

CSCE 475/875 Multiagent Systems
Handout 7: Learning and Communication
September 12, 2017

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

Introduction

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

The 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. There is a number of features that can be applied to classifying the 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. (a) The relevance of involvement and (b) role played during involvement. With respect to 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. With respect to the role, an agent may act as a “generalist” in so far as it performs all learning activities (like centralized learning, but centralized learning precludes interaction!) or it may act as a “specialist” – learning a particular activity.

Goal-specific features. (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. The first feature leads to the important distinction between learning that aims at an improvement with respect to a single agent (e.g., its motor skills or inference abilities) and 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). The second feature leads to the important distinction between conflicting and complementary learning goals.

The learning method.

- rote learning (i.e., 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 (i.e., 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 (i.e., 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 (i.e., solution-preserving information of knowledge and skills from a solved to a similar but unsolved problem)
- learning by discovery (i.e., discovering and gathering new knowledge and skills by making observations)

The learning feedback. The learning feedback indicates the performance level achieved so far.

- supervised learning (i.e., the feedback specifies the desired activity of the learner and the objective is to match this desired action as closely as possible),
- reinforcement learning (i.e., the feedback only specifies the utility of the actual activity of the learner and the objective is to maximize this utility),
- unsupervised learning (i.e., 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)

Learning and Communication

Learning and communication are related to each other:

- (1) *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.
- (2) *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.

Both lines of research are related to the following issues:

- (1) What to communicate (e.g., what information is of interest to the others)

- (2) When to communicate (e.g., what efforts should an agent investigate in solving a problem before asking others for support)
- (3) With whom to communicate (e.g., what agent is interested in this information, what agent should be asked for support)
- (4) 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)

Reducing Communication by Learning

Broadcasting is costly. Direct communication paths are not always known.

The primary idea underlying *addressee learning* is 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. With the help of the acquired knowledge, tasks can be assigned more directly without the need of broadcasting their announcements to all agents.

Suppose each task has its specification. As such, one could establish similarity between each pair of tasks. Suppose that agent a has previously asked agent b to help with task T_i successfully. Given a new task T_j , if T_j is very similar to T_i , then it makes sense for agent a to approach agent b again.

Improving Learning by Communication

Agents cannot be assumed to be omniscient without violating realistic assumptions. In general, agents have incomplete information about:

- (1) the environment in which it is embedded and the problem to be solved
- (2) other agents
- (3) the dependencies among different activities and the effects of one's own and other agents' activities on the environment and on potential future activities.

Two forms of improving learning by communication:

- (1) **Learning based on low-level communication**, that is, relatively simple query-and-answer interactions for the purpose of exchanging missing pieces of information (knowledge and belief) – shared information, and
- (2) **Learning based on high-level communication**, that is, more complex communicative interactions like negotiation and mutual explanation for the purpose of combining and synthesizing pieces of information – shared understanding