Computers, Environment and Urban Systems 48 (2014) 73–85

Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

journal homepage: www.elsevier.com/locate/compenvurbsys

Incorporating Bayesian learning in agent-based simulation of stakeholders' negotiation

Majeed Pooyandeh *, Danielle J. Marceau

Geocomputing Laboratory, Department of Geomatics Engineering, University of Calgary, 2500 University Drive NW, Calgary, AB T2N 1N4, Canada

ARTICLE INFO

Article history:
Received 7 November 2013
Received in revised form 2 July 2014
Accepted 12 July 2014
Available online 15 August 2014

Keywords:
Agent-based model
Stakeholders' negotiation
Bayesian learning
Fuzziness
Land development

ABSTRACT

This paper describes the incorporation of a Bayesian learning algorithm into an agent-based model designed to simulate stakeholders' negotiation when evaluating scenarios of land development. The objective is to facilitate reaching an agreement at an earlier stage in the negotiation by providing the opportunity to the proposer agent to learn his opponents' preferences. The modeling approach is tested in the Elbow River watershed, in southern Alberta, Canada, that is under considerable pressure for land development due to the proximity of the fast growing city of Calgary. Five agents are included in the model respectively referred to as the *Developer* agent, the *Planner* agent, the *Citizen* agent, the *Agriculture-Concerned* agent, and the *WaterConcerned* agent. Two types of land development scenarios are evaluated; in the first case, only the geographical location is considered while in the second case, the internal land-use composition is also varied. The *Developer* agent that is equipped with the Bayesian learning capability attempts to approximate its opponents' fuzzy evaluation functions based on the responses he receives from them at each round of the negotiation. The results indicate that using this approach, an agreement can be reached in fewer number of negotiation rounds than in the case where the Developer agent selects the subsequent offers based merely on its own utility. The model also indicates how the satisfaction of each agent evolves during the negotiation. This information is very useful for decision makers who wish to consider stakeholders' perspectives when dealing with multiple objectives in a spatial context.

1. Introduction

Global urbanization and the resulting concerns about land sustainability have generated an urgent need for examining scenarios of land development (Wu, 1996). These scenarios are images of future land-use patterns if certain land development regulations were to be adopted by decision makers (Xiang & Clarke, 2003). Different "what if" scenarios based on stakeholder inputs and feedback facilitate the investigation of possible land development patterns without bearing the costs of implementing them (Van Noordwijk, Tomich, & Verbist, 2003). This process of incorporating multiple views and coping with pluralistic wishes requires negotiation (Forester, 1999). Negotiation is a complex decision-making process where each party autonomously represents its viewpoints and interacts with the other parties to resolve conflicts and reach an agreement while attempting to maximize all parties' payoffs (Choi, Liu, & Chan, 2001; Jennings et al., 2001). It typically involves a combination of objective facts along with values and emotions and can be highly deviated from rationality due to individual and competitive biases (Bazerman & Moore, 2008). Computer models can facilitate human negotiation by processing a wide range of alternatives and examining their outcomes in the presence of biases (Oliver, 1997).

While land development scenarios have been in practice for years, it is only in the past two decades that the employment of computer models for creating and evaluating them has become possible. These models vary from GIS functionalities (Almeida et al., 2005; Batty & Xie, 1994; Hilferink & Rietveld, 1999; Joao & Walsh, 1992) to sophisticated computational approaches, such as agent-based modeling (ABM) in which the spatial capabilities of GIS are combined to Artificial Intelligence techniques (Benenson & Torrens, 2003; Ligmann-Zielinska & Jankowski, 2010; Matthews, Gilbert, Roach, Polhill, & Gotts, 2007). Software agents as autonomous problem solving entities can support the automation of complex negotiations by negotiating on the behalf of stakeholders and providing adequate strategies to achieve realistic, win–win agreements (Rahwan, Kowalczyk, & Pham, 2002).

Agent-based modeling (ABM), which has roots in Artificial Intelligence, possesses outstanding features for simulating and testing scenarios to support decision making (Mensonides, Huisman, &

* Corresponding author.
E-mail addresses: mpooyand@ucalgary.ca (M. Pooyandeh), dmarceau@ucalgary.ca (D.J. Marceau).

http://dx.doi.org/10.1016/j.compenvurbsys.2014.07.003
0198-9715/

Dignum, 2008). It employs a bottom-up approach in which the interactions of the individual decision makers are simulated (Bone, Dragicevic, & White, 2011). In ABMs, entities of the system being investigated are represented as autonomous individual agents that are intelligent and purposeful and act based on their own interests, values and goals (Matthews et al., 2007). They are aware of their environment, can communicate with each other and adapt their behavior (Beck, Kempener, Cohen, & Petrie, 2008). This modeling approach is particularly adapted to deal with situations where the agents seek their own benefit in the usage of a limited common resource and where a solution needs to be reached to ensure the sustainability of this resource (Marceau, 2008). The capability of these models to connect heterogeneous individual behaviors to collectively emerging patterns makes them suitable for modeling land development scenarios, which requires considering a pluralistic standpoint towards the problem in hand (Lempert, 2002).

Agent-based automated negotiation refers to negotiation conducted with computer agents using artificial intelligence techniques in which two or more agents multilaterally bargain resources for mutual intended gain (Beam & Segev, 1997). A computer agent is situated in some environment and is capable of flexible problem solving behavior to fulfil a specific purpose (Jennings et al., 2001). It has been demonstrated that negotiating agents may obtain significantly improved outcomes compared to results achieved by humans (Jonker et al. 2012). Different agent-based negotiation models have been proposed (Lopes, Wooldridge, & Novais, 2009). *Game-theoretic models* are particularly interesting in the context of land development. In these models, the parties choose a strategy to maximize the negotiation outcome by an iterative exchange of proposals. If the preference information of a player is known to all other players, then the game is one with *complete information*; otherwise it is called a game with *incomplete information* (Ausubel, Cramton, & Deneckere, 2002). In a multi-objective negotiation regarding shared environmental resources such as land, dealing with incomplete information is typically the case.

In the absence of complete information, learning techniques can be used by the agents to acquire knowledge about the other agents' preferences or changes in the environment. Incorporating learning techniques in negotiation offers two main advantages (Gerding, van Bragt, & La Poutre, 2000). First, an agent can adjust its own negotiating strategies to obtain better deals based on its previous negotiation experiences. Second, learning can be used to update expectations regarding other parties' strategies. A suitable conflict management approach such as a negotiation must foster learning among the parties (Lee, 1994). This is vital to the sustainability of decisions in any natural system (Daniels & Walker, 1996). The elements of such systems need to adapt to changing environments and such adaptation is done through learning. While it is an inherent feature of human decision making process (Daniels & Walker, 1996), a computer model which attempts at simulating such decision making needs to accommodate learning as well. Learning in this context not only improves the negotiation outcomes, but also provides insights into the possible avenues for agreement in real world negotiations. Small changes in learned behaviors can often result in unpredictable changes in the resulting macro-level emergent properties of the multi-agent group as a whole (Panait & Luke, 2005).

Due to the semi-cooperative nature of land development, in which agents compete over a resource but also attempt to perform a common task, the notion of learning is particularly significant. While learning is a missing component in many real world negotiations of land development (Forester, 1999), a simulation model aimed at improving such negotiations need to explicitly incorporate learning. The stakeholders as users of a negotiation support system equipped with learning capability can investigate the evolution of opinions among the opponents that results from the learning capability. They can understand the significance of learning the opponents' perspectives and how it enhances the negotiation outcomes.

Several learning approaches have been used in agent-based negotiation to facilitate the agreement among agents (Panait & Luke, 2005; Weiß, 1996). They aim at obtaining a better performance in the future based on the experiences gained in the past (Alpaydin, 2004; Kulkarni, 2012). One of the popular learning approaches in agent-based negotiation is Reinforcement Learning (RL). In RL, a numerical performance measure representing an objective is being maximized (Szepesvári, 2010). At each iteration, the agent takes an action that changes the state of the environment; such transition is communicated to the agent through a scalar reward called *reinforcement signal* that evaluates the quality of the transition (Kaelbling, Littman, & Moore, 1996). The study of Bone and Dragićević (2010) is a good example of the use of RL to improve the negotiation results in a multi-stakeholder agent-based forest management model. However, a common issue with RL is to find a balance between *exploration* that consists in taking sub-optimal actions to discover new features, and *exploitation* that involves using the knowledge currently available about the world (Coggan, 2004). Each action must be repeated several times to obtain a reliable estimate of its expected reward (Kulkarni, 2012). Generalization is another issue in RL in which a function approximator such as neural network is needed to generalize between similar situations and actions (Boyan & Moore, 1995; Sutton, 1996).

Other learning techniques have been employed in agent-based negotiation. Choi et al. (2001) used a genetic algorithm to enable an agent to learn its opponents‘ preferences based on the counter-offers received during the previous rounds of negotiation. This approach requires a large number of rounds to obtain meaningful results. Carbonneau, Kersten, and Vahidov (2008) used a neural network to predict the opponents‘ negotiation moves in electronic negotiations. Other than the requirement for a large number of negotiation rounds, the generalizability of the approach presented in this study is also limited.

A promising approach to deal with the issue of learning in agent-based negotiation is Bayesian learning in which the probability of a hypothesis is updated based on acquired evidence. In other words, the posterior probability distribution of a hypothesis is computed conditioned to the evidence obtained through new data. It has been demonstrated that Bayesian learning provides the opportunity to learn an opponent‘s evaluation function in a fewer negotiation rounds in comparison with a no-learning scenario (Hindriks & Tykhonov, 2008). Moreover Bayesian learning is not data intensive and can yield noticeable results in a reasonable number of negotiation rounds (Kotsiantis, 2007). Domingos and Pazzani (1997) performed a large-scale comparison of the Bayesian approach with a number of algorithms for decision tree induction, instance-based learning, and rule induction on standard benchmark datasets, and found it to be superior in comparison with other learning schemes, even on datasets with substantial feature dependencies. They also concluded that the Bayesian approach can be a better method than most powerful alternatives when the sample size is small. This is important in land management negotiation, in which the number of land development scenarios that a developer can propose is limited and therefore the search space, i.e. the number of possible alternatives, is not large. In such a case, the learning needs to be accomplished in a few rounds of negotiation. This is in contrast with cases where several negotiation rounds must be completed for the learning to be achieved.

One of the initial attempts to incorporate Bayesian learning in agent negotiation was made by Zeng and Sycara (1998) who devel-