A Multi-context BDI Recommender System: from Theory to Simulation

Amel Ben Othmane*, Andrea Tettamanzi†, Serena Villata†, Nhan Le Thanh†

*Université Côte d’Azur, ADEME, Inria, CNRS, I3S France

amel.ben-othmane@inria.fr

†Université Côte d’Azur, CNRS, Inria, I3S, France

{andrea.tettamanzi, serena.villata, nhan.le-thanh}@unice.fr

Abstract—In this paper, a simulation of a multi-agent recommender system is presented and developed in the NetLogo platform. The specification of this recommender system is based on the well known Belief-Desire-Intention agent architecture applied to multi-context systems, extended with contexts for additional reasoning abilities, especially social ones. The main goal of this simulation study is, besides illustrating the usefulness and feasibility of our agent-based recommender system in a realistic scenario, to understand how groups of agents behave in a social network compared to individual agents. Results show that agents within a social network have better collective performance than individual ones. The utility and the satisfaction of agents is increased by the exchange of messages when executing intentions.

I. INTRODUCTION

Nowadays, recommender systems must cope with an increasing demand of complexity, for example, an application for recommending routes in a traffic scenario should deal with different contextual information (e.g., information about the user location) and other non-logical components linked to human behavior like desires, beliefs or emotions. For this reason, multi-agent systems are considered as suitable alternatives for modeling and simulating this kind of real-world scenarios, where different entities autonomously interact in a dynamic and uncertain environment. In particular, one of the most popular agent architectures, the Belief-Desire-Intention (BDI) model, seems to be particularly suitable for this task. Under this model, the mental state of the agent is composed by a set of beliefs, desires and intentions that consist of informational, motivational, and deliberative states, respectively.

Recently, the Artificial Intelligence (AI) community is putting much effort on the investigation and evaluation of recommender systems based on intelligent agents. Such a kind of systems have been applied so far in different fields such as health-care [1], tourism [2] and traffic and transportation [3]. A complete taxonomy of recommender agents can be found in Montaner *et al.* [4]. However, few works only combine BDI agents and recommender systems. Among them, we refer in particular to the approach of Casali *et al.* [5] who propose a BDI recommender agent in the tourism domain.

In this paper, we propose a multi-agent simulation to evaluate the overall behavior of the multi-context BDI recommender system we presented in [6], where different strategies are

applied. The proposed framework aims at recommending a plan for a user taking into account different contexts. For this purpose, we give an overview of the different theories used to define contexts, and explain how all those contexts are relied together to define the whole behavior of the system. In order to evaluate the goodness of the proposed recommendation, we compare the performance of the system with two different strategies, namely the *solitary agent* strategy, and the *social agent* strategy.

The reminder of this paper is organised as follows: Section II provides an overview of our multi-context BDI formal framework, highlighting the main features of the system. The simulation and the results are discussed in Section III. Finally, some conclusions are drawn.

II. THE MULTI-AGENT FRAMEWORK

The specification of our agent model is based on Multi-Context Systems (MCS) [7], to allow for a separation of the definitions of the different formal components or units. A MCS is defined as a group of interconnected units $\{\{C_i\}_{i \in I}, \Delta_{br}\}$ where each context is defined as a tuple $\langle L_i, A_i, \Delta_i \rangle$ where L_i , A_i and Δ_i are the language, axioms, and inference rules, respectively. Δ_{br} is a set of bridge rules, i.e., rules of inference, which relate formulas in different units. A bridge rule is of the form:

$$C_1 : \phi, C_2 : \psi \rightarrow C_3 : \theta$$

and it can be read as: if the formula ϕ can be deduced in context C_1 , and ψ in C_2 , then the formula θ is to be added to the theory of context C_3 .

The advantage of adopting MCS is illustrated below, where we use MCS for the specification of our extended BDI agent based on [7]. As visualized in Figure 1, our multi-context BDI agent is defined as follows:

$$Ag = (\{BC, DC, GC, SC, PC, IC, CC\}, \Delta_{br})$$

where BC , DC , GC represent respectively the Belief Context, the Desire Context and the Goal Context which model an agent mental attitude. PC , IC and CC are functional contexts that represent respectively the Planning Context, the Intention Context and the Communication Context. SC is for the Social Context, and it models social influence between agents. We first extended the classical BDI model with others

Fig. 1. The Multi-context BDI Agent Model

contexts. The Goal Context, for example, is introduced based on [8], where goals are considered as a list of desires that, besides being logically consistent, are also maximally desirable. Second, in order to represent and reason about graded notions of beliefs, desires and goals, we use the classical propositional language with additional connectives as language L_i , following [8], [9], where uncertainty reasoning is dealt with possibility theory [10]. The behavior of these contexts is handled by means of internal deduction rules Δ_i and axioms L_i which contain axioms from classical propositional logic, and from necessity and possibility measures of possibility theory. The detailed formalization of each context can be found in [6].

Concerning the Social Context, we assume that an agent has the tendency to be socially influenced by other agents to adopt a certain mental attitude if it has similarities with the latter without the need to be in an explicit social relationship with it, e.g., to have the same goals or to be in the same location. Consequently, if an agent a_i is similar to another agent a_j , a direct link is created with the latter. Between them, we consider a trust relationship and the trustworthiness of a_i towards agent a_j about an information ϕ is interpreted as a necessity measure $\tau \in [0, 1]$.

For specifying and reasoning over plans, i.e., the Planning and Intention contexts, we propose to adopt the 5W (Who, What, Why, When, Where) vocabulary¹, which is relevant for describing different concepts and constraints related to plans and allows spatial and temporal reasoning over plans and intentions.

On the one hand, the behavior of each context is handled by axioms and inference rules. On the other hand, the overall behavior of the system is handled by bridge rules like Rule (2) linking GC to DC , and expressed as follows:

$$(2) \models GC : G(a_i, \phi) = \delta_\phi \rightarrow DC : D^+(a_i, \phi) = \delta_\phi$$

It can be read as follows: if an agent a_i has as goal ϕ with a satisfaction degree δ_ϕ in a GC then it positively desires ϕ

with the same degree δ_ϕ in a DC . For more details about the proposed model, we refer the reader to [6].

To show the applicability of our multi-agent BDI framework, an experimental model is presented and evaluated in the next section using the NetLogo Platform.

III. THE SYSTEM SIMULATION

In agent-based systems with spatial reasoning and social behavior, a visual output is needed to display the agents' moving and interaction in two or three dimensional spaces. The Netlogo graphical user interface offers the possibility to design agents with different shapes and positions. Each agent in the simulation is a multi-context BDI agent whose behavior is described in the previous section. An agent represents a user with different desires and beliefs that are randomly initialized. The aim of the simulation is to compute a recommendation based on a user initial set of beliefs and desires, and to see how our agent will adapt the recommendation in two cases:

- the agent is part of a social network (social agent), i.e., it has links with other agents similar to it,
- the agent is considered as a solitary agent, i.e., it has no interaction with other agents.

Plans are a list of activities that consist of moving from one destination to another. Each destination contains some rewards that an agent will obtain when it reaches that destination. The amount of rewards for each agent is random. Once rewards are gained, an agent will broadcast information about the number of remaining rewards in the correspondent destination to similar agents. These agents will decide to accept or not this recommendation according to the trust degree in the sender, and whether there is any information in their knowledge base (desire or belief base) that contradicts this one. If an agent decides to accept the recommendation, then it adds this information to its desire base, and then trigger the recalculation of its intentions according to the updated desire base.

A. Experimental Setup

Table I summarises the parameters that can be varied for different use cases. As shown in Figure 2, agents are initially randomly distributed in the space (patches in NetLogo). They also have different profiles (desires, trust degrees, positions, etc...). Links are also created randomly between agents according to an initial link number defined by the user at the beginning of the simulation (on the left-hand side of Figure 2).

We used Netlogo 5.3.1 version to implement our simulation. For the BDI behavior and the communication context we use two available NetLogo libraries [11], one for BDI-like agents and the other for ACL-like communication, allowing the development of goal-oriented agents that communicate using FIPA-ACL messages. We developed the rest of the behavior of the agents using the NetLogo language with some extensions. The objective of the simulation is to assess the effects of these agents on the system as a whole (and not only to assess the effect of individual agents on the system).

¹<http://ns.inria.fr/huto/5w/>

Fig. 2. The User interface of our multi-agent simulation in NetLogo. The person icon represents an agent which represents a user. Flags represent destinations in which agents can go. Labels represent an agent intention which consists of two elements: the name, mapped to a NetLogo command, and a done-condition, mapped to a NetLogo reporter. Intentions are stored in a stack, and are popped out when to be executed. If the done-condition is satisfied, the intention is removed and the next intention is popped out consecutively. The figure shows also, on the right side, how the graphs are updated dynamically as the program runs.

TABLE I
THE SCALE AND DISTRIBUTION OF PARAMETERS IN THE SIMULATION.

Parameters	Scale	Distribution
Number-of-agents	0-100	Random
Desires	0-50	Random
Beliefs	0-100	Random
Intentions	0-10	Random
Links	0-100	Random
Gain	0-50	Random

B. Experimental Results and Discussion

The model and experimental data were analysed using the RNetLogo extension [12]. Once the experiment is set up, each agent will have a list of random desires, beliefs are empty at the beginning. According to these desires and the aforementioned behavior, an agent calculates the recommendation which has a plan as output. This plan will become the agent's intention, and the agent will execute it. In the case of a solitary agent, it will execute its plan without any change. Only a new belief from an external source that does not contradict the agent initial belief can make it change its intention. In the other case, i.e., a social agent, similar agents will communicate a set of proposed recommendations with the aim to influence the others to change their beliefs or desires. If the recommendation is accepted, the agent will recalculate its intentions based on the recommendation, and it will follow a new plan. Metrics such as utility or satisfaction are calculated using the following equations:

$$\text{utility}(p) = \frac{\sum_{i \in G_S} g(i)}{\sum_{j \in D_{\text{initial}}} d(j)}$$

where G_S is a set of goals satisfied by a plan p , and D_{initial} is a set of initial desires of an agent.

$$\text{satisfaction-degree}(p) = \max\{G(\phi_i), i \in [0, n]\}$$

where n is the number of goals satisfied by a plan p . The utility measure estimates how much the user needs (desires) match the recommendation (plan). The satisfaction degree, as its name suggests, computes the user satisfaction about a recommendation based on its initial degrees of desires.

The mean gain of the agents is also reported, and results are showed in Figure 3. We can see that agents within a social context, i.e., agents that communicate in order to influence each other, accrue more gain most of the time in comparison with those without a social context. These results show that a social population could have a greater social welfare than a non social one when agents have similar interests.

For comparison, we calculate the average satisfaction degree and utility over time for 50 agents in the case of individual and social agents. One may expect that the probability of gaining utility will increase with exchanging messages. Figure 4 confirms this expectation. It shows that utility augments considerably within social agents compared to the utility within individual agents. We notice that the average utility is the same over time for individual agents. We can deduce that exchanging ones beliefs and desires increases, on average, the agents utility.

In Figure 5, we can see the average satisfaction of the agents about the recommendations they received (plans). The average is higher within social agents than within individual ones. We can conclude that agents get more satisfaction collectively from exchanging information than alone.

These results provide for agents further motivation to engage in communications with similar trustworthy agents and support our modeling choices. It is also interesting to note

Fig. 3. Mean gain of agents with and without a social context.

Fig. 4. Mean utility of agents with and without social context.

Fig. 5. Mean satisfaction degree of agents with and without social context.

how communities of agents (e.g., agents with similar interests) likely to be influenced are more efficient collectively than individual agents. However, it is also interesting to see how the system will behave if some malicious agents communicate incorrect information. An interesting approach is presented in [13], where a score pair (trust, distrust) is used. We are currently studying the possibility of merging the two approaches.

IV. CONCLUSIONS

In this paper, we have presented an agent-based simulation of the framework we proposed in [6]. The purpose of the simulation was to evaluate agent behaviors adopting two different strategies (the social and the individual strategy) in order to infer the quality of recommendations. Results show that agents achieve a better performance collectively when they are in “communities”, i.e., agents with shared interests (thus similar to each other), then when they are acting as solitary agents. We believe that the issues of trust and recommendation are tightly related. Results show that exchanging beliefs or desires with trustworthy agents can improve the whole performance of agents. However, we ignore how the framework will behave when errors are introduced in the communicated information. For that reason, we need a mechanism for checking information reliability and updating the trust value, accordingly. We believe that approaches like [13] can be used to extend our agent model with further reasoning abilities and, consequently, to deal with information reliability as well. Our ongoing work, then, is to expand the multi-agent simulation with those abilities. For future work, we will investigate how to extend the agent framework with temporal and spatial reasoning to be more representative of real-world applications.

REFERENCES

- [1] B. Lopez, S. Aciar, B. Innocenti, and I. Cuevas, “How multi-agent systems support acute stroke emergency treatment,” in *IJCAI Workshop on Agents Applied in Health Care*, 2005, pp. 51–59.
- [2] M. Batet, A. Moreno, D. Sánchez, D. Isern, and A. Valls, “Turi@: Agent-based personalised recommendation of tourist activities,” *Expert Systems with Applications*, vol. 39, no. 8, pp. 7319–7329, 2012.
- [3] B. Chen and H. H. Cheng, “A review of the applications of agent technology in traffic and transportation systems,” *Intelligent Transportation Systems, IEEE Transactions on*, vol. 11, no. 2, pp. 485–497, 2010.
- [4] M. Montaner, B. López, and J. L. De La Rosa, “A taxonomy of recommender agents on the internet,” *Artificial intelligence review*, vol. 19, no. 4, pp. 285–330, 2003.
- [5] A. Casali, L. Godo, and C. Sierra, “A tourism recommender agent: from theory to practice,” *Inteligencia artificial: Revista Iberoamericana de Inteligencia Artificial*, vol. 12, no. 40, pp. 23–38, 2008.
- [6] A. B. Othmane, A. G. Tettamanzi, S. Villata, M. Buffa, and N. Le Thanh, “A multi-context framework for modeling an agent-based recommender system,” in *8th International Conference on Agents and Artificial Intelligence (ICAART2016)*, 2015.
- [7] S. Parsons, N. R. Jennings, J. Sabater, and C. Sierra, “Agent specification using multi-context systems,” in *Foundations and Applications of Multi-Agent Systems*. Springer, 2002, pp. 205–226.
- [8] C. da Costa Pereira and A. G. Tettamanzi, “An integrated possibilistic framework for goal generation in cognitive agents,” in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 1239–1246.
- [9] —, “Syntactic possibilistic goal generation,” in *ECAI*. Citeseer, 2014, pp. 711–716.
- [10] D. Dubois and H. Prade, “Possibility theory: qualitative and quantitative aspects,” in *Quantified representation of uncertainty and imprecision*. Springer, 1998, pp. 169–226.
- [11] I. Sakellariou, P. Kefalas, and I. Stamatopoulou, “Enhancing netlogo to simulate bdi communicating agents,” in *Artificial Intelligence: Theories, Models and Applications*. Springer, 2008, pp. 263–275.
- [12] J. Thiele, “R marries netlogo: Introduction to the netlogo package,” *Journal of Statistical Software*, vol. 58, no. 1, 2014.
- [13] P. Victor, C. Cornelis, M. De Cock, and P. P. Da Silva, “Gradual trust and distrust in recommender systems,” *Fuzzy Sets and Systems*, vol. 160, no. 10, pp. 1367–1382, 2009.