Toward Acquiring a Human Behavior Model of Competition vs. Cooperation

David V. Pynadath, Ning Wang, and Chirag Merchant
Institute for Creative Technologies
University of Southern California
Los Angeles, CA
pynadath@usc.edu, nwang@ict.usc.edu, merchant@ict.usc.edu

ABSTRACT

One of the challenges in modeling human behavior is accurately capturing the conditions under which people will behave selfishly or selflessly. Researchers have been unable to craft purely cooperative (or competitive) scenarios without significant numbers of subjects displaying unintended selfish (or selfless) behavior (e.g., Rapoport & Chammah, 1965). In this work, rather than try to further isolate competitive vs. cooperative behavior, we instead construct an experimental setting that deliberately includes both, in a way that fits within an operational simulation model. Using PsychSim, a multiagent social simulation framework with both Theory of Mind and decision theory, we have implemented an online resource allocation game called "Team of Rivals", where four players seek to defeat a common enemy. The players have individual pools of resources which they can allocate toward that common goal. In addition to their progress toward this common goal, the players also receive individual feedback, in terms of the number of resources they own and have won from the enemy. By giving the players both an explicit cooperative goal and implicit feedback on potential competitive goals, we give them room to behave anywhere on the spectrum between these two extremes. Furthermore, by moving away from the more common two-player laboratory settings (e.g., Prisoner's Dilemma), we can observe differential behavior across the richer space of possible interpersonal relationships. We discuss the design of the game that allows us to observe and analyze these relationships from human behavior data acquired through this game. We then describe decision-theoretic agents that can simulate hypothesized variations on human behavior. Finally, we present results of a preliminary playtest of the testbed and discuss the gathered data.

ABOUT THE AUTHORS

David V. Pynadath is the Director for Social Simulation Research at the University of Southern California's Institute for Creative Technologies. He serves as co-chair of the annual Workshop on Plan, Activity, and Intent Recognition. He has published papers on social simulation, multiagent systems, teamwork, plan recognition, and adjustable autonomy. He has developed multiagent systems for applications in social simulation, virtual training environments, automated personal assistants, and UAV coordination. He is the co-creator of PsychSim, a unique multiagent social simulation framework that combines Theory of Mind and decision theory.

Ning Wang is a research scientist at the Institute for Creative Technologies of the University of Southern California and the chief scientist at Curious Lab, a consulting service that specializes in assessment of training and simulation technology. Dr. Wang's research is in the design of intelligent virtual characters that motivate students to engage in learning activities. Dr. Wang has extensive experience in the design and assessment of virtual agent based training simulations and educational games. She is one of the pioneers in socially intelligent pedagogical agents—virtual animated characters with human-like behavior that facilitate learning. Her work on the pedagogical agents' socially intelligent feedback lays the groundwork for the design of trust-inducing explanations for the robots in human-robot teams.

Chirag Merchant is a research programmer at the Institute for Creative Technologies of the University of Southern California. He has worked as the lead developer on multiple military training environments, experimental social simulations, and other applications of human-computer interaction in virtual environments.

Toward Acquiring a Human Behavior Model of Competition vs. Cooperation

David V. Pynadath, Ning Wang, and Chirag Merchant
Institute for Creative Technologies
University of Southern California
Los Angeles, CA
pynadath@usc.edu, nwang@ict.usc.edu, merchant@ict.usc.edu

INTRODUCTION

As our world grows more populous and more interdependent, the solutions to our problems can rarely be carried out by any one individual, or even by any one group. Collaboration across people and agencies has become critical for effective strategies for disaster response (e.g., Kapucu, N., Arslan, T., & Demiroz, F., 2010), homeland security (e.g., Hocevar, Thomas, & Jansen, 2006), etc. Examining the historical record of such collaborations illustrates the factors that lead such diverse entities to work together to successfully achieve their common goal, as well as the factors that lead to miscoordinated behavior among the separate entities (Johnson, McLaughlin, & Christensen, 1982). Therefore, simulating such domains (e.g., for analysis, training, decision support) requires an ability to also simulate this range of collaborative outcomes.

There is a great deal of research on how to build the necessary computational models of the conditions under which people will behave selfishly or selflessly. Researchers have been unable to craft purely cooperative (or competitive) scenarios without significant numbers of subjects displaying unintended selfish (or selfless) behavior. Subjects famously display cooperative behavior in the Prisoner's Dilemma game, despite the fact that the prescribed dominant behavior is to do the opposite (Rapoport & Chammah, 1965). Such experimental observations have informed efforts to model the motivations behind—to literally "rationalize"—such deviations from expectations of competition or cooperation (e.g., Margolis, 1984). Researchers in agent-based simulation have similarly attempted to operationalize the varying motivations underlying these deviations by embedding them within an autonomous decision-making framework (e.g., Axelrod, 1997).

In this work, we seek to create an experimental testbed within which we can study various permutations on this theme of competition vs. cooperation. We expand the common two-player games (e.g., Prisoner's Dilemma, dictator game, ultimatum game) into a four-player game that enriches the behavior space to allow different competitive and cooperative relationships to emerge within the same scenario. Furthermore, rather than placing participants in a game where competition is the "rational" strategy and watching the emergence of cooperative behavior, we instead create a cooperative game while also planting the seeds for competitive behavior. Our aim is thus an experimental setting that is rich enough to engender both cooperation and competition within a common task, while still being of manageable enough complexity to fit the gathered data within an operational simulation model.

Using PsychSim (Marsella, Pynadath & Read, 2004; Pynadath & Marsella, 2005), a multiagent social simulation framework with both Theory of Mind and decision theory, we have implemented an online resource allocation game called "Team of Rivals", where four players seek to defeat a common enemy. The players have individual pools of resources which they can allocate toward that common goal. However, in addition to their progress toward this common goal, the players also receive individual feedback, in terms of the number of resources they own and have won from the enemy. By giving the players both an explicit cooperative goal and implicit feedback on potential competitive goals, we give them room to behave anywhere on the spectrum between these two extremes.

We discuss the design of the game that allows us to observe and analyze these relationships from human behavior data acquired through this game. We then describe decision-theoretic agents that can simulate hypothesized variations on human behavior. We then present preliminary results from small-scale experiments and discuss how those results will inform our future studies using this testbed.

GAME DESIGN

The fundamental concept underlying our experimental testbed is collaborative resource allocation. The overall goal is thus to allocate resources to achieve a common objective. However, we plant the seeds of competition by giving the players individual pools of resources that they have control over and by giving them individual credit within the achievement of the common objective.

Team of Rivals Game

We designed the resource allocation task as the defeat of a common enemy by a four-player team in a fictionalized world. We divide the world into 42 territories, 38 of which are initially owned by this enemy, while each of the other 4 is owned by one of the players (with the four territories being spread out as far as possible on the map). It is important to note that ownership is given to individual players, not to the overall team. Therefore, the individual players may have personal incentives regarding which of them owns the territory, even though this individual ownership does not affect the overall team goal.

The players can allocate resources toward the conquest of enemy-owned territories, and the team wins the game if all enemy-owned territories are captured within 20 rounds. Each enemy-owned territory has a fixed number of defending resources, ranging from 5 to 20 across different territories, and observable by the players. On each turn, the players can allocate any number of resources (or no resources) from their available pool to any enemy-owned territory. The probability that the enemy loses the territory is inversely proportional to the number of resources the player team allocates to it. More precisely, if there are r_E defending resources and r_P player resources, the probability that the player team wins is $r_P/(r_P + r_E)$. Thus, the more resources allocated, the greater the chance for team success. At the same time, there is an element of chance that gives the team a chance to lose, regardless of how many resources they allocate. Likewise, they have a chance to win, regardless of how much the enemy defenders outnumber their allocation.

If the player team wins, they come one step closer to achieving their goal of winning all of the territories. However, as already mentioned, the newly won territory's ownership is given to a single player. If only one player allocated resources toward the conquest of the territory, then that player is granted ownership of the territory. If multiple players allocated resources, then the player who allocated the most resources is granted ownership. Ties between two or more players are broken randomly according to a uniform distribution.

Figure 1 shows the state of the game during the first round. The gray territories (both light and dark gray) are owned by the enemy. Each of the four players is designated by a color (red, green, blue, and yellow). The territories won by each player are displayed in the color of that player. The players can allocate resources to only those territories that border a territory owned by any player. The interface denotes such available territories with a dark gray background, while territories that are inaccessible are light gray. It is important to note that any player can allocate resources to any dark gray territory in Figure 1, *not* just to territories that border their own. This decision increases the opportunities for multiple players to allocate resources to the same territory, either to cooperate and jointly win the territory, or to compete and try to "outbid" each other for ownership. It also allows us to measure the degree to which people are influenced by adjacency, even when it has no inherent value.

The information bar at the top of the screen is in the color of the given player (i.e., Figure 1 is the view for the Green player). The colored circles contain the number of resources currently owned by each player. Each player receives 5 resources at the beginning of each turn. Thus, they start the game with 5 resources with which to make their first-round allocations. Any resources allocated toward enemy territories are lost. Any unallocated resources are carried over to the next turn. Each enemy territory has a fixed value (ranging from 5 to 12 resources) that are given to the player who takes ownership of the territory. Therefore, the player who allocates the most resources not only gains the territory, but also the increase in resources that comes with it. The players start each round with a minimum of 5 resources, but may have many more based on such gains.

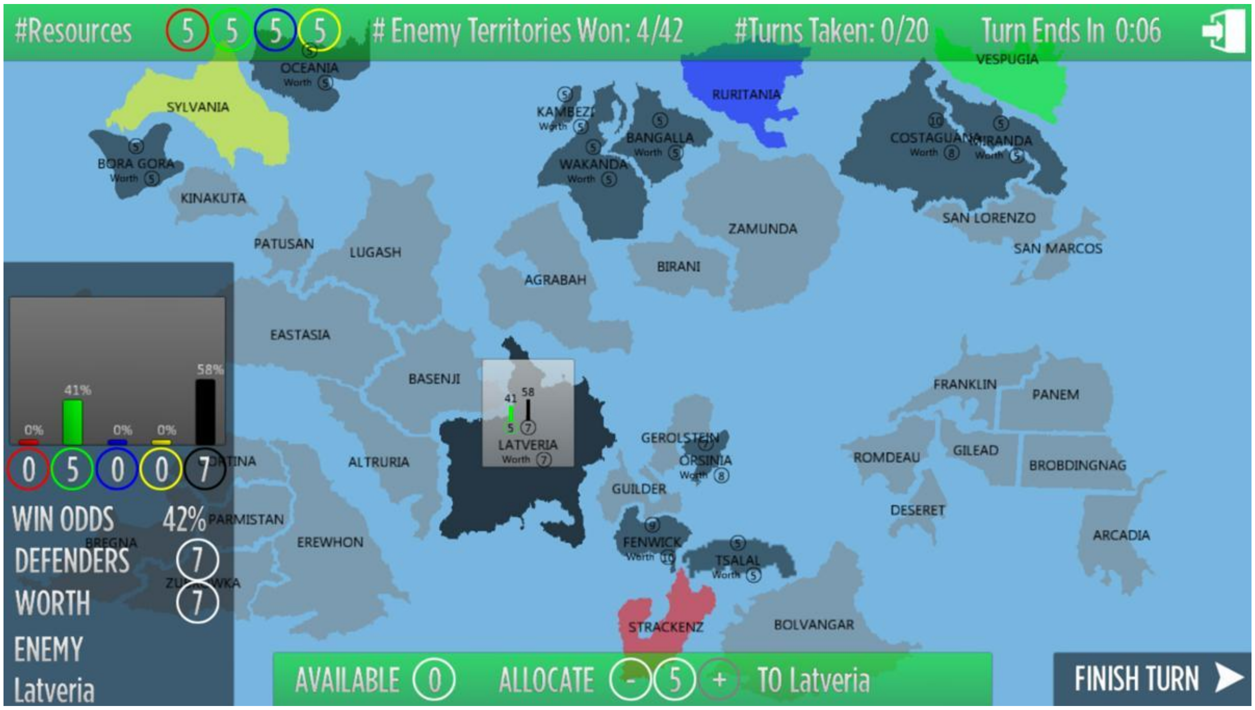

Figure 1. Team of Rivals game interface during Green's first turn of the game.

Each round lasts at most one minute. During a round, the players concurrently select territories and allocate resources to them. As they do so, they get feedback on the probability of their winning the territory (in the bottom left corner of Figure 1). Displaying this probability provides the player with real-time information on how their allocation of resources will affect their chance of winning, removing the challenge of assessing the likelihood of winning the territory across possible allocations. By observing how each participant weighs that likelihood in his or her final decision, we can gauge their risk attitudes. More precisely, by comparing their allocations against the expected "utility" (e.g., gain in territories and resources), we can measure their degree of risk aversion.

More importantly for our purposes, each player's current allocation also appears on the *other* players' interfaces as well. Thus, all of the players get almost instantaneous updates about their teammates' allocation decisions. This information allows dynamic interaction among the team during the round, even without any explicit communication. The players are all free to use this information in any way they choose. In other words, they may observe that a teammate is invading a certain country and therefore choose to contribute resources to assist. Alternatively, they may choose to invest their resources elsewhere so as to increase their chances of gaining a territory for themselves. The players may change their allocation as many times as they can during each round's one minute of game play.

At the end of each round, the players see a table showing all of their teammate's final allocations (as well as their own). The simulation then computes the outcomes of all of the allocations and updates the state of the game. We currently perform the probabilistic outcomes using a fair random-number generator, but it would also be possible to manipulate the outcomes at this stage instead (e.g., to observe how the team reacts to consistent overestimates of win probabilities). The players then return to the Figure 1 screen for the next round, with the updated display of their resources and territory ownership.

The current implementation employs a context of military conquest, but it is important to note that we can change this context without changing any of the underlying game mechanics. For example, we could alternatively describe the 38 enemy-owned territories as being affected by a virus outbreak instead. The players would then be cooperating in seeking to eradicate the virus by allocating resources (e.g., medical personnel, health supplies) to the territories. Players would still be incentivized to seek individual goals, because of the "credit" they would receive in clearing territories. Thus, even without changing the underlying game mechanics, we can explore the impact of the task description on cooperative vs. competitive behavior in such settings.

Implementation

The game's front-end (displayed in Figure 1) was developed for the Unity 3D Web player, allowing us to include anyone with Web access as a potential participant. This Web client communicates with our hosted server, which maintains the state of the game, collects player moves from the clients, and sends back the resulting game states. The server functionality was implemented using the Django Web framework. Django was a natural choice for the server implementation, as it is built upon the Python programming language shared by our social simulation framework, PsychSim (described in the following section).

AGENT MODELS

PsychSim

We build both the game itself as well as the agent models of player strategies using the multiagent social simulation framework, PsychSim (Marsella, Pynadath, & Read, 2004; Pynadath & Marsella, 2005). PsychSim represents individuals and groups as autonomous agents that integrate two multiagent technologies: recursive models and decision-theoretic reasoning. Recursive modeling gives agents a *Theory of Mind* (Whiten, 1991), to form complex attributions about others, incorporate such beliefs into their own behavior, and enrich the explanations provided to the user. In addition, the agents employ boundedly rational, decision-theoretic reasoning to quantitatively assess risk/reward tradeoffs. Thus, these agents represent a decision-making model that generates behavior by reasoning from a declarative representation of their goals and beliefs. We have used PsychSim to model a range of cognitive and affective biases in human decision-making and social behavior (e.g., Pynadath & Marsella, 2013; Pynadath, Si, & Marsella, 2014). PsychSim thus provides a natural basis for our current exploration of the interplay of cooperation and competition in human decision-making.

Another motivation behind the use of PsychSim is its successful application within multiple simulation-based training environments. The Tactical Language Training System (TLTS) is an interactive narrative environment in which students practice their language and culture skills by talking to non-player characters built upon PsychSim agents (Si, Marsella, & Pynadath, 2005). We also used PsychSim's mental models and quantitative decision-theoretic reasoning to model a spectrum of negotiation styles within ELECT BiLAT training system (Kim et al, 2009). UrbanSim used a PsychSim-driven simulation to put trainees into the role of a battalion commander trying to maintain stability, fight insurgency and crime, reconstruct the civic infrastructure and prepare for transition in a densely populated urban environment (McAlinden et al., 2009). SOLVE used PsychSim agents to populate a virtual social scene where people could practice techniques for avoiding risky behavior (Klatt et al., 2011; Miller et al., 2011).

We have also used PsychSim to build experimental testbeds for studying interactions between a single person and a single agent. In one such testbed, we gathered data on human behavior within the simulation of a wartime negotiation (Wang, Pynadath, & Marsella, 2015). The results of that experiment allowed us to quantify the ways in which people would be motivated by affective concerns (e.g., avoiding war) that were not explicitly incentivized by the game itself. More recently, we used a PsychSim agent to autonomously generate behaviors for a simulated robot that teamed with a person, in a study of trust within human-robot interaction (Wang et al., 2015). We build upon PsychSim's capability for such experimental use in the expanded interaction of the current investigation.

PsychSim represents the game-level decision-making problem facing the players as a Partially Observable Markov Decision Problem (POMDP) (Kaelbling et al., 1998). However, the game itself is completely observable, in that there are no hidden states. The mental states (e.g., preferences) of each player are unobservable to the others, but in this investigation we simplify the problem by allowing all of the agents to know the true preferences of the others (this obviously does not apply to the human participants who play the game). We can thus model the problem as a Markov Decision Problem instead (MDP). An MDP is a tuple $\langle S,A,P,R \rangle$, with S the set of states, A the set of actions for the agent, P the transition probability representing the effects of the actions on the states, and R the reward function that expresses the player's preferences. We describe each of these components in more detail within the Team of Rivals game context.

S, State

The state of the world represents the evolution of the game state over time. We use a factored representation (Boutilier et al., 1999) that allows us to separate the overall game state into orthogonal features that are easier to specify and model. There are features corresponding to aspects of each territory:

- **Owner:** The player who owns this territory, or else "Enemy" if it is still enemy-owned.
- **Occupants:** The number of defensive resources in the territory (irrelevant if not enemy-owned).
- **Value:** The number of resources conveyed to the owner upon winning this territory (irrelevant in subsequent rounds).
- **Invaders:** The total number of resources allocated by the player team to this territory.

There are also features corresponding to aspects of each player's current state as well:

- **Resources:** The number of resources currently available for the player to allocate.
- **Territories:** The number of territories currently owned by the player.
- **Value:** The total value of the territories currently owned by the player.

There is also a global state feature representing the current **Round** of game play.

A, Actions

The actions available to each player are allocations of resources to territories. The set *A* for Player *X* thus contains actions of the form "Player *X* allocates ⟨number⟩ resources to ⟨territory⟩". The ⟨territory⟩ must be one that is accessible from a team-owned territory. A player may choose multiple such actions, each with a different accessible ⟨territory⟩. The total ⟨number⟩ of resources across all a player's action choices must be no more than his or her current number of resources. A player may also choose no such actions, e.g., to save up resources for the next turn.

P, Transition Probability

The transition probability represents the effect of the players' chosen actions on the game state. We give only a small number of examples here, but for the most part, *P* represents the rules of the game in a straightforward manner. For example, on the completion of each round (i.e., for any set of player actions), the value of **Round** increases by 1 with 100% probability. Similarly, the number of **Invaders** of a territory is equal to the sum of resources allocated to it across the actions chosen by all four players, again with 100% probability.

The transition probability of the **Owner** of a territory is more complicated, in that it must capture the game mechanics of winning a territory. As described in the previous section, we want the probability of any player winning the territory to decrease with the number of defensive resources present. Therefore, the probability that the **Owner** of a territory is still "Enemy" after a given turn is the ratio of the number of **Occupants** of the territory to the sum of the **Occupants** and **Invaders**.

The probability that the players win the territory is 1 minus this probability. Again, as described in the previous section, we set the **Owner** to be the player who allocates the most resources to the territory with 100% probability (in the case when the Enemy loses). Ties are broken by a coin flip from a uniform distribution. This game mechanic gives an even stronger incentive for the players to allocate more resources, as the one who allocates the most gets the spoils of victory. At the same time, players allocating fewer resources to a territory than another will help contribute to the team goal (by reducing the probability of enemy victory), but will have nothing to gain individually.

R, Reward Function

The reward function allows us to capture these varying motivations, similar to the work on modeling both selfishness and altruism, mentioned in the opening section of this paper. An agent's reward function is a quantitative assessment of the current game state in terms of its desirability to that agent. An agent who completely adopts the team goal would simply have a reward that decreases with the number of enemy-owned territories. An agent who is also interested in gaining territories for itself would also have a component of its reward that increases with the number of territories it owns. The following section presents other possible models of human motivations in this game.

Agent Decision-Making

Having specified the game within the PsychSim language, we can apply existing POMDP algorithms to autonomously generate decisions for the individual player agent (Kaelbling et al., 1998). Such algorithms enable the agent to consider candidate allocations of its resources, weigh the relative likelihoods of winning various territories based on such allocations, and compute an expected reward gain (or potentially lose) for each such allocation. It can then choose the allocation that maximizes this expected reward, although we have freedom to modify this selection if so desired to better capture the types of reasoning exhibited by people. Importantly, this algorithm can autonomously generate behavior without any additional specification, allowing us to observe differences in behavior that result from varying modeling parameters (e.g., the relative importance of team vs. individual goals).

PRELIMINARY RESULTS FROM PILOT STUDIES

The testbed aims to provide an experimental platform to study how people decide when to pursue competitive vs. cooperative goals, and to collect human interaction data to build computational models of the decision-making process for the agents. The research objective is based on the assumption that the interactants of the testbed favor one of two abstract goals: maximizing their personal gain, or maximizing the collective gain of the team. In particular, we hypothesize that individuals play the Team of Rivals game in order to "gain as many territories/resources for oneself", "help the team defeat the enemy", or some combination of the two. To verify our assumption, we conducted a small-scale pilot study to understand what are the goals players have when they are playing the Team of Rivals game.

Pilot Study

We recruited eight participants from a university of a major United States metropolitan area. Of the eight participants, two are research staff on the project, two are graduate students majoring in computer science, and four are U.S. Army cadets. Participants formed two groups of four players. Each group played one game. During the game, after each round, participants filled out an in-game survey regarding the reasons behind their actions during the last turn. The in-game questionnaire has four open-ended questions. The second question listed below was shown as three questions, each regarding one of the other players (represented by color) on the survey.

> *"I (Blue) chose to Attack a country on my own / Attack a country together / Not attack because…"*
>
> *"Green/Red/Yellow chose to Attack a country on his/her own / Attack a country together / Not attack because…"*

After the game, participants filled out a post-game survey about the strategies used during the game. Each group then participated in a post-game interview.

> *"I (blue) chose my actions in the game in order to…"*
>
> *"Green/Red/Yellow chose his/her actions in the game in order to…"*

We collected all of the goals that participants had and what goals they thought others had during the game and after the game, including surveys and interviews. Data on the goals participants themselves had during the game confirmed our hypothesis on competitive and cooperative goals. In particular, the goals participants reported included "I wanted to get as many territories for myself" (competitive) and "I want to help the team win the game" (cooperative).

However, these two goals were not held constant; instead, they evolved as the game went on. For example, participants reported that they had the goals for gaining territories for themselves until towards the end of the game, at which time they switched to act to help the team win the game. This is largely due to circumstances in the game. For example, if a player did not have enough resources to attack a country on his or her own and did not have enough time (e.g., not enough turns left in the game) to accumulate resources, then he/she may choose to cooperate.

Additionally, if there were fewer than 4 territories left, then the 4 players may choose to cooperatively attack the enemy. The evolution of competitive/cooperative goals throughout the game reflected the general theme that the goals participants had were highly context-dependent.

Participants also reported other goals we did not hypothesize. Some participants reported that “I want to help Blue/Green/Red/Yellow (another player)” and “I attacked country X to make the connected country available for attack in the next round.” Some participants reported the goal of following the norm “I attacked the enemy because others were doing the same thing.” Some participants reported goals of challenging oneself, e.g. “I attacked the enemy alone because I want to get a territory by myself.” Some participants reported both competitive and mischievous goals such as “I want to steal country X from Blue/Green/Red/Yellow (by allocating more resources to country X than Blue/Green/Red/Yellow at the last minute).”

The goals participants reported reflected *where* they choose to attack as well. Some participants reported “I attacked Country X because they defeated me several times in a row”, or “… to “build my own kingdom” (e.g. owning several contiguous countries), or “… because I want to own countries all over the world”, or “… because I want to spread my resources around (and gamble with my luck”. When participants did not attack, their goals were clear, namely “to accumulate more resources for future attacks”.

The goals participants inferred from other players about the other players’ goals were largely similar to the goals they had for themselves.

Modeling the Observed Player Goals

Having collected these self-reported motivations from our participants, one possible next step is to operationalize them within our PsychSim agents’ reward functions. For example, the incentive to “steal” a territory indicates that players perceived reward, not only from getting the territory for themselves, but also from preventing another player from getting it. In other words, this particular player had a component of reward that decreased as the territories owned by the teammate increased.

Some other incentives were map-related, such as attacking a country to make adjacent countries available for allocation on the subsequent round. This decision is one form of heuristic that guides short-term decision-making in a way that increases long-term rewards. It remains to be seen whether such decisions emerge naturally from the POMDP’s long-term strategizing, or whether this additional heuristic incentive must be added explicitly to more accurately capture human decision-making. On the other hand, there is no in-game incentive for owning contiguous territories, so that observed behavior indicates a new reward component for us to add to our agent models.

Once we have such a set of candidate reward functions, we can then instantiate them within the PsychSim agents and have them autonomously generate in-game behaviors in place of our human participants. By experimenting with different team compositions (i.e., a different combination of the four players’ rewards), we can explore the impact of individual motivations on the overall team outcome in simulation. This methodology allows us to extrapolate from hard-to-collect human behavior data to easy-to-collect simulated results from systematic exploration of the space of teams.

DISCUSSION

Although two games are obviously much too small a sample from which to draw many conclusions about the underlying human behavior, it is clear that our “Team of Rivals” testbed is capable of engendering a broad space of cooperative vs. competitive strategies. The diversity of motivations expressed by even our limited set of participants encourages us to believe that the game design is rich enough to give people the freedom to express their own individual decision-making processes. At the same time, the constrained nature of the game still allows us to translate these individual motivations into an operational language of agent-based simulation. Thus, we see the game design as successful as the basis for a methodology for gathering human behavioral data that can inform simulation models of competition and cooperation.

The next obvious step is to gather these behavioral data on a much larger scale. We are currently exploiting the Web-based nature of our platform to set up a study that will use participants recruited from Amazon Mechanical

Turk. Our previous work successfully drew from such a pool to quickly gather data on a large number of games with participants from multiple countries (albeit currently restricted to only English-speaking participants) (Wang et al., 2015). Analyzing the collected data will give us insight into the competitive vs. cooperative behavior of broader demographics of participants as well. By building agent-based models from the gathered data, we will be able to generate simulation libraries that capture a much wider range of cooperative/competitive decision-making than is currently possible.

Our testbed also gathers data on the *dynamics* of the players' allocations, not just their final decisions. We can therefore study how players' decisions change as they observe their teammates' decisions. Such direct observations of behavior can reveal more of the players' collaborative attitudes than self-report in surveys. Anecdotal evidence certainly suggests that these dynamics are a rich source of such indicators.

Beyond analysis of data gathered using the current game setup, there are numerous variations of the game that we can use to explore additional aspects of the phenomenon. For example, it is trivial to manipulate the parameter settings (e.g., starting number of territories, resources, defenders, etc.). We might expect that making the overall team task more challenging might lead to more cooperation, so as to avoid overall mission failure. As already mentioned, we can also completely recast the game context itself (e.g., responding to a disease outbreak instead of conquering the world). Such a context shift might also increase the emergence of cooperative behavior.

We can enrich these behaviors even more by adding a communicative component to the in-game interaction. For example, we could allow players to make commitments to each other about allocating resources jointly (i.e., coalition formation). The capability for communication opens up a much broader behavior space, one that might overwhelm our ability to analyze the resulting data. Nevertheless, expanding the game to include some constrained communicative or signaling acts would add another interesting dimension to the studies we can support.

In conclusion, our PsychSim-based "Team of Rivals" testbed offers a novel platform to gather human behavioral data on competitive and cooperative behavior. The game design grants the players enough freedom to pursue a large variety of strategies, while still being constrained enough to support translation of the data into agent-based simulation models. The multiplayer aspect of the game goes beyond the typical dyadic interactions of such experimental platforms. Our limited playtesting provides preliminary evidence that our platform meets the requirements of our investigation, encouraging us to pursue larger-scale studies and build validated models of human competitive/cooperative behavior. In previous work, we have translated data covered in such a platform into simulation models capable of autonomously generating human-like decision-making behavior (Wang et al., 2012). Our implementation thus provides researchers of such models with a novel experimental testbed to both gather and analyze behavioral data for input into simulation models that can support analysis and training in domains where cooperation and competitive exist simultaneously.

ACKNOWLEDGEMENTS

This work was sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). The content does not necessarily reflect the position or policy of the Government. No official endorsement should be inferred.

REFERENCES

Axelrod, R. M. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton University Press.

Boutilier, C., Dean, T., & Hanks, S. (1999). Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, *11*(1), 94.

Hocevar, S. P., Thomas, G. F., & Jansen, E. (2006). Building collaborative capacity: An innovative strategy for homeland security preparedness.

Johnson, H. W., McLaughlin, J. A., & Christensen, M. (1982). Interagency collaboration: Driving and restraining forces. *Exceptional Children*, *48*(5), 395-399.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, *101*(1), 99-134.

Kapucu, N., Arslan, T., & Demiroz, F. (2010). Collaborative emergency management and national emergency management network. *Disaster Prevention and Management: An International Journal*, *19*(4), 452-468.

Kim, J. M., Hill Jr, R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., Marsella, S., Pynadath, D. & Hart, J. (2009). BiLAT: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*,***19(***3):289-308.

Klatt, J., Marsella, S., & Krämer, N. C. (2011, January). Negotiations in the context of AIDS prevention: an agent-based model using theory of mind. In: *Proceedings of the International Conference on Intelligent Virtual Agents*, pp. 209-215.

Margolis, H. (1984). *Selfishness, altruism, and rationality*. University of Chicago Press.

Marsella, S.C., Pynadath, D.V., and Read. S.J. (2004). PsychSim: Agent-based modeling of social interactions and influence. In *Proceedings of the International Conference on Cognitive Modeling*, pp. 243-248.

McAlinden, R., Gordon, A., Lane, H. C., & Pynadath, D. (2009). UrbanSim: A game-based simulation for counterinsurgency and stability-focused operations. In *Workshop on Intelligent Educational Games at the 14th International Conference on Artificial Intelligence in Education,* pp. 41-50.

Miller, L. C., Marsella, S., Dey, T., Appleby, P. R., Christensen, J. L., Klatt, J., & Read, S. J. (2011). Socially optimized learning in virtual environments (SOLVE). In *Interactive Storytelling*, pp. 182-192.

Pynadath, D.V. and Marsella, S.C. (2005). "PsychSim: Modeling theory of mind with decision- theoretic agents." In Proceedings of the International Joint Conference on Artificial Intelligence, pp. 1181-1186.

Pynadath, D.V. and Marsella, S.C. (2013). "Socio-Cultural Modeling through Decision-Theoretic Agents with Theory of Mind" In *Advances in Design for Cross-Cultural Activities*, Ed. by Denise M. Nicholson & Dylan D. Schmorrow, pp. 417-426.

Pynadath, D.V., Si, M., and Marsella, S.C. (2014). "Modeling theory of mind and cognitive appraisal with decision-theoretic agents." In *Social emotions in nature and artifact: emotions in human and human-computer interaction*, pp. 70-87.

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's dilemma: A study in conflict and cooperation* (Vol. 165). University of Michigan press.

Si, M., Marsella, S.C., and Pynadath, D.V. (2005). THESPIAN: an architecture for interactive pedagogical drama. In: *Proceedings of the International Conf. on Artificial Intelligence in Education* , pp. 595–602.

Wang, N., Pynadath, D. V., & Marsella, S. C. (2012). Toward automatic verification of multiagent systems for training simulations. In *Proceedings of the International Conference on Intelligent Tutoring Systems,* pp. 151-161.

Wang, N., Pynadath, D.V., Marsella, S.C. (2015). Subjective Perceptions in Wartime Negotiation. *IEEE Transactions on Affective Computing*, **6**(2):118-126.

Wang, N., Pynadath, D. V.,Shankar, S., U. K.V., and Merchant, C. (2015). Intelligent agents for virtual simulation of human-robot interaction. In: *Proceedings of the International Conf. on Human-Computer Interaction*, to appear.

Whiten, A. (Ed.). (1991). *Natural theories of mind: Evolution, development and simulation of everyday mindreading*. Oxford: Basil Blackwell.