

CSCE 475/875 Multiagent Systems
Handout 8: Learning and Teaching
 September 12, 2017
 (Based on Shoham and Leyton-Brown 2011)

Introduction

1. The capacity to learn is a key facet of **intelligent** behavior.
2. Learning is **complex**: (1) the interaction between learning and teaching, (2) the settings in which learning takes place and what constitutes learning in those settings, (3) the learning of the other agents will be impacted by the learning performed by individual agents, and (4) the yardsticks by which to measure this or that theory of learning in multiagent systems.
3. The **simultaneous** learning of the agents means that every learning rule leads to a dynamical system, and sometimes even very simple learning rules can lead to complex global behaviors of the system.
4. In the context of multiagent systems one *cannot* separate the phenomenon of *learning* from that of *teaching*; when choosing a course of action, an agent must take into account not only what it has learned from other agents' past behavior, but also how it wishes to influence their future behavior.

Example 1. Stackelberg game. Player 1 must teach player 2. (*Note: Why can't player 2 teach player 1? Who has a dominant strategy? Is there a unique Nash equilibrium?*)

	<i>Player 2 plays L</i>	<i>Player 2 plays R</i>
<i>Player 1 plays T</i>	1,0	3,2
<i>Player 1 plays B</i>	2,1	4,0

Example 2. Pure coordination game. Which driver should be the teacher? Either driver could play the teacher with equal success. However, if both decide to play teacher and happen to select uncoordinated actions (Left, Right) or (Right, Left) then the players will receive a payoff of zero forever. *Is there a learning rule that will enable them to coordinate without an external designation of a teacher?*

	<i>Driver 2 Left</i>	<i>Driver 2 Right</i>
<i>Driver 1 Left</i>	1,1	0,0
<i>Driver 1 Right</i>	0,0	1,1

When we speak about learning strategies, these should be understood as complete strategies, which involve learning in the sense of **choosing actions** as well as **updating beliefs**. One consequence is that learning in the sense of "accumulated knowledge" is *not always* beneficial. In the abstract, **accumulating knowledge never hurts, since one can always ignore what has been learned.**

Example 3. The Game of Chicken. The principle of the game is that while each player prefers not to yield to the other, the worst possible outcome occurs when both players do not yield. (Note: The “Watch out: I’m crazy” policy – less knowledge is more, in terms of teaching the other player?)

	<i>Player 2 Yield</i>	<i>Player 2 Dare</i>
<i>Player 1 Yield</i>	2,2	1,3
<i>Player 1 Dare</i>	3,1	0,0