Learning and Communication

(Based on Chapter 6 of G. Weiss, (Ed.), Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, MIT Press, 1999.)

Leen-Kiat Soh

What is Learning?

The acquisition of new knowledge and motor and cognitive skills and the incorporation of the acquired knowledge and skills in future system activities, provided that this acquisition and incorporation is conducted by the system itself and leads to an improvement in its performance

Differencing Features

- Degree of decentralization
- Interaction-specific features
- Involvement-specific features
- Goal-specific features
- Learning method
- Learning feedback

Learning in agents vs. offline learning?

Degree of Decentralization

- Concerning distributedness and parallelism
- One extreme is that a single agent carries out all learning activities sequentially
- The other extreme is that the learning activities are distributed over and parallelized through all agents in a MAS

Interaction-Specific Features

When working as a team, what do you learn? What interactions?

- Concerning interactions required for realizing a decentralized learning process
  - *the **level of interaction*** (ranging from pure observation over simple signal passing and sophisticated information exchange to complex dialogues and negotiations),
  - *the **persistence of interaction*** (ranging from short-term to long-term),
  - *the **frequency of interaction*** (ranging from low to high),
  - *the **pattern of interaction*** (ranging from completely unstructured to strictly hierarchical: peer-to-peer, broadcast, etc.), and
  - *the **variability of interaction*** (ranging from fixed to changeable) as some learning requires only minimal interaction, some maximal.

Involvement-Specific Features

- Concerning (a) **the relevance of involvement** and (b) **role played during involvement**
- Relevance: there are two extremes
  - the involvement of an agent is not a condition for goal attainment because its learning activities could be executed by another available agent as well; and
  - the learning goal could not be achieved without the involvement of exactly this agent
- Role
  - an agent may act as a "generalist" in so far as it performs all learning activities
  - or it may act as a "specialist" – learning a particular activity.

Scenarios where every agent MUST learn?

Scenarios where agents should learn to become generalists, or specialists?

Goal-Specific Features

- Concerning (a) the **type of involvement** that is tried to be achieved by learning and (b) the **compatibility of the learning goals** pursued by the agents.
- Type of Involvement
  - learning that aims at an improvement with respect to a single agent (e.g., its motor skills or inference abilities)
  - learning that aims at an improvement with respect to several agents acting as a group (e.g., their communication and negotiation abilities or their degree of coordination and coherence)
- Compatibility of Learning Goals
  - Conflicting or complementary?

Scenarios where agents learn to improve group?

Scenarios where agents have conflicting learning goals? (Think about pursuit games)

Learning Methods

Which one is most common in agent learning?

Complexity

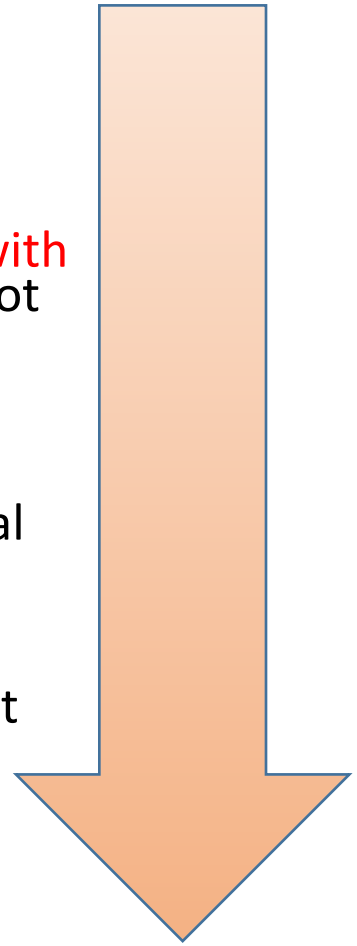

- **Rote Learning**
  - direct implantation of knowledge and skills without requiring further inference or transformation from the learner, like primary/elementary school
- **Learning from Instruction and by Advice Taking**
  - operationalization—transformation into an internal representation and integration with prior knowledge and skills—of new information like an instruction or advice that is not directly executable by the learner
- **Learning from Examples and by Practice (or by Doing)**
  - extraction and refinement of knowledge and skills like a general concept or a standardized pattern of motion from positive and negative examples or from practical experience
- **Learning by Analogy**
  - solution-preserving information of knowledge and skills from a solved to a *similar* but unsolved problem
- **Learning by Discovery**
  - discovering and gathering new (novel) knowledge and skills by making observations

Learning Feedback

- Concerning **how to give (what) feedback on performance level achieved** so far
- **Supervised Learning**
  - Specifies the desired activity of the learner and the objective is to match this desired action as closely as possible
- **Reinforcement Learning**
  - Specifies the utility of the actual activity of the learner and the objective is to maximize this utility,
- **Unsupervised Learning**
  - no explicit feedback is provided and the objective is to find out useful and desired activities on the basis of trial-and-error and self-organization processes

Supervised learning: Often offline

RL a popular approach in MAS. Why? Environment too complex, can't prescribe the desired activities that would lead to goal states

Unsupervised Learning: Pattern recognition

Learning and Communication

Learning and Communication

- **Learning to communicate**
  - Learning is viewed as a method for reducing the load of communication among individual agents
    - communication usually is very slow and expensive, and therefore should be avoided or at least reduced whenever this is possible

Learning to communicate: scalability

- **Communication as learning**
  - Communication is viewed as a method for exchanging information that allows agents to continue or refine their learning activities
    - learning is inherently limited in its potential effects by the information that is available to and can be processed by an agent

Communication as learning: Speech acts

Learning and Communication 2

- **What** to communicate
  - e.g., what information is of interest to the others
- **When** to communicate
  - e.g., what efforts should an agent investigate in solving a problem before asking others for support
- **With whom** to communicate
  - e.g., what agent is interested in this information, what agent should be asked for support
- **How** to communicate
  - e.g., at what level should the agents communicate, what language and protocol should be used, should the exchange of information occur directly—point-to-point and broadcast—or via a blackboard mechanism

What, when, with whom: can be learned at real-time

How: more difficult, meta-learning, might have to be learned offline

Example | Reducing Communication by Addressee Learning

- Broadcasting is costly; direct communication paths are *not* always known
- ***Addressee Learning***
  - to reduce the communication efforts for tasks announcement by enabling the individual agents to acquire and refine knowledge about the other agents' task solving abilities
    - Suppose that agent *a* has previously asked agent *b* to help with task T_i successfully
    - Suppose that now agent *a* encounters a new task T_j
    - If T_i is very similar to T_j, then *a can decide to approach b* again, instead of asking help from all its neighbors

Do we do this?

Example | Improving Learning by Communication

When do we share information?
When do we share understanding?

- Agents cannot be assumed to be omniscient without violating realistic assumptions. In general, agents have incomplete information about:
  - the environment in which it is embedded and the problem to be solved
  - other agents
  - the dependencies among different activities and the effects of one own's and other agents' activities on the environment and on potential future activities.
- Two forms
  - **Learning based on low-level communication**: relatively simple query-and-answer interactions for the purpose of exchanging missing pieces of information (knowledge and belief) – shared information
  - **Learning based on high-level communication**: more complex communicative interactions like negotiation and mutual explanation for the purpose of combining and synthesizing pieces of information – shared understanding

Connection to MAS?

Single agent learning: can ignore other agents, and focus on learning about the environment

Multiagent learning: more complex but often necessary in order to improve performance

Learning is powerful, adds autonomy to agents! But it is complex

Silly Question: In a 1-fish competition, each contestant can only submit one fish. Each contestant has 4 hours to fish. Each contestant is accompanied by a referee. The referee's job is to measure the fish for the contestant, accept one and only one submission, and deliver the submitted fish to the authority. The contestant with the longest fish wins the competition. *What would you do as a contestant?*