



A reinforcement learning approach to improve the argument selection effectiveness in argumentation-based negotiation

Ariel Monteserin*, Analía Amandi

ISISTAN Research Institute, CONICET – UNCPBA, Campus Universitario, Paraje Arroyo Seco, 7000 Tandil, Argentina

ARTICLE INFO

Keywords:

Reinforcement learning
Argument selection
Argumentation-based negotiation
Autonomous agents

ABSTRACT

Deciding what argument to utter during a negotiation is a key part of the strategy to reach an expected agreement. An agent, which is arguing during a negotiation, must decide what arguments are the best to persuade the opponent. In fact, in each negotiation step, the agent must select an argument from a set of candidate arguments by applying some selection policy. By following this policy, the agent observes some factors of the negotiation context (for instance, trust in the opponent and expected utility of the negotiated agreement). Usually, argument selection policies are defined statically. However, as the negotiation context varies from a negotiation to another, defining a static selection policy is not useful. Therefore, the agent should modify its selection policy in order to adapt it to the different negotiation contexts as the agent gains experience. In this paper, we present a reinforcement learning approach that allows the agent to improve the argument selection effectiveness by updating the argument selection policy. To carry out this goal, the argument selection mechanism is represented as a reinforcement learning model. We tested this approach in a multiagent system, in a stationary as well as in a dynamic environment. We obtained promising results in both.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

In multiagent systems, negotiation is a fundamental tool to reach an agreement among agents with conflicting goals. The essence of the negotiation process is the exchange of proposals. Agents make proposals and respond to proposals in order to converge on a mutually acceptable agreement. However, not all approaches are restricted to that exchange of proposals. Several approaches to automated negotiation have been developed. One of them is the argumentation-based approach (see e.g. (Amgoud, Dimopoulos, & Moraitis, 2007; Geipel & Weiss, 2007; Kraus, Sycara, & Evenchik, 1998; Ramchurn, Jennings, & Sierra, 2003; Rahwan et al., 2003; Sierra, Jennings, Noriega, & Parsons, 1998)). In argumentation-based approaches, agents are allowed to exchange some additional information as arguments, besides the information uttered on the proposals. Thus, in the context of the negotiation, an argument is seen as a piece of information that supports a proposal and allows an agent (a) to justify its position of negotiation, or (b) to influence the position of negotiation of other agents (Jennings, Parsons, Noriega, & Sierra, 1998).

In contrast to agents without an argumentative ability, an argumentative agent, in addition to evaluating and generating proposals,

must be able to evaluate, generate and select arguments (Ashri, Rahwan, & Luck, 2003; Rahwan et al., 2003). Argument evaluation processes incoming arguments and updates the agent's mental state as a result. Argument generation and selection are related to the production of outgoing arguments. When the agent has to argue during a negotiation, it first generates a set of candidate arguments, for example by using explicit rules (Kraus et al., 1998; Ramchurn et al., 2003), then the agent selects what argument to utter by applying a selection policy. The agent usually *observes* the context of the negotiation and decides which type of argument to utter by following the argument selection policy. This policy takes into account several factors of the negotiation context: trust in the opponent (Ramchurn et al., 2003), agreement urgency, authority relation with the opponent (Sierra et al., 1998), expected utility and argument strength (Kraus et al., 1998), among others.

The argument selection is considered as the essence of the strategy in argumentation-based negotiation (Rahwan et al., 2003). Therefore, the success of the negotiation depends on the effectiveness of this mechanism. In the literature, the selection policy is generally represented as a set of explicit rules that determines which negotiation factors should be observed and what type of argument should be uttered in each context. Nevertheless, these policies do not take into account the process of learning new rules or updating existent ones. Because of the constant appearance of new factors, opponents and types of agreement in the negotiation context, learning is essential. In addition, opponents are heterogeneous, thereby,

* Corresponding author.

E-mail addresses: ariel.monteserin@isistan.unicen.edu.ar (A. Monteserin), analía.amandi@isistan.unicen.edu.ar (A. Amandi).

we do not think that all opponents, in the same context, will respond to the same arguments in the same way.

To solve this limitation, we present a novel approach to improve the argument selection effectiveness in the context of argumentation-based negotiation. Our approach uses the well-known Reinforcement Learning mechanism (Sutton & Barto, 1998; Kaelbling, Littman, & Moore, 1996) to update the argument selection policy of an agent. Reinforcement learning (RL) is a problem faced by an agent that must learn behaviour through interaction with a dynamic environment. We select this mechanism because naturally emulates the way in which people learn to select arguments in the real life. Moreover, the argument selection mechanism can be naturally modelled as a reinforcement learning problem.

In this work, the argument selection policy is represented by a set of preferences. These preferences determine how suitable it is to utter an argument in a given context. Each preference is composed by the argument, the set of factors that describe the negotiation context (trust, authority role, urgency and utility, among others), and a preference level described by two values: support and confidence. These preferences are structured in a hierarchy. At the top levels of the hierarchy are situated the most general preferences and in the low levels, the most particular ones. Initially, this hierarchy is empty, but new preferences are added as the agent gains experience by arguing in different negotiations. This structure allows us to add new factors to the preferences dynamically and make them more specific. In addition, we update the support and confidence values of each preference taking into account the success or failure of the argument uttered (for example, an argument is successful when it is accepted or not rebutted by the opponent).

We have tested our approach in a multiagent system in which the agents must negotiate with other agents to reach an agreement. We have obtained promising results. We compared the argument effectiveness between an agent selecting the arguments following a static selection policy and an agent using the reinforcement learning approach to learn and update the preferences for argument selection. This comparison was made in a stationary environment as well as in a dynamic one. In a stationary environment, we found that the effectiveness of the first agent was between 30% and 45%, whereas the second agent started at 0%, increased logarithmically and reached a final effectiveness of 70%. In the dynamic context, the agent using the RL approach also obtained the best argument effectiveness.

The paper is organized in the following way. Section 2 introduces concepts and related work in the area of argumentation-based negotiation. Section 3 presents how the argument selection mechanism is modelled as a reinforcement learning problem. In Section 4, we define the argument selection policy and the RL algorithm to train it. Section 5 presents the experimental results. Finally, in Section 6, concluding remarks and future work is described.

2. Argument selection in argumentation-based negotiation: background and related work

In accordance with the work of Rahwan, Sonenberg, and McBurney (2005), there are two major strands in the literature on argumentation-based negotiation: (a) attempts to adapt dialectical logics for defeasible argumentation by embedding negotiation concepts within these (Amgoud, Parsons, & Maudet, 2000; Parsons, Sierra, & Jennings, 1998); and (b) attempts to extend bargaining-based frameworks by allowing agents to exchange rhetorical arguments, such as promises and threats (Amgoud & Prade, 2005; Kraus et al., 1998; Sierra et al., 1998). Our work is situated in the second strand.

As we have introduced earlier, in an argumentation-based negotiation approach, agents can exchange arguments (particularly, rhetorical arguments) in order to justify their proposals, to persuade their opponent, and to reach an expected agreement. In

contrast to agents without this argumentative ability, an argumentative agent must be able to (a) evaluate incoming arguments and update its mental state as a result; (b) generate candidate outgoing arguments; and (c) select an argument from the set of candidate arguments (Ashri et al., 2003). An argument is a set of one or more meaningful declarative sentences known as the premises along with another meaningful declarative sentence known as the conclusion. There are several types of rhetorical arguments that an agent can generate during a negotiation. Three general argument types are defined in the literature on argumentation-based negotiation: appeals (Amgoud & Prade (2004) define them as explanatory arguments), rewards and threats (Kraus et al., 1998; Sierra et al., 1998). Appeals are used to justify a proposal; rewards to promise a future recompense; and threats to warn of negative consequences if the counterpart does not accept a proposal. Moreover, we can define several subtypes of appeal by varying their premises: past promise, counterexample, prevailing practice and self-interest, among others.

In this work, we focus on the selection of arguments. Rahwan et al. (2003) consider argument selection as the essence of the strategy in argumentation-based negotiation. Argument selection is concerned with selecting the argument that should be uttered to a counterpart from the set of candidate arguments generated by the argument generation process. Once the candidate arguments have been generated, the argument selection mechanism must apply some policy, in accordance with the agent's mental state, to select the best argument. Argument selection policies are diverse. Kraus et al. (1998) define that the candidate arguments are ordered by their severity, then they select the weakest, taking into account appeals as the weakest argument and threats as the strongest argument. Ramchurn et al. (2003) define rules for argument selection by observing the trust in the opponent and the expected utility of the proposal. For example, they state that if the trust is low and the utility is high then the agent should send a strong argument, but if the trust is high and the utility low, then it should utter a weak one. In the work of Sierra et al. (1998), several authority roles among agents are taken into account to generate and evaluate arguments. Moreover, other factors influence the negotiation process and they should be taken into account during the argument selection. For instance, the time available to reach the agreement influences directly the negotiation process, affecting the agent behaviour in different ways: the agent can be patient or impatient. Thus when the agent is patient, it gains utility with time and when the agent is impatient, it loses utility with time (Fatima, Wooldridge, & Jennings, 2004). Other works analyse the information that composes each argument. Schroeder (1999) chooses the shortest argument in order to reduce the target to counter-argue. Amgoud and Prade (2004) assign a strength to each argument in accordance with the beliefs with which it was built. All these works establish different factors and rules to select the best argument. However, they define static policies for argument selection. That is, they do not define how to learn and update the selection policy nor how to integrate different factors or incorporate new ones as the agent gains experience.

Additionally, the design of negotiation strategies has been studied from several perspectives (Baek & Kim, 2007; Carbonneau, Kersten, & Vahidov, 2008; Lin, Chen, & Chu, 2011). Particularly, Rahwan, McBurney, and Sonenberg (2003) determine that a negotiation strategy may be defined as a rule or algorithm which specifies what the agent should utter and when, in a particular negotiation interaction. In that direction, Rahwan et al. identify some factors that may influence the design of the strategy. Among these factors, we can stress: goals (what goals the agent wants to achieve from undertaking a negotiation), counterparts (the nature of the other participants), and resources (the time and the resources available for the agent), among others. Therefore, the argument selection

Download English Version:

<https://daneshyari.com/en/article/10321972>

Download Persian Version:

<https://daneshyari.com/article/10321972>

[Daneshyari.com](https://daneshyari.com)