Gravel Grant Proposal
Simulator for Social Agent Behavior

Steven Damer
Department of Computer Science and Engineering
Master's Student
898 Raymond Ave. Apt 3
St. Paul, MN, 55114
Phone: (651)917–2557
Email: damer@cs.umn.edu

December 12, 2003

Abstract

My research interest is social behavior in computer agents. I will extend the functionality of the simulator for discrete games with interagent communication I have written, use it to develop socially aware agents, and then attempt to integrate those agents into an environment where they can interact with human players.

Applicant Name: Steven Damer
Applicant Signature:

Faculty Mentor Name: Maria Gini
Faculty Mentor Signature:

1 Introduction

I'm interested in social behavior for computer agents. There are many situations where there are opportunities for agents to do better by cooperating, but it is not possible to assume that agents you encounter will cooperate with you [1, 2]. I believe that social behavior is a mechanism which agents can use to take advantage of opportunities to cooperate, without unduly opening themselves up to the risk of another agent taking advantage of them.

Many situations in which agents can interact include the possibility of cooperation, but the complexity of the environment makes creating social agents difficult. Furthermore, any social agents which you do develop are very difficult to re-use in other environments. I want to develop a system where social agents can be created and evaluated in isolation from the environment. In order to accomplish this, I want to abstract potential agent interactions in terms of their results on the affected agents. For example, an interaction which can occur in a Massively Multiplayer Online RolePlaying Game (MMORPG) is one agent attacking another. Instead of worrying about things like what weapon or spell was used, or where the attack took place, agents will only be aware that one agent performed an action with a negative effect on the other agent.

Frequently computer games offer players the opportunity to interact with entities which are supposed to be intelligent. Unfortunately due to limitations in the entities coding they often do not act in a realistic manner. A better simulation of social behavior by computer driven entities in a game will result in better suspension of disbelief on the part of the player, and therefore a better game experience.

My goal is to create a system which can be used to support many areas of research into social behavior in agents. When complete this system will support exploration of many kinds of social communication, in many different contexts. Using simple abstractions of the game data, this system will be useful for developing more effective and realistic social behaviors for computer driven entities in games.

2 Description of System

For my project in CSci 8551 - Intelligent Agents I have written a simulator which can easily be extended with new agents and new games [4]. It can run tournaments between groups of agents with opportunities for communication between agents. I believe it will form a useful test platform for exploring social behavior, and developing social agents.

The simulator consists of 5 modules. The ExperimentServer, GameServer, and GossipServer modules provide the services to process the agent environment. The GameFactory and PlayerFactory modules are interfaces which provide a framework in which to create custom agents and games.

2.1 The ExperimentServer

The ExperimentServer module will read a set of tournaments to run from a configuration file. Each tournament is specified by a number of rounds, a type of game to play, and a set of agents to compete. In each round, games are run between sets of agents determined by the type of game in the tournament. Between rounds, the agents in the tournament have the opportunity to communicate with the other agents in the tournament.

2.2 The GameServer

The GameServer module will take a set of agents and a game, and run that game in a series of turns, with a subset of the agents simultaneously choosing a move each turn. Between turns agents have the opportunity to communicate with other agents in the game.

2.3 The GossipServer

The GossipServer module runs communication between agents for the GameServer and ExperimentServer modules. Currently communication is implemented by passing Strings between the agents. I am planning to extend this module to process the strings as XML documents, to allow for easier processing of communication

on the part of the agents. I am also planning to add support for passing messages back to the game, which would allow agents to make statements which can be confirmed as true or false by the game, and allow agents to pre-commit to a move in the following turn.

2.4 The GameFactory

The GameFactory module provides instances of games to the TournamentServer for the TournamentServer to pass to the GameServer. A game is played in a series of turns, with a subset of all the agents in the game making a move each turn. A game also provides data about the state of the game to the players, and the interface includes some optional functions which will allow agents to choose a move without needing to reason about the underlying game. Currently, the only game I have implemented is a uniformly distributed random normal form game with 2 players with a configurable number of moves [3].

2.5 The PlayerFactory

The PlayerFactory module provides instances of players to the TournamentServer. Player memory is persistant between games in a tournament, but not between tournaments. Each player must implement functions to receive data about a game, to receive messages from other agents, to send messages to other agents, and to choose a move in a game. Currently I have 2 players implemented, both of which do not reason about their moves, and use functions provided by the game to select their moves.

3 Social Communication Goals

I want to use my simulator to explore social communication between agents. Since full social communication is extremely complex, I have a number of relatively simple steps which I will implement.

3.1 Communication about moves

The first and simplest stage of communication is communication about moves in the game. Using this agents will be able to propose moves to make in the next turn and accept proposals.

3.2 Communication about trust

Building on that stage, we can include the concept of trust. Agents will be able to communicate about which agents they have found trustworthy, and which agents lied to them. An extension of this stage is to include promises of future communication about trust in the negotiation about moves. For example, an agent could say 'If you do move 1, I will tell everyone that you are very trustworthy' or 'If you don't do move 2, I will tell everyone that you are a cheater.'

3.3 Communication about debt

Another elaboration we can add to the first stage is the concept of debt. Since games are frequently asymmetric, agents find themselves in the position of making a sacrifice to their personal utility so that the group of agents playing the game can prosper. Agents will be able to express the idea that another agent is in debt to them, or that they are in debt to another agent. Agents can use this to negotiate about moves and about which agent should make a sacrifice for the group. For example, an agent might say 'If you do move 1, I will owe you 10 units' or 'If I do move 2, you will owe me 10 units.'

3.4 Signalling deception

As communication gets more complex, the issue of deception comes up. One thing I find intellectually perplexing about humans is that when humans lie, they frequently send some signals which indicate they are lying. I would like to explore this by giving agents the opportunity to include an indicator of deception with

each message over a noisy channel, and observe what effect this has - specifically, if voluntarily signalling deception is a useful practice.

3.5 Communication about relationships

Finally, we can implement communication about social relationships. Agents will be able to describe social relationships in terms of membership, and the obligations held by the members of the relationship. This could allow agents to communicate more efficiently about debt and retribution. For example, if agents encounter each other infrequently, it might be useful to be able to form groups to collectively manage obligations. You might see a message like 'If you do move 1, I will revoke your debt to agent 3 because me and agent 3 are in a relationship.' or 'I will do move 2 to punish you for lying to agent 3 about the move you were going to make in a previous game.'

3.6 Full social communication

Ultimately I would like to implement full social communications between agents, so that agents can create relationships on their own, and potentially lie about relationships they are in. We might see a message like 'I would like to form a group with you to collectively punish people who take advantage of us' or 'You owe me 5 units because agent 3 refused to revoke my debt to him because he said you weren't in a relationship with him that would allow you to revoke a debt on his behalf.'

4 Other Improvements to the System

I will also implement some more games so that players have the opportunity to use social communication in a richer environment. Currently the uniformly distributed random normal form game I have implemented allows complete knowledge of the environment. All agents know all payoffs for every agent, and what moves were chosen by every agent. I want to add options to hide some or all of that information. I would also like to add a Split the Dollar game occuring after the normal form game, which would allow agents to settle any debt between them resulting from an asymmetric payoff for the moves chosen. It would be interesting to remove the simultaneous move aspect of the game, and allow one agent to move after observing the choice made by the other agent. More complex games with a persistant state and a greater number of turns would also be interesting to implement, but I want to be careful to avoid taking attention away from the social communication aspect of the simulation.

5 Applications

Once I have implemented socially behaving agents in the simulator, I would like to attempt to use those behaviors in the context of a more elaborate game, preferrably one including human players. Two games which I currently have in mind are MMORPGs and computerized Diplomacy. I would be very interested in hearing about any other games to which people think my approach could usefully be applied.

5.1 MMORPGs

A MMORPG is a computer game played by a large number of people, each controlling a single character in a simulated environment which also contains many places to go and hostile monsters and treasure. To implement my approach within a MMORPG, it is necessary to abstract out common player interactions, such as attacking, healing, buffing, and trading goods. It would also be good to examine more sophisticated interactions, such as giving aid (or failing to) in a fight, or providing information about the game world. One significant advantage in MMORPGs is that it should be possible to provide human players with a simple interface to allow them to communicate with computer players. The largest difficulty is to transform a continuously changing game world into a discrete series of agent interactions. Once implemented, I think that socially aware agents could greatly add to the realism of the NPCs in the world, since they would be more difficult to take advantage of, and could take actions on their own.

5.2 Diplomacy

Diplomacy is a seven player strategy game simulating World War 1. The rules are very simple and deterministic. Since players start with equal forces, and the victory condition is control of half the map, the key to victory is to negotiate alliances to conquer your neighbors. A successful player is one who can negotiate good alliances, but is also capable of betraying their alliance partners when the time is right. A well played game of Diplomacy is a shifting web of deception and betrayal. There are a number of servers which run Diplomacy games on the web through email. Furthermore, there is work out there towards automating the strategic and tactical choices made by a player in the game. [1] Unfortunately, it is probably impossible to create an agent which can participate in the full game, due to the difficulty of processing email from the human players. Fortunately there are already many games run where there is no communication between players, and the only interaction is by moves made on the board. Given a social agent, it should be possible to combine it with a program to solve the strategic and tactical problems to allow it to play in one of those games (there are some conventions of social meaning to certain moves in those games, but due to the discrete nature of the moves, those social meanings should be easy to translate). Furthermore, Diplomacy players would probably be open to running some games with a well specified communication format which would allow agents to participate fully.

6 Timeline

My first goal is to finish the implementation details of the simulator. I anticipate that this can be completed in a month. Implementing the various types of messages, and agents which can use them is my next goal. I think I can complete 3.1 in 2 weeks, 3.2 in 1 month, 3.3 in 1 month, 3.4 in 1 month, 3.5 in 1 month, and 3.6 in 2 months. At that point I will have a system which can be used to create and test socially aware agents, and I will be ready to look at integrating them into an environment which includes human interaction. I will also use the system to explore the interactions between agents implementing various different types of social behaviors. I will also be able to examine the value that social behaviors bring to agents.

7 Faculty Mentor and Budget

My faculty mentor for this project will be Maria Gini. I will use the grant funds to support myself while I do work on the project, but I anticipate that the opportunity to get suggestions and feedback from other members of GRAVEL will be more valuable to me.

8 Conclusion

I look forward to working with the GRAVEL project in the future. Computer games are an important application of my research interest in artificial social intelligence, and I am very interested in learning what other researchers are doing in that area.

References

[1] R. M. Axelrod. *The evolution of cooperation.* Basic Books, 1984.

[2] Robert Axelrod. *The complexity of cooperation.* Princeton University Press, 1997.

[3] Michael Bowling and Manuela Veloso. Multiagent learning using a variable learning rate. *Artificial Intelliegence*, 136:215–250, 2002.

[4] Gerhard Weiss, editor. *Multiagent systems. A modern approach to Distributed Artificial Intelligence.* The MIT Press, 1999.

[1] http://devel.diplom.org/Zine/F2003M/Norman/ai_article.html