

Dynamics of Social Interactions in a Network Game

Eunkyung Kim, Luyan Chi, Yu-Han Chang, Rajiv Maheswaran

Information Sciences Institute, University of Southern California

4676 Admiralty Way, Marina del Rey, CA 90292

Email: eunkyung@usc.edu, lchi@usc.edu, ychang@isi.edu, maheswar@isi.edu

Abstract—We characterize the social dynamics of human players in a transactional online game. We introduce two new approaches for understanding temporal behavior: the introduction of entropy dynamics, where we measuring the change in the entropy of certain descriptive distributions over the course of a game, and the use of clustering to discover temporal dynamics in subpopulations of experimental subjects. Experiments were conducted in the context of the Social Ultimatum Game, a multi-agent multi-round extension of a classical game-theoretic domain, using two distinct populations. We show that we are able to extract temporal behavior from the use of entropy dynamics and identify unique subpopulation behavior with respect to generosity, reciprocity and fairness.

INTRODUCTION

Social dynamics are at the core of human behavior, and many researchers in social computing have proposed new ways of characterizing aspects of social behavior. Social network analysis is a general term that describes this area, but much of the work in social network analysis focuses on static aspects of social networks. Understanding temporal dynamics is an emerging and important topic. Additionally, links in the network typically only represent simple aspects of a relationship, e.g., existence of a friendship or agreement on an issue, rather than potentially taking on many possible values in a multi-dimensional space. There are many real-world interactions that would clearly be better modeled by enabling the expression of additional structure, such as business transactions or favors between people over time.

In this paper, we investigate social dynamics in the context of an online game, where transactions are made between various parties in an unconstrained network over the course of many rounds of game play. These transactions can be thought of as business offers, favor requests, or any other activity between two parties. We are interested in characterizations of the social behavior in such settings over time. Here, we are viewing the network at a transactional level, where actions between entities create momentary links. We might not assign a permanent link or weight to that relationship. Parties learn from each of these transactions, and thus future actions are influenced by what is learned in past transactions.

There are three main challenges to characterizing social dynamics in this kind of domain: 1) we need new methods to handle multi-dimensional vectors over time, 2) there may be multiple subgroups within a population that exhibit different behavior, but it is difficult to extract the patterns that characterize them from the data, and 3) observed actions vary and cannot necessarily be correlated with observable agent

states, since much of the decision-making and learning is internal to each subject. To solve 1) and 3), we extend existing characterizations of behavior by computing these features over time — both by computing values at each time period, and by computing average values and slopes of values within time intervals. Time intervals enable us to partly deal with the variability of the data. To further characterize this variability, we introduce a notion of entropy dynamics: measuring the change in the entropy of certain descriptive distributions over the course of a game.

To solve 2), we first transform our data by computing these temporal measures on a sequence of time intervals for each game, and apply clustering to the resulting vector to separate the data. We can then recompute the original measures on the separated data, which enables us to characterize groups of individuals that exhibit similar temporal behaviors in their game play. Without first separating the data by clustering, these behaviors often do not manifest themselves in the overall measures calculated across all the players (i.e., in the non-separated dataset), either because it is a behavior associated with only a minority population and thus does not significantly affect the overall mean, or aspects of the behavior in the separated groups may be correlated and in some cases average out to unremarkable values, or other factors.

Our analysis is carried out in the context of the Social Ultimatum Game (SUG), an extension of a classical game-theoretic domain, the Ultimatum Game, where two participants have the potential to share an endowment. In the classical game, one player proposes a split and the other can either accept or reject it. It is a two-player one-shot interaction.

In the Social Ultimatum, this is extended to an environment where each player plays in a society where they repeatedly have the opportunity to split an endowment and also the have the ability to choose the recipient, enabling them to be both the proposer and recipient multiple times. We investigated this game through an online computational interface where each player was represented by an avatar and interacted with avatars for the other players. Experiments were conducted in two subject pools: one used staff, masters students and undergraduate at a U.S. university; another used doctoral students and faculty at an international conference. We show how our approaches were able to draw out unique temporal pattern differences both between the subject pools and within the subject pools.

Our choice to investigate an extension of the Ultimatum Game was motivated by its long history and because it is a

leading example of where traditional game-theoretic reasoning fails to predict human behaviors [9], [19], [11]. In particular, we are interested in multi-agent domains where humans make sequential decisions over time. The Social Ultimatum Game more closely fits assumptions that an individual has choice over whom to interact with and that an individual persists in a society where there are multiple opportunities to interact with others, while maintaining the socio-cultural issues that have been identified by other researchers in the one-shot Ultimatum Game. Understanding behavior in contexts such as the Social Ultimatum Game can yield insights into how humans interact but also into how to build software agents that can replicate or facilitate particular types of behavior. This investigation into human social behavior has practical applications for training in virtual environments [21], large-scale social simulation [3], and adversarial modeling [1].

RELATED WORK

Economists and sociologists have proposed many variants and contexts of the Ultimatum Game that seek to address the divergence between the “rational” Nash equilibrium strategy and observed human behavior, for example, examining the game when played in different cultures, with members of different communities, where individuals are replaced by groups, where the players are autistic, and when one of the players is a computer. In general, most individuals in western industrial cultures play fairly instead of being selfishly rational. Interestingly, isolated non-industrialized cultures, people who have studied economics, groups, autists, and playing with a computer all tend to lead to less cooperative behavior [9], [19], [17], [12], [2], [6]. However, even within supposed similar cultural groups (e.g., non-industrialized cultures), there exists significant variability in behavior [10].

Recently, there has also been other work attempting to model human behavior in the realm of social computation by observing convergence and other properties in relationship graphs over time, with applications to multi-agent systems, social networks and other domains [14]. Research on identification and development of social networks includes analyzing event-driven growth [20] and inferring social situations by interaction geometry [7]. Other work has described algorithmic methods to discover temporal patterns in networked interaction data [13].

There is a long literature on time-series metrics [18], however, these metrics do not capture the temporal causality patterns that are key to evaluating human behaviors in our domain. Stochasticity in social networks has been investigated in the context of friendship networks [22]. Prisoner’s Dilemma, another classical game-theoretic domain, has been studied in a graphical setting with simulated agents [16]. To help gain insight into behavior, several visualization approaches have been proposed [15], [8]. These approaches all begin the investigation into the area that is the focus of this paper: characterization of social interaction dynamics over time, where the interaction linkages between players at each time period are defined on a more complex action space.




Fig. 1. The Social Ultimatum Game: In each round k of the game, each player m receives an endowment e and makes an offer $q^{m,n}$ to player n as to how much of the endowment they are willing to give to the recipient. Then, every recipient with an offer, makes decisions $d_{m,n}^k$ as to whether to accept each offer they have received. Acceptance means that the endowment is split as proposed. Rejection means both players receive nothing for that offer. This interaction is repeated for K rounds, where players can choose new offer recipients, make offers with different values, and make different acceptance decisions over time.

THE SOCIAL ULTIMATUM GAME

The classical Ultimatum Game is a two-player game where P_1 proposes a split of an endowment $e \in \mathbb{N}$ to P_2 where P_2 would receive $q \in \{0, \delta, 2\delta, \dots, e - \delta, e\}$ for $\delta \in \mathbb{N}$. If P_2 accepts, P_2 receives q and P_1 receives $e - q$. If P_2 rejects, neither player receives anything. The subgame-perfect Nash or Stackelberg equilibrium has P_1 offering $q = \delta$ (i.e., the minimum possible offer), and P_2 accepting, because a “rational” P_2 should accept any $q > 0$, and P_1 knows this. Yet, humans make offers that exceed δ and reject offers greater than the minimum.

To represent the fact that people operate in societies of multiple agents and repeated interactions, we introduce the Social Ultimatum Game. An illustration and description of the game is shown in Figure 1. Players, denoted $\{P_1, P_2, \dots, P_N\}$, play $K \geq 2$ rounds, where $N \geq 3$.

In each round k , every player P_m receives a new endowment e and chooses a recipient r_m^k and makes them an offer $q_{m,n}^k$ (where $n = r_m^k$). Each recipient P_n then considers the offers they received and makes a decision $d_{m,n}^k \in \{0, 1\}$ for each offer $q_{m,n}^k$ to accept (1) or reject (0) it. If the offer $q_{m,n}^k$ is accepted by P_n , then P_n receives $q_{m,n}^k$ and P_m receives $e - q_{m,n}^k$, where e is the endowment to be shared. If the offer $q_{m,n}^k$ is rejected by P_n , then both players receive nothing for that particular offer in round k . Thus, P_m ’s reward in round k , U_m^k , is the sum of the offers they receive and accept (if any are made to them) and their portion of the proposal they make, if their proposal was accepted:

$$U_m^k = (e - q_{m,n}^k)d_{m,n}^k + \sum_{j=1 \dots N, j \neq m} q_{j,m}^k d_{j,m}^k \quad (1)$$

The total rewards for P_m over the game is the sum of per-round winnings, $U_m = \sum_{k=1}^K U_m^k$.

APPROACH

Traditional analysis of the Ultimatum Game has focused on measures such as the average offered amount and the average rejection rate given different offer amounts. In the Social Ultimatum Game, we begin by examining these measures,

however, the structure of the game provides the opportunity to look at more complex features of behavior over time. Because players can adapt as they learn about each other over the course of the game, we can examine concepts such as generosity, reciprocity, exploration, exploitation and stability.

For each round k , and each player m , there is a tuple that characterizes the offer that the player made, the offers that the player received, and the acceptance decisions for all those offers:

$$(q_{m,n}^k, \{q_{n,m}^k\}, d_{m,n}^k, \{d_{n,m}^k\})$$

Each *player trace* of behavior in a 20-round Social Ultimatum Game for a particular player m is a sequence of these tuples where $k = 1, 2, \dots, 20$. A game trace can be characterized by the union of these tuples over all players in that game.

Human behavior is notoriously difficult to characterize on an individual basis. It can be noisy and dependent on many unobservable states. We thus focus primarily on measures that can be aggregated over multiple time periods in order to characterize general trends in behavior. We define windows or intervals over the game periods, so that we can measure the change in these metrics over time. A *window* $w(k, \tau)$ is the set of rounds $\{k, k+1, \dots, k+\tau-1\}$.

We define a set of metrics, or features, that describe each of the traces we observe. These features will define a vector that describes each of the traces collected in our experiments. Each of these features capture an aspect of human behavior in a multi-round transactional game. They should find wide applicability in games other than Social Ultimatum, where multiple agents must interact over many rounds, and where there are values or target players assigned to the transactions.

Since some of our experiments involved software agents, we will often restrict our measures to cases where the human is executing the action of interest. For example, for offer distributions, we restrict our attention to offers that are made by human players. Below we describe each of these features, which are defined per player, per time period:

- Offer amounts
- Offer distribution entropy
- Recipient distribution entropy
- Likelihood of reciprocation
- Acceptance of offers received by value

We describe each of these features in detail below. We first discuss some temporal features that are indexed by round. Each of these features enables us to transform a player trace into a feature trace v^F of length at most $K = 20$, for each feature F . Some of the vectors v^F will have length less than K because the features are defined over windows of game rounds.

Offer Amounts. We define a feature trace for each player m to be $v_{(m)}^O$. Each entry of $v^O(m, k)$ corresponds to the offer amount made by player m in round k , i.e., $v^O(m, k) = q_{m,n}^k$.

Offer Distribution Entropy. There may be variability in the offer amounts made by a single player over the course

of a game. We introduce the notion of entropy dynamics to capture the changes in this variability over time. For example, we may observe that a player makes widely differing offer amounts early in the game, and then settles down to a much narrower range, or even to a single value, later in the game. Here, the offer distribution entropy would be high early on, and low towards the end of the game.

First, we assign an offer distribution to the random variable $O_{w(k,\tau)}$ by using the offer values over the window $w(k, \tau)$. This yields a probability mass function over the possible offer values, with weights corresponding to their frequency of occurrence within this window. The offer distribution entropy is then $H(O_{w(k,\tau)})$. We normalize the standard information theoretic entropy so that the value is bounded above by 1. For our specific experiment and analysis, we choose the window length to be $\tau = 5$ and the number of possible offer values is 11 ($\{0, 1, \dots, 9, 10\}$), so we have:

$$v^{H(O)}(k, 5) = -H(O_{w(k,5)}) / \log(1/11)$$

for $k = 1, 2, \dots, 16$.

Recipient Distribution Entropy For each player, we assign a recipient distribution to the random variable $R_{w(k,\tau)}$ by using the recipients over the window $w(k, \tau)$. This yields a probability mass function over recipients, representing the likelihood that a player made an offer to each of the four other players. The recipient distribution entropy is then $H(R_{b(k,\tau)})$. Again, we normalize this such that the value is bounded above by 1. We choose the window length to be $\tau = 5$ and because we are focusing on 5-player games, we have:

$$v^{H(R)}(k, 5) = -H(R_{w(k,5)}) / \log(1/4)$$

for $k = 1, 2, \dots, 16$.

Reciprocation. We measure the degree to which players respond to offers by reciprocating with an offer in the next time period. We define a reciprocation measure for each round $k \geq 2$ for each player which has value is 1 if the player received one or more offers in the previous round and chose one of the those players as the recipient in the current round, and has value is 0, in any other case. Note that, the value is 0 if the player received no offers in the previous round. We can thus define a length-19 feature trace for each player trace. This measures if the player to a reciprocal action in each round.

Clustering

Our experiments described in Section involved two separate subject pools. By looking deeper into the data, we can also understand internal differences within each subject pool, in addition to the differences across the two subject pools. To uncover such differences, we cluster the traces within each subject pool.

For clustering, we first transform each feature trace into a feature vector of 8 dimensions. The feature traces are broken up into 4 equal segments for length-20 and length-16 traces and into segments of length $\{4,5,5,5\}$ for the reciprocity traces which are length-19. The first four dimensions of the feature

vector are the average value of the feature in the corresponding segment. The second four dimensions of the feature vector are the average slopes of the feature value changes within the corresponding segment.

We then apply standard k -means clustering on this data. For most of our results, we choose $k = 2$. We confirmed the choice of number of clusters by evaluating different values of k using the Davies-Bouldin index and the Dunn index, as well as verification with silhouette graphs and visual confirmation [4], [5]. Larger values of k did not yield noticeably better clustering on these criterion.

In addition, we consider several non-temporal features of the player traces.

Acceptance of Offers Received by Value. This measures the likelihood by which offers of a certain value received by a player were accepted. For each player m and some condition on offer value, we count the occurrences where the player received an offer of that value and the number of times that offer value was accepted. We then create a feature vector for each player $v^A = [A(1) A(2) A(3) A(4) \bar{A}(5)]$ where $A(q)$ is the likelihood that a player will accept an offer of value q and $\bar{A}(q)$ is the likelihood that a player will accept an offer of value q or higher.

We also consider average offer value, average offer distribution entropy and average recipient distribution entropy, which are the averages of the temporal features for each player over all rounds.

EXPERIMENTS

In our experiments, players interacted with each other exclusively through the game’s interface, shown in Figure 2, on provided iPads. Once a player logs in, they receive a randomly chosen avatar and screen name. The avatars and screen names that players see for other players are consistent within a game, but are not the identity (i.e., avatar and screen name) that those players see of themselves. This was done to prevent players from identifying the actual human beings that they are playing with and to aid in preventing collusion.

Data was collected from two distinct pools of human subjects. The first pool was recruited from undergraduates, masters students and staff at the University of Southern California, henceforth referred to as “University”. In each round, every player is given the opportunity to propose a “\$10” split with another player of their choosing. Games ranged from 20 to 50 rounds. A conversion rate of 10 SUG dollars to 25 U.S. cents was used to pay participants. This works out to \$5 per 20 rounds per player in an egalitarian social-welfare maximizing game. Each game typically lasted 15-20 minutes. The subjects participated in organized game sessions and a typical subject played three to five games in one session. Between three and seven players participated in each game. In this experimental setup, we collected data from 27 human subject games. For this paper, we focus on the eight 5-person 20-round games in the dataset, i.e., 40 player-game traces. This is in order to have a common comparison with the subsequent experiment pool.

The second subject pool was recruited from attendees of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), composed primarily of doctoral students and faculty in computer science and is henceforth referred to as “Conference”. As in the first pool, in each round, every player is given the opportunity to propose a “\$10” split with another player of their choosing. Here, all games involved 5 players and 20 rounds.

The reward structure here was different: participants were told that only the player with the highest reward obtained in any single game would receive a cash prize (worth approximately U.S. \$35, in local currency). Participants could play more than once but most only played once. An additional complexity is that participants were told that the players that they saw in the game could be human or a software agent. They were not told how many would be software agents or what strategy the software agents would be using.

There were 24 games in this setup: 16 games with three humans and two software agents and 8 games with four humans and one software agent. In total, we collected 80 human game traces. The software agents were of two types:

- *Round-Robin*: This agent picked three players at random and would rotate among them as recipients of their offer in a round-robin manner. The offer was always an even split of the endowment (\$5-\$5).
- *Tit-for-Tat*: In the first round, or any subsequent round where the agent did not receive any offers in the previous round, this agent would pick a player at random and offer them a \$7-\$3 split, i.e., the software agent would keep \$7, and the recipient would get \$3. In any round where the agent received an offer in the previous round, it would choose the agent who made it the highest offer and make the same offer to that agent. If multiple agents made the same (best) offer, the agent picked a recipient from among them randomly.

Both agents accepted all offers. The games with two agents had one agent of each type. The games with one agent had only the Tit-for-Tat agent.

RESULTS

The results of the experiments are discussed for the various features mentioned in Section . The first collection of figures shows a particular feature indexed either by round or by a round where a window began. Error bars in each figure display the standard error of the feature’s mean value at that time index. In each subfigure of the clusters within each pool, the number in brackets indicates the number of player traces that fall within that cluster.

Offer Values

Our first analysis looks at the offer values made by human players over time. The first subfigure of Figure 3 shows the average values of offers made by human players in both the University and Conference experiments indexed by the round the offer was made. We see that the University players were uniformly more generous in every single round when

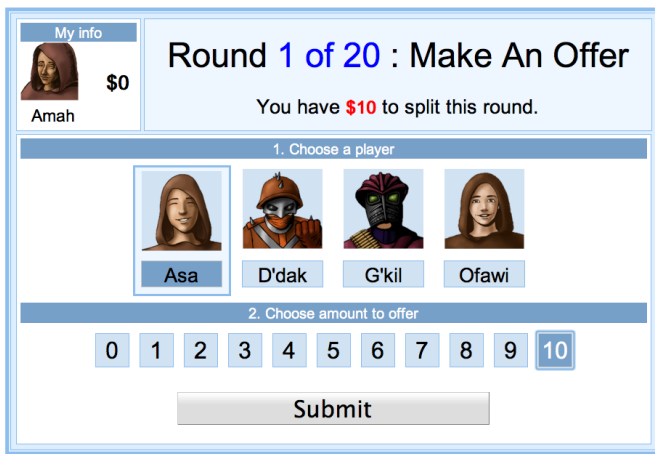


Fig. 2. The Social Ultimatum Game Interface




Fig. 3. Offer Value

compared with the Conference players. Furthermore, the University players maintained their levels of generosity throughout the game, with perhaps a small dip in the last three rounds. The Conference players, on the other hand, had a steady decline in generosity, where their average offer in the final round was 58% of the average offer in the first round (as compared to 93% for the University players).

The second and third subfigures show two clusters within the University and Conference subject pools respectively. We see that the first cluster for each pool are very similar in terms of offer values. One difference is that in the Conference setting the offer values for the high offer value group decreases quickly in the last three rounds. Another difference is that in




Fig. 4. Offer Value Entropy

the University setting the high group is 75% of the population whereas in the Conference setting it's only about half of the population.

The second cluster for the University setting keeps a steady offer value at a lower level. The second cluster for the Conference setting, on the other hand, has a very different behavior. These players have a steady gradual decline in offer value, decreasing below \$2 in the 6th round. This cluster makes up about half of the Conference pool subjects. Thus, the decline shown in the overall Conference pool consists of the combination of two different decline behaviors (the sharp decline of the high group, and the slow continuous decline of the low group). We note that the sharp jump in offer value in round 12 for the low cluster in the University pool was primarily caused by a single player deviating significantly from his other offers to make an offer of \$10 in that round and the relatively small size of the cluster.

Offer Value Entropy

Offer value entropy indicates the uncertainty of the player with regard to offer value. Alternately, it indicates the degree of exploration the player is performing with respect to offer value in the given window. The first subfigure of Figure 4 shows the average entropy of offer values made by human players in both the University and Conference experiments. We see that the University players had a relatively steady entropy level while the entropy levels slowly decreased for the Conference players. More significantly, the entropy values are contained in a very tight band between 0.172 and 0.302. This indicates that there was very little uncertainty about offer value. For reference, an

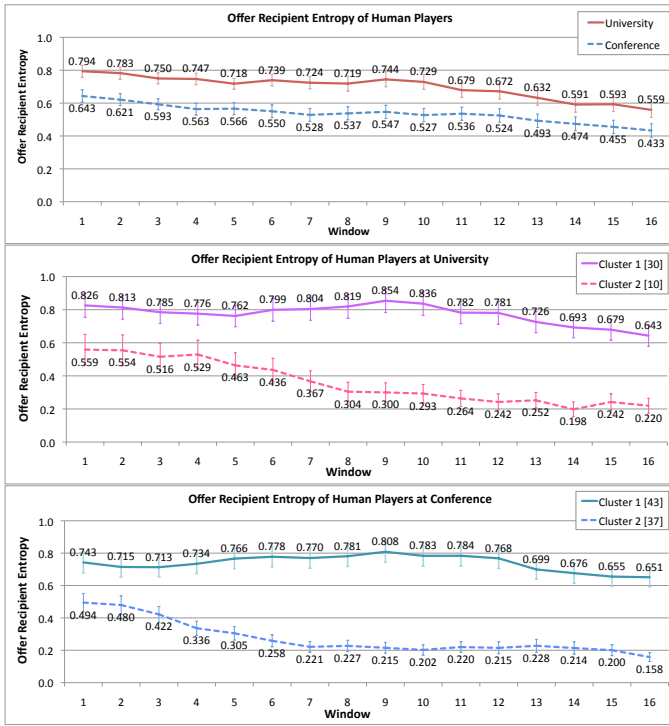


Fig. 5. Offer Recipient Entropy

offer value distribution that has equal support among two offer values would lead to a normalized entropy of 0.289. Thus, as a group, most players did not change their offer values very much over time. The second and third subfigures show that clustering does not add much explanation and verifies that most players had low offer value entropy.

Offer Recipient Entropy

Offer recipient entropy indicates the degree to which a player is alternating among recipients in the given window. The first subfigure of Figure 5 shows the average entropy of the offer recipient for human players in both the University and Conference experiments. As in the case of offer value entropy, the window index refers to the round where the window (of length five) started. We see that the University players had significantly higher offer recipient entropy than the Conference players, uniformly across all rounds: In fact, 75% of the University players' average entropies are larger than the largest Conference player's entropy, and 69% of the Conference entropies are smaller than the smallest University entropy. However, they both share a similar profile of decrease over time. The entropy in the last window compared to the first window is 70% for the University players and 67% for the Conference players. We note that, as opposed to offer value entropy, the entropies of the offer recipient remain high, i.e., players do not converge to offering to a single player. As a reference, alternating evenly between three recipients would lead to a normalized entropy of 0.8 and alternating evenly between two recipients would lead to a normalized entropy of 0.5. This indicates that players in both games prefer to explore

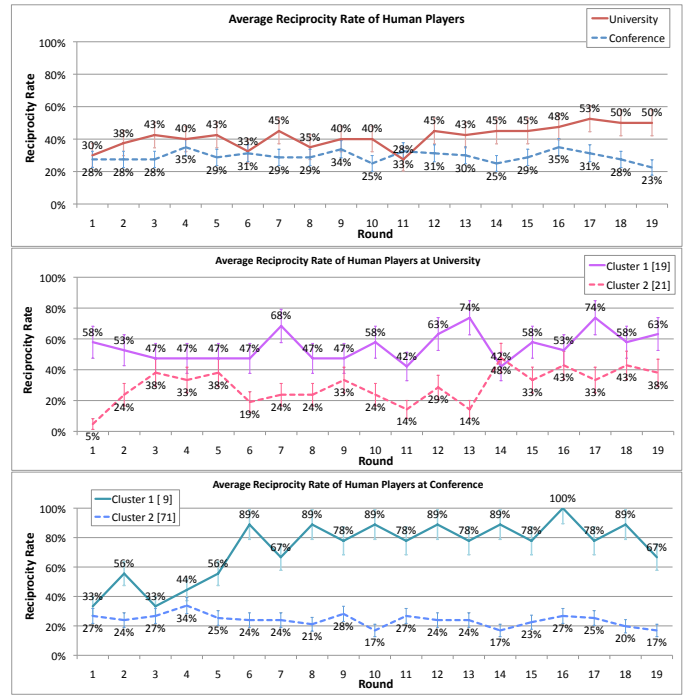


Fig. 6. Reciprocity

over recipients significantly more than exploring over value.

Looking at the clusters in the second and third subfigures of Figure 5 contributes a finer characterization of behavior. We see that both pools had a collection of players who had very high recipient entropies and another collection of players that had very low recipient entropies. The differences in recipient entropies for the clusters within a subject pool was high for both experimental settings. The Conference pool had a much larger percentage of players who had low entropy explaining the lower overall average entropy when comparing between the subject pools.

Reciprocity

The reciprocity feature traces show a player's ability or desire to enter into a stable partnership with another player. The first subfigure of Figure 6 shows the reciprocity percentage for all players in each experimental pool over the rounds of the game. This graph shows that the Conference pool has a steady reciprocation rate of about 30%, while the University pool starts at 30% and slowly climbs to 50%.

The second and third subfigures show the clusters for each pool. While the University clusters show two groups that each have slight ascents in reciprocation over time, the Conference cluster reveals a small group (9 out of 80) who were able to form stable relationships very quickly (by round 6) and maintain them until the end of the game. This characterization shows the value of combining clustering and temporal analysis in identifying subgroups with unique behavior.




Fig. 7. Offer Acceptance

Offer Acceptance

Offer acceptance characterizes a player's willingness to accept a split based on the value of the offer received. This has been interpreted to reflect a player's notion of fairness as determined by the threshold below which they will not accept offers. The first subfigure of Figure 7 shows the acceptance rate for various offer values, averaged over all players in each experimental pool. Values over \$5 were represented as one dimension due to their near universal acceptance. The first chart shows that both populations have similar acceptance profiles except for a large difference for \$1 offers, where University participants would mostly reject the offer, while Conference participants would mostly accept the offer.

In the second and third subfigures, which show the clusters for each pool, we see that there is important diversity within each population. In the University pool, there is a significant difference in participants' willingness to accept a \$2 offer based on cluster, where 18% (7 out of 40) of players would choose to usually reject it (95% rejection rate). Similarly in the Conference pool, there a significant percentage of the population (24%, or 19 out of 80) who usually reject \$1 offers as opposed to the remaining players who accept them at an 80% rate. For a profit maximizing agent, the knowledge of the existence of these smaller subpopulations within the pool can alter the design of the agent's strategy.

Correlations Across Features

In addition to the temporal analyses above, we examine the correlation between different characterizations of the game data. In Figure 8, we plot heatmaps that illustrate various dependencies: (1) between a player's average offer value (over




Fig. 8. (Above) Heatmap representing distribution of Conference player traces across average offer value and average offer distribution entropy; (Below) Heatmap representing distribution of Conference player traces across average offer value and average recipient distribution entropy. The numbers are the percentage of traces that fall within that cell.

all rounds of a game) and the same player's offer distribution entropy, and (2) between a player's average offer value and the player's recipient distribution entropy. Both figures show data collected from the Conference subject pool. We can see that there are two rough clusters of player traces. The main cluster corresponds to players whose average offer amount ranges between \$3 and \$5. These players also have higher recipient distribution entropy and higher offer value distribution entropy. They made a wider variety of higher offers to a larger number of different partners. The second, smaller cluster has lower average offer amounts, between \$1 to \$2, and has significantly lower recipient distribution entropy and offer distribution entropy. In these traces, the players have likely entered into stable reciprocal relationships where the partner was willing to accept low offer amounts.

CONCLUSION

This paper represents a first step towards a framework for characterizing social network behavior where links between

the actor nodes can represent multi-dimensional transactions in time. Such a framework would be useful for understanding many kinds of real social interactions, from business relationships to behaviors in MMOGs (massively multiplayer online games), where data is becoming increasingly available. It would also be a useful tool for creating agent models for realistic simulation and training systems based on a better understanding of real interaction data.

We argue that for this kind of domain, it is particularly important to consider features over time, either by looking at feature values at each specific time period, or by averaging or calculating slopes of the values over windows of time. Furthermore, clustering these traces is important. The clustering enables us to characterize groups of behavior that would otherwise not be apparent from a graph of average values (and standard errors) over time. For example, without clustering, we would not have seen that there is a small group of Conference subjects that achieve stable reciprocal relationships very quickly. The data in this cluster is not apparent once it is merged together with the data from the other cluster, if we are only examining means over time. Furthermore, by taking time into account, we can observe that the differences in the changes in reciprocation rates over time.

We see that the Conference subjects with high reciprocation rates actually only achieve those high rates after an initial phase where their reciprocation rates are not substantially higher than any of the other players. These kinds of temporal changes are important characteristics of human play, and we only observe them by choosing to examine features over time, and clustering together similar traces.

Applied to the Social Ultimatum Game, which is a good example of an abstract game that models social interactions over time, these tools produce a variety of interesting results and expose aspects of human play that are not necessarily intuitive. In the future, we will extend our analysis to cluster the traces using a larger set of features, such as incorporating different types of feature-based behavior traces into a single input vector for clustering. We will also investigate better clustering strategies such as transductive clustering techniques that allow us to seed clusters with representative behavior traces.

Application of these methods to other social interaction datasets will also be interesting. Furthermore, we will investigate scaling the game and the analyses up to larger number of players and to larger number of time periods. This will enable us to study a variety of different situations corresponding to realistic scenarios, such as when the number of players is so large that a single player cannot possibly interact with all the other players even once. Steady state behaviors will also be more easily verified by increasing the number of rounds.

Acknowledgements. This material is based upon work supported by the AFOSR Award No. FA9550-10-1-0569.

REFERENCES

- [1] David Carmel and Shaul Markovitch. *Learning Models of Intelligent Agents*. AAAI Press, 1996.
- [2] C. Robertand P. J. Carnevale. Group choice in ultimatum bargaining. *Organizational Behavior and Human Decision Processes*, 72(2):256–279, 1997.
- [3] Paul Davidsson. Agent-based social simulation: A computer science view. *Journal of Artificial Societies and Social Simulation*, 2002.
- [4] D. L. Davies and D. W. Bouldin. A cluster separation measure. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 1(224–227), 1979.
- [5] J. Dunn. Well separated clusters and optimal fuzzy partitions. *Journal of Cybernetics*, 4(95–104), 1974.
- [6] R. H. Frank, T. Gilovich, and D. T. Regan. Does studying economics inhibit cooperation? *The Journal of Economic Perspectives*, 7(2):159–171, 1993.
- [7] G. Groh, A. Lehmann, J. Reimers, M.R. Friess, and L. Schwarz. Detecting social situations from interaction geometry. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pages 1–8, Aug. 2010.
- [8] Jeffrey Heer and Danah Boyd. Vizster: Visualizing online social networks. In *IEEE Information Visualization (InfoVis)*, pages 32–39, 2005.
- [9] Joseph Henrich. Does culture matter in economic behavior? ultimatum game bargaining among the machiguenga. *American Economic Review*, 90(4):973–979, 2000.
- [10] Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean Ensminger, Natalie Smith Henrich, Kim Hill, Francisco Gil-White, Michael Gurven, Frank W. Marlowe, John Q. Patton, and David Tracer. “economic man” in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences*, 28:795–815, 2005.
- [11] Joseph Henrich, Steven J. Heine, and Ara Norenzayan. The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3):61–83, 2010.
- [12] E. Hill and D. Sally. Dilemmas and bargains: Theory of mind, cooperation and fairness. Working paper, University College, London, 2002.
- [13] Hsun-Ping Hsieh and Cheng-Te Li. Mining temporal subgraph patterns in heterogeneous information networks. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pages 282–287, Aug. 2010.
- [14] Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. In *Proceedings of the National Academy of Science*, 2010.
- [15] C.K.-S. Leung and C.L. Carmichael. Exploring social networks: A frequent pattern visualization approach. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pages 419–424, Aug. 2010.
- [16] Lingzhi Luo, N. Chakraborty, and K. Sycara. Prisoner’s dilemma in graphs with heterogeneous agents. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pages 145–152, Aug. 2010.
- [17] Alan G. Sanfey Mascha van’t Wout, René S. Kahn and André Aleman. Affective state and decision-making in the ultimatum game. *Experimental Brain Research*, 169(4):564–568, 2006.
- [18] Theophano Mitsa. *Temporal Data Mining*. CRC Press, 2010.
- [19] H. Oosterbeek, R. Sloof, and G. van de Kuilen. Cultural differences in ultimatum game experiments evidence from a meta-analysis. *Experimental Economics*, 7(2):171–188, 2004.
- [20] Baojun Qiu, K. Ivanova, J. Yen, and Peng Liu. Behavior evolution and event-driven growth dynamics in social networks. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pages 217–224, Aug. 2010.
- [21] Jeff Rickel. Intelligent virtual agents for education and training: Opportunities and challenges. In Angélica de Antonio, Ruth Aylett, and Daniel Ballin, editors, *Intelligent Virtual Agents*, volume 2190 of *Lecture Notes in Computer Science*, pages 15–22. Springer Berlin / Heidelberg, 2001.
- [22] Tom A.B. Snijders, Gerhard G. van de Bunt, and Christian E.G. Steglich. Introduction to stochastic actor-based models for network dynamics. *Social Networks*, 32(44–60), 2010.