Chapter XLIX
Agents for Multi-Issue Negotiation

John Debenham
University of Technology, Australia

ABSTRACT

This chapter describes a generic multi-issue negotiating agent that is designed for a dynamic information-rich environment. The agent strives to make informed decisions by observing signals in the marketplace and by observing general information sources including news feeds. The agent assumes that the integrity of some of its information decays with time, and that a negotiation may break down under certain conditions. The agent makes no assumptions about the internals of its opponent—it focuses only on the signals that it receives. Two agents are described. The first agent conducts multi-issue bilateral bargaining. It constructs two probability distributions over the set of all deals: the probability that its opponent will accept a deal, and the probability that a deal should be accepted by the agent. The second agent bids in multi-issue auctions—as for the bargaining agent, this agent constructs probability distributions using entropy-based inference.

INTRODUCTION

This work is based on the assumption that when an intelligent agent buys a hat, a car, a house, or a company, she does so because she feels comfortable with the general terms of the deal. This "feeling of comfort" is achieved as a result of information acquisition and validation. Negotiation is as much of an information acquisition and exchange process as it is an offer exchange process—one feeds off the other.

The generic multi-issue negotiation agent Π draws on ideas from information theory. Game theory (GT) tells us what to do, and what outcome to expect, in many well-known negotiation situations, but these strategies and expectations are derived from assumptions about the internals of the opponent. Game theoretic analyses of bargaining are founded on the notion of agents as utility optimizers in the presence of complete and incomplete information about their opponents (Muthoo, 1999).

Two probability distributions form the foundation of both the offer evaluation and the offer making processes. They are both over the set of all deals and are based on all information available to the agent. The first distribution is the probability that a deal is acceptable to Ω. The second distribution is the probability that a deal will prove to be acceptable to Π—this distribution generalizes the notion of utility.

Π may not have a von Neumann-Morgerstern utility function. Π makes no assumptions about the internals of Ω in particular whether it has a utility function. Π does make assumptions about: the way in which the integrity of information will decay, preferences that its opponent may have for some deals over others, and conditions that may lead to breakdown. It also assumes that unknown probabilities can be inferred using *maximum entropy probabilistic logic* (McKay, 2003) that is based on random worlds (Halpern, 2003). The maximum entropy probability distribution is "the least biased estimate possible on the given information; i.e., it is maximally noncommittal with regard to missing information" (Jaynes, 1957). In the absence of knowledge about Ω's decision-making apparatus, Π assumes that the "maximally noncommittal" model is the correct model on which to base its reasoning.

A *preference relation* is an assumption that Π makes about Ω's preferences for some deals over others—for example, that she prefers to pay a lower price to a higher price. A *single-issue preference relation* assumes that she prefers deals on the basis of one issue alone, independent of the values of the other issues. A preference relation may be assumed prior to the negotiation, or during it based on the offers made. For example, the opponent may display a preference for items of a certain color; Faratin, Sierra, and Jennings (2003) describe a basis for ordering colors. The preference relations illustrated here are single-issue orderings, but the agent's reasoning operates equally well with any preference relation as long as it may be expressed in Horn clause logic.

THE MULTI-ISSUE NEGOTIATION AGENT Π

The integrity of information decays in time. Little appears to be known about how the integrity of information, such as news-feeds, decays. One source of information is the signals received by observing the behavior of the opponent agents both prior to a negotiation and during it. For example, if an opponent bid $8 in an auction for an identical good two days ago, then my belief that she will bid $8 now could be 0.8. When the probability of a decaying belief approaches 0.5, the belief is discarded.

Agent Architecture

Incoming messages from all sources are time-stamped and placed in an "In Box" X as they arrive. Π has a knowledge base K and a belief set B. Each of these two sets contains statements in L. K contains statements that are generally true, such as $\forall x(Accept(x) \leftrightarrow \neg Reject(x))$—that is, an agent does one thing or the other. The belief set $B = \{\beta_i\}$ contains statements that are each qualified with a *given sentence probability* $B\{\beta_i\}$, which represents an agent's belief in the truth of the statement. These sentence probabilities may decay in time.

A *deal* is a pair of commitments $\delta_{\Pi:\Omega}(\pi, \omega)$ between an agent Π and an opponent agent Ω, where π is Π's commitment and ω is Ω's commitment. $D = \{\delta_i\}_{i=1}^{D}$ is the deal set—the set of all possible deals. If the discussion is from the point of view of a particular agent, then the subscript "Π:" may be omitted, and if that agent has just one opponent, the "Ω" may be omitted as well. These commitments may involve multiple issues and not simply a single issue such as

Related Content