Abstract

This thesis studies a class of problems where rational agents can make suboptimal decisions by ignoring a side effect that each individual action brings to bear on the common good. It is generally believed that a mutually desirable strategy can be enforced as a stable outcome for rational agents if the imminent threat exists that any deviator from the strategy will be punished. This thesis expands this understanding, arguing that rationally bounded agents can learn to self-organize to stabilize on mutually beneficial outcomes *without* the explicit notion of threat. As an approach to demonstrate this capability, a double-layered multi-agent learning algorithm, known here as IMPRES (implicit reciprocal strategy learning), has been developed.

In game theory, it is generally assumed that the players (agents) of a game are of equal ability. This thesis takes a contrasting view. The foundation of this work is inspired by the concept of \bounded rationality", where some agents may have more privileges than others either because they are exposed to different parts of information in the environment, or because they simply have higher computational power. Based on this intuition, this thesis investigates how agents can boost their performance by utilizing the notion of social learning - learning from one another in an agent society.

Theoretical and empirical results show that the IMPRES agents learn to behave rationally as if they are in a virtually optimal Nash equilibrium of a repeated game. To my knowledge, IMPRES is the first algorithm that achieves this property in games involving more than two players under imperfect monitoring.