Negotiations in the Context of AIDS Prevention: An Agent-Based Model Using Theory of Mind

Jennifer Klatt[1], Stacy Marsella[2], and Nicole C. Krämer[1]

[1] University Duisburg-Essen, Department of Social Psychology: Media and Communication, Forsthausweg 2, 47058 Duisburg, Germany
{jennifer.klatt,nicole.kraemer}@uni-due.de
[2] University of Southern California, Institute for Creative Technologies, 12015 Waterfront Drive, Playa Vista, CA 90094, USA
marsella@ict.usc.edu

Abstract. For the purpose of an AIDS prevention game, a model was developed that focuses on training safe sex negotiations. Non-player characters in the game are socially intelligent agents that are equipped with a Theory of Mind that allows them to reason about the mental processes and behavior of others. The underlying model for the negotiation about safe sex between player and agent was implemented in multi-agent simulation software. It consists of two agents who have different goals of either safe or unsafe sex, actions to achieve these goals, and the wish to come to an agreement. The model was evaluated for the agent-agent conversation to test the basic functioning.

Keywords: Virtual agents, negotiation modeling, Theory of Mind, AIDS prevention.

1 Introduction

AIDS is one of the major health threats of this and the last century. More than half of the newly infected persons are men who have sex with men (MSMs) [1]. Currently, the best way to fight AIDS is to prevent it from being transmitted, so it is crucial to achieve a behavioral change in people. New media interventions have been shown to improve intentions towards safe sex in comparison to other intervention methods [2]. Recently a web-distributed health intervention was developed to achieve behavioral change. In the "Solve-It" game the user meets new people in a social setting, can go home with one of them, and negotiate safe sex. The focus is on young adult MSMs due to that population's higher infection rate. Since the real life situation involves aroused states that influence decision-making, the prevention needs to create a realistic atmosphere, and especially affective states (e.g. arousal). By providing appropriate stimuli, the non-conscious biases of the real situation will be replicated [3].

Modeling of a realistic negotiation requires an opponent with human-like behavior. Theory of Mind [4] is the basis for countless acts in social behavior: concepts as empathy, irony, a white lie, deception, or hints can be performed using the ability to reason about others [5]. Also, it is necessary to be able to negotiate with others. In

H. Högni Vilhjálmsson et al. (Eds.): IVA 2011, LNAI 6895, pp. 209–215, 2011.

order to be able to come to an agreement or to even persuade the interaction partner and change his behavior, the negotiator needs assumptions on the current beliefs of the interaction partner and how they can be changed [6]. Also, ToM is a prerequisite for negotiations regarding the emotions that occur. Humans usually perceive and decode others' emotions during the negotiation which influences their behavior. The current emotional state of the negotiator himself is another aspect that influences his behavior, strong emotions lead to different behavior than weak emotions [5].

One example for teaching social skills with using agent models is FAtiMA that was developed to teach children how to deal with bullying by interacting with agent-driven characters. The project has been adapted to simulate a form of theory of mind, double appraisal [7], where the agents consider the results of their action as if the action was done to them. The ELECT BiLAT application [8] is explicitly designed to teach negotiation skills, but in a cross-culture military context. The underlying simulation software used for the model is the PsychSim social simulation tool [9; 10], a multi-agent system whose agents are equipped with a Theory of Mind. However, in BiLAT fixed values are used for rewards, unlike the model we describe here.

The focus of the work discussed lies in the negotiation phase of the Solve-It game. An important aspect for the game is to create a non-player negotiation partner that employs humans-like tactics. To that end, we have used autonomous agents that are equipped with social intelligence, a Theory of Mind [4], to drive the negotiation partner's behavior. This intelligence includes the ability to put one-self into someone else's shoes and foresee future moves of both the user and itself. So the agent needs to make assumptions about the user's goals, feelings, and what action he will choose to achieve this goal. Therefore, Theory of Mind like capacities are necessary to enable an updating of the model on the user during interaction, allowing a form of user modeling that can support adaptive behavior as the user proceeds through the game. In this paper, we discuss the model we developed, evaluate its behavior in a series of simulation experiments and discuss the results.

2 Description of the Model

The model used in the Solve-It game consists of two agents, Adam and Bart, who negotiate whether they are going to have safe or unsafe sex. They make various negotiation moves, offering, counter-offering and complimenting in an effort to persuade the other to accept their offer. The model was developed in the PsychSim multi-agent system [9; 10]. PsychSim uses a decision-theoretic framework for quantitative modeling of multiple agents that allows agents to reason about tradeoffs in achieving their goals. PsychSim agents are equipped with a Theory of Mind that allows them to reason about beliefs and behavior of other agents. To realize an agent model for negotiating safe sex in PsychSim, we specified the agent's *states*, *goals*, *goal weights*, *actions*, *action dynamics,* and *beliefs* that define how the agent will behave.

States: provide a description of the world. The states we model are *Offered, Risky*, *Attractive*, *Arousal*, and *FeelingSafe*. *Offered* keeps track of which offer has currently been made. *Risky* is a characteristic of the agent about whether he has a tendency to engage in unsafe sex. *Attractive* is the degree of physical attractiveness. *Arousal* and

FeelingSafe are emotional states that describe the contradictory feelings that are involved with having safe or unsafe sex: to get aroused but also to feel safe.

Actions: change the values of states. The actions of the agents include: *OfferUnsafe*, *offerCondom*, *compliment*, *accept* and *rejectNegotiation*. *OfferUnsafe* and *offerCondom* are actions to offer safe or unsafe sex. *Compliment* is an action where one agent makes a compliment to the other one with the idea to create a better atmosphere between them. If one agent agrees to an offer that he did not make himself, he can use the action *accept*. *RejectNegotiation* can be chosen when one agent wants to end the conversation. Actions change states through action dynamics. For example the action *compliment* changes the values for accepting unsafe sex on *Arousal* and *FeelingSafe*. To influence the partner, the agent who wants to reach an unsafe sex agreement, Adam, compliments Bart and will become more attractive and less risky. This does not mean that he really is safer, but the complimented agent, Bart, will perceive him as a less risky and a more attractive choice than before. This has the consequence that he is more willing to accept an unsafe offer for two reasons: First, he will believe there is a higher reward for *Arousal* which is caused by the increased *Attractive* value and second, the smaller penalty for *FeelingSafe*, caused by the lower *Risky* value. In essence, the compliment changes Bart's view of the payoff, the benefit, and cost, of having sex with Adam.

Goals: states that need to be either maximized or minimized. For example, an agent can want to maximize his *FeelingSafe* or his *Arousal*. Agents can also have multiple goals and preferences or weights over which of these goals are most important. For example, in most of the examples below, Adam weights maximizing *Arousal* 70% and 30% for *Feeling Safe*. Bart's default values are 76% for *Feeling Safe*, 24% for *Arousal*. In the model's decision theoretic framework, the reward the agent receives for maximizing *Arousal* will be proportional to the value of the *Arousal* and the weight the agent gives that goal. Agents can also have goals of maximizing another agent's goals. This goal strongly influences the behavior.

Beliefs: consist of models of all agents (including itself), representing their states, beliefs, goals, and actions. So Adam can have a belief about Bart's goals e.g. his preference for maximizing *FeelingSafe*. Beliefs have a bounded recursive structure: Adam can have a belief about what Bart believes about his goal.

The conversation between agents is structured in steps: one step means both agents take one negotiation turn. The agents are able to think several steps ahead, so they can reason about their and the other one's future actions. How much they think ahead is defined by the *horizon* which can be set to a value individually for the agents. In the case of the experiments conducted with the model, the agents usually have a horizon of two steps per agent in the future (so four conversational turns in total) they can consider for choosing their first action. The recursive assumptions about the mental models of the agents are called the *belief depth* and are per default set to three levels. The mental models of oneself and the other agent can be correct, incorrect or use a probability distribution over possible mental models of the other that is updated during the conversation according to the other agent's actions (using a Bayesian update procedure). However, the default setting is a correct model.

One possible ending of the negotiation is the acceptance of an offer (action *accept*). Then the reward or penalty is added to the goal states *Arousal* and *FeelingSafe*: The

more attractive the partner is the more arousing sex is with him. Therefore, a higher value is added to the state *Arousal* when they accept the offer. *FeelingSafe* is dependent on how risky the agent is: If the offer of safe sex is accepted, *FeelingSafe* is increased. However if the offer is unsafe, it will be decreased according to the level of the other one's risk value. This means that a person who is perceived as unsafe will make us feel less safe when we have unsafe sex with him. Another possibility is the action of rejecting the negotiation (*rejectNegotiation*). No reward follows this action, but also no penalty is given which can be the most compelling option.

The described design of the model was chosen for several pedagogical reasons: First of all, we chose the extreme positions for the agents to have the maximum negotiating training. So Bart prefers *FeelingSafe* and Adam prefers *Arousal*. *Arousal* will be most increased when they agree on unsafe sex, whereas the *FeelingSafe* state is maximized for safe sex; unsafe sex even decreases its values. Also the agent with the unsafe sex goal is intended to be hard to negotiate with in order to enhance training effects. Second, the negotiation always reaches an agreement so the user experiences the rest of the game. A third aspect contains the variation of the states *Risky* and *Attractive* and the goal weights: in future versions of the game factors like drinking alcohol will change the preferences for *Arousal* or *FeelingSafe*, their models about the other one and finally, as a result, the actions: They will perceive others as less risky and more attractive, their arousal will go up and their horizon reduces. Therefore, the likelihood for unsafe sex increases.

3 Experiments

Experiments were run to validate the model used in the Solve-It game. A first set of experiments showed the basic behavior of the model. In a standard conversation with one agent who wants safe sex and the other wants unsafe sex, the outcome will be unsafe sex. This is due to the action *compliment* that changes the likelihood of the agent with the goal of safe sex to agree to unsafe sex (an example for the standard conversation can be found in Table 1, lines 1 and 2). First experiments tested the manipulation of the goals: When both agents want unsafe sex they agree on unsafe sex and the same is true for safe sex. So the basic functioning of the model is given.

To explore the behavior of the agents in a more cooperative negotiation, we conducted a second set of experiments where each agent could have an additional goal of maximizing the other's goals, additional to their goals to maximize *Arousal* and *Feeling safe*. We set the value for this goal weight to three different values: 20% as a moderate weight, 60% as a weight that makes helping the other one to achieve his goals more important than the own goals and 80% as an extreme position. Since all combinations with 20% resulted in the standard conversation, they will not be discussed here. For an overview see Table 1.

Results show that Adam, the one who most wants to maximize *Arousal*, keeps negotiating in the same way when he does not care for Bart's goals at all, when he cares for it 60%. Only when he cares for him 80% he behaves differently. A change in Bart's goal weight affects his behavior stronger: 60% of caring for Adam's goals is enough for him to compliment Adam because this changes Adam's values in case of unsafe sex. Adam receives higher values for his *Arousal* then, so Bart helps him even

more to maximize this state. When Adam starts, he offers unsafe sex, Bart compliments him, but this time Adam chooses to *rejectNegotiation*, because this increases his value for *Feeling Safe*. If in a human conversation somebody offers unsafe sex, receives a compliment and replies with rejection, he as well makes himself feel safer because he realizes that the other one tries to manipulate him. When Bart cares for Adam 80%, he behaves the same when he starts. But when Adam starts, he just directly accepts: unsafe sex.

When both agents care equally much for each other, for 60% the conversation proceeds as follows: Bart starts with complementing, then Adam compliments, too. They have the same reason for doing so: They want to maximize the other's goal state: Bart wants to maximize Adam's *Arousal* and Adam helps to maximize Bart's *FeelingSafe*. Then Bart offers unsafe sex and Adam agrees. When Adam starts, he instantly offers safe sex (because this helps best to let Bart achieve his goal) and Bart accepts. This is due to the fact that they care more for each other than themselves. With a goal weight of 80% they agree always on unsafe sex. A reason for this is the structure of the goals: In this case both agents care very much for the goals of the other one. So Adam believes that it is Bart's goal to achieve Adam's goal which makes Adam offer again his favorite action. So with 80%, a very high level of caring for the partner, the situation can be compared to not care about the other at all.

Table 1. Overview over results of varying how much they care for each other. Abbreviations used: *OC: offerCondom, OU: offerUnsafe, C: compliment, A: accept, RN: reject negotiations.*

Caring for partner	Outcome	Turn 1	Turn 2	Turn 3	Turn 4
Adam 60%	Unsafe	Bart *OC*	Adam *C*	Bart *OU*	Adam *A*
	Unsafe	Adam *C*	Bart *OU*	Adam *A*	
Adam 80%	Safe	Bart *OC*	Adam *A*		
	Unsafe	Adam *C*	Bart *OU*	Adam *A*	
Bart 60%	Unsafe	Bart *C*	Adam *OU*	Bart *A*	
	Unsafe	Adam *OU*	Bart *C*	Adam *RN*	Bart *A*
Bart 80%	Unsafe	Bart *C*	Adam *OU*	Bart *A*	
	Unsafe	Adam *OU*	Bart *A*		
Both 60%	Unsafe	Bart *C*	Adam *C*	Bart *OU*	Adam *A*
	Safe	Adam *OC*	Bart *A*		
Both 80%	Unsafe	Bart *C*	Adam *OU*	Bart *A*	
	Unsafe	Adam *C*	Bart *C*	Adam *OU*	Bart *A*

In mental models the agents have assumptions about the other one's recent states, beliefs, and goals. The goals are very important in this model because they define which action will be chosen. We ran experiments with three different possible mental models with Adam and Bart: a correct one (so the agent has correct assumptions about the other one's goals), a false one that assumes that the other one has the same goals like himself and a probability distribution over alternative models. The latter is implemented as a probability of 50% for the correct model and 50% for the false model. It adapts the probability of the two models according to the actions of the other agent. All combinations of the three mental model types were evaluated (see Table 2).

False mental models. In the case that Adam has a correct model and Bart has the false one, the standard conversation occurs. However, when Adam's model is wrong and Bart has a correct one, Bart offers condoms. Then Adam offers unsafe, because he thinks that Bart is like himself and will accept this offer immediately. Now Bart does not expect Adam to accept another offer of condoms but to reject the negotiation; therefore he rejects the negotiation himself. Hence, under these conditions they do not reach an agreement, regardless of who starts the conversation. When both agents have a wrong mental model both offer their favorite option and then the first agent does not expect that another offer would lead to success, so he rejects the negotiation. Hence, only when Adam has a correct mental model, an agreement is reached: unsafe sex.

Probability distribution over alternative mental models. When Adam has the correct mental model and Bart uses the probability distribution over possible mental models the standard conversation occurs. When Adam uses the probability distribution or when both use it, a different conversation results: Bart starts with *offerCondom*, which is accepted by Adam. When Adam starts and offers unsafe sex, Bart rejects the negotiation, essentially to avoid being manipulated, for he believes that if he did not do this, Adam would compliment him and he would finally agree to his offer. Unfortunately for Bart, Adam still compliments after this and Bart agrees. If someone in a human conversation asks for unsafe sex and the partner wants to go away he might as well try to change the partner's mind by saying something nice.

The values of 50% for the false model and 50% of the correct model are adapted during the conversations. Even with the low number of only two turns per agent, they reach values up to 70% for the correct model and 30% for the false model. This change of probabilities for the mental models would impact future interactions.

Table 2. Overview over results of varying the mental model of each other. Abbreviations used: *OC: offerCondom, OU: offerUnsafe, C: compliment, A: accept, RN: reject negotiations, PD: probability distribution*.

Mental model	Outcome	Turn 1	Turn 2	Turn 3	Turn 4
Adam correct,	Unsafe	Bart *OC*	Adam *C*	Bart *OU*	Adam *A*
Bart wrong	Unsafe	Adam *C*	Bart *OU*	Adam *A*	
Adam wrong,	No agreement	Bart *OC*	Adam *OU*	Bart *RN*	
Bart correct	No agreement	Adam *OU*	Bart *RN*		
Both wrong	No agreement	Bart *OC*	Adam *OU*	Bart *RN*	
	No agreement	Adam *OU*	Bart *OC*	Adam *RN*	
Adam correct,	Unsafe	Bart *OC*	Adam *C*	Bart *OU*	Adam *A*
Bart *PD*	Unsafe	Adam *C*	Bart *OU*	Adam *A*	
Adam *PD*,	Safe	Bart *OC*	Adam *A*		
Bart correct	Unsafe	Adam *OU*	Bart *RN*	Adam *C*	Bart *A*
Both *PD*	Safe	Bart *OC*	Adam *A*		
	Unsafe	Adam *OU*	Bart *RN*	Adam *C*	Bart *A*

4 Conclusion and Outlook

In order to create game characters that are challenging to negotiate with, we implemented an agent-based model of negotiation with ToM capacities. Computer

experiments were run to evaluate the behavior of the agent model. Most experiments conducted revealed an intuitive and logical behavior that shows that the model in general is functioning well.

These various computational experiments are just a first step in evaluating the model. It needs to be also tested for a) their pedagogical appropriateness for the game and b) the subjective impressions of human subjects. As noted earlier, the negotiation dialog paths generated by the model were incorporated in the current version of the Solve-It game. Those behaviors were vetted by the design team as realistic and 50 distinct paths were incorporated into the game. The Solve-It game will be tested and evaluated in large scale longitudinal clinical trial. In addition, we plan to test human subjects directly interacting with the agents in a negotiation to reveal if and how much the agents' behavior is perceived as realistic human behavior.

Various negotiation models use Theory of Mind approaches to make models of negotiations more realistic. Here most approaches use dynamics that are not dependent on emotions. The context of safe sex negotiations is a highly emotional topic that could not be addressed with only rational aspects. The model described in this paper tries to combine handling of emotions with general structures of negotiations. However, the model described only addresses a social negotiation that deals with a safe sex context. Future research could develop a general approach to modeling emotions and social negotiations without a specific context.

Acknowledgments. We thank the National Institute for Mental Health (NIMH) and the German Academic Exchange Service for their support of this project.

References

1. Centers for Disease Control and Prevention. HIV among Gay, Bisexual and Other Men Who Have Sex with Men (MSM) (2010)
2. Noar, S.M., Pierce, L.B., Black, H.G.: Can computer-mediated interventions change theoretical mediators of safer sex? A meta-analysis. Hum. Comm. Res. 36, 261–297 (2010)
3. Bechara, A., Damasio, H., Tranel, D., Damasio, A.R.: Deciding advantageously before knowing the advantageous strategy. Science 275, 1293–1294 (1997)
4. Baron-Cohen, S.: Autism: the empathizing-systemizing (E-S) theory. Annals of the New York Academy of Sciences 1156, 68–80 (2009)
5. Sally, D.: Social maneuvers and theory of mind. Marq. L. Rev. 87, 893–902 (2004)
6. Sycara, K.P.: Persuasive Argumentation in Negotiation. Theo. a. December 28, 203 – 242 (1990)
7. Aylett, R., Louchart, S.: If I were you: double appraisal in affective agents. In: Proc. of the 7th Int. Conf. on Autonomous Agents and Multiagent Systems, pp. 1233–1236 (2008)
8. Kim, J., Hill, R., Durlach, P., Lane, H., Forbell, E., Core, M., Marsella, S., Pynadath, D., Hart, J.: BiLAT: A Game-Based Environment for Practicing Negotiation in a Cultural Context. Int. J. of Art. Int. in Education 19, 289–308 (2009)
9. Marsella, S.C., Pynadath, D.V., Read, S.J.: PsychSim: Agent-based modeling of social interactions and influence. In: Proc. of 6th Int. Conf. on Cognitive Modeling, pp. 243–248 (2004)
10. Pynadath, D.V., Marsella, S.C.: PsychSim: Modeling Theory of Mind with Decision-Theoretic Agents. In: Proc. of the Int. Joint Conf. on Artificial Intelligence, pp. 1181–1186 (2005)