On Learning Negotiation Strategies by Artificial Adaptive Agents in Environments of Incomplete Information

Jim R. Oliver

INSEAD, France,
jroliver@ebay.com

Abstract. Automated negotiation by artificial adaptive agents (AAAs) holds great promise for electronic commerce, but non-trivial, practical issues remain. Published studies of AAA learning of negotiation strategies have been based on artificial environments that include complete payoff information for both sides of the bargaining table. This is not realistic in applied contexts. Without loss of generality, we consider the case of a seller who knows its own preferences over negotiation outcomes but will have limited information about the private values of each customer. We propose a learning environment that takes advantage of partial information likely to be available to the vendor. General strategies are learned for a group of similar customers – a market segment – through a simulation approach and a genetic learning algorithm. In addition, we systematically further relax constraints on the opponent's preferences to further explore AAA learning in incomplete information environments.

1 Introduction

Researchers have demonstrated that artificial adaptive agents (AAAs) can learn effective negotiation strategies. But the experimental environments in which these agents learn are also artificial and depart from what is plausible in natural settings. In particular, the environments completely specify the payoff structure for the opposing sides, but in applied contexts each side keeps information private or even misrepresents it. One side never completely knows the other. Still, we have some hunches about our opponents. Wouldn't it be nice if there were a way to take advantage of what is known about our opponent and incorporate it systematically and effectively into a learning environment for AAAs? This paper examines the potential of and perils of incomplete information and the impact on AAA performance.

1.1 The Challenge of Negotiation

Bargaining and negotiation always have, and always will, play a critical role in commerce. Even within a firm, allocation decisions usually have a negotiation component. Despite its importance and prevalence throughout history,

human negotiation performance falls significantly short of optimal. Experiments, field studies, and our common experience demonstrate that even in simple negotiations people often reach sub-optimal agreements, thereby "leaving money on the table" [Cam90]. Economic efficiency would be greatly increased if the cost of negotiation could be decreased or superior agreements could be effected. Computer science and information systems researchers have approached this problem in two ways. One stream of research has sought to provide tools that support human negotiators (e.g. [RS97]), while another stream seeks to build systems that can negotiate by themselves. Our interest here is in the latter approach, specifically systems of Artificial Adaptive Agents (AAAs) based on evolutionary computation.

The essence of negotiation is two or more parties trying to arrive at a single agreement from a set – often large – of potential agreements. Part of the challenge and the opportunity of negotiation arises from differences between the parties as to how they value possible outcomes.

Negotiation can be viewed as a *search* process. Optimization is also often thought of as a search process, but it differs in important ways from negotiation search. Optimization seeks the maximum or minimum of a known (computable), static, search space. In negotiation, each side has private information, and neither knows the other's exact utility function. The situation is further complicated because both sides have incentive to misrepresent their preferences. Thus, each party has differing views of the search space, which change over time as the bargaining session unfolds. Finding an optimal agreement in this dynamic environment is extremely challenging because both sides are in competition, but must jointly search the space of possible agreements.

Despite the challenging nature of negotiation, researchers pursuing automated negotiation have had provocative successes. A multitude of approaches have been taken, and the results have been reviewed extensively elsewhere (e.g. [Bic01]; [Far98]); thus we will only provide brief background. [CSLC01], [TWL00], and [Oli96a] used genetic algorithms [Gol89] to evolve agents that can negotiate in complex, large, multi-dimensional search spaces. [Oli96a] used simple strategies that capture common decision making processes, but lack the expressive power to capture and react directly to a particular dynamic bargaining path. [TWL00] extended the work by using Finite State Machines (FSMs) to represent strategies. Another stream of research is represented by [DKL96] who use genetic programming techniques to evolve agents that can play coalition games.

1.2 Two Problems: Learning Curves and Knowing Your Opponent

The promise of Artificial Adaptive Agents learning to negotiate is exciting, but many practical problems remain to be overcome before the approach is viable for commercial transactions between firms. We turn our attention to just one of these difficulties, a specific applied problem that is important and