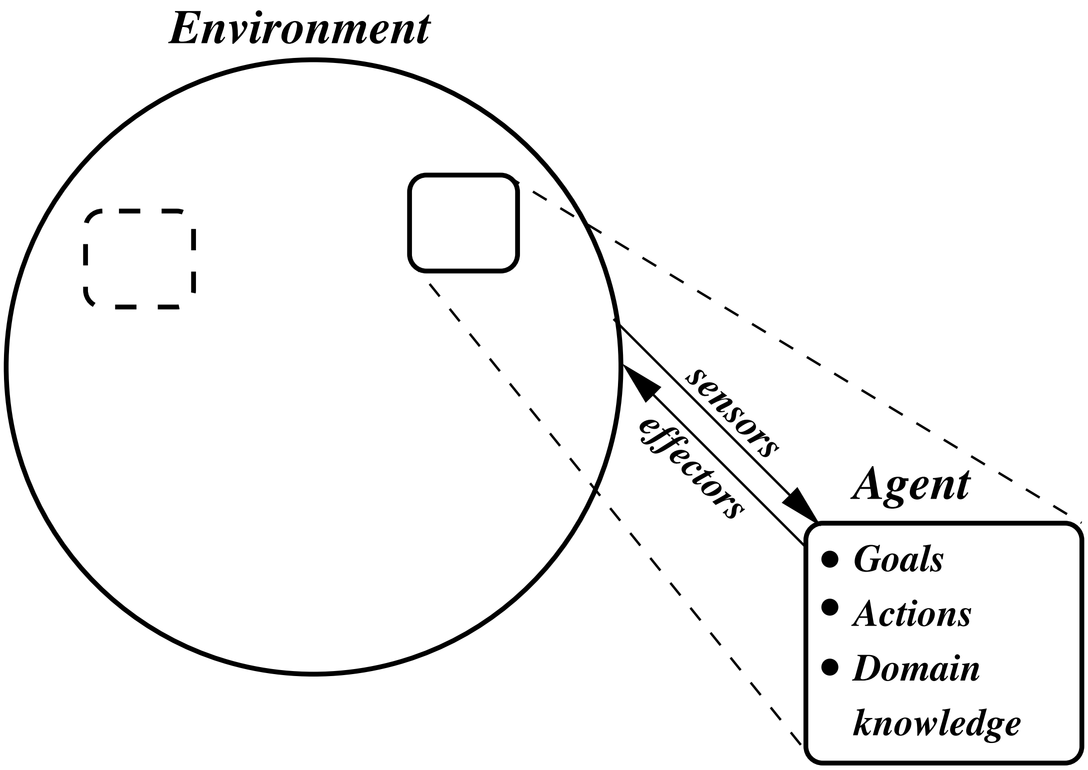

- Agent models itself, the environment, and their interactions.
- If other agents exist, they're considered environment.

Multiagent Systems

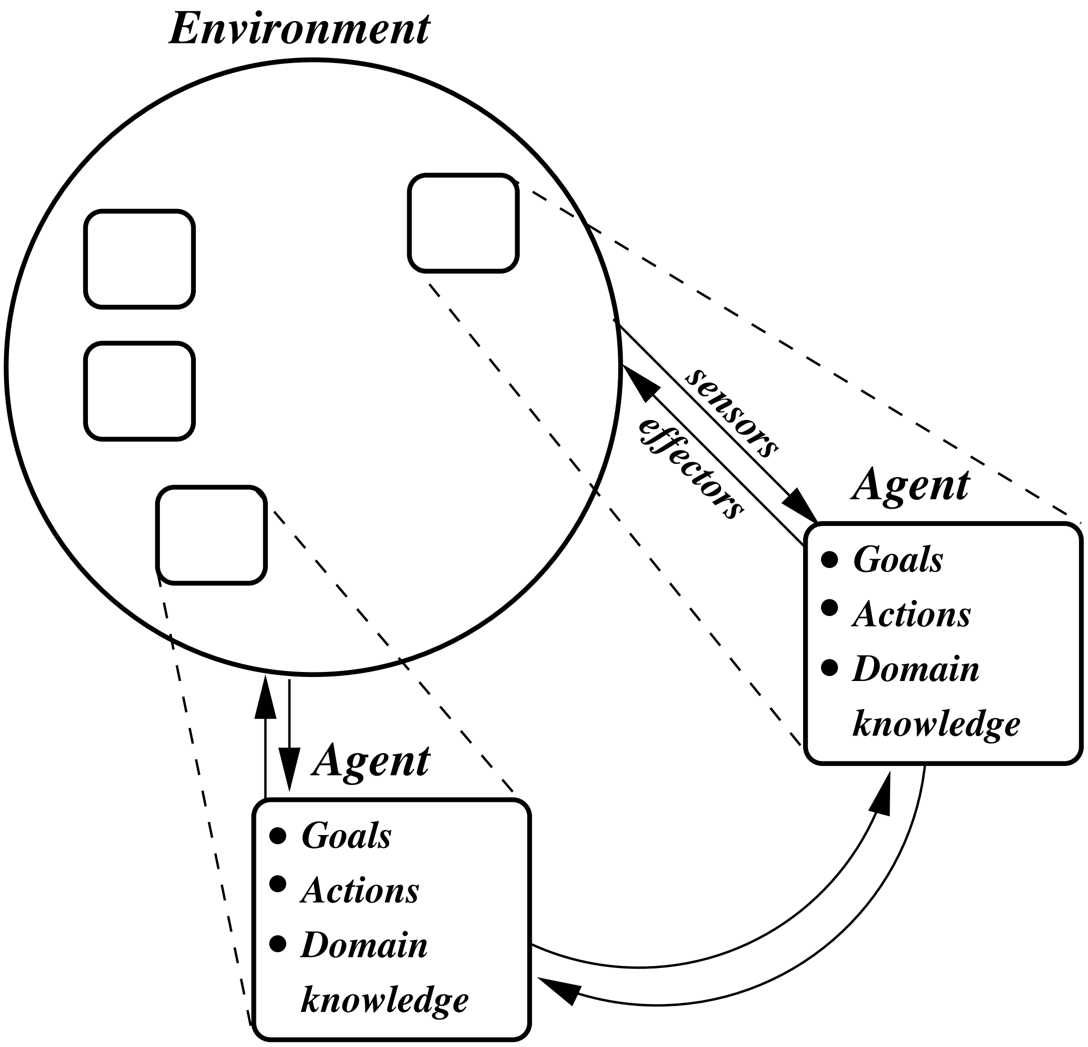

- Agents model each others' goals and actions.
- Possible direct interaction.
- Dimensions: degree of agent **heterogeneity**; degree of **collaboration**; **communication** capabilities.

Homogeneous Multiagent Pursuit

- One **identical** agent per predator.
- Possibly limited information about other agents' internal states.

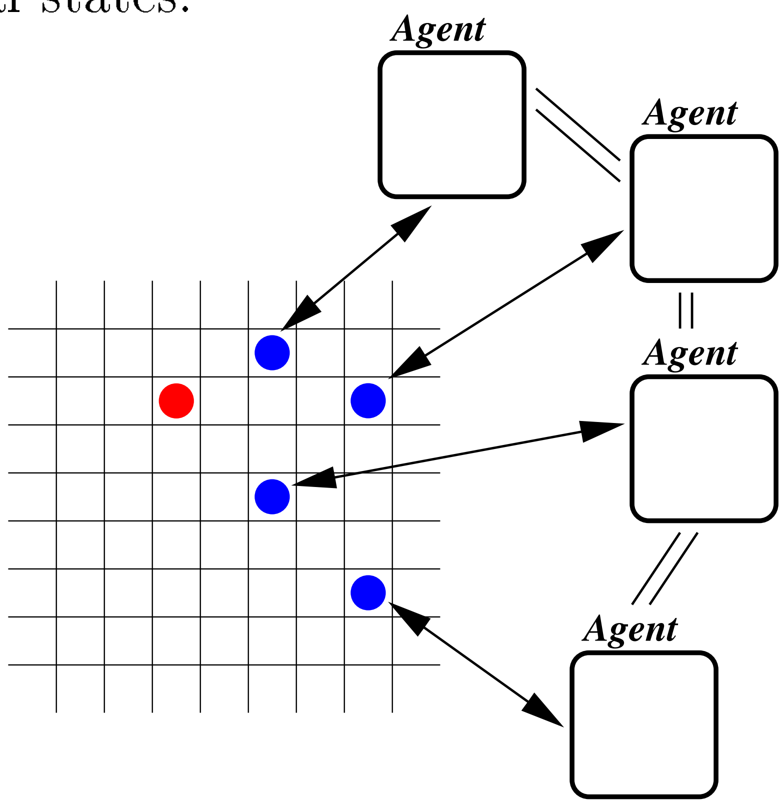

- Each agent to closest capture position. *Stephens and Merx*
- Predict other agents' moves using **Recursive Modeling**. *Vidal and Durfee*
- **Game theory**: each agent maximizing own utility leads to coordination. *Levy*
- Claim that game is easily solved with greedy decisions (diagonal, hexagonal games): attract to prey, repel other predators. *Korf*

Homogeneous Agents

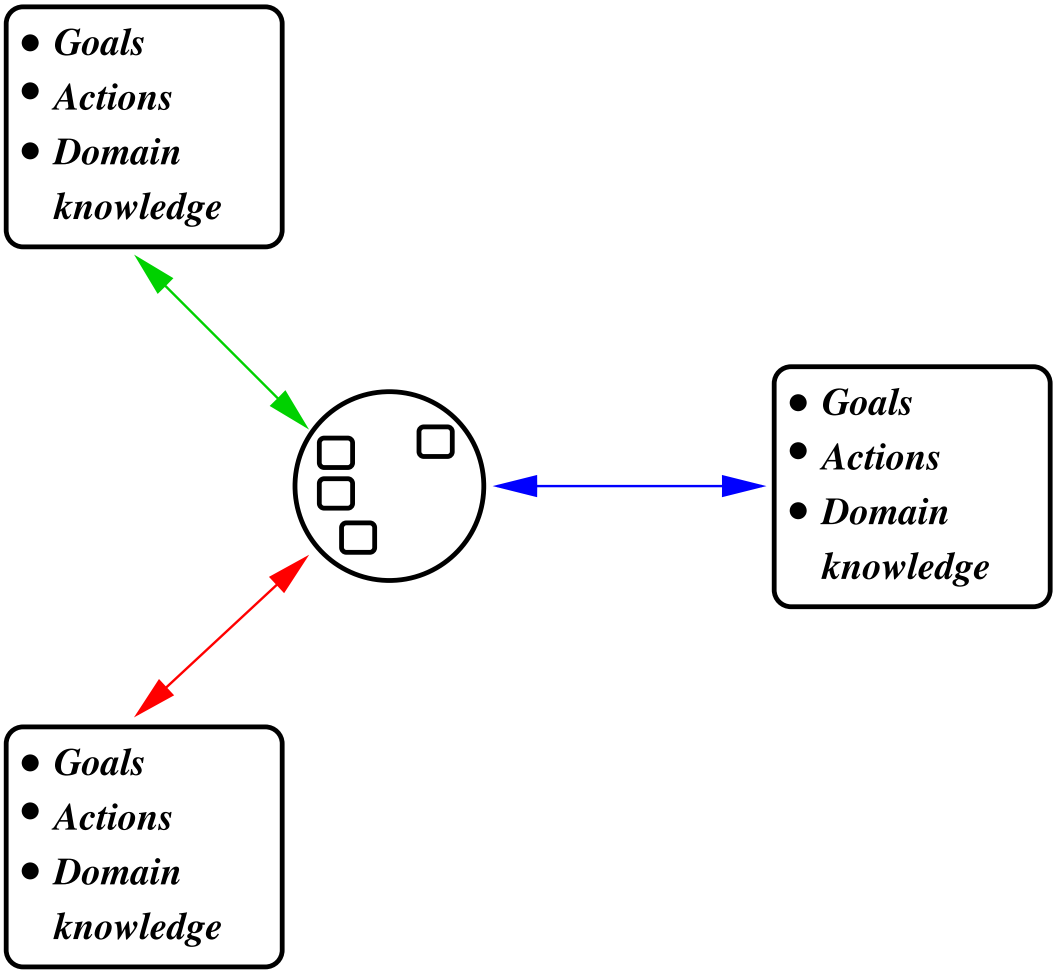

- Only sensors and effectors differ.
- Necessary condition.
- For now assume no direct interaction.

Homogeneous Agents

Issues

- Reactive vs. deliberative agents.
- Local or global perspective.
- Modeling of other agents' states.
- Dynamic environment.
- Knowledge of team?
- How affect others?

Techniques

- Reactive Behaviors for Formation maintenance. *Balch*
- (limited) Recursive Modeling Method (RMM). *Durfee*
- Stigmergy. *Holland*
- Q-learning for behaviors like foraging, homing, etc. *Mataric*

Learning opportunities

- Enable others' actions
- Sensor data $\rightarrow$ Other agent's sensor data

Heterogeneous Multiagent Pursuit

- Prey also modeled as an agent.
- Goals and actions differ among agents.

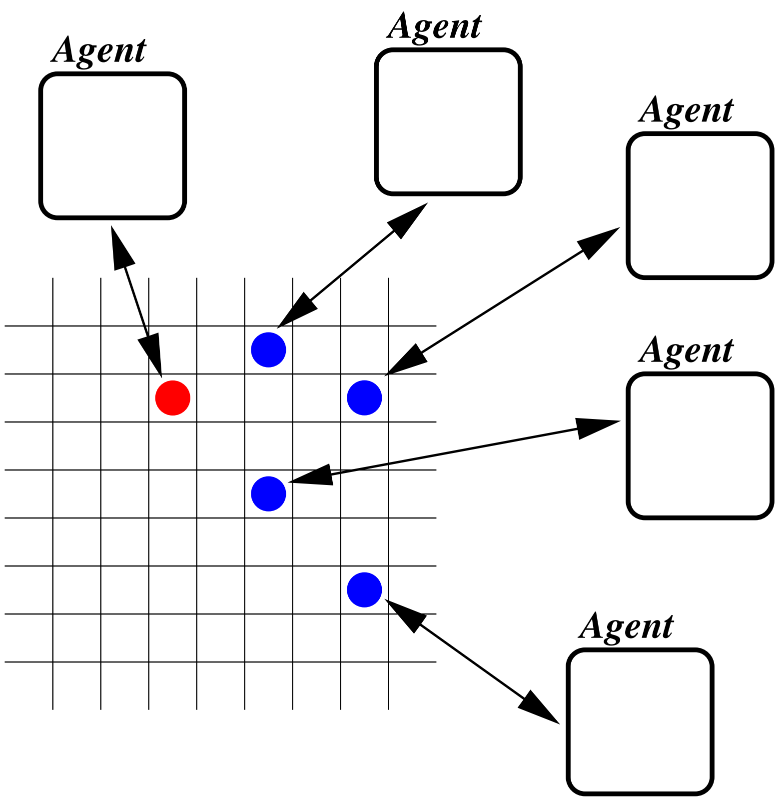

- Evolve teams of predators: implicit communication. *Haynes and Sen*
- Model others by set of default rules, then store negative cases and update behaviors (no co-occupation). *Haynes and Sen*
- **Evolve prey behaviors** (linear always succeeds?). *Haynes and Sen*

Heterogeneous Agents

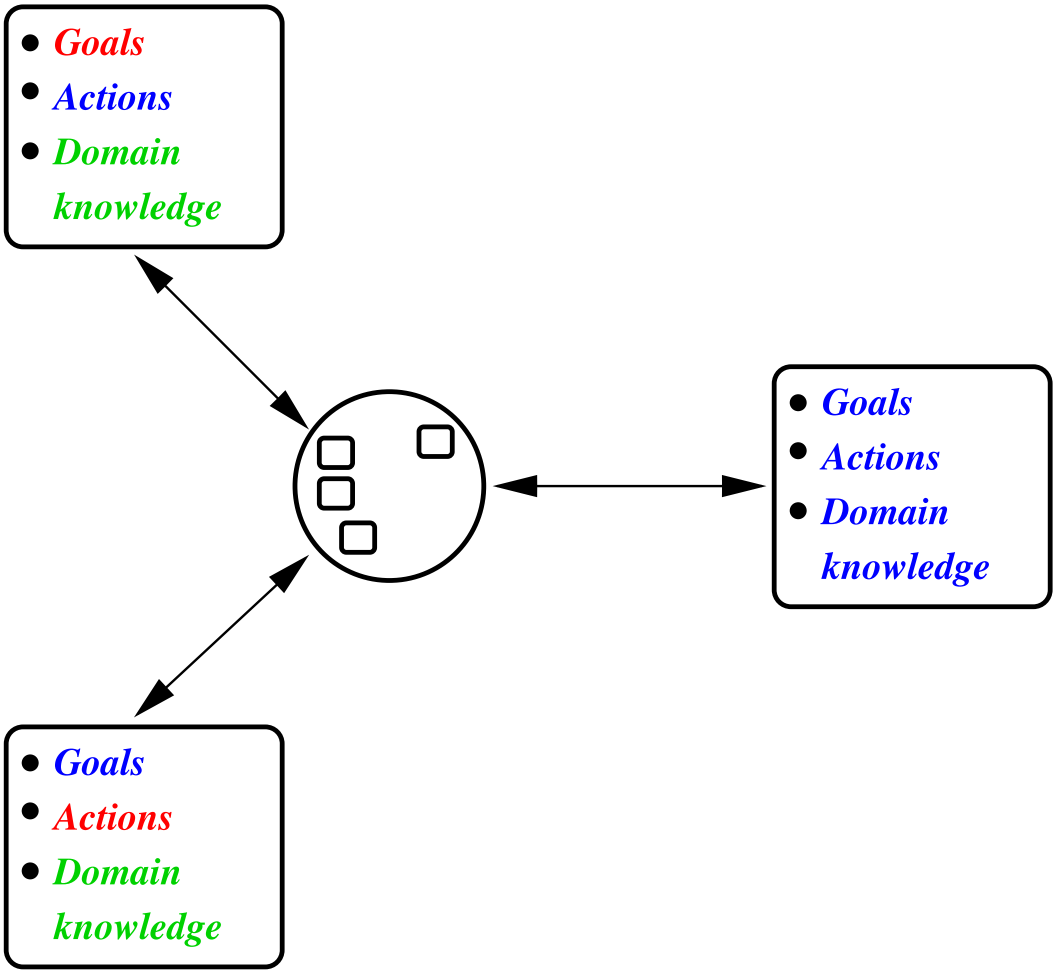

- Goals, actions, and/or domain knowledge may differ.
- Still assume no direct interaction.

Heterogeneous Agents

Issues

- Benevolence vs. competitiveness.
- Stable vs. evolving agents ('arms race').
- Modeling of others' goals, actions, and knowledge.
- Resource management (interdependent actions).
- Social conventions.
- Roles.

Techniques

- Game theory, iterative play. *Rosenschein*
- Minimax-Q. *Littman*
- Competitive co-evolution. *Haynes and Sen/Grefenstette/Rosin*
- Belief/Desires/Intentions (BDI) model: OASIS. *Rao and Georgeff*
- Deduce intentions through observation. *Huber and Durfee*
- Autoepistemic reasoning (ignorance). *Permpoontanalarp*
- Model as a team (individual $\rightarrow$ role). *Tambe*
- Social reasoning: depend on others for goal ($\neq$ game theory). *Sichman and Demazeau*
- GAs to deal with Braes' paradox (more resource $\rightarrow$ worse). *Arora and Sen*
- Local resource decisions (blissful ignorance). *Lesser*
- Multiagent RL for adaptive Load Balancing. *Schaerf, Shoham, and Tennenholtz*
- Focal points/Emergent conventions. *Rosenschein/Woolridge*
- Design agents play different roles. *Lesser*

Learning opportunities

- Credit/blame in competitive scenarios.
- Behaviors that blend well with team.
- Prediction of others' actions.
- Dynamic role assumption.

Communicating Multiagent Pursuit

- Predators can communicate.

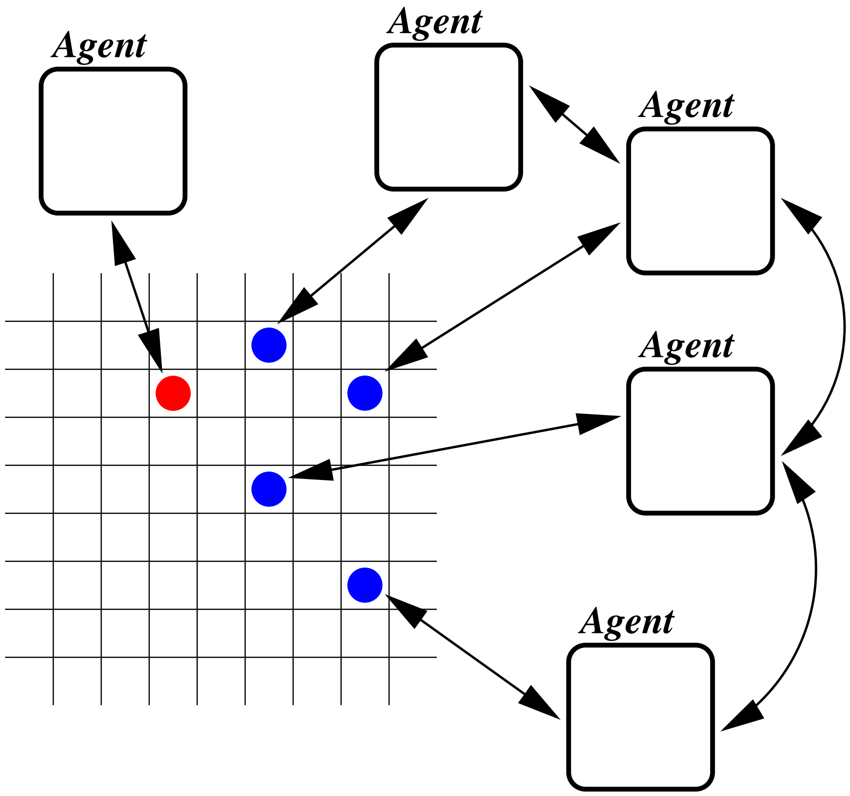

- Predators use Q-learning and share either sensation, episodes, or policies (predators, scouts, limited view). Share sensation for joint task. *Tan*
- Choose capture position to approach one at a time: farthest first. *Stephens and Merx*
- One agent controls all. *Stephens and Merx*

- Exchange data; exchange data and goals; or give up control: tradeoff between communication costs and better decisions. *Benda et al.*
- Start autonomously, then communicate, then negotiate, then control. Cost increases as freedom decreases. *Osawa*

In the extreme, we get back to the single agent scenario

Communicating Agents

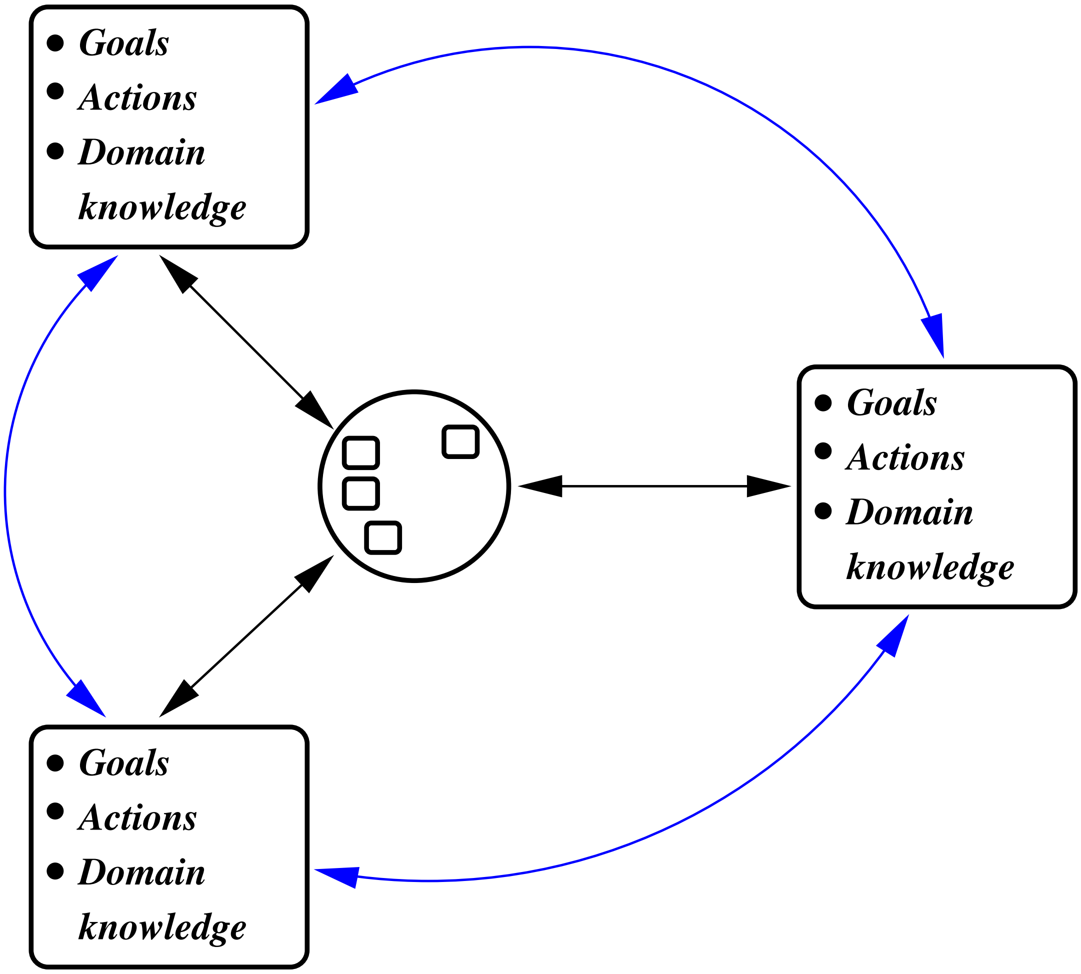

- Information can be transmitted directly among agents.
- Blackboard or point-to-point communication.

Communicating Agents

Issues

- Understanding each other.
- Planning communicative acts.
- Benevolence vs. competitiveness.
- Resource management
  (schedule coordination).
- Commitment/decommitment.

Techniques

- Language protocols: KIF for content (*Gensereth and Fikes*), KQML for message format (*Finin et al.*).
- Speech acts. *Lux and Steiner/Levesque*
- Bayesian learning in negotiation: model others. *Zeng and Sycara*
- Training other agents' Q-functions (track driving). *Clouse*
- Contract nets for electronic commerce. *Sandholm and Lesser*
- Market-based systems. *Huberman and Clearwater*
- GPGP. *Decker*
- Internal, Social, and Collective (role) commitments. *Castelfranchi*
- BDI commitments only over intentions. *Rao and Georgeff*
- Commitment states (potential, pre, and actual) as planning states. *Haddadi*
- Coalitions. *Zlotkin and Rosenschein*

Learning opportunities

- Evolving language.
- Effects of speech acts on global dynamics.
- Communication utility.
- Commitment utility.

Domain Issues

- Number of agents.
- Amount of time pressure (real time?).
- Cost of communication.
- Cost of failure.
- User involvement.
- Dynamically arriving goals?
- Environmental uncertainty: *Decker*
  - a priori in the domain;
  - in the actions of other agents;
  - in outcomes of an agent's own actions.

Examples:

WWW, design, planning, entertainment, games, air-traffic control, air combat, personal assistants, load-balancing, leg control, soccer.

Open Issues For Pursuit

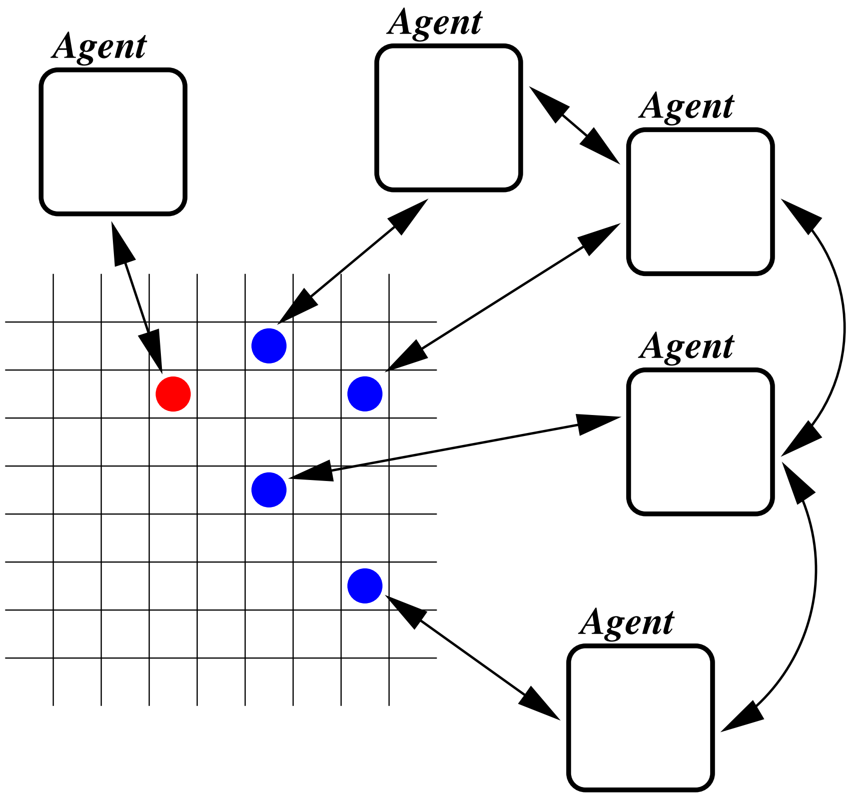

- Co-evolution of predators and prey.
- Combine collaborative and adversarial reasoning in predators.
- Switching among behaviors based on perceived behavior of opponent(s).
- Behaviors based on predicted future.

Where There's Work To Be Done
(A partial list)

- Combine collaborative and adversarial reasoning.
- Dynamic role assumption.
- Dynamic (sub)goals.
- Real-time planning.
- Layering of behaviors.