Coordination

José M Vidal

Department of Computer Science and Engineering University of South Carolina

September 7, 2005

Abstract

We analyze the problem of coordination in game-theoretic terms [Vlassis, 2003, Chapter 4] and the emergence of conventions [Shoham and Tennenholtz, 1997].

Coordination Game

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

- Perhaps Alice and Bob should coordinate.
- In general, a game might have many Nash and Pareto equilibriums.
- We can say that coordination is the process via which agents agree on an equilibrium.
- A simple algorithm is for all agents to
  1. Determine all the equilibria.
  2. Order them.
  3. Pick the first.

Social Conventions

- A social convention constraints the agents to only take certain joint actions.
- That is, it eliminates boxes in the payoff matrix.
- An equilibrium can be found faster.
- But, it might eliminate good equilibria.

Learning Social Conventions

- IDEA: Let agents learn social conventions (like the pigs).

Highest Cumulative Reward (HCR) rule

An agent switches to a new action iff the total payoff obtained from that action in the latest m iterations is greater than the payoff obtained from the currently-chosen action in the same time period.

Settings

Prisoner's Dilemma

	c	*d*
c	3,3	0,5
d	5,0	1,1

Coordination Game

	a	*b*
a	1,1	-1,-1
b	-1,-1	1,1

Guaranteed Convergence

Theorem

The HCR update rule guarantees eventual emergence of coordination and of cooperation, that is, rational conventions in the respective games.

Experimental Results

- Convergence is reduced as the update frequency is reduced.
- Erasing their memory too often prevents convergence, but system is resilient to occasional memory loss.
- Shorter memory windows m are better.
- Try out in my NetLogo implementation.

Roles

- A role limits the actions an agent can take.
- By assigning roles we effectively reduce the size of the payoff matrix. Equilibria calculations are easier.
- Each role assignment represents a (possibly) different matrix, with new equilibria.
- Role assignment can be very hard to do optimally.

Coordination Graphs

- Assume that the global payoff is the sum local payoff functions as in

$$u(s) = f_1(s_1, s_2) + f_2(s_2, s_3) + f_3(s_3, s_4)$$

- This can be drawn as 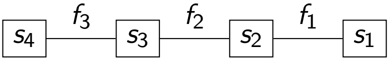

- We can now find the best strategy via iterative maximization.

Iterative Maximization

- What is the best strategy?
- Eliminate agent 1:

$$\max_{s} u(s) = \max_{s_2, s_3, s_4} \{ f_3(s_3, s_4) + \max_{s_1} (f_1(s_1, s_2) + f_2(s_1, s_3)) \}$$

Iterative Maximization

- Resolve the inner max and let that value be $f_4(s_2, s_3)$, we then have

$$\max_s u(s) = \max_{s_2,s_3,s_4} \{f_3(s_3, s_4) + f_4(s_2, s_3)\}$$

so agent 1 has been eliminated.

Iterative Maximization

- We can eliminate agent 2 by defining $f_5(s_3)$ to replace the old f_4, we then have

$$\max_{s} u(s) = \max_{s_3, s_4} \{ f_3(s_3, s_4) + f_5(s_3) \}$$

and have thus eliminated agent 2.

Iterative Maximization

- We eliminate agent 3 with the function $f_6(s_4)$ so we are left with
$$\max_s u(s) = \max_{s4} f_6(s_4)$$
which agent 4 can solve by itself, choosing a_4^*.
- This strategy can then be plugged into 3's formula so it can calculate its strategy, and so on, until 1.
- This is faster than computing all equilibria in the full game.

Conclusion

- Social conventions and roles are easy to implement and minimize communication and computations.
- A more sophisticated technique is negotiation, which requires communication and reasoning.

Shoham, Y. and Tennenholtz, M. (1997).
On the emergence of social conventions: modeling, analysis, and simulations.
Artificial Intelligence, 94:139–166.

Vlassis, N. (2003).
A concise introduction to multiagent systems and distributed AI.
Informatics Institute, University of Amsterdam.
http://www.science.uva.nl/~vlassis/cimasdai.